

# DEIXIS



Department of Energy  
Computational Science  
Graduate Fellowship

2002 - 2003



DEIXIS - THE DOE CSGF ANNUAL  
DEIXIS - THE DOE CSGF ANNUAL



Department of Energy  
Computational Science  
Graduate Fellowship



# DEIXIS

Department of Energy  
Computational Science Graduate Fellowship

## TABLE OF CONTENTS

2002 - 2003

### Editor

Shelly Olsan

### Design

Juls Design, Inc.

### Contributors

Jacob Berkowitz  
William J. Cannon  
Victor D. Chase  
Peter Gwynne

*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

Copyright 2002 by the Krell Institute. All rights reserved.



PAGE 6



4

### Practicum Experiences

#### 4 An Enlightening Experience

##### 4 Diem-Phuong Nguyen

*Modeling Fire Leads to Exciting Research*

##### 8 Benjamin Keen

*Developing a Mathematical Toolbox*

PAGE 31



PAGE 13



12

### Department of Energy Lab Research

#### 12 Supernova Express

*Lawrence Berkeley National Laboratory*

#### 16 Simulating the Ultrasmall

*Sandia National Laboratories*

#### 20 Future Climate

*Oak Ridge National Laboratory*

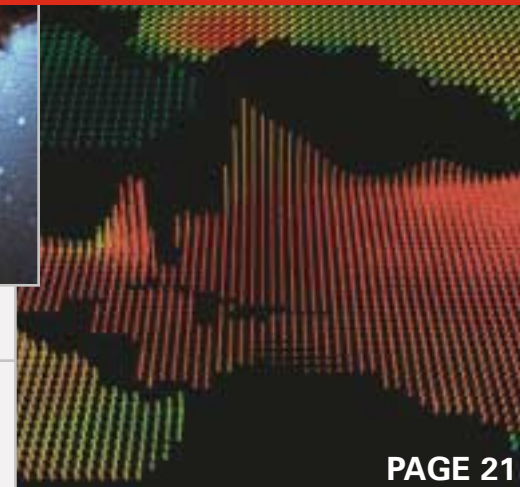
#### 24 Housing Lasers

*Lawrence Livermore National Laboratory*

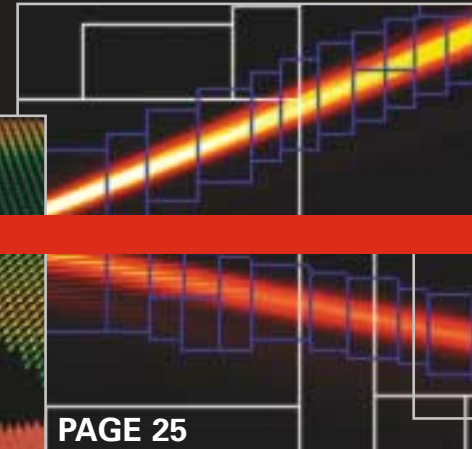
#### 28 Working a Beat

*Brookhaven National Laboratory*

PAGE 21



PAGE 25



44

### Alumni Profiles

44 Kenneth Comer  
*Procter and Gamble*

45 John Dolbow  
*Duke University*

46 Marcus Martin  
*Sandia National Laboratories*

47

### Alumni Directory

52

### Fellows Directory

*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 supported by the Office of Science and the Office of Defense Programs.

PAGE 38



PAGE 43



# An Enlightening Experience

**THE DEPARTMENT OF ENERGY** Computational Science Graduate Fellowship (DOE CSGF) program is one of a kind. This program not only requires a broad Program of Study that branches across three major disciplines (science/engineering, applied mathematics, and computer science) and offers generous support to its participants, but it also provides the opportunity to work in a DOE national laboratory on a three-month practicum assignment.



Benjamin Keen and Diem-Phuong Nguyen, DOE CSGF fellows, pose at the 2002 annual fellows conference.

**THE** practicum is designed not only to complement the fellows' research efforts, but also to expose fellows to the exciting world of ideas and scientific research that is happening in our nation's labs.

Two bright stars in the DOE CSGF program, Diem-Phuong Nguyen and Benjamin Keen, happened upon the fellowship by accident. Diem received an e-mail and Ben heard about it from a faculty member. Both decided to "give it a shot."

That chance has paid off for both of these fellows, who were selected from a highly competitive pool of candidates and have gone on to become bright stars in the program.

Having no experience in or with the DOE labs, Ben and Diem weren't sure what to expect when it came time to complete their respective practicums. But, after 12 weeks of living, learning, and lighting it up in the labs, they both became respected collaborators and found themselves right at home.

## Modeling Fire Leads to Exciting Research

DIEM-PHUONG NGUYEN

University of Utah | Lawrence Livermore National Laboratory | Story by William J. Cannon

Anyone who missed the movie *Backdraft* need only turn on the nightly news for a reminder of how unpredictable and unforgiving fire can be. Our collective horror has been kindled by images of Colorado, Arizona and other Western states burning and by stories of terrorists' bombs calculated to inflict maximum damage.

All of this has stoked a sense of urgency among combustion researchers, who have spent years studying the fundamentals of explosions and the path of fire. Diem-Phuong Nguyen's interest in the behavior of fire was sparked somewhat late, during a lecture in graduate school.

"I was in chemical engineering, just kind of hanging around," says Nguyen, who is finishing her doctorate this year at the University of Utah in Salt Lake City. "I really didn't have a research project. Then I went and heard a talk about the Accelerated Strategic Computing Initiative,"

better known as ASCI. "Philip Smith (chairman of Utah's chemical and fuels engineering department) was the speaker. He went through some sophisticated combustion simulations, and I was hooked. I didn't anticipate any of this."

Nguyen's group, Combustion and Reaction Simulations (CRSIM), headed by Smith, is developing software that simulates large-scale so-called pool fires of up to 20 meters in diameter. CRSIM belongs to a larger group, Center for the Simulation of Accidental Fires and Explosions (CSAFE), which simulates high-energy explosions. CSAFE is one of five university-based centers that are part of ASCI, funded by the Department of Energy.

"We look at accident scenarios," she says. "Say jet fuel catches on fire. We simulate that because, for obvious reasons, you can't do it experimentally. Each explosion is different – each model has ramifications for how to design better safety systems for activities like transporting combustibles. The real concern is transportation accidents, both airline and ground, but our methods apply to forest fires, too. We work on keeping fires from spreading. We can design facilities better by putting up barriers that

redirect fires from what we're trying to protect — if we know how they spread."

That's a big "if." Combustion is a complex process of fluid flow dynamics, fire structure, turbulence and other physics combined with thousands of chemical reactions. As much progress as there has been numerically simulating accidental fires and explosions, no one has come up with a reliable way to calculate both the turbulent conditions inside and surrounding an inferno and the fire's complex chemistry, involving as many as 200 species of highly reactive atoms called free radicals. Free radicals spawn the massive chain reactions that characterize combustion, converting fuel and oxygen into other gases and, not least, soot.

"Soot is an integral part of all fires, even jet fuel, because the soot radiates and strongly affects the temperatures and pollutants coming off from the fire," Nguyen says. "Remember that luminous fireball and the billowing black cloud of smoke during the aircraft explosion and fire on September 11? Both contain large amounts of soot. The luminous part is the soot particles radiating; smoke comes when the soot particles have cooled down."

Image 1



Both images are snapshots (one time frame) of a transient 3-D 1 meter heptane pool fire simulation completed at the University of Utah. These images show the soot volume fraction in the fire. Image 1 simulation uses traditional equilibrium calculations for chemistry. Image 2 uses the intrinsic lower dimensional manifold (ILDm) method for the chemistry calculation. ILDM is what Nguyen is working on as part of her dissertation.

Image 2



The real concern is transportation accidents, both airline and ground, but our methods apply to forest fires, too.

*These images are snapshots of a transient 3-D simulation of a 10 meter heptane pool fire completed by Nguyen's research group at the University of Utah. This simulation was computed on 1000 processors (LANL machine) for 36 hours and the images are showing temperature. Although this simulation uses only equilibrium, Nguyen is currently coupling the ILDM chemistry to the same code so that she can produce a detailed chemistry simulation of the same large scale pool fire. The detailed chemistry should greatly improve the structure and temperature predictions of the fire.*

*The arrows in Image 2 and 3 are velocity vectors. They represent the direction of the fluid flow (in this case the air and fuel and hot gases). The size of the arrow indicates the magnitude of the flow.*

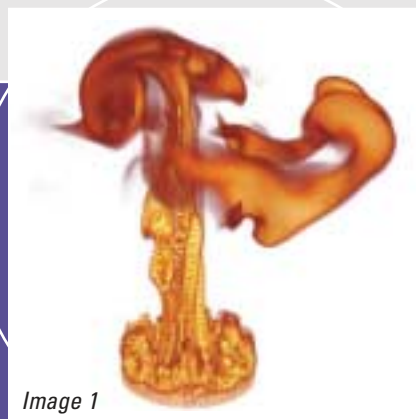


Image 1



Image 2

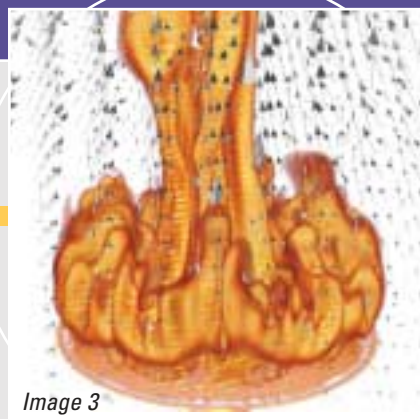


Image 3

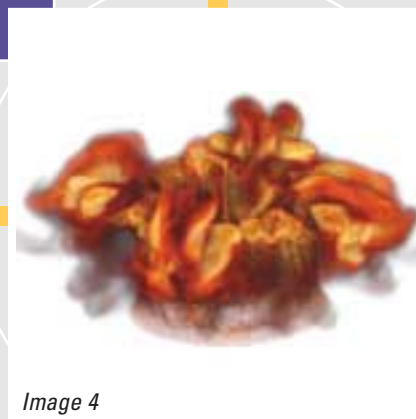


Image 4

### Flaming Rubik's Cube

So-called computational fluid dynamic codes, which factor in both flow and detailed chemistry that are equally important in figuring out what fire will do next, occupy enormous swaths of computing space and time. Codes must model the flow characteristics or chemistry that unfolds in each region of the fire. Then it is up to the investigator to assemble the pieces and deduce the course of a fiery event.

"An actual solution of the system of turbulence and chemical equations without modeling would be computationally prohibitive," Nguyen says, "taking thousands of years given our current computing power."

Through a standard technique called the intrinsic lower dimensional manifold (ILDm) method, an investigator can whittle a thousand or more reactions down to three that are representative of the whole chemistry during an event. "You can do a mathematical analysis to determine which reactions are dominant, which reactions are most important at a given time under given conditions," Nguyen says.

A simulation that tracks the evolution of a fire through time must account for temperature fluctuations, flame structure, soot, pollutants and other factors that affect the life of a fire at 1-, 2-, 3-second intervals, and so on — or "time steps," as Nguyen calls them, each a snapshot of the fire after it is ignited. Each time step requires an accompanying series of what are called Navier-Stokes equations, having to do with mass conservation, momentum, energy and chemical reactions. And for a realistic picture, they must be solved simultaneously.

Suppose Nguyen's imaginary fire ignited from a pool of liquid fuel 1 meter in diameter — about the size of a campfire. To model the growth of the fire with precision, she must "set up a computational domain larger than the fire" — a 5 by 5 by 5 meter cube enclosing the campfire pit. Each factor calculated can be highly variable, depending on when it happens and where. For example, the fire will be hotter in the center than at the edges. Nguyen pictures the 5 meter box as a see-through Rubik's cube, with little red-hot squares at the center of the cube, fading into less-bright hues as her mind's eye moves toward the surrounding medium, air. In fact, that's exactly what the model does — subdivides the big cube into 100 boxes in each direction, 100 by 100 by 100.

This allows calculation of such things as a mean temperature for each box — the way a weatherman would assign separate temperatures to, say, Cleveland and Columbus. Just as crime statistics,

**Right off, Nguyen could see that the fellowship was unusual not only for the amount of support it provided but also because it required a three-month practicum at one of the Department of Energy's national labs.**

parking problems and smog vary from city to city, conditions inside each box vary in their physics and chemistry and thus must be separately calculated.

"That means," Nguyen notes, "that we solve for temperature, velocities, pollutants and so forth in every single box — that's a million boxes. So for each time step, we solve a million sets of equations. You can see why we need a supercomputer: Our total number of calculations equals a million times the number of time steps in the overall simulation."

Each of those million boxes constitutes a "grid cell," together forming a "mesh." Nguyen has introduced an innovation into the mix to reduce the size of the mesh, and hence the calculation time. She incorporated a "subgrid" scale chemistry model that, in effect, shrunk the number of little boxes from a million to 125,000 without losing the ability to capture the detailed chemistry and physics in each tiny cell. Like thousands of local weathermen working independently, "the little squares become decoupled, independent of each other, so that you can calculate the chemistry on many different processors, and that saves time."

Even at that, these codes require as many as a thousand processors running continuously for a day and half to yield six seconds of data.

### More Livermore, Anyone?

Nguyen's entry into the DOE CSGF program was as inevitable as e-mail. "I didn't know anything about the fellowship program. I got an e-mail forwarded to me about it and was told to go to a website. I got an application and submitted it three days later. Now there are posters all over the department."

Right off, Nguyen could see that the fellowship was unusual not only in the amount of support it provided but also because it required a three-month practicum at one of the Department of Energy's national labs.

"You get to choose anything you want to, at any of the labs. It's a nice opportunity. There's no overhead; the DOE takes care of everything."

She met with scientists at Sandia, Los Alamos and Lawrence Livermore laboratories, choosing the latter for reasons not entirely scientific: "I really wanted to spend some time in the San Francisco Bay area."

And it didn't hurt, either, that Livermore was home to the Center for Applied Scientific Computing, which put her "in touch with different computing systems, supercomputers and a variety of software that give you a chance to explore aspects you may not have explored in school."

### 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 towards a PhD at a United States university.

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 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.

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*.

Would she do it again? Yes, definitely. And literally. “The first experience was so positive, I went back for a second one.”

The most valuable part of a practicum in computational science shares a term from computing but has nothing to do with any particular computational apparatus: “networking,” Nguyen says. “It’s great exposure to the scientific community, to other minds, to people who have ideas that are different from yours. People working in the combustion community have different points of view. The practicum fosters new creative thinking; it’s an ongoing education.”

And possibly “a little scary,” Nguyen confesses. “Everyone is so smart.” As if the indigenous scientists aren’t overwhelming and impressive enough, the labs often bring in special guests, famous Nobel laureates. “I got to meet Edward Teller.”

Would she do it again? Yes, definitely. And literally. “The first experience was so positive, I went back for a second one.” To avoid being tied to one computational flow code, she spent her second summer testing her subgrid reaction models and coupling them to codes developed at Livermore. “Our group has had half a dozen

DOE CGSF graduate students work with us during the summer over the past four to five years, and we have found that they are all very good and enthusiastic,” says William Henshaw, a Lawrence Livermore applied mathematician and Nguyen’s practicum adviser.

What about a third practicum for Nguyen? No — as much as she would have liked another, she is too busy finishing up her degree. But she hopes that she is not finished with the national labs.

Before her practicum experiences, she hadn’t really considered a career at a national lab. The whole idea of working at a government lab sounded dull, programmed, one-dimensional. “But I found it was such an open community, with everyone encouraging you and great resources, which include not only the equipment but also the people.”

To which she adds, without pause, “I would be interested in working at a national lab. Unlike a university, you don’t have to worry about teaching. I know now that I want to concentrate on research.”

## Developing a Mathematical Toolbox

BENJAMIN KEEN

University of Michigan | Lawrence Berkeley National Laboratory | Story by Victor D. Chase

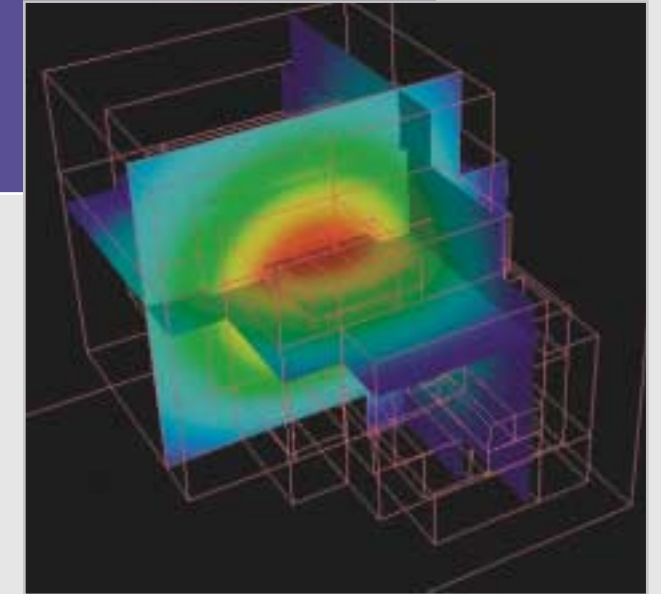
His 6-foot 7-inch build goes a long way towards making him “a very striking physical presence,” says Phil Colella, a senior scientist at Lawrence Berkeley National Laboratory (LBNL), as he describes Ben Keen. Keen, a 26-year-old grad student in mathematics at the University of Michigan, spent the summer of 2001 working in Colella’s lab in Berkeley, California, while fulfilling the practicum requirement of his DOE Computational Science Graduate Fellowship.

In keeping with his physical appearance, Keen also has a “striking and confident personality,” says Colella, who served as Keen’s practicum advisor. This self-assurance allowed Keen to interact with the other scientists in Colella’s group just as though he were one of them, which “is unusual for someone his age, and at his point in their career,” said Colella. “He behaved as a peer of the others in the group.” As for his part, Keen clearly enjoyed being at LBNL. “It’s a nice collegial atmosphere,” he says. “Very university-like except that you’re not at a university.”

### A Computational Toolbox

Keen’s practicum mission was to work on the development of an algorithm that is part of a larger software “toolbox” designed to provide scientists involved in

The Applied Numerical Algorithms Group at Lawrence Berkeley National Laboratory has developed a series of adaptive mesh refinement (AMR) tools for solving problems involving partial differential equations. This image shows how two of these AMR tools, Chombo and ChomboVis, can be used to solve a Poisson problem and visualize the results. Such a problem may involve determining the equilibrium state of electrical fields or temperature distributions.



computer simulation with an array of standard mathematical tools that they can turn to and adapt to their specific needs. The intent is to avoid having to reinvent the mathematical wheel each time one is needed.

“It’s a set of tools that would allow us to easily build custom software that on one hand would help someone doing combustion engineering, and on the other hand the same set of tools would allow us to help high energy physicists who are building particle accelerators,” says Colella. “Even though those are very different applications, the underlying mathematics is very similar; consequently the same software tools can be reused in different ways to support those different applications.”

The umbrella programs under which this work falls have the flavor of alphabet soup. First comes SciDAC, for DOE’s Scientific Discovery through Advanced Computing program. While the development of widely applicable computational tools has been going on for almost a decade, it has gained steam within the past year, after becoming a core SciDAC program. Then, under SciDAC, there is ISIC, for Integrated Software Infrastructure Centers. One of those centers is the Applied Partial Differential Equation Center, or APDEC, which is Colella’s group.

And within APDEC there is Chombo, which is not an acronym, but a Swahili word aptly referring to “variously a tool, a container, or a generally useful implement,” explains Colella. The algorithm that Keen was involved in is part of Chombo.

The fact that the algorithm has no name of its own may be indicative of the fact that algorithms do not seem to get their due when credit is divvied up for the ongoing massive leaps in computer power and speed. “In the last 10 years a lot of people have come up with really excellent algorithms that are every bit as responsible as [improved hardware] for increasing computer speed,” says Keen. Yet they get little credit.

“It would be too strong to say that algorithms get zero credit among those who know, but certainly to the general public, the idea is that the ‘big computers’ are the big win — but it’s really the ‘smart algorithms.’ It’s not increasing computer speed so much as allowing people to do larger simulations faster,” he says.

### A Mathematical Middleman

The algorithm that Keen was helping perfect during his practicum relates to work being done to represent complex boundaries in geometry. “Things like the interior of a cylinder in an automobile engine,” says Colella.

While the representation of the boundaries is a separate effort, it can be analogized to the creation of a weirdly shaped cake pan.

“What I was involved in was making tools for baking a cake in that pan,” says Keen, the tools being partial differential equations.

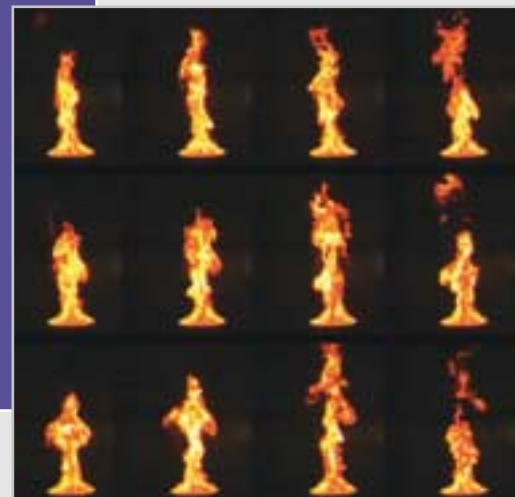
This was a perfect assignment for Keen, who sees himself as a mathematical middleman, otherwise known as an applied mathematician. “There are people who work on purely theoretical tools, and then there are people more on the engineering side, who try to apply numerical schemes to specific problems. I stand between the two — from the more theoretical points of view down to the more practical, so I think of myself as a tool builder,” says Keen.

Keen compares his vision of himself in mathematics and computational science to the role of a process chemist, which happens to be what his father is in Charleston, West Virginia, where Keen grew up.

“Even though those are very different applications, the underlying mathematics is very similar.”



These images are photos taken from a high speed video (500 fps) as part of an experiment at the University of Utah. They are images of a burning pool (12 inch diameter) of Jet A fuel. The different frames are 80 ms apart. This experiment is an actual fire similar to the one illustrated by the simulation images on page 9.



Reference: Ciro, W. “Heat Transfer at Interfaces of a Container of High-Energy Materials Immersed in a Pool Fire”. PhD Thesis (In Progress), University of Utah, 2002.

*ChamboVis utilized to solve Poisson problems and visualize the results.*

“A theoretical chemist might come up with a new system of synthesizing a compound. The purely practical person would be the chemical engineer who looks at making a million pounds of this for the lowest possible cost with the minimal amount of waste.

“But, there’s a process chemist who stands between these two chemists to move the technology along. That person takes something from a lab to a pilot plant to a commercial plant. That is where I stand. I see myself as a process guy who moves something from a small demonstration that still needs to have some work done on it so the average person can take it off the shelf and use it,” he says.

## SCOPE OF PROGRAM

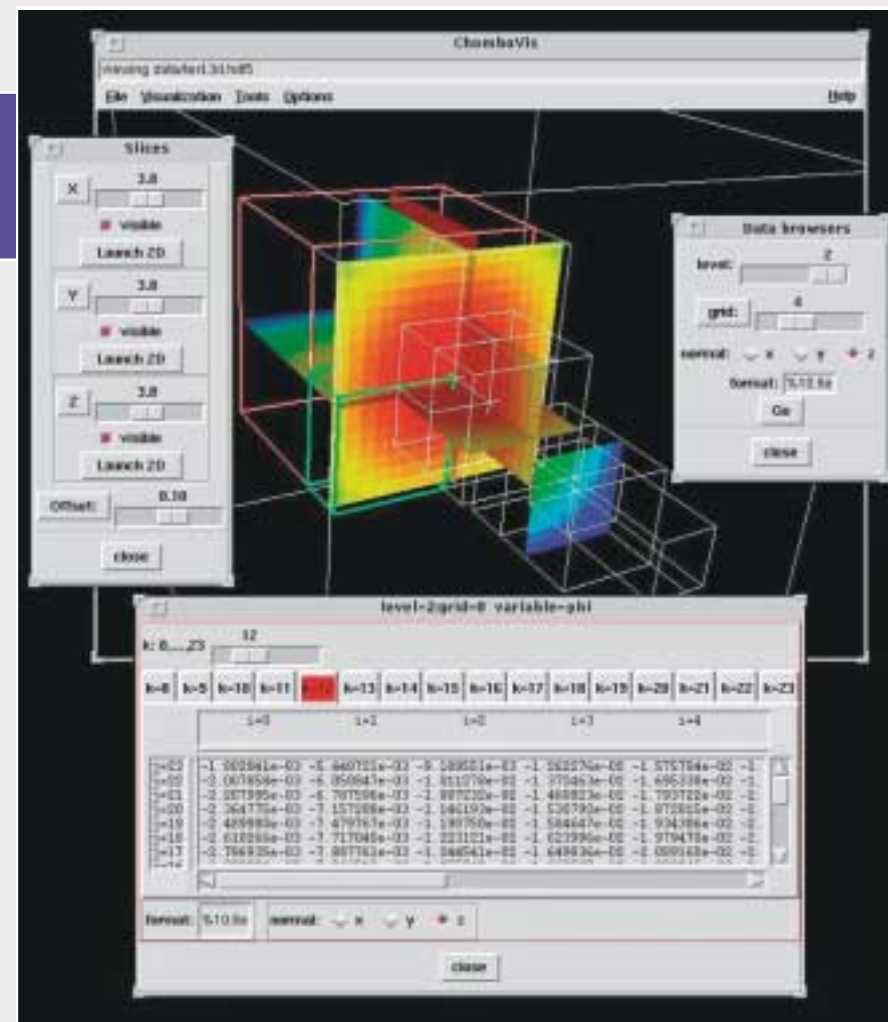
Since its inception, the DOE CSGF program has supported 164 students in over 50 universities all over the U.S. Currently it supports nearly 75 students in 22 states: Arizona, California, Colorado, Georgia, Illinois, Iowa, Massachusetts, Michigan, Minnesota, New Jersey, New Mexico, New York, North Carolina, Oregon, Pennsylvania, Rhode Island, Tennessee, Texas, Utah, Virginia, Washington, and Wisconsin.

### Of Water and Air

In 1998, Keen graduated *summa cum laude* from the University of Southern California and took up residence at University of Michigan, where he had been awarded a Regents Scholarship to pursue a PhD in mathematics and computational science, which he hopes to receive in 2003. On the way, he has picked up a Master’s in applied mathematics.

His dissertation will deal with inviscid compressible flows, which pertains to high-speed flows possibly containing shock waves.

“Both water and air exhibit viscosity and compressibility,” explains Keen. “In the kind of flow I’m interested in — flow around a supersonic aircraft,



for example — one neglects the effects of viscosity because they’re relatively small and pays attention to the effects of compressibility because they’re very important. Likewise, someone interested in the flow of water around a submarine would take the opposite point of view. In either case, the same basic ideas go toward building the partial differential equations one finally solves, but the equations naturally turn out to be very different, and the numerical tools one uses are likewise very different.”

Computational science enters the mix when this technology is applied to the computer modeling of things such as air passing around an aircraft wing. The computer simulation technology that Keen is working on could be used, for example, to test a new wing design through simulation, rather

than having to go through the more costly and slower process of building a prototype and testing it in a wind tunnel. Similar simulations could be used to examine a wide range of fluid flows, from what goes on inside turbo machinery to a nuclear explosion.

### The Fellowship

Keen applied for the DOE Computational Science Graduate Fellowship (DOE CSGF) in late 1999, at the suggestion of Robert Krasny, a faculty member in the University of Michigan mathematics department. “I really didn’t expect to get it,” he says, “but it seemed worth a shot.”

While Keen found the application process straightforward, he did find the requirement that he fill out a

program of study, including a list of classes he intends to take, different from the requirements of other grant applications he had completed. This is because the DOE CSGF requires more classwork than do most other fellowships, to provide the interdisciplinary training needed to work in computational science.

Taking the time to fill out the application for the DOE CSGF proved to be worth a lot more than a shot for Keen. He received a positive response to his application within a month of submitting it. Since then, “I’ve ended up taking far more classes than other people in my department,” he says. In fact, he adds, “I’m probably just a few classes away from qualifying for a Master’s in aerospace engineering.” Which is all to the good, as he sees it, since his

program for some 10 years as a member of various panels and as a professor at the University of California at Berkeley, Colella knew that a practicum is a fellowship requirement. When he and Keen struck up a conversation, the similarities of their work soon became apparent.

### “It Will Compute”

During his 12-week stay at LBNL, Keen found that as an “outsider living in the house” of the LBNL researchers who created the algorithm, he was able to bring a fresh perspective to their work and turned up several problems in the software.

Of the perspective he brought to the party, Keen modestly notes that

**Despite Keen’s modesty, Colella points out that the kind of work he did during his practicum is particularly difficult because it deals with reconciling the ideal circumstances of mathematical theory with the imperfect world of mechanical systems.**

professional function as an applied mathematician will require working closely with engineers. “I’ve got to be able to speak to engineers,” he says, “so I can’t just take math classes.”

Keen found his practicum match through his University of Michigan academic advisor, Smadar Karni, who invited Colella to speak at a conference she had organized. Having been involved in the DOE Computational Science Fellowship

“It’s not that I had anything to offer concerning the architecture, etc., but just that ‘hey, here’s an outside user fitting a number of the qualities of the eventual users of this software’ — so the things that I find difficult/confusing others might too, and the capabilities that I might want as part of implementing this algorithm might be things that other folks need too. “I was taking a first user pass through some code, and blowing up some

particular landmines that everyone knew were going to be in there,” Keen explained. “It didn’t require a deep understanding on my part, just an ability to be careful about asking the right kind of questions. So it wasn’t too hard for them to say ‘Oh, you need to do this such and such a way’, or ‘Ah, this is messed up, let me fix it.’ ”

Despite Keen’s modesty, Colella points out that the kind of work he did during his practicum is particularly difficult because it deals with reconciling the ideal circumstances of mathematical theory with the imperfect world of mechanical systems. And, he says, “that’s something Ben took to very quickly. He was able to identify problems with the algorithm we had proposed, and working with him we’ve come up with solutions that make it more robust. As a result, it will compute something sensible and correct under almost any circumstances.”

And even though Keen finished his practicum in 2001, he continues to serve as a consultant to Colella’s group. Members will, on occasion, send e-mail messages to him asking what he thinks of an issue, such as how errors are reported to users. And, says Colella, “Even though he’s in Michigan, and has other things on his mind, he responds in a very thoughtful way.”

Keen also hopes to use Chombo in his own research. “It’s very good to be able to make use of this tool,” he notes.

Undoubtedly the highest compliment of all came when Colella was asked if he would hire Keen as a full-time member of his group. “Wouldn’t hesitate for a second,” he shot back, even before the question was complete.

## DOE CSGF HIGHLIGHTS

- > **Payment of all 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](http://www.krellinst.org/csgf)

By Jacob Berkowitz

# Supernova Express

**A WHITE DWARF IS EXPLODING.** This carbon-oxygen star, slightly smaller than our sun, has been gradually accreting hydrogen, helium and other atoms from a neighboring star. This siphoning-off has doomed it. As the white dwarf grows in mass, its temperature and density rise without bound as it approaches the Chandrasekhar limit. Then the carbon atoms ignite.

Trillions of high-energy protons careen through the star's supercharged atmosphere. The result is a runaway thermonuclear explosion that will completely obliterate the star.

Peter Nugent watches this Type Ia supernova on his computer screen at Lawrence Berkeley National Laboratory's National Energy Research Scientific Computing Center (NERSC). The explosion process wasn't exactly what he was expecting. Which is fine, because he created it.

A staff scientist with NERSC's scientific computing group, Nugent is at the forefront of a new realm of supercomputer-based computational astrophysics. He grew up looking at the stars in his Easton, Pennsylvania, backyard through a Newtonian telescope given to him by his grandfather. Now he's using an IBM RS/6000 SP, the world's second-most-powerful unclassified supercomputer, to virtually model solar phenomena that would otherwise be impossible to test experimentally in the lab.

## The Dawn of a New Era of Supernova Science

When Nugent graduated with a PhD in physics from the University of Oklahoma in 1996, he was on the cusp of a revolution in Type Ia supernova research.

Three years earlier, the Carnegie Institution of Washington's Mark Phillips demonstrated that these supernovae are remarkably accurate beacons of cosmic distance. Since they all explode with about the same amount of mass, in a near-uniform fashion, their spectra and light curves represent 'standard candles' which vary with distance and time.

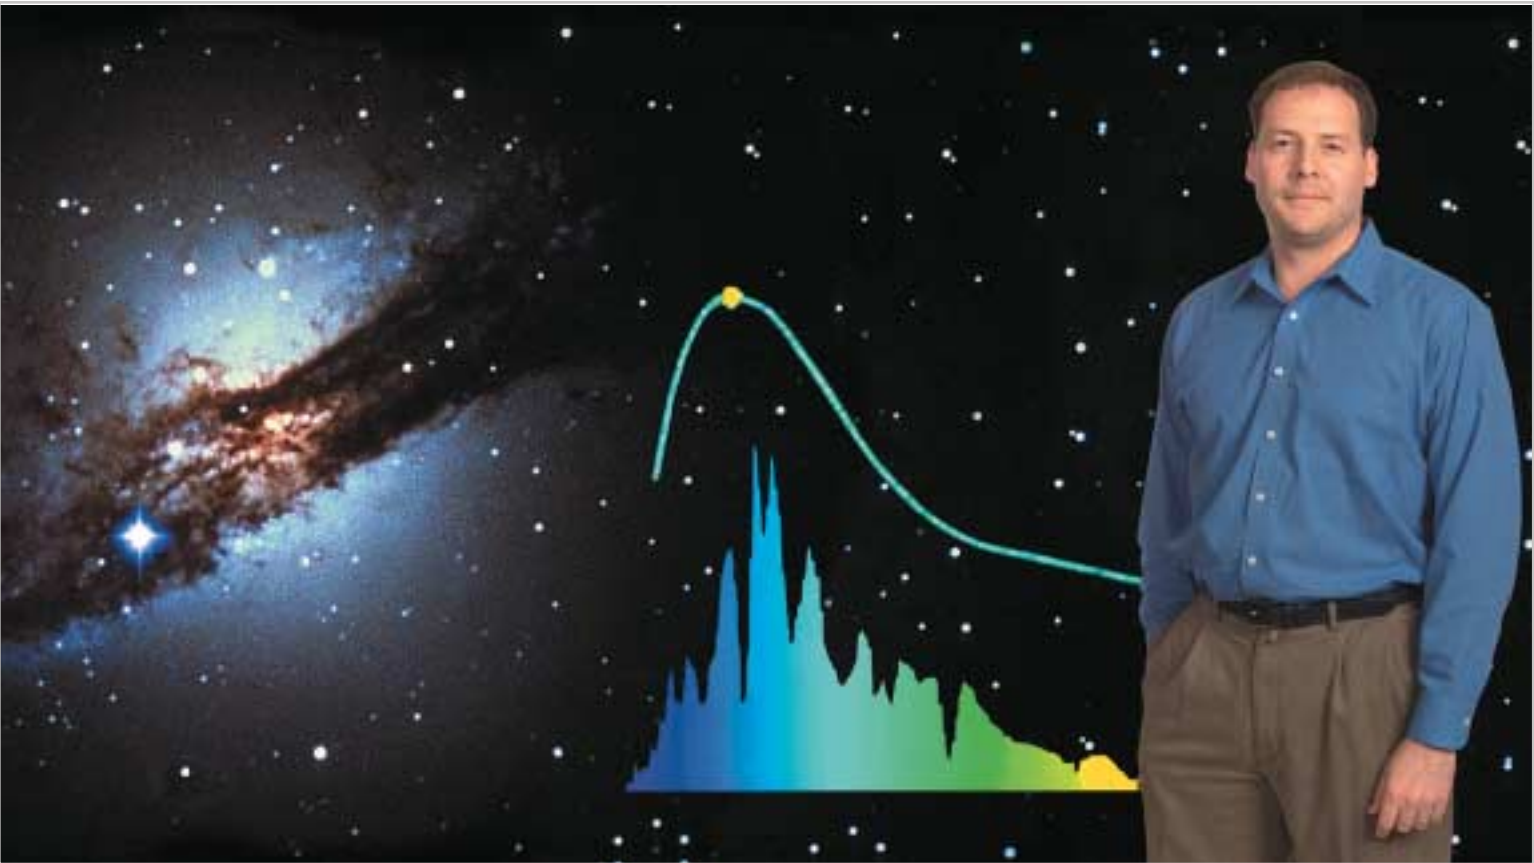
The new understanding of the value of Type Ia supernovae to cosmology was quickly bolstered by two major technological advances in astronomy: the launch of the Hubble Space Telescope and the development of the charged coupling device (CCD). As a space-based telescope, the Hubble is able to capture the fainter light of ancient high-redshift supernovae, unobservable from Earth-based observatories. As digital imaging systems, CCDs enable electronic computational subtraction of two images taken at different times.

This confluence of theoretical and technological advances led to the creation of the Supernova Cosmology Project (SCP), led by Saul Perlmutter of Berkeley Lab's physics division. The project's goal is to use high-redshift Type Ia supernovae in exploring a bevy of cosmological questions. The group pioneered the automatic detection of supernovae through using CCDs.

Nugent's doctoral thesis, on the modeling of Type Ia supernova spectra, made him a perfect addition to the Supernova Cosmology Project.

"When I graduated, the Project didn't have anyone to interpret their data and push the envelope by suggesting better ways to take measurements," says Nugent, who was hired as the Project's first theoretician, starting out as a post-doctoral researcher in Berkeley Lab's physics division.

The project's goal is to use high-redshift Type Ia supernovae in exploring a bevy of cosmological functions.



Peter Nugent studies supernovae to answer basic questions about the nature and destiny of the universe.

"There's a great synergy between observations and theory in this field. For me it's a very exciting one because there's a lot of contact between the theorists and the observers. What we do with the modeling is create a theoretical spectrum and then compare that model with observed data. We then frame questions in ways that people can do data analysis and say, 'The models are telling us that if these parameters change then we'd expect to see this signal,'" says Nugent.

This joint observational and computational modeling approach has already led to what the journal *Science* called the "Breakthrough of the Year for 1998". Using dozens of supernovae at distances greater than four billion light years, the SCP and their competitors, the High-Z Supernova Search Team, were able to determine that the universe is expanding at an accelerating rate.

It's an observation that implies the existence of a mysterious, anti-gravity property of space called "dark energy".

Two years later, using data from SN 1999ff, the most distant Type Ia supernova ever found, Nugent and Adam Reiss of the Space Telescope Science Institute were able to roughly determine when the expansion of the universe began to accelerate.

## Modeling the Insanely Big

Like genomics and climate modeling, creating virtual supernovae deals in the realm of the mind-numbingly large and complex. In short, it's a job for a supercomputer.

Modeling the spectrum of a Type Ia supernova means recreating a combination of millions of spectral lines. Each line represents an atomic

type within the exploding star's rapidly evolving atmosphere. These include sulfur and silicon, with hundreds of thousands of lines representing iron in various transition states. The spectral lines are in turn strongly dependent on temperature and density in the various zones of the star's atmosphere.

Mimicking the observational data, the models begin at about one second after ignition and continue to 20 days, at which time the supernova reaches peak brightness.

"The memory requirements are incredibly large," Nugent says. "The atomic data we feed in to do our calculations is about a gigabyte of data. It's handled differently for 100 different zones of the atmosphere, so that makes for more than 100 gigabytes of I/O. And then we model each of the atoms."

Peter Nugent graduated from Bowdoin College in 1990 and attended the University of Oklahoma, where he received his PhD in Physics (concentration in Astronomy) under Eddie Baron and David Branch in 1996. From there he started a post-doc at Lawrence Berkeley National Laboratory (LBL), working for the Supernova Cosmology Project under Saul Perlmutter. In 1998 the group shared the journal *Science* 'Breakthrough of the Year Award' for the discovery of the accelerating universe. In 2001, Peter became a scientist at the National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory and that summer was an integral member of the group that discovered the most distant supernova to date. He has recently been appointed to the Hubble Space Telescope User's Committee and is a lead member of the Supernova Acceleration Probe and Supernova Factory.

**Further Reading**  
*K-corrections and Extinction Corrections for Type Ia Supernovae*, P. Nugent, A. Kim & S. Perlmutter, Publications of the Astronomical Society of the Pacific, 2002, vol 114, pg. 803.

*The Farthest Known Supernova: Support for an Accelerating Universe and a Glimpse of the Epoch of Deceleration*, A. Riess, P. Nugent et al. Astrophysical Journal, 2001, vol 560, pg. 49.

*Measurements of Omega\_M and Omega\_Lambda from 42 High-Redshift Supernovae*, S. Perlmutter et al., Astrophysical Journal, 1999, vol. 517, pg. 565.

**Contact:**  
Peter Nugent  
nugent@panisse.lbl.gov

Moving to 3-D

Until recently, his supernova models were created in one dimension. Now the work is about to get a lot harder.

“We’re just about to go to 3-D with our code, and this will cube the time and memory requirements and exhaust the amount of computing time I have to run a couple of simulations. So I want to make sure I’m doing this in the best possible fashion,” Nugent says.

The move to 3-D models is part of the SCP’s efforts to move beyond spectral analysis to gain even more detailed understanding of these stars’ behavior.

One key element of this is study of the polarization spectra. As Nugent puts it, when a Type Ia explodes, is it aspherical like a football or relatively round like a basketball?

Addressing this question entails modeling very small polarization differences (in the range of half of one percent) in the flux spectrum. And this in turn requires a model that involves a much greater number of photons, and requires more memory, in order to produce a statistically significant result.

To make the move to 3-D, Nugent is drawing on NERSC’s computational science expertise to create ever-smarter simulations.

And it just so happens that one of the best ways to model a supernova is to physically mimic it within the computational hardware.

“We divvy up our atmosphere between the processors. The way you can think of it is that the photons are making their way through the supernova atmosphere by traveling from processor to processor,” says Nugent.

With more than 3,000 processors, NERSC’s IBM RS/6000 SP is ideal for this modeling, which requires rapid communication between processors and large amounts of memory. The processors are distributed among 184 compute nodes with 16 processors per node. Each node has a common pool of between 16 and 64 Gbytes of memory.

Yet even with this computational workhorse, there’s still the question of how to best divide up the atmosphere to maximize the computational efficiency.

“You could have a very low-density chunk of your atmosphere where there are very few interactions. So if a photon comes in, it basically leaves immediately, you do very little with it. So you’d like to have just a few processors handle all those parts of the atmosphere,” he says. Nugent is working with NERSC computational scientists Osni Marques and Tony Drummond to assess questions of load balancing, as well as exploring ways to use dynamic allocation of the processors for on-the-fly reallocation of processing power.

Dynamic reallocation could play a particularly important role when it comes to creating 3-D visuals. Nugent is collaborating with NERSC’s science visualization staff to finesse both the processing and data representation of these models.

“At some point we’re going to get simulations that will take days to run and we may want to go into the middle of a simulation and look at the data in a visual way and ask: ‘Is this going anywhere near the result we want? And if not, let’s just stop it right now and start down another path,

so we don’t waste all our time on the supercomputer.’ So they would help us interact with the supercomputer,” says Nugent.

The role of computational supernova simulations will be increasingly important as the SCP shoots for its next goal: SNAP — the SuperNova Acceleration Probe. This multi-agency, multi-institution proposal, led by Saul Perlmutter and Michael Levi of Berkeley Lab’s physics division, is for the creation of a space-based, 2-meter telescope with the largest and most sensitive CCD ever.

Unlike even the biggest ground-based telescopes — such as the Keck 10-meter telescope in Hawaii — currently used to detect and take high-redshift supernova spectra, SNAP will be able to detect and take spectra of the faint light (especially in the infrared) from the most distant and ancient high-redshift supernovae. The data would be essential to helping us understand the nature of dark energy, says Nugent. It will help not only to confirm that the so-called cosmological constant is nonzero, but also to answer questions as to whether it has changed over time.

According to Nugent, theoretical modeling will help guide what SNAP actually snaps images of. The plan is for the probe to take images in batch mode and probably find thousands of supernovae. The question would be which ones to focus on so as to capture their spectra.

Nugent plans to use simulations to create a comparative grid of both Type Ia and Type II supernova data. It will enable observers to use the initial handful of data they receive to select the supernovae they want to follow.

“Right now we’re sitting at a point very much like Einstein at the turn of the last century. We’re trying to detect things that we have no idea what they are — dark energy. It has huge importance for understanding the universe and basic physics. Getting that data into a form so that theorists can understand it will come about through supercomputing.”

WHY THE DOE FINDS ITSELF AT HOME IN THE STARS

>> *Supernova research brings together the really big and the really small. And in the process it demands an astronomical amount of collaboration, says computational astrophysicist Peter Nugent.*

“There has always been a tie between astrophysical phenomena and work in nuclear physics. With both there’s a whole range of energies, temperatures and densities that are inaccessible in the laboratory,” says Nugent, a staff scientist with the Department of Energy’s National Energy Research Scientific Computing Center (NERSC) based at Lawrence Berkeley National Laboratory.

Nugent notes that physicist Hans Bethe is perhaps best known for his role in the Manhattan Project. However, Bethe’s 1967 Nobel Prize in physics was awarded for his contribution to the theory of nuclear reactions, especially as it related to his discoveries concerning energy production in stars.

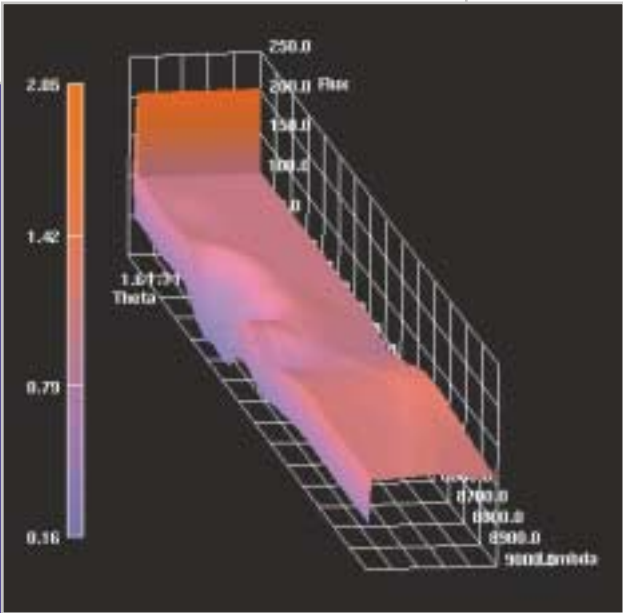
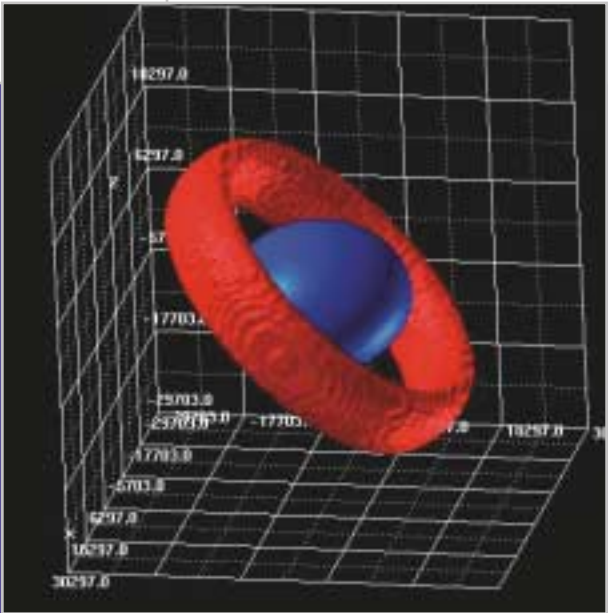
The intersection of elementary-particle and nuclear physics with observational and computational astrophysics means that supernova research requires a high degree of interdisciplinary cooperation. A collaborative

approach is also motivated by the need to use a variety of telescopes and staff to identify, track and record spectra from supernovae during the course of their visible lives (about 3 months from Earth-based telescopes).

In 1999, Nugent was one of 32 authors — from the observers to the data processors and interpreters — on a supernova journal article for which he had provided the computational analysis.

Says Nugent: “With these cosmological observations, astronomy and particle physics are coming together in a way in which, if you understand cosmology, then you certainly have great insights into fundamental particle physics. Because we’re pushing back to exactly what came out of the Big Bang and what are the equations which govern the physics of the universe. In that, many people believe there will be great insights into particle physics.”

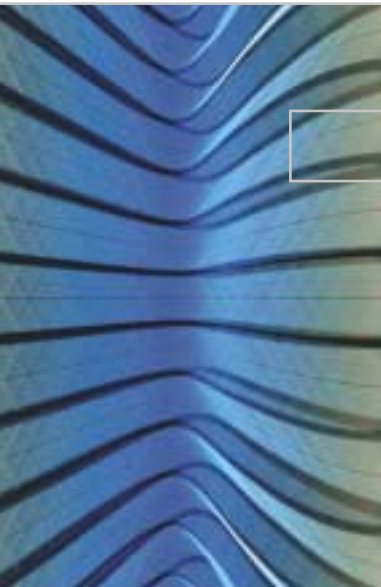
Here we present a spectrum synthesis calculation of a supernova atmosphere surrounded by a toroid. The layout of the atmosphere is presented on the right while at the left we have a graph of the flux vs. wavelength vs. viewing angle. As the viewing angle shifts towards the toroid, the strength of the absorption increases dramatically. Data that confirm such a model would for the first time put strong constraints on the progenitors of Type Ia supernovae.



By Peter Gwynne

# Simulating<sub>the</sub> Ultrasmall

**COMPUTER SCIENTISTS HELP** to ensure the safety, security, and reliability of the nation's stockpile of nuclear weapons by developing better ways to design and analyze devices of microscopic size. The Cold War may have ended a decade ago, but a stockpile of nuclear weapons remains an essential element of national defense.



>> Astrophysics

>> Nanoscience

>> Climate Change

>> Laser Physics

>> Computational Medicine

>> Radiation Transport

>> Computational Biology

>> Materials Science

To maximize the stockpile's ability to discourage attacks by enemies, particularly in these troubled times, engineers must ensure that the weapons remain in reliable working condition. They must also refurbish them continually to incorporate the latest technological advances.

Those tasks face a key limitation. Like a majority of other nuclear nations, the United States refuses to test its weapons directly by exploding them in the atmosphere or underground. So the guardians of the stockpile – most of them in the Department of Energy's national laboratories – are placing greater reliance on simulation techniques to keep the weaponry up to date.

The Microsystems and Engineering Sciences Applications (MESA) project at Sandia National Laboratories in Albuquerque, New Mexico, is designed to provide a critical element of such simulation work. It will facilitate the design, production, and testing of weapons components called microsystems. "Integrated microsystems will streamline the number of weapons components, reduce production maintenance costs, and minimize the number of fault points in the weapons system," explains Sandia scientist Rodney Schmidt.

Schmidt has particular interest in one facet of microsystems known as MEMS, for microelectromechanical systems. These are gears, motors, sensors, and other parts no larger than grains of pollen and made entirely from the silicon compounds used in computer chips and integrated circuits. They have potential uses in fields as diverse as consumer electronics and medicine.

Sandia has developed pioneering technologies that enable the production, on computer chips, of intelligent MEMS systems that know where they are and what is happening around them. This work, says Schmidt, "is an important element of our stewardship over maintaining the continued safety, security, and reliability of the stockpile."

Not only are MEMS components minuscule in size, but many of them have extremely complex shapes. So it's no surprise that designing and fabricating the parts can present extraordinary challenges.

## The Power of Computer Simulation

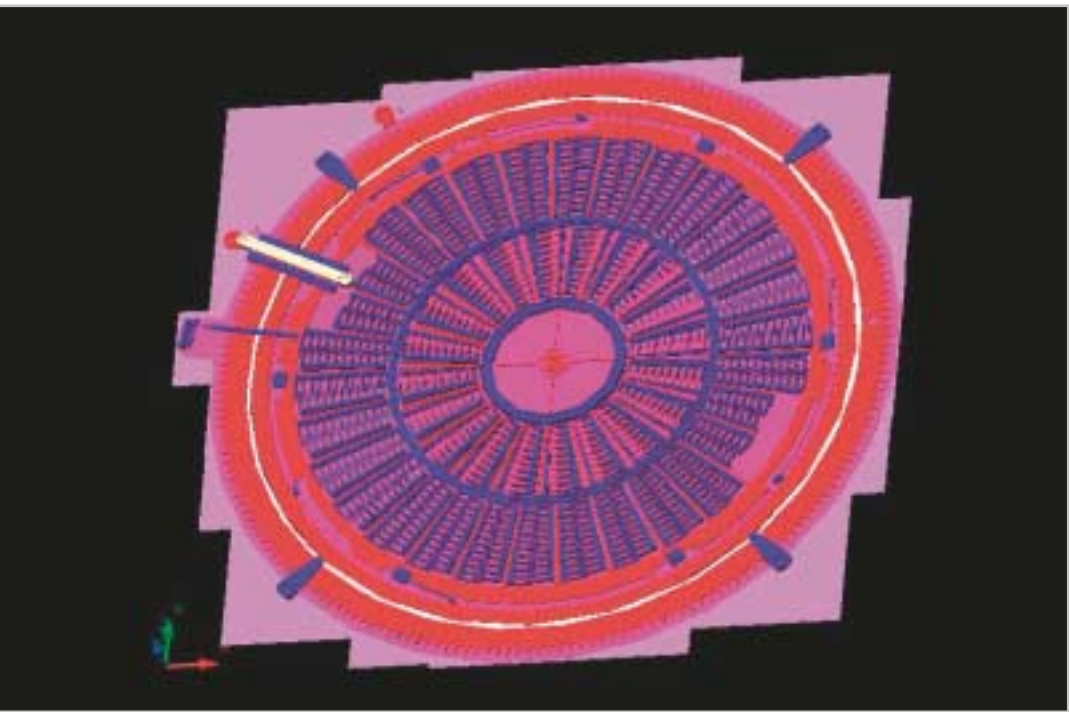
Experiments in fabrication facilities to test the designs of MEMS devices are a big problem because good measurements are very difficult to make at extremely small scales; the experiments typically require a long lead time to set up and conduct, and as a result the tests cost a great deal of money. Fortunately, a cheaper alternative is emerging: using computers to simulate design and performance.

That's where people like Schmidt come in. As leader of a Sandia project called Computational Technology for MEMS, he works closely with MEMS scientists and engineers to determine how computer simulation can best help them. The project is part of Sandia's Science-based Modeling

and Simulation Roadmap for MEMS, an effort that aims to coordinate and tie together all the different types of MEMS computational work needed at Sandia.

"The MEMS project is part of the MESA vision, recognizing that Sandia's computer science center needs to be active in contributing to the MEMS work and to satisfy the computational needs of the people in the MEMS center," Schmidt explains. "Computation has become an essential element in our vision of the national laboratories. We believe that the development and application of computational tools in support of MEMS R&D is fundamental to advancing the state of the art in MEMS design, fabrication, and performance analysis. We expect that computational tools will impact all aspects of MEMS work at Sandia in the future."

What specific mission does Schmidt's group have? "Our goal is to identify and develop the computational tools, algorithms, and models most important to making simulation science a key element of Sandia's MEMS R&D program," he explains.

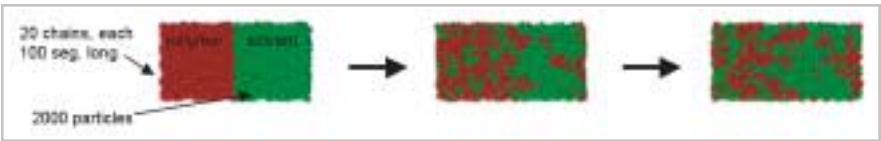


## Fabrication and Design Challenges

Many of the microfabrication technologies used to produce MEMS devices stem from the techniques that the microelectronics industry has developed to create electronic semiconductor devices. The three dominant technologies are surface micromachining, bulk micromachining, and LIGA. Surface micromachining forms 3-dimensional structures by deposition on and etching of thin films. Careful construction of the lithographic masks that control these steps and the application of a final selective "release" etch permits the

creation of complicated free-standing movable parts. In bulk micromachining, wet chemical etching (using various acids or hydroxide mixtures) and dry etching (using highly directional plasma techniques) allow deep, high aspect ratio features to be sculpted into the bulk of materials such as silicon, quartz, and glass. LIGA, jokingly referred to as Long Incomprehensible German Acronym, stands for Lithographie Galvanoformung und Abformung, or lithography, electroplating, and molding. The process forms small 3-D structures by mold fabrication followed by electro-deposition.

Computationally generated solid model of a Torsional Ratchet Actuator (TRA) with 13,833 vertices. The diameter of the outer gear is 825 microns.



Molecular Dynamics simulation of polymer dissolution in LIGA.

"Integrated microsystems will streamline the number of weapons components, reduce production maintenance costs, and minimize the number of fault points in the weapons system."

# MEMS

## ROD SCHMIDT

Rod Schmidt received his Masters degree from the University of Utah (1984) and PhD from the University of Minnesota (1988), both in mechanical engineering. He has worked at Sandia National Laboratories since the fall of 1987, where his research interests in computational modeling and simulation have been applied in areas ranging from hypothetical severe nuclear reactor accidents to turbulent reacting flow to microsystems.

### Further Reading

C. Jorgensen, D. Melander, R. Schmidt, S. Plimpton, *Development and performance of a PVM based parallel geometric modeler for MEMS*, Proceedings of the MSM2002, 5th International Conference on Modeling and Simulation of Microsystems, San Juan, Puerto Rico, April 22-25, 2002.

Additional information on MEMS work at Sandia can be found at [www.mems.sandia.gov](http://www.mems.sandia.gov).

### Contact:

Rod Schmidt  
[rschmi@sandia.gov](mailto:rschmi@sandia.gov)

To create more effective and more complex MEMS, designers must incorporate an increasing number of process steps and material layers into the fabrication processes. For example, in surface micromachining, the current standard at Sandia is a process known as SUMMIT V (for Sandia Ultra-Planer Multilevel MEMS Technology) which, after several hundred different steps, yields 5 distinct polysilicon levels. Designing a part for manufacture with this process requires more than defining the final geometry directly. Rather, a designer must create a set of up to 14 two-dimensional lithographic masks that, combined with the specifics of the process, create the final MEMS product. Not surprisingly, this indirect method of designing parts presents a challenge that becomes more difficult, time consuming, and subject to error as the complexities of the parts and processes increase.

These difficulties make the use of accurate, robust, and fast design tools essential. Of particular importance is the ability to automatically generate an accurate 3-D solid-body representation of the MEMS device that will result from a particular set of 2-D masks. Designers can then visually evaluate

and numerically analyze the solid model for correctness. Working with a model in this way improves designers' ability to recognize problems and achieve early success with production.

### Developing Collaboration

The MEMS technology team doesn't work alone on its projects. "As a group doing applied research we bridge the gap between the ground-breaking work of computer scientists and applied mathematicians working on massively parallel computer systems and the applied engineers and scientists who have real-life problems that need solutions today. Right now we're working closely with the people who are designers in the surface micromachining area. One of our goals is to further develop collaboration with fabrication specialists," Schmidt explains. "One of our projects involves working closely with surface chemists to understand better the detailed physics associated with the etching and deposition processes. We have to understand their problems and they have to understand what we can do so that together we can come up with the best solutions. We have easy access to each other."

The group also has access to the powerful computing facilities of the national laboratory, as well as a host of world class specialists in computational mathematics, algorithms, and parallel computing. "Sandia brings the ability to leverage massively parallel computing to solve challenges of this type," Schmidt says. "Some of the things we can do are unique. The national labs have particular strengths, including large numbers of talented people who can collaborate in an environment in which different people can provide diverse resources in an interdisciplinary situation."

Challenging technical problems of national import, state of the art computational resources, and a multi-talented multidisciplinary staff provide an atmosphere that Sandia hopes will continue to attract the best and the brightest scientists and engineers. To boost the attraction, the Department of Energy offers computational science graduate fellowships to qualified individuals.

The MEMS program has already signed up its first DOE CSGF fellow. Matthew J. McNenly from the University of Michigan is a student who will spend the summer of 2002 working with staff member Michael McGallis in the Non-continuum Transport and Microscale Science Department.

### Exciting Opportunity

How did Schmidt get into this line of work? "My manager invited me to lead the project when it was initiated," he recalls, noting that he hadn't been an expert in MEMS up to the time of his selection. He had previously carried out applied research in two main areas: modeling turbulent reacting flow and developing computer codes to simulate the physics of severe accidents in nuclear reactors. The fact that he is now working on computer simulations for MEMS, he says, "reflects the Sandia environment. You can work in an area for five or six years and then move to something different and exciting when the opportunity arises."

His team doesn't work in isolation at the Lab. Members frequently collaborate with scientists in industry. That benefits everyone involved. "To show that components going into weapons are safe and reliable, we need broad experience of how they perform in a range of environments," Schmidt explains. "Incorporating the components in commercial items unrelated to nuclear weapons is thus important to obtaining that experience. This helps us to understand that the fundamental technology is working reliably."

The computational tools and techniques that Schmidt's group is developing have broad potential application and will help to stimulate applications of the technology in civilian areas. For example, one promising field is bioMEMS, a technology that will lead to products such as biosensors, micropumps, and microlabs on chips with direct applications in medicine. As Schmidt explains, "All of us at the labs recognize and appreciate that the technology being developed in support of national security concerns often has spin-offs into a wide range of non-defense related areas that can have a direct impact on improving our quality of life."

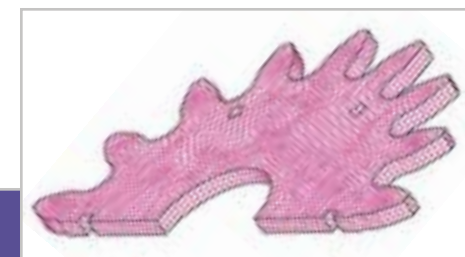
Microsystem geometries are complex and need to be understood in detail for design and qualification.



a. (above) Actual device cross-section.



b. (below) Idealized model generated from design masks.



c. Computational mesh of the gear teeth portion.

## FIVE SEPARATE PROGRAMS

>> Schmidt's group is currently working on five distinct tasks that focus on different computational problems. They are:

### > A 3-D parallel process geometry modeler.

This is a computer code written for multiple central processing units that automatically converts design layouts of two-dimensional masks into solid body representations of the three-dimensional components that the masks will produce. The goal is to provide fast, accurate, and robust simulations that will permit designers to recognize and correct potential design problems.

### > A "part-to-art" MEMS design tool.

This project aims to revolutionize the design of surface micromachining parts by allowing the designer to create 3-D designs directly, without having to define 2-D masks. Schmidt's group is developing algorithms and computer software that will automatically generate the 2-D mask sets from a 3-D solid body representation created by the designer.

### > A level-set based feature scale modeling code.

Many questions remain about fundamental aspects of the surface chemistry at the extremely small scale of MEMS devices during etching and deposition. This work, called ChISELS, for Chemically Induced Surface Evolution with Level-Sets, tracks the evolving surface topology while simultaneously modeling the complex surface chemistry that drives the process.

### > MEMS specific automatic meshing development.

To analyze and test the performance and behavior of a new MEMS design, a computational mesh that represents the device's complex geometry must be generated. However, accurate results depend strongly upon the quality of the mesh used in the simulation. This work is developing new approaches for automatically generating high quality meshes.

### > Molecular dynamics for LIGA.

This effort seeks to explore the viability of using what is called "molecular dynamics simulations" to understand better the physical processes that control the dissolution of polymers.

By Jacob Berkowitz

# Future Climate

**“IT’S ALWAYS BEEN THE CASE** that the people with the largest computer have been able to produce the best weather forecasts,” says John Drake, speaking on the phone from his office at the Department of Energy’s Oak Ridge National Laboratory in Tennessee.

He’s been listening intently to weather forecasts for the past several weeks. The southeastern seaboard has been inundated with heavy late-winter rains. Yet, beyond the immediate question of whether he should be wearing hip waders to work, weather forecasts are small computational potatoes to this computer scientist and climate modeler. While a long-range weather forecast is seven days, Drake’s pioneering parallel supercomputer climate models are currently pushing out forecasts for the next *two centuries*.

These parallel climate models developed by Drake, in close collaboration with climatologists from the National Center for Atmospheric Research (NCAR), in Boulder, Colorado, and other national labs, are providing an unprecedented peek at the fine details of future climates. This supercomputer-generated information isn’t just about average global temperatures, but increasingly about regional details, such as the mean temperature in central Iowa in 2075. It’s crucial climate forecasting that is helping inform U.S. and international policy on energy use as it relates to global climate change.

**The Parallel Supercomputers that Could**

In the late 1980s, Drake’s background in applied mathematics with a focus on computationally demanding fluid dynamics problems made him the perfect candidate to work in the quickly developing field of climate modeling.

“Climate is the granddaddy of fluids problems. The general circulation of the atmosphere and the oceans are fluids problems,” notes Drake, now the Computational Climate Dynamics Group Leader at Oak Ridge.

However, it is a problem that severely outpaced the computing power of the day.

“The climate models of ten years ago were simplistic by the standards of today. That’s not because people didn’t know what to put in them. They just couldn’t afford to do the computation. So they made simpler models,” recalls Drake. “In the last three to four years we’ve gone from having the power only to simulate

the atmosphere in a climate setting to being able to create coupled general circulation models, ones that combine the atmosphere with oceans, sea ice and land surfaces. That’s a big change in complexity and our understanding of what constitutes climate.”

Back in the dawning days of the Internet, the challenge of producing better climate models pushed Drake and his colleagues to explore a new computational approach: massively parallel computers.

“This was in about 1989, just when early Intel chips were getting cheap and it became obvious that you could string a lot of these together and compete with a supercomputer,” Drake says.

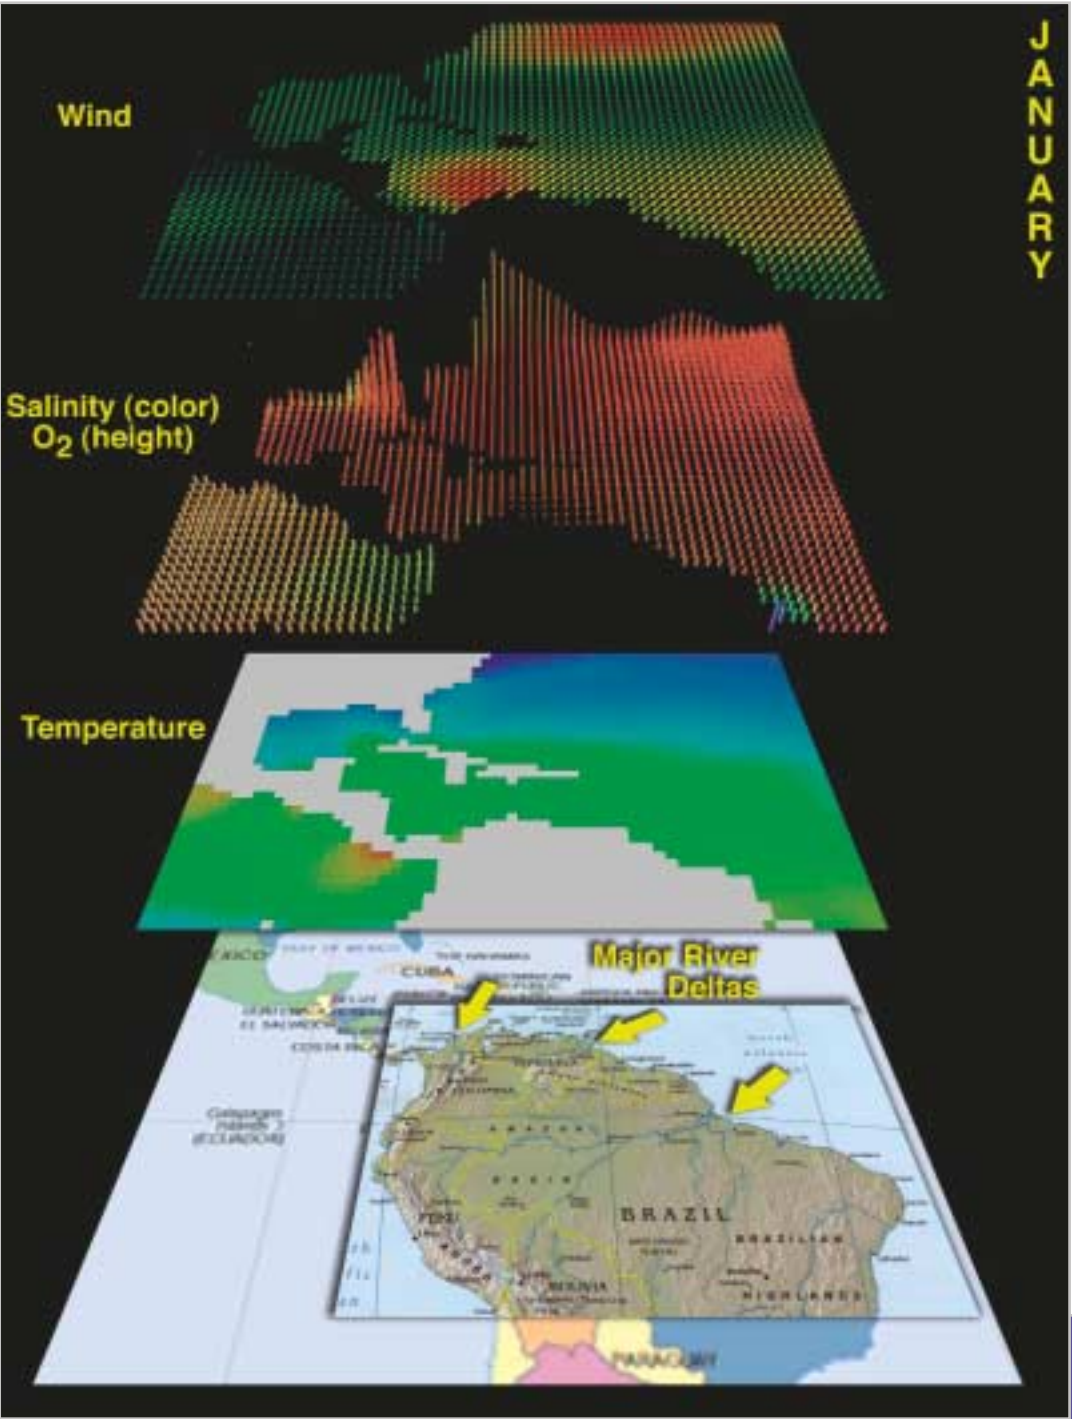
At the time, notes Drake, massively parallel computers were largely a computer-science novelty to climatologists. Most were accustomed to running sophisticated vector-based models on CRAYs (though some of these models, including a pioneering one at NCAR, were parallel vector applications). But to Drake —

working at the borderlands of mathematics, supercomputing and climatology — distributed-memory massively parallel systems offered an intriguing advantage: their structure mimicked the highly interactive nature of climate.

“The realization with massively parallel computing was that if you had lots of computer chips you could break a calculation up into several pieces. You could divide and conquer. So instead of one long vector, you had a bunch of short vectors, and the memory is spread out,” recalls Drake. His Oak Ridge lab had one of the first of what he calls “toy” massively parallel computers.

However, the distributed-memory architecture of parallel computers also presented a major new problem. How would climate models deal effectively with message passing, the sharing of information between processors? After all, when processors are sharing information they’re not performing calculations.

In the early 1990s, the DOE created the CHAMMP (Computer Hardware, Advanced Mathematics, and Model Physics) program to directly address this problem. The program was a collaboration between staff computational scientists at DOE’s Oak Ridge and Argonne National Laboratories, and climatologists from the NCAR.



This example, based on University of Wisconsin–Milwaukee Comprehensive Ocean–Atmosphere Data Set (UWM-COADS) observational data, uses visualization to illustrate regional climate dynamics in the Caribbean. Winter and summer data from 1980 are visualized and compared to show how measured physical properties change seasonally. Such dynamics and their simulation are widely viewed as critical to the improvement of computational climate models applied to long time scale studies (e.g. 200+ years). As part of a national climate modeling initiative, computer simulations of these and other features will be incorporated into global climate models to study their impact.

Its goal was straight forward — was it possible to adapt a climate model and all its code to a massively parallel computer? The answer, as all climatologists now know, is a resounding yes. Parallel climate models are the gold standard for climate simulation.

But the question resulted in a still-simmering computational controversy over how to maximize the efficiency and accuracy of these parallel models.

There are two main contending approaches: grid point approximations and transform methods. While transform methods are the traditional approach to climate models, the grid point method is more computationally efficient, since it relies on more local mathematics. Drake, an advocate of the transform method, counters that speed isn't everything. The two methods have subtle differences in the results they produce.

“There was a ‘bake-off’ last year; we did evaluations of three different algorithms, comparing a grid point model with two versions of transform models, all of which had had significant investment to make parallel algorithms. The chemists like one and the meteorologists like another,” Drake says.

For example, tropical regions are better represented with the transform model, while the grid point model more faithfully conserves atmospheric chemicals, making it chemists’ preferred choice. One solution that Drake and others are exploring is to combine the best elements of both approaches with focus grid transform models.

The “More” Challenge

The lock-step march of advances in computational and climate science is now presenting a whole new palette of problems for climate modelers. Call it the “more” factor. Climate modelers like Drake are faced with *more* data, *more* complex climate models and ever *more* intricate parallel computing hardware. It's a phenomenon that's fostering a greater-than-ever level of interdisciplinary collaboration in developing climate simulations.

Working in partnership with Los Alamos Laboratory's Bob Malone, Drake is co-leading a new collaborative model development project as part of the DOE's Scientific Discovery through Advanced Computing (SciDAC) initiative. It's the largest and most diverse DOE climate modeling project to date, bringing together six national laboratories and 20 staff scientists.

One of the team's first tasks is to overcome the programming hurdle presented by the latest generation of parallel supercomputers, the ten-million dollar IBM p690. Making message passing seem like a relatively easy problem, its Power4 chip adds a deep memory hierarchy to the mix.

“Not only do you have to deal with the fact that memory is now distributed across several processors, you have to deal with the fact that within a processor there are different levels of memory,” says Drake. “We’re developing a programming paradigm and rehashing the algorithms to match this kind of computer so that we’ll be able to take advantage of

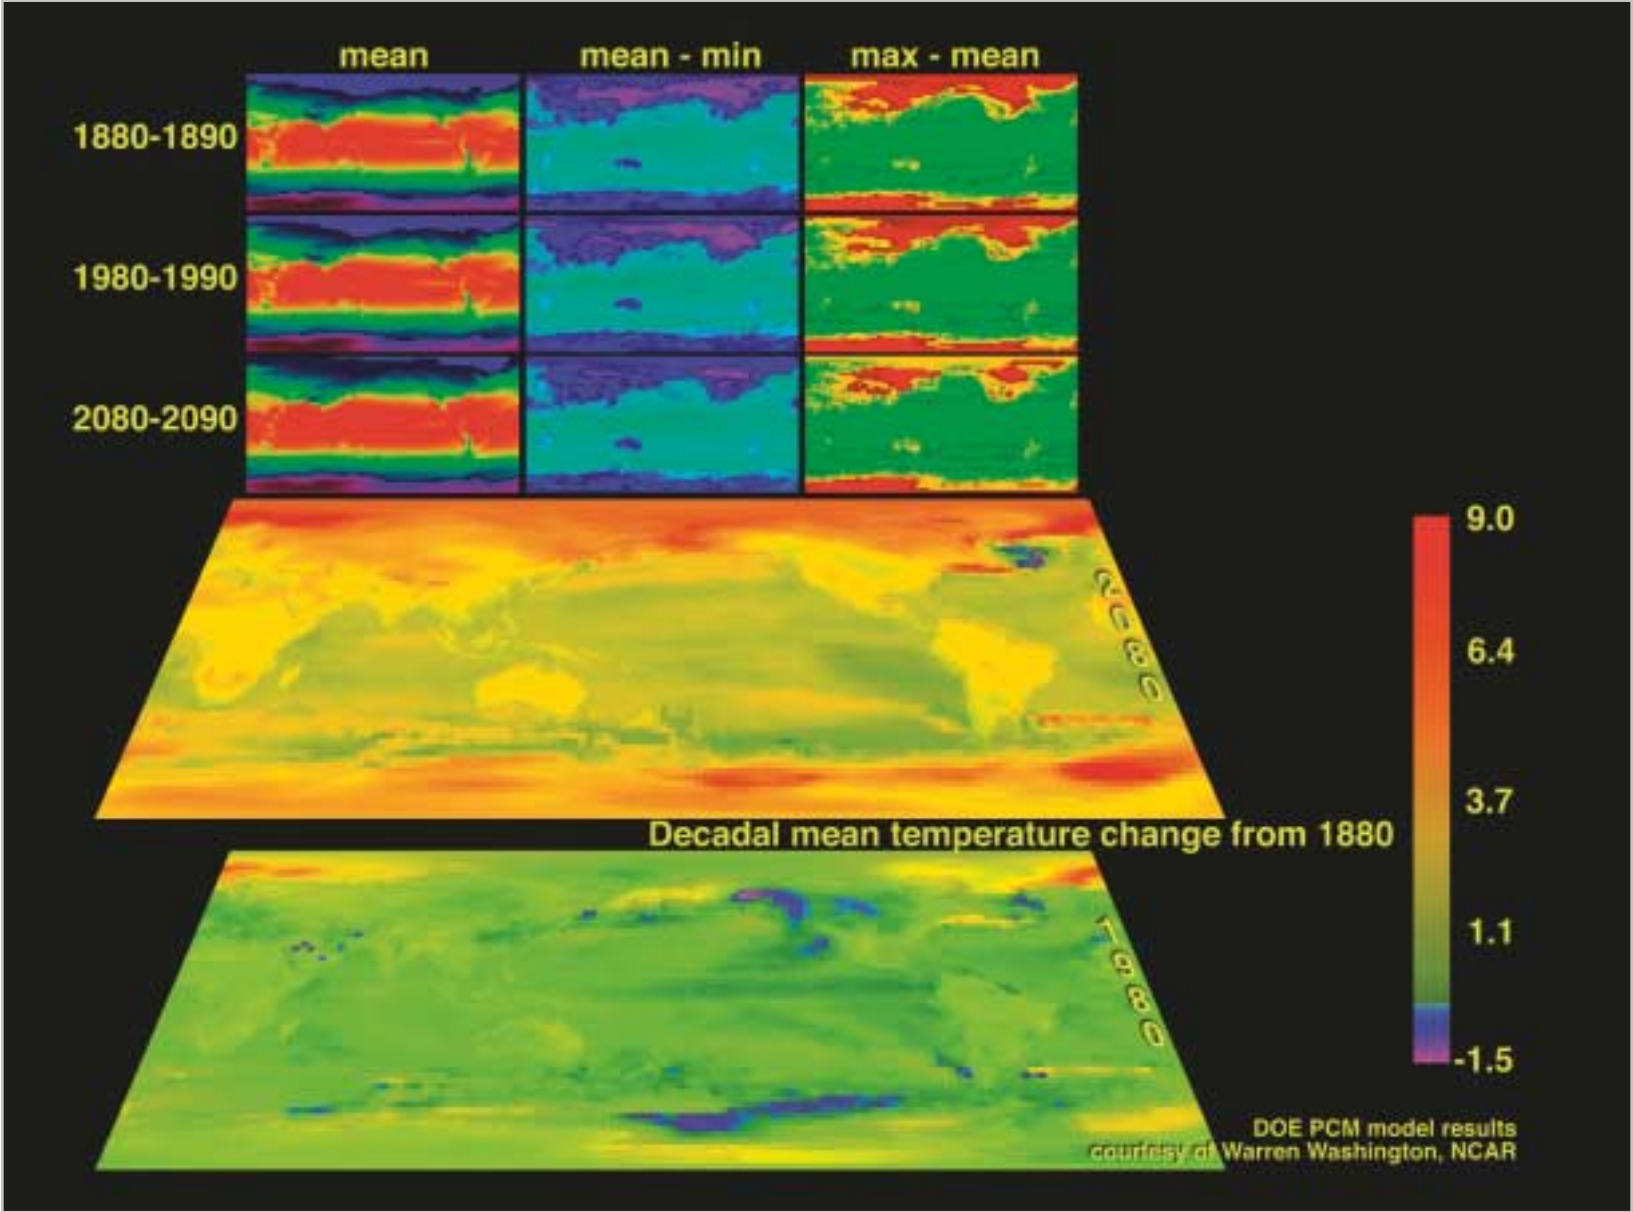
its speeds. The issue here really is memory speeds. The processing speed is going up like Moore's Law, but the effective distance to memory is not keeping pace.”

The combination of more complex hardware architecture and more complex models is also spurring Drake and Malone's team to grapple with software development issues. These are particularly relevant for recent global coupled models. These models integrate the climatic roles of atmosphere, oceans, sea ice, land surfaces and, increasingly, other factors such as volcanic activity and solar variability.

“One of the things that we’ve tried to do recently is bring modern software practices into this situation,” notes Drake. “It used to be that the problems were small enough, and the codes were simple enough, that individual scientists could in fact know everything that's in the code. But now with this complexity that's no longer the case. You really have to work in teams.”

Addressing the triumvirate of “more” problems holds the promise of opening the way for the next generation of climate simulations, ones that will provide ever greater accuracy and resolution.

The accuracy of global climate models is a major scientific — not to mention political — issue. Using a parallel climate model that incorporates many of Drake's algorithms, NCAR's Warren Washington recently created a simulation of historical data (1870 to present) with a 95 percent confidence interval. Drake is now working with him to shepherd through Oak Ridge's



and NCAR's massive parallel computers what might well be the most definitive U.S. models yet for the impacts of global warming.

Unlike past simulations, Washington's current climate forecasts are based on ensembles of multiple runs of a model. This is the way that weather forecasts are constructed. However, only in the past several years has there been the computing capacity to perform ensemble models for multi-century global circulation climate models.

While it's our grandchildren and their descendants who will ultimately assess the accuracy of these multi-century models, 15-years later Drake is still

working to improve the quality of his models for the ultimate in fluid dynamics problems.

Says Drake, “When people start to ask questions about the implications of climate change, one of the things they quickly get to is what's going to happen to the rivers, vegetation, and agriculture on a local level. Not just east coast or west coast. You'd like to be able to have 30-kilometer resolution over the U.S. Currently we cannot provide that kind of information. But I'm hoping that new algorithms and increased computing power will provide a way to achieve this increased resolution and accuracy.”

Decadal climate variations are important indicators of long-term climatological factors. Comparisons or decadal averages clearly show polar warming trends, especially in the north.

JOHN DRAKE

John Drake received his PhD in mathematics from the University of Tennessee in 1991. Since 1999 he has been the Group Leader of the Climate Dynamics Group in the Computational Sciences Section at Oak Ridge National Laboratory. His current research interests include numerical methods and parallel algorithms for climate dynamics, numerical solution of partial differential equations and integral equations, with particular regard to supercomputing applications.

**Further Reading**  
*A Vorticity-Divergence Global Semi-Lagrangian Spectral Model for the Shallow Water Equations*, J.B. Drake, D.X. Guo, ORNL/TM-2001X216, November 2001.

*Statistical Downscaling of the United States Regional Climate from Transient GCM Scenarios*, William M. Putman, J.B. Drake and G. Ostrouchov, in the Proceedings of the AMS 12th Conference on Applied Climatology, May 2000.

**Contact:**  
John Drake  
drakejb@ornl.gov

The lock-step march of advances in computational and climate science is now presenting a whole new palette of problems for climate modelers. Call it the “more” factor.

By Victor D. Chase

# Housing Lasers



**IT IS AS BIG** as a professional sports stadium and will eventually house 192 powerful lasers, all of which will be focused on a BB-sized capsule sitting inside a small tin can-like structure. When fired off in unison, the lasers’ powerful light beams will crunch the tiny capsule, causing the atoms inside to fuse together. The result will be a colossal burst of energy from the same nuclear fusion process that takes place in the Sun and the other stars in the universe.

In this instance, however, the venue will not be the fiery interior of a star, but the more tranquil surface of planet Earth. The specific location is the National Ignition Facility (NIF), now under construction at Lawrence Livermore National Laboratory (LLNL), Livermore, California.

The earthbound fusion process that will take place in NIF, which is scheduled for completion in 2008, has been successfully created in smaller, predecessor laser facilities at LLNL, known as SHIVA and NOVA, but more energy had to be used to create those fusion reactions than was created by them. The intent is to reverse that equation in NIF, so that the process produces more energy than it consumes. If the process is successful, nuclear fusion could provide an endless source of energy, since the fuel contained in the target capsule consists of the hydrogen isotopes deuterium and tritium, and hydrogen is the most abundant element in the universe.

Another reason for the construction of NIF falls under the heading of “stockpile stewardship,” which refers to maintaining the nation’s nuclear weapons capability. Although the Comprehensive Nuclear Test Ban Treaty has not been ratified by Congress, the U.S. has complied with the spirit of the treaty and has not tested nuclear weapons since 1992. NIF will be used to replicate the intense pressures and temperatures that occur during a nuclear explosion without creating an actual explosion.

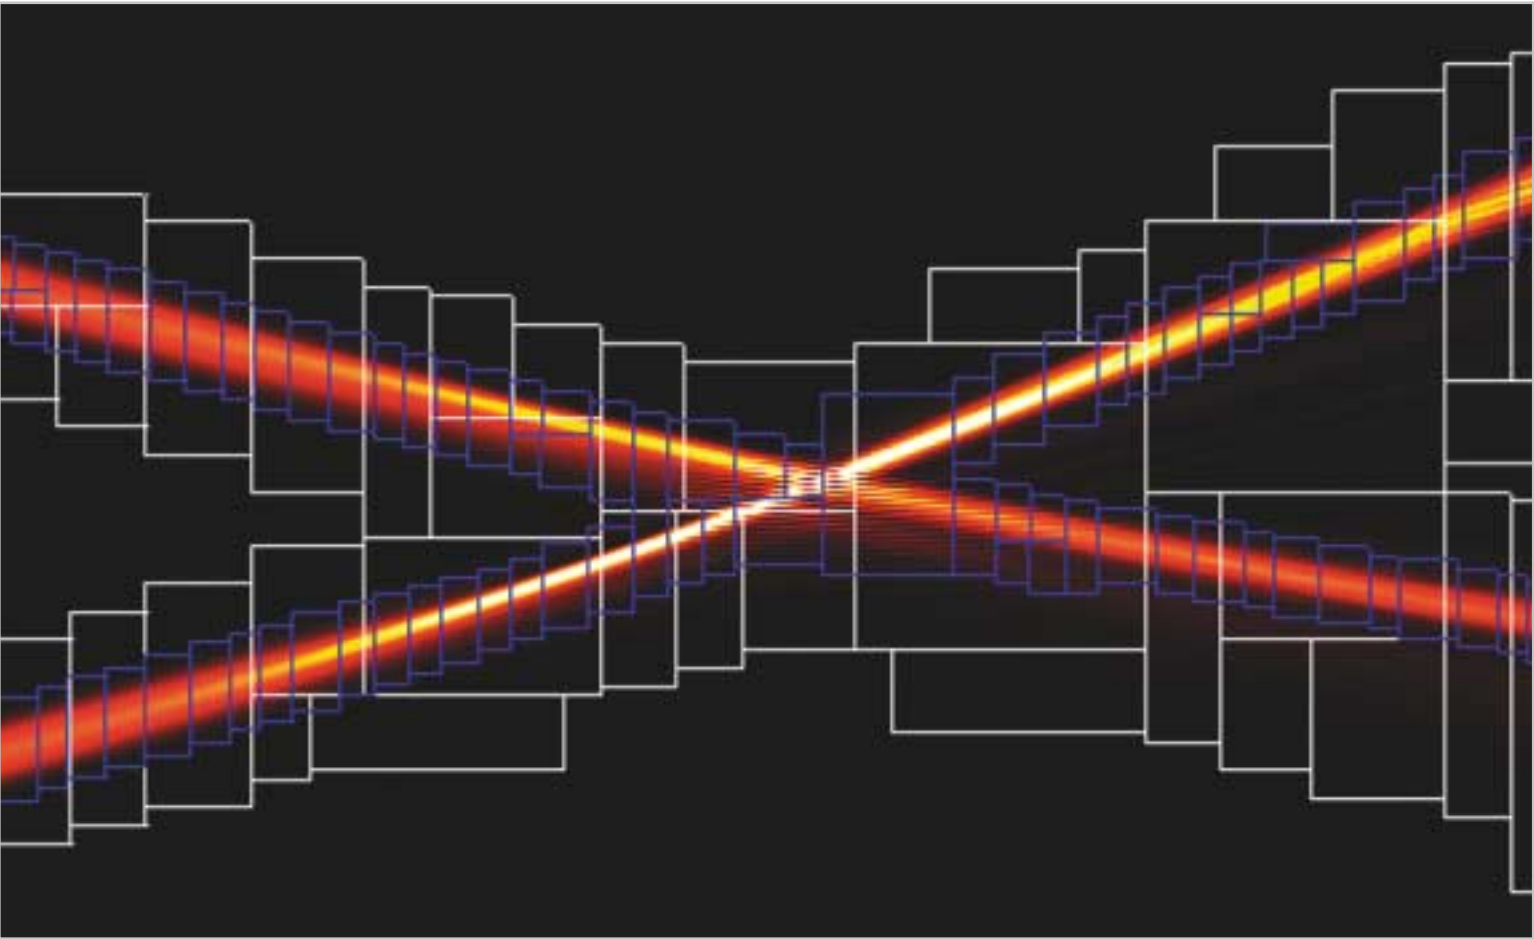
“There’s a large effort to try to compensate for not being able to test weapons anymore by using computational methods, doing numerical simulations, as well as getting experimental information from facilities like NIF,” explains Milo Dorr, a LLNL mathematician and staff scientist.

## Of X-Rays and Plasma

Of course, recreating the energy-producing process that takes place in the stars is an extremely complex process that requires, among other things, a tremendous amount of computational simulation work long before laser ever strikes target.

Complicating matters is the fact that the actual nuclear fusion process that will take place inside NIF isn’t as direct as it may seem, in that it is not the laser beams themselves that create pressure on the fuel capsule. Rather, pressure is created by the X-rays that are generated when the laser beams hit the inside of the can containing the capsule, called a hohlraum, as well as the capsule.

In addition to generating X-rays, the high temperatures created by the laser beams ionize the fuel and the



hohlraum materials. The result is plasma, which, instead of being composed of atoms, is a mix of positively charged ions and negatively charged electrons. Such matter is found in the atmospheres of stars, and in thermonuclear reactions. Unfortunately, the plasma, in turn, can have a deleterious effect as the laser light passes through it.

“The plasma could bend a beam in a different direction. It could even reflect some light backward out of the hohlraum, which is undesirable because that sends the energy back towards very expensive optics,” explains Dorr. Or, energy can cross from one laser beam into another, which can cause an undesirable pressure pattern on the fuel capsule. Hence, ensuring

that the energy in the laser beams actually ends up striking its intended target requires an extensive understanding of what happens when light attempts to pass through the plasma.

## Simulating Plasma and Light

Developing such an understanding by actually creating a plasma and firing laser beams into it is not always a practical solution, since the facilities that can produce these conditions are few and the expense of doing so is great. So scientists have turned to the next best thing, namely, simulating plasma and laser light interactions by use of computer models to augment the empirical data from experiments.

That’s where Dorr and his group of mathematicians, physicists, and other scientists come in.

Among them is Jeff Hittinger, a postdoctoral fellow who received a PhD in aerospace engineering and scientific computing from the University of Michigan with the aid of a DOE Computational Science Graduate Fellowship (DOE CSGF). Another DOE CSGF Fellow, Edward McKay Hyde, worked on the computer simulation project at LLNL to fulfill the practicum requirement of the DOE CSGF.

*An ALPS adaptive mesh calculation of laser intensity for crossed beams in a plasma flow. The white boxes indicate the first refined level and the blue boxes indicate the second and finest refinement level.*

**NIF will be used to replicate the intense pressures and temperatures that occur during a nuclear explosion without creating an actual explosion.**

## MILO DORR

Milo R. Dorr is a member of the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory (LLNL). He holds a BS degree in mathematics from Grove City College, and MA and PhD degrees in mathematics from the University of Maryland. Since joining LLNL in 1983, Milo has developed numerical methodologies for neutron and charged particle transport, with emphasis on scalable algorithms for implementation on massively parallel computers. He has also developed adaptive mesh refinement algorithms and codes for solving systems of equations describing plasma processing of materials and laser plasma interactions. Milo is a member of the Society for Industrial and Applied Mathematics and has authored or co-authored approximately 30 scientific publications.

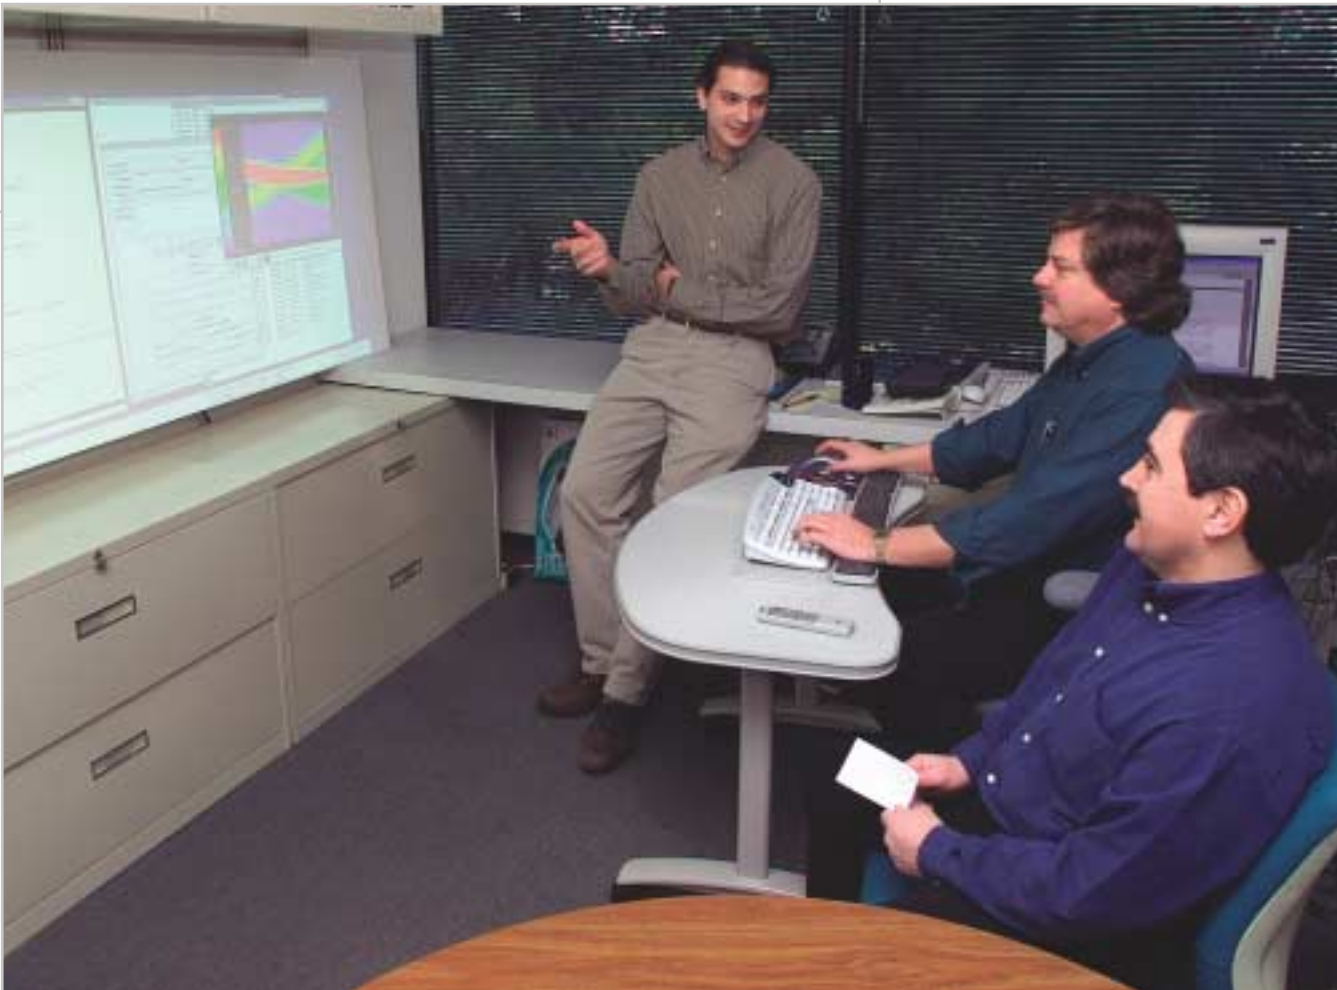
### Further Reading

*Simulation of Laser Plasma Filamentation Using Adaptive Mesh Refinement*, M.R. Dorr, F.X. Garaizar and J.A.F. Hittinger, J. Comput. Phys., 177 (2002), pp. 233-263.

**General information on laser-driven fusion and the National Ignition Facility can be found at:** [www.llnl.gov/nif](http://www.llnl.gov/nif).

### Contact:

Milo Dorr  
Milodorr@llnl.gov



*Jeff Hittinger, Milo Dorr and F. Xabier Garaizar (left to right) discuss ALPS.*

During his practicum at DOE's facility in Los Alamos, New Mexico, Hittinger worked on computer modeling of radiation hydrodynamics, which deals with the effects of light propagation through a fluid, such as air or a plasma, under conditions of intense radiation. Although different from the NIF project he is now engaged in, the subject is closely related. And, in fact, he obtained his postdoctoral assignment at LLNL through contacts he made while at Los Alamos.

"My group leader at Los Alamos was actively involved in the DOE CSGF program," he said. "He and members of his research group moved from Los Alamos to Livermore and put me in contact with a member of the ALPS team," said Hittinger, using the acronym for Adaptive Laser Plasma Simulation, the NIF project in which he and Dorr are currently involved.

## Conserving Computer Power

The purpose of ALPS is to improve the efficiency of an existing computer model that simulates the complex laser/plasma interaction in the hohlraum more efficiently. Dorr hopes to accomplish this through ALPS by focusing computational power precisely where it is needed, when it is needed, rather than have the calculations look at the entire picture all of the time.

Currently, modeling work is parceled out among a cluster of parallel computers that run simultaneously to solve a problem. Combined, these machines are equivalent to 512 individual workstations, and can run 680 billion multiplications per second. Such computer power doesn't come cheap, and whatever can be done to make its use more efficient is valuable.

**"The basic idea is to use the computer resolution only where you need it and, therefore, save on the total cost of the calculation."**

terms of how much memory I have to use in the computer and how long I have to run the calculation," says Dorr. Besides, it turns out that it is not necessary to look at all the cells all of the time.

"In certain parts of the plasma there may be a lot of things going on, so I need to use a lot of grid cells in that particular area, whereas elsewhere there may not be any laser light at all, so I don't need as much grid resolution," explains Dorr. Hence, the purpose of this project is to turn that uniform grid of cells into an adaptive grid — as in the Adaptive in ALPS — that can adapt in concert with the needs of the modeling process.

"Suppose I'm only interested in something that's going on in the middle of this rectangle. Then I don't need to have such high resolution in the outer part of the rectangle," says Dorr. "What I'd like to do instead of using one billion mesh cells everywhere, is concentrate my mesh cells in the middle of the problem where there are fine detailed features that need to be resolved, and in the outer portion of the region, where I don't care as much about what's going on, I could have a much coarser resolution."

The result could be the need to look at a million rather than a billion mesh cells. "The basic idea is to use the computer resolution only where you need it and, therefore, save on the total cost of the calculation," concludes the LLNL mathematician. With the ALPS program, all of this would be accomplished automatically,

with the computer determining, on its own, when and where to go into a coarser mode of operation. "Part of the research is to try to understand how to do that automatically," says Dorr.

## A Two-Way Street

That computer modeling can be used to simulate something as complex as laser/plasma interaction is predicated on an understanding of the physical rules that govern what happens in such environments. "It's a matter of understanding what those rules are, and being able to represent them as a set of equations," says Dorr. "There is a sequence of steps that starts off with the modeling of basic plasma physics, in other words, writing down a set of equations describing the physics. The next step is designing algorithms to solve those equations and implementing them on large computers."

To verify the calculations, experiments are sometimes run at the University of Rochester's Omega laser facility, a much smaller sibling of the in-the-making NIF. But, since there are more potential experiments to be run in Omega than time and money allow, ALPS and Omega are entering into a two-way relationship in which Omega tests ALPS calculations, and ALPS calculations help determine which experiments Omega should run.

"What we're hoping to do is use our computational model to try to winnow out some cases, so we can tell the experimental folks, 'We've seen this behavior, these types of problems, so you may want to design an experiment in this regime. On the other hand, we ran all these other problems and nothing really showed up so you may not want to waste any time probing in that part of the parameters'," says Dorr.

Although ALPS has been an ongoing program for about three years, according to Dorr, "this is an open-ended research area. As people start doing more experiments on NIF itself, I think there will be an even greater demand for computational simulation support, not only to design the experiments, but also to understand what happened when they did them. I think this is a growth area."

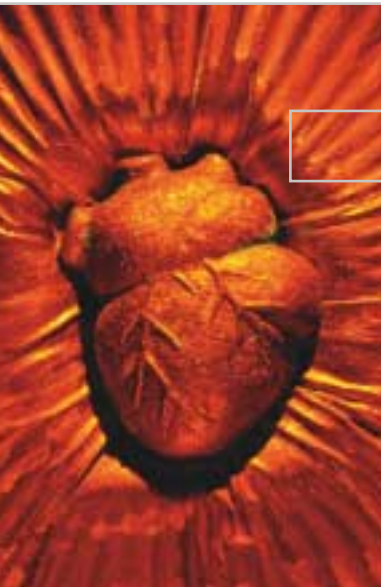
## NATIONAL IGNITION FACILITY (NIF) FAST FACTS

- >> NIF is a **192-beam 1.8 MJ laser** for creating conditions of extreme temperatures and pressures in the laboratory.
- >> NIF has **missions** in national security, energy, and science research.
- >> NIF is a nationwide **collaboration** between government, national laboratories, and industry.
- >> NIF technology requirements enable cutting-edge **laser and optics** developments.

Source: [www.llnl.gov/nif](http://www.llnl.gov/nif)

By Jacob Berkowitz

# Working a Beat



**IN THE FALL OF 2001** an article was published in *Cardiac Electrophysiology Review* that summarized the current work on computer modeling of atrial fibrillation, the debilitating quivering of the heart’s upper chambers. All things being equal, there was probably about as much chance that Wonho Oh would ever see the article as that he would get hit by lightning.

Except for one thing — the article’s concluding statement that “the final limitation of simulations is the inadequacy of available computer resources to handle all the details.”

Not long after the article appeared, the Brookhaven National Laboratory mathematician and supercomputing scientist received a call from Flavio Fenton, the Director of Electrophysiology Research of the Heart Institute at the Beth Israel Medical Center in New York City.

That call began Oh’s ongoing journey to help develop better simulations of the heart’s electrical activity, especially arrhythmia, or irregular beating, using the power of parallel supercomputing. It’s a pursuit that puts Oh (who has no training in biology) at the forefront of research into an ailment that accounts for one-third of all cardiac deaths in the United States.

Oh was introduced to heart electrophysiology because Flavio Fenton’s rabbit heart ventricle simulations were slow enough to give an eager researcher a heart attack. Working on a single processor workstation, it took about three hours to model a single second of electrical activity in a rabbit’s ventricles — the heart’s two lower chambers and pumping workhorses. At this rate it would be impossible to achieve Fenton’s goal: modeling fibrillation and arrhythmia in the human atria and ventricles. Human hearts are at least four times larger in diameter than a rabbit’s, but the increase in computing power required is even greater, reflecting the cubic rise in volume of the larger heart.

### Model Complexity

To study fibrillation, Fenton needed to significantly increase the complexity of his models — another factor that would require more computing power. He needed to create bi-domain models that accurately reflect the fact that the heart’s ion-mediated electrical current travels into and out of cells.

Earlier heart electrophysiology models treated this bioelectrical system as a mono-domain one, with electrical activity limited to cells’ interiors.

“When you go from mono-domain to bi-domain you need to solve not only a partial differential equation but also a Poisson equation, and this involves more computing power. The bottleneck is the Poisson equation. It couples the partial differential equations for the intra- and extra-cellular environments. Unless you have a fast algorithm, the same method for the mono-domain can be at least ten times slower with the Poisson,” says Fenton, who is also a visiting research scientist in the physics department at Hofstra University in Hempstead, New York. Interestingly, it was Oh’s experience working with geologists studying flow through porous media that prepared him to help cardiologists.

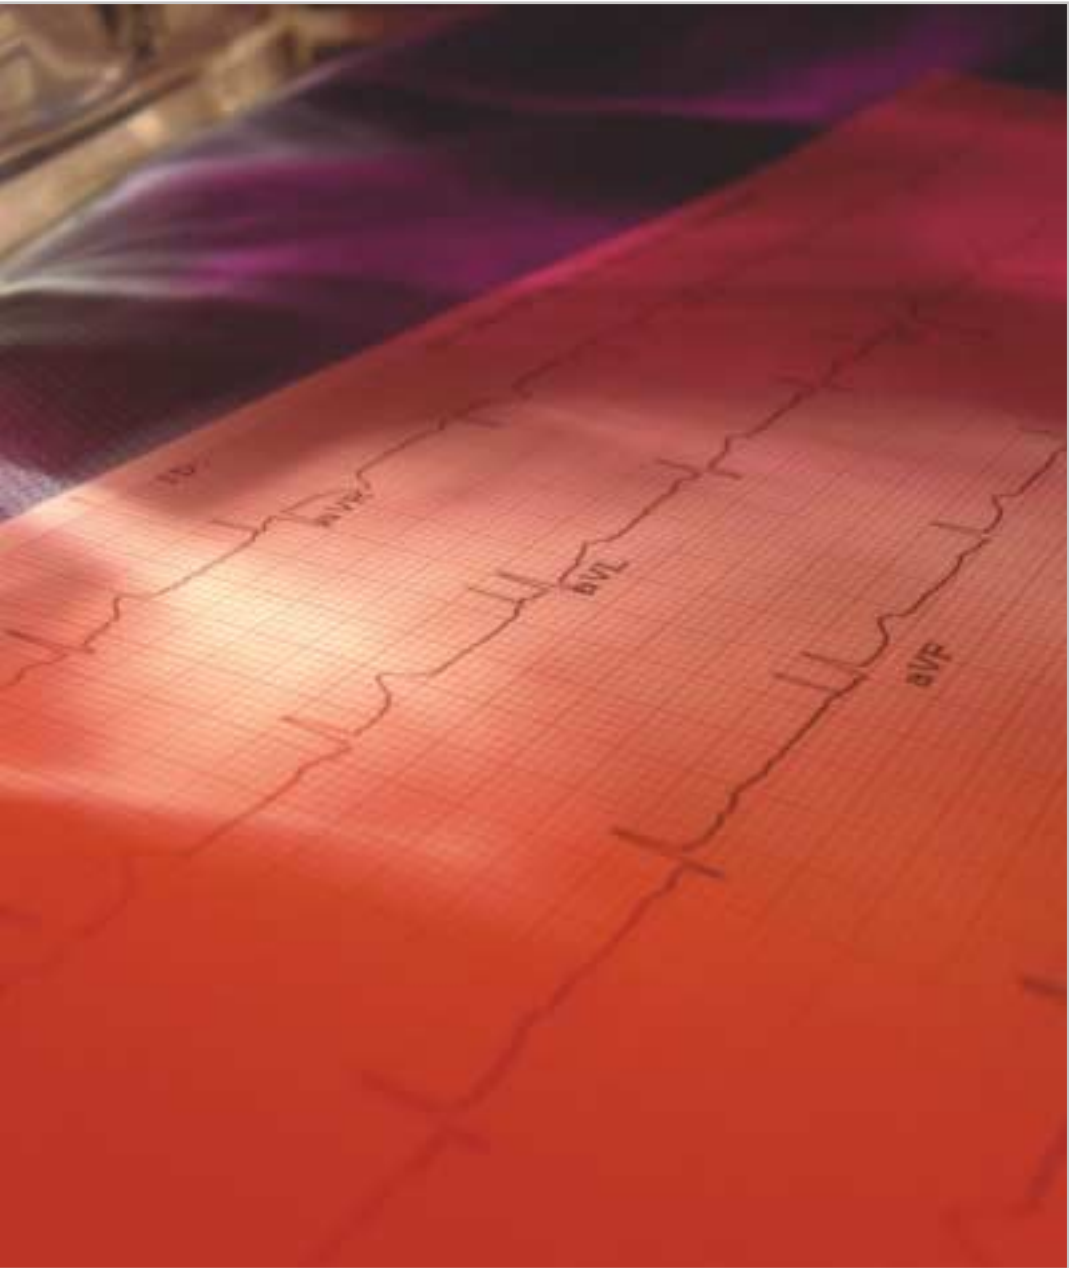
“I was solving oil reservoir equations,” says Oh, who’s also a professor at State University of New York in Stony Brook. “These involved an elliptic partial differential equation that looks similar to the bi-domain equation. The major difference is that the heart problem includes a time derivative, it changes with time.”

### Electrical Pump

While most of us know the heart as a pump, fewer of us think of the large, essential muscle in our chests as an *electrical* pump. Yet, it is electrical activity that regulates the heart’s normal, steady rhythm of between 60 and 90 contractions a minute. In a healthy heart, the cardiac rhythm originates in a small area located in the upper part of the right atrium, called the sinoatrial node. The impulses from the heart’s natural pacemaker travel as an electrical wave through the atria, and then — after a hundredth-of-a-second delay that allows blood to enter the ventricles — through the cells of the ventricles. This electrical stimulation causes the cells to contract forcefully, pumping blood through the body.

Arrhythmia occurs when, as a result of structural or dynamical problems, one of these electrical waves is locally blocked and breaks, turning the wave into a potentially deadly spiral of electrical activity.

“When these spiral waves are produced in the heart, they have a much faster frequency than your normal pacemaker,” says Fenton. “Since they rotate much faster, they take control and your heart starts pumping much faster — called tachycardia. Then the spirals tend to break into multiples. And when



“When you go from mono-domain to bi-domain you need to solve not only a partial differential equation but also a Poisson equation, and this involves more computing power.”

COLLABORATORS

**Wonho Oh** received his PhD in applied mathematics from the State University of New York at Stony Brook. He is currently an Assistant Professor in the Department of Applied Mathematics and Statistics at the State University of New York at Stony Brook, as well as a Research Associate in the Center for Data Intensive Computing at Brookhaven National Laboratory.

**Flavio H. Fenton** received a BS degree in physics from Universidad Nacional Autonoma de Mexico (UNAM) in 1990 and MS and PhD degrees in physics from Northeastern University in 1992 and 1999, respectively. Since 1999 he has worked as a visiting research scientist at Hofstra University and, since 2001, as Director of Electrophysiology Research at the Heart Institute, Beth Israel Medical Center, New York. His current research interests focus on models of cardiac cellular electrical activity and computer simulations of arrhythmias.

**Publications**  
*A Critical Analysis of Raleigh-Taylor Growth Rates*, W. Oh, J. Glimm, J. Grove, X.Li and D. Sharp. Journal of Computational Physics, 169 No. 2, pp. 652-667, 2001.

*Multiple mechanisms of spiral wave breakup in a model of cardiac electrical activity*, F.H. Fenton, E.M. Cherry, H.M. Hastings, and S.J. Evans, to appear in Chaos, Vol. 12, No. 3, 2002.

*Real-time computer simulations of excitable media: JAVA as a scientific language and as a wrapper for C and FORTRAN programs*, F.H. Fenton, E.M. Cherry, H.M. Hastings, and S.J. Evans, BioSystems 64, 73-96, 2002.

**Contact:**  
Wonho Oh      Flavio Fenton  
woh@bnl.gov    Flavio.H.Fenton@hofstra.edu

you have a few of these spirals all over your heart, each one at a very high frequency and out of phase from the others, then the whole heart is just quivering and not pumping anymore.”

About ten percent of all Americans over the age of 65 have some level of chronic atrial fibrillation, a condition that results in weakness and numerous long-term circulation problems. Ventricular fibrillation always results in death if not treated within minutes.

Heart Modeling

In the early 1980s, Beth Israel Medical Center cardiologist Dr. Steven Evans began collaborating with a mathematician to numerically model the heart’s electrophysiology. In 1998, he recruited Fenton, who was then working in high energy physics, to further pursue the computational modeling of this electrical activity.

The pursuit of heart electrophysiology modeling is driven by the fact that cardiologists are already well aware of the potential to treat arrhythmia. Patients with abnormal heart rhythms can undergo what has become a standard diagnostic procedure known as an electrophysiology study. A catheter is inserted into a patient’s femoral vein in the groin, and then guided up into the heart. The doctor then stimulates the heart with electrical signals to make it beat at various rates and to observe any irregularities. Based on this test, the cardiologist can use a variety of drugs, surgical techniques, or a small, automatic defibrillator worn by a patient, to prevent arrhythmia or to convert irregular rhythms back to normal.

The problem is that even given these medical techniques, cardiac arrhythmias are still a major killer, and their cause and dynamics — the breaking of the electrical wave and spiral wave behavior — still remain largely a mystery.

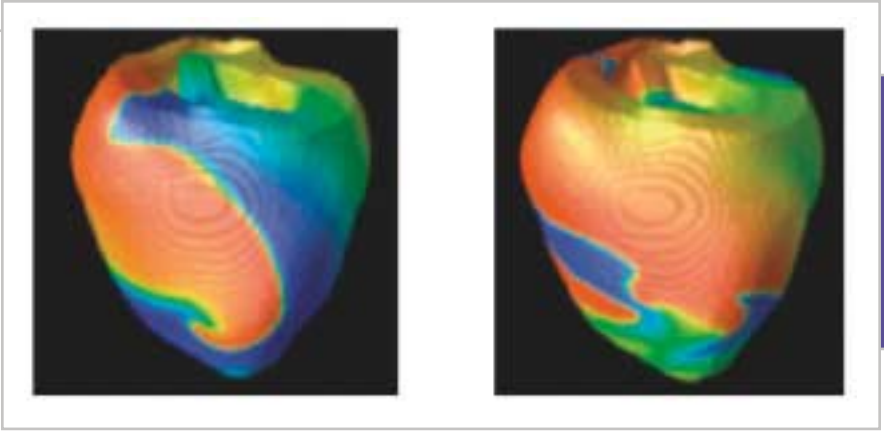
The Code

Working with experimental data, Brookhaven’s Oh is methodically developing the parallel super-computing code and algorithms that will enable Fenton, Evans and other researchers to minutely model atrial and ventricular activity with greater precision and speed than ever previously achieved.

Initially he’s developing a three-dimensional slab model, one that will act as a validation and testing ground for building the parallel code.

“When you stimulate a part of the slab tissue, you expect a certain behavior of the electrical potentials, as seen experimentally. At this point, I have a basic set of code and I have been testing it in a serial fashion in a single processor with a small tissue sample to see whether the resulting potentials look reasonable. Qualitatively it’s working, the wave is propagating, but I still need to confirm it quantitatively,” Oh says.

After this initial phase is completed, the electrophysiologists will get what they are waiting for — movement of the model to Brookhaven’s Galaxy cluster parallel supercomputer. The Galaxy cluster consists of 77 Pentium III dual processor nodes, each with one gigabyte of memory. The nodes currently communicate through a Message Passing Interface (MPI) library, though Oh says he plans to use OpenMP format as well as MPI in the near future. (The OpenMP



Single spiral (VT), multiple spirals (VF) on simulated rabbit ventricles.

The blue shows tissue that is electrically polarized, and red, depolarized. (i.e., tissue at about -90 mV for blue and about 15 mV for red.)

format allows multiple threads on a single node and eliminates unnecessary message passing between the processors on a single node.)

“The final goal is to make my code work for parallel machines so that I can get the result very quickly,” says Oh. “The real challenge is that you want to have a scalable algorithm, so that once you have numerous processors available, the computing speed grows at the same rate as the number of processors.”

Lifesaving Knowledge

It is this speed at the bi-domain, human ventricles or atria level that will significantly advance the theoretical understanding and, it is hoped, treatment of arrhythmia, says Fenton.

“The speed of parallel supercomputing allows for a parameter search, one that explores the whole range of possible conditions, such as ischemia (reduced blood flow) to the levels of sodium ions. You need to be able to do a lot of different simulations with many different parameters, and you can’t do it if it takes a day to do one simulation — then you don’t get anywhere.”

Getting somewhere in science is always important, but never more so than when the prospect of lifesaving knowledge is on the horizon. Sophisticated new simulations of fibrillation and arrhythmia in human hearts hold the promise of providing the *in-silico* pre-screening of new anti-arrhythmia drugs, more effective defibrillators, and improved surgical techniques to treat arrhythmia.

All of this is being aided by a computer scientist who at the start of the project had to reach for a dictionary to understand cardiac terminology.

Concludes Oh, “Simulations are not going to replace experimental data. But numerical simulations are highly reproducible — they’re not dependent on the conditions present in an experiment which make it difficult to reproduce them. Once you know that what you’re simulating is identical to the experiments, then you can do more replicates and do them faster without harming anything.”

“Simulations are not going to replace experimental data. But numerical simulations are highly reproducible — they’re not dependent on the conditions present in an experiment which make it difficult to reproduce them.”

By Peter Gwynne

# Attila the Code

**A TEAM AT THE LOS ALAMOS** National Laboratory has developed a new computer model that predicts how radiation travels through space. Potential applications range from better design of nuclear power plants to improved delivery of medical therapies.



>> Astrophysics
>> Nanoscience
>> Climate Change
>> Laser Physics
>> Computational Medicine
>> Radiation Transport
>> Computational Biology
>> Materials Science

“There’s a new field in physics, beyond experimental and theoretical physics,” says John McGhee, a member of the Computer and Computational Sciences Division at the Los Alamos National Laboratory. “It has to do with numerical modeling and simulation.” It also has very much to do with the work currently under way in McGhee’s division, which involves entities known as transport codes.

Transport codes are computer models designed to solve radiation transport equations — the complex mathematical formulations that model the ways in which various types of radiation, including x-rays, gamma rays, protons, neutrons, and infrared light, move through space and matter. McGhee and a small team of co-workers have developed and continue to perfect a transport code they call Attila. The Attila code can be applied to solve the transport equations to model the behavior of radiation in such applications as determining dosage rates for radiation therapy, protecting the crews in orbiting spacecraft, shielding nuclear reactors, calculating heat transfer via infrared radiation in industrial processes as diverse as steel manufacturing and glass making, and even modeling the progress of forest fires.

### Why Attila?

Why the name Attila? It has nothing to do with any Hun-like characteristics of the code. “I opened my Webster’s Collegiate Dictionary, looked at names, and grabbed one that was short and easy to type,” McGhee recalls. “I started in the As.”

One characteristic that the code shares with the notorious Hun is sheer power. It can help design “anywhere you have a radiation source,” says McGhee. “It involves a lot of physics, a lot of math, and a lot of computer science.” The code provides a complete solution of the radiation transport equations in terms of spatial dependence, angular dependence, and energy dependence. As a result it permits scientists and engineers to deal with such radiation effects as dosage rates, rates of energy transfer, the deposition of energy in two-dimensional and three-dimensional regions, and thermal loading — all issues of value in a variety of scientific and engineering applications.

How does Attila differ from other transport codes? “Traditionally there were two ways of solving transport

problems,” says team member Todd Wareing. The Monte Carlo method takes a statistical approach but can suffer severe statistical errors, and it can also consume prohibitively large amounts of computer time. Deterministic solutions, in contrast, are based on using discrete values for the variables involved in radiation transport — variables such as the location, direction of motion, and speed of the transported radiation. Historically, the deterministic solution method involved modeling radiation flow through a discrete “mesh” that can be thought of as a collection of small rectangular boxes. However, says Wareing, “limiting the transport solutions to a rigid, rectangular mesh can make it very difficult to model complex geometries.”

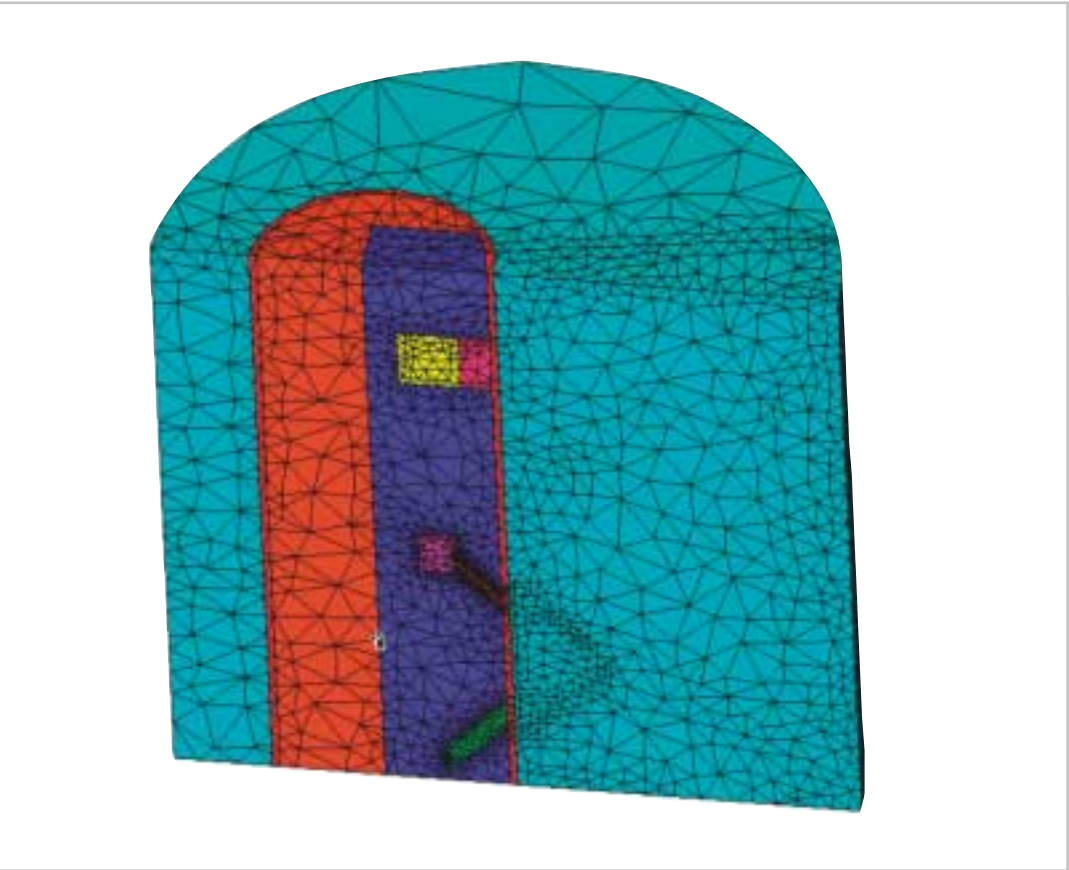
### A Deterministic Approach

What the Los Alamos team has done with Attila is develop a deterministic approach to solving transport equations that avoids the oversimplifications imposed by rectangular meshes. “Attila is, as far as we know, the first unstructured tetrahedral mesh transport code,” says Wareing. In geometrical terms,

adds Shawn Pautz, another member of the Los Alamos team, “Attila is based on a tetrahedral element” rather than a rectangular, or orthogonal, one. That added complexity gives Attila a great deal more flexibility in modeling realistic three-dimensional problem geometries.

The team is now extending its work by developing new codes based on even more complex meshes containing a variety of element types. In addition, Pautz aims to apply parallel computing to a follow-on to the Attila code. This involves using several different computers simultaneously to work on the same problem. “Work was previously done in our group for orthogonal meshes but it was not immediately obvious how to extend that work to unstructured meshes,” he says. “We’ve come up with an approach that seems to scale quite well.” Parallelization demanded some on-the-job education. “It was another big step in that we had to learn other fields with which we were not familiar,” Pautz recalls.

By current standards, the code doesn’t make excessive computational demands. “Attila doesn’t require any special computing power,” says McGhee. “Attila can solve realistic problems using only a common desktop workstation. The technology that Attila represents combines in a



very synergistic way with advances in computing power, memory prices, and storage capacity to provide a unique numerical simulation capability that wasn’t available before.” The specific needs for hardware, memory, and disk space are determined by the size of the particular problem to be solved.

Attila offers two key advantages over conventional methods of solving transport problems: speed and flexibility. In some cases the code can produce solutions to radiation transport problems in as little as one tenth to one quarter of the time that traditional Monte Carlo solutions need. It also allows its users arbitrarily to reduce the number of computational elements required for solutions. That also reduces the need for computing time, without significant loss in accuracy of the results.

### Attila in Action

Attila promises fruitful applications in three general fields. First is radiation transport calculation for the commercial nuclear industry. By effectively calculating the distribution of neutrons — elementary particles that are fundamental components of atomic nuclei — and gamma rays — a form of high-energy radiation — the code permits nuclear engineers to calculate the protection and shielding needed by nuclear power plants, and to check that those components of existing plants are working effectively. Similarly, Attila promises to help designers of futuristic fusion reactors to tame the power of the hydrogen bomb in entirely safe structures. Tracking nuclear radiation also offers value to developers of other forms of energy. Geologists searching for new

Computational mesh for a gamma density well-logging tool

One characteristic that the code shares with the notorious Hun is sheer power.

“There is a lot of computer science, numerical analysis, and software engineering required in addition to the knowledge of the physical processes and mathematics involved to create a simulation tool like Attila.”

## COLLABORATORS

**John M. McGhee** is a member of the technical staff at Los Alamos National Laboratory. He holds a PhD in Nuclear Engineering and specializes in numerical modeling and simulation with applications to neutral and charged particle radiation problems.

**Shawn D. Pautz** is a member of the technical staff at Los Alamos National Laboratory. He has degrees in Nuclear Engineering from Arizona State University and Texas A&M University. His professional interests are in the field of neutral and charged particle transport methods development, particularly in software engineering, code verification, and parallelization strategies.

**Todd Wareing** received a PhD in Nuclear Engineering and Scientific Computing from the University of Michigan. He is currently a team leader at Los Alamos National Laboratory. Todd’s technical work involves the development of advanced spatial differencing schemes and efficient iterative solution techniques for solving the radiation transport equation on both structured and unstructured grids.

### Further Reading

*Coupled Electron-Photon Transport Methods on 3-D Unstructured Meshes*, Wareing, T.A., Jim E. Morel and John M. McGhee, Trans. Am. Nucl. Soc. Vol 83, ANS Winter Meeting, Washington D.C., 12-16 Nov, 2000. (LA-UR-00-3126).

*Discontinuous Finite Element Sn Methods on 3-D Unstructured Grids*, Wareing, Todd A., John M. McGhee, Jim. E. Morel and Shawn D. Pautz. Nuclear Science and Engineering, Vol 138, No. 2, July 2001 (LA-UR-99-6340).

### Contact:

John McGhee  
mcghee@lanl.gov

Shawn Pautz  
pautz@lanl.gov

Todd Wareing  
wareing@lanl.gov

sources of fossil fuels use nuclear well-logging tools to characterize the strata associated with reservoirs of oil and gas. Attila can help them to analyze their findings faster and more efficiently.

In addition, the code works with infrared radiation, which transfers heat. That gives it a wide range of potential uses in the design of industrial processes large and small. Engineers can apply the code to understand heat transfer in large-scale designs of steel furnaces, large boilers, and glass manufacturing plants. At a smaller scale, automotive engineers can use it to determine how heat travels under the hoods of cars of new design. And at an even more minute level, electronic engineers can apply Attila to understand how heat builds up in integrated circuits.

Heat transfer studies don’t apply only to industrial processes. One significant area of potential application has value in a natural phenomenon that frequently causes extensive damage: forest fires. Scientists who try to understand the development of those fires may be able to apply Attila to refine their calculations.

A third major application of Attila involves the calculation of radiation dose rates for people and materials. Because Attila can calculate the passage of radiation through any type of material and predict the amount of radiation delivered to any particular location, it has ready application in determining the amount and profile of radiation therapy that doctors should apply to cancers and other medical conditions.

The same sort of calculation can help designers of manned spacecraft to determine the type of shielding necessary to protect astronauts and sensitive instruments on spacecraft from the radiation always present in space.

## Commercial Prospects

Development of Attila started as a semi-commercial enterprise. The vehicle was a Cooperative Research and Development Agreement (CRADA), a type of government-sponsored project that brings together national laboratories and commercial organizations to help develop promising research into usable technologies. In 1995, when the Attila team had just started its investigation of tetrahedral mesh transport codes, it coordinated a CRADA in association with the oil and gas industry. “They use nuclear well-logging tools to determine the density and porosity of rock formations. They put a radiation source in the rock and measure the radiation reflected from the formation,” explains McGhee. At the time, the Monte Carlo approach represented the only way to interpret the measurements. The CRADA enabled the Los Alamos team to apply Attila to the problem. The result: “We now have an accurate model for an important well-logging tool known as a gamma-gamma density tool,” says Wareing. “Our work was the first to apply transport codes to this type of tool.”

The team hopes to move Attila to the marketplace. “It’s being evaluated for commercial use,” says McGhee. “We

have hopes that Attila will be in use in several commercial markets within the next year or so.”

Meanwhile the team continues to perfect Attila and makes plans for follow-on projects, benefiting from the specialized technical facilities available at the Los Alamos National Laboratory. “We knew that our parallelization technique would work on four or eight processors, but the effects we wanted to see would have to occur at up to hundreds and even thousands of processors,” says Pautz. Current parallel prototype algorithms can efficiently utilize as many as 250 processors. “In any other place I would have had to hunt for computing time,” Pautz continues. “Here I used leftover computer time from other teams.”

The scientists at the national laboratory have proved just as important as equipment in helping to develop the Attila transport code most efficiently. “You start out with a graph outlining the tasks that need to be done,” explains Pautz. “None of

us was familiar with graph theory. But there were people right around me whom I could ask. What otherwise might have taken weeks or months for me to solve took less than a week.” McGhee agrees. “Theoretically we could have accomplished the work we’ve done over the Internet,” he says. “But it’s certainly a lot more convenient to have everybody co-located — to have experts ready at hand.”

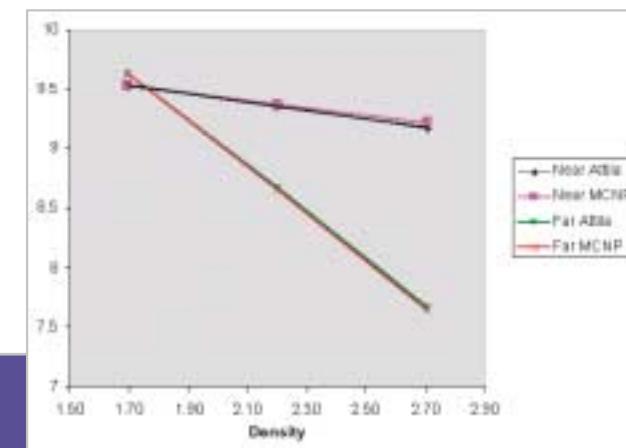
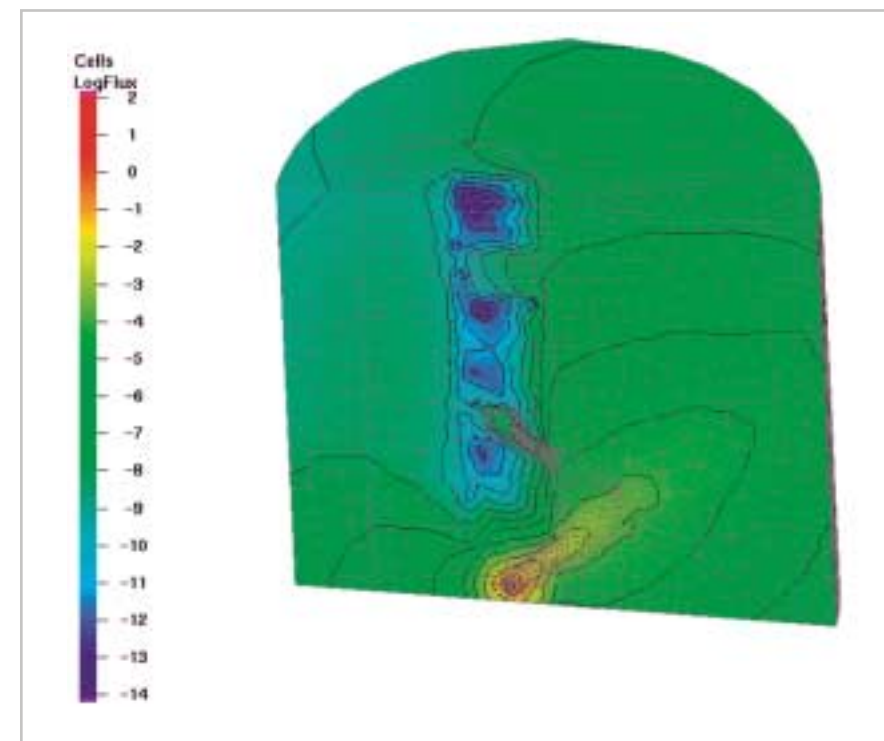
## Other Attractions

Los Alamos offers additional attractions independent of science and technology. “There are a lot of benefits to living here,” McGhee says. “We are near the city but do not experience the smog and congestion associated with large metropolitan areas. We have opportunities for outdoor recreation.”

The team has also taken the opportunity to expand its scientific understanding. McGhee, Pautz, and

Wareing all have PhDs in nuclear engineering. But these days, says McGhee, “we often think of ourselves as computational physicists. There is a lot of computer science, numerical analysis, and software engineering required in addition to the knowledge of the physical processes and mathematics involved to create a simulation tool like Attila.” In Wareing’s view, the national laboratory’s facilities and people, along with the cutting-edge work in transport numerical methods, represent a strong attraction for scientists.

Graduate students have the opportunity to share in the excitement. In fact, a student from the University of New Mexico is already working on charged particle enhancements for Attila. That’s a significant problem because many transport codes can deal only with neutral particles such as neutrons. Plainly, the work offers challenges beyond the norm.



Left:  
Color contour plot of Scalar Flux in gamma density well-logging tool

Right:  
Logarithm of collision rate versus bulk density of rock formation of the gamma density well-logging tool

By William J. Cannon

# Designer Cells

**THE BODY IS OFTEN COMPARED TO A MACHINE.** Many scientists have liberated themselves from the metaphor in favor of fresh, raw literalism: the body *as* machine. The human body is made of parts that are self-contained machines interacting with one another. Molecular-level changes lead to cellular ones; cellular to tissue, and so on to organs and organisms, providing whatever the body needs to plug along.

The tiny machines and the processes by which they harness, transform and transport energy, the biochemical and molecular components from which they are built and rebuilt, the network that links them and keeps them working together, environmental insults that trigger developmental errors that cause them to fail — all follow the rules of physics and chemistry and can therefore be modeled mathematically and computationally. The real test becomes whether the cell works according to a *design* whose complexity can be defined in the most efficient mathematical and computational terms possible.

## The Virtual Cell Project

At the Department of Energy’s Pacific Northwest National Laboratory (PNNL) in Richland, Washington, David Dixon and his team are building what he calls the “virtual cell”. Under the auspices of PNNL’s Biomolecular Systems Initiative, they are five years into an intense and massive data-churning,

analysis and modeling effort to enable predictions of real cell behavior; information that may one day assist experimentalists in fashioning working cells and cell parts of their own design. This work is made possible by the massive hardware, data storage and networking systems and software developed by PNNL’s Theory, Modeling & Simulation Directorate in the William R. Wiley Environmental Molecular Sciences Laboratory (EMSL), a DOE national user facility located at PNNL.

“We’re taking advantage of all of the biological data coming in, including genomics and proteomics to identify the machinery of the cell, including proteins and protein complexes that do the work in cells and pass instructions essential to life,” says Dixon, associate director for Theory, Modeling & Simulation in the EMSL. “We want to understand the location of all these things in space and how they change in time so we can understand how the cell processes information and how it communicates with other cells.”

Much of what we currently know about the structural details of cellular machinery, proteins and DNA is based on optical imaging, other microscopies, crystallography and nuclear magnetic resonance. For instance, the tinker-toy model of DNA’s double-helix structure arose from interpreting its crystal portrait. Today protein crystals and the amino acids and peptides from which they are made are depicted as colorful, 3D computer renderings that look like a dancer festooned with ribbons — sans the dancer. Through observation and experiment and associated ribbons, it is possible to capture snapshots of a particular dance — and to see what effect removing or handicapping particular dancers or adding others will have on a performance. The virtual cell project will help define what happens when our metaphoric dancers speed up or slow down, or if one or two let go of their ribbons or rearrange them between the frames of crystal snapshots. Perhaps more importantly, why are the ribbons and dancers behaving this way?

## Cellular Simulations

The virtual-cell approach is to traditional biology as digital film is to a series of stills in a slide projector. Molecular biologists can deduce the activity aplenty in the slide-show approach, and Dixon notes that

the models developed by his team are built on data gleaned from experimental biology. Experiments tease out a cell’s functional biology, inform models and offer a way to check a simulation’s verisimilitude. The simulation adds a sort of narrative fluidity to the frozen-moment stillness of experimental results and can reveal a biology that may be invisible to experimentalists. A life-or-death bit of biology may take but a nanosecond, yet it’s crucial that everything that takes place in that sequence is accounted for if a model is to prove useful.

For instance, Haluk Resat, a member of Dixon’s group, has run simulations of the signaling agent epidermal growth factor, or EGF. Cell signaling is a sort of wireless communication that controls all complex functions in the hierarchical network, from protein to cell to, in this case, skin. EGF works by binding with an EGF receptor, a portal on the cell’s membrane. This binding sets off a cascade of chemical reactions inside the cell that produce signals that make their way to the nucleus, where protein-building genes are expressed. Resat’s model integrated what is known about EGF to fine-tune the cascade reaction and, in the process, to quantify not only which proteins are expressed but where inside the cell the processes that lead to protein



Ribbon diagram representation of the complex of the protein Ras with the Ras-binding domain of Raf, which is an essential component of the EGFR and other signaling pathways that regulate cell differentiation and proliferation.

expression occur — information that was otherwise invisible between slides in the projector.

“Because we’ve run the simulations for a long time, we see different features than you would in experimental snapshots. The models must account for movement through time so they can predict how a cell will respond to a given perturbation and to various chemical reactions.” The resulting signals disperse throughout the cell, leading to further biochemical reactions, which can lead to proper function or to disease, depending on the nature of those reactions and where they take place.

Today protein crystals and the amino acids and peptides from which they are made are depicted as colorful, 3D computer renderings that look like a dancer festooned with ribbons — sans the dancer.

## DAVID A. DIXON

David A. Dixon is Associate Director for Theory, Modeling & Simulation in the William R. Wiley Environmental Molecular Sciences Laboratory (EMSL) at the Pacific Northwest National Laboratory (PNNL) in Richland, WA. He leads one of the largest groups of computational chemists in the world (~40 permanent staff) and is responsible for the Molecular Sciences Computing Facility, part of the EMSL DOE User Facility. He is also the leader of the Virtual Biology effort at PNNL. He has a BS from Caltech in chemistry (1971) and a PhD from Harvard (1976) in physical chemistry. Prior to coming to PNNL, Dr. Dixon was on the faculty at the University of Minnesota Chemistry Dept. for 6 years and was a member of the research staff at DuPont's Central Research & Development, Wilmington, DE, for more than 12 years. Dr. Dixon's awards include a Harvard Junior Fellowship, a Sloan Fellowship, a Dreyfus Teacher-Scholar Fellowship, and the 1989 Leo Hendrik Baekeland Award presented by the American Chemical Society. He is a Fellow of the American Association for the Advancement of Science and of the American Physical Society and an Adjunct Professor, Chemistry, at the University of Utah. Dr. Dixon's main research interest is the application of numerical simulation techniques to chemical and biological problems, with a focus on electronic structure theory. Dr. Dixon has published more than 385 papers on a wide range of topics.

### References

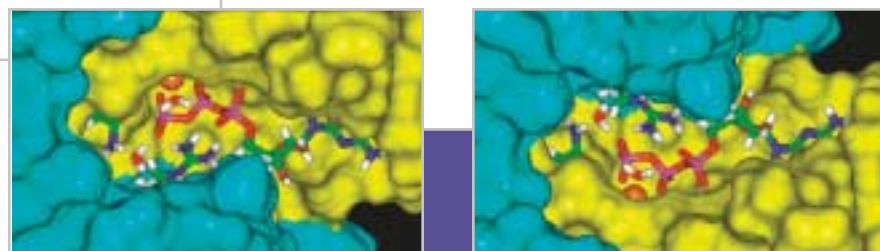
*The arginine finger of RasGAP helps Gln61 align the nucleophilic water in GAP-stimulated hydrolysis of GTP*, H. Resat, T. P. Straatsma, D. A. Dixon, and J. H. Miller, Proc. Nat'l. Acad. Sci., USA, 2001, 98, 6033.

*An Algebraic-Combinatorial Model for the Identification and Mapping of Biochemical Pathways*, J.S. Oliveira, C.G. Bailey, J.B. Jones-Oliveira, and D.A. Dixon, Bulletin of Math Bio., 2001, 63, 1163.

*Probability-Weighted Dynamic Monte Carlo Method for Reaction Kinetics Simulations*, H. Resat, H. S. Wiley, and D. A. Dixon, J. Phys. Chem. B, 2001, 105, 11026.

### Contact:

David Dixon  
David.Dixon@pnl.gov



*Interface between the proteins in the p21ras (yellow): RasGAP (blue) complex. Close up look at the structure of the active site where hydrolysis of GTP takes place. Parts of the proteins are removed for clarity. Shown as sticks are the GTP molecule, waters involved in the hydrolysis reaction, and the important amino acid residues p21ras:Gln61 (on the left) and RASGAP:Arg789 (at the right).*

In a paper published last year in the *Proceedings of the National Academy of Science*, simulations by Resat, T. P. Straatsma, J. Miller and Dixon, as part of the virtual cell project, pointed to a chemical reaction involving an “arginine finger” structure in a group of so-called molecular switches near membrane surfaces. The switches mediate the chemistry that regulates cell proliferation and differentiation. This protein, called Ras, is of special interest because a mutant form of the protein is found in nearly a third of all tumors.

“The simulations showed that biology has to conform to the rules of chemistry and physics, in that charge balance — an equal number of positive and negative charges — in the active site of the protein plays an important role,” Dixon says. Without charge balance, essential chemistry cannot take place, and “the protein does not function properly, the switch does not work and you can get rapid cell division leading to cancer.”

Their simulation suggested the arginine finger plays a role — unseen by those studying the Ras crystal structure — in refereeing reactions involving water, or hydrolysis, that are dampened in the mutant forms of Ras. Their work has

forced molecular biologists to look more closely at the arginine finger and offer a much more refined, alternative, experimental-based explanation of hydrolysis in the Ras complex.

### Modeling Techniques

As in digitizing visual information, data compression is a key issue in computational biology — that is, what is the least amount of information-space required to generate the greatest fidelity and resolution? How can a program yield the most cogent explanation of what is happening in real life without occupying supercomputers for decades?

The challenge is akin to ransacking mountain-size haystacks of genomic, proteomic and other raw data for a needle or two of salient information about not only the reactions of a given protein but also cellular networks and the entire dynamics of a biological system. The task requires nothing short of new mathematical approaches such as “Petri nets” developed by PNNL's Joe Oliveira, C.G. Bailey, Janet B. Jones-Oliveira and Dixon. The Petri nets allowed them to come up with a novel way to map the information flow through biochemical pathways.

### Other Applications

In another recent test of mathematical honing, Resat joined Dixon and PNNL colleague Steve Wiley to streamline a venerable method called Dynamic Monte Carlo simulation for calculating the properties of physical systems. Their method incorporated computational efficiencies into a sophisticated probability scheme that accounts for chemical reactions and physical processes that occur at widely varying times.

“We’re developing techniques that will work with a wide range of cells,” Dixon says. Besides the work in mammalian cells discussed above, “we’re also interested in modeling microbial cells as part of the DOE’s Genomes to Life initiative that began with the Human Genome Project.” This effort speaks directly to DOE’s core missions: basic science, energy research and environmental remediation. Microbes have odd appetites; on the menu of some might be the greenhouse gas methane, or toxic metals in the ground.

“In a practical sense, we’re interested in the question of how to modify parts of a cell to deal with nuclear waste — or even how to get a cell to *not* do something,” Dixon says. “You want to control cells so that they will chew up things that they should. If you want to prevent the mobility of uranium or plutonium, for example, you don’t want to oxidize another metal such as chromium that makes the toxic material in the ground more mobile.”

When enough simulations are run and verified by experimental data, will the result be a massive database of blueprints for synthesizing cells? With good chemistry and engineering it’s certainly possible that whole designer cells would follow, though Dixon’s group will continue to concentrate on how the cellular machinery acts and how best to modify that behavior.

**“The simulations showed that biology has to conform to the rules of chemistry and physics, in that charge balance — an equal number of positive and negative charges — in the active site of the protein plays an important role.”**

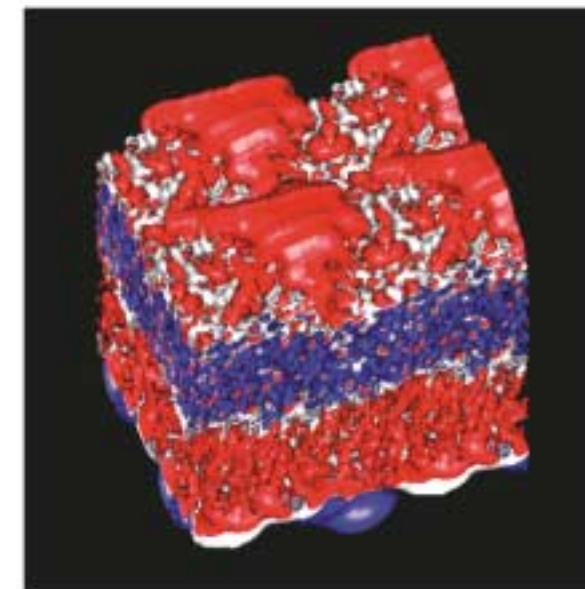


Image 1

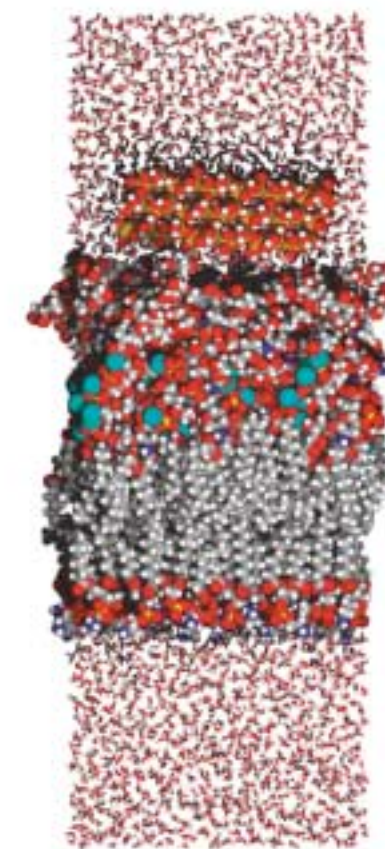


Image 2

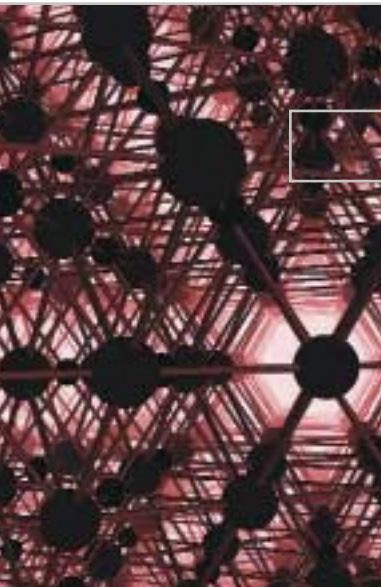
*1) Electrostatic potential isocontours at -150 mV (red), 0 mV (white), and +150 mV (blue) for a side view of the rough lipopolysaccharide membrane of Pseudomonas aeruginosa, after 1 ns of molecular dynamics simulation.*

*2) Space filling representation of the rough lipopolysaccharide membrane of Pseudomonas aeruginosa in contact with a goethite mineral fragment, with water molecules displayed using a stick representation.*

By Victor D. Chase

# Meeting Materials

**EFFICIENCY, STRENGTH, AND LONGEVITY**, all highly desirable attributes of life, are equally applicable to the world of inanimate, manmade objects. And just as medical researchers strive to bring these characteristics to the living, so too do materials scientists endeavor to make them a reality for artifacts.



>> Astrophysics
>> Nanoscience
>> Climate Change
>> Laser Physics
>> Computational Medicine
>> Radiation Transport
>> Computational Biology
>> Materials Science

And just as advanced technology has given new tools to those medical researchers, so too has it furthered the quest of materials scientists, primarily through the power of computer simulation.

Most of the inanimate objects in question are made of either metallic or ceramic materials, both of which are actually conglomerations of minuscule crystals. What looks like a solid aluminum spar of an aircraft wing, for example, is really a mass of microscopic crystals. And the spaces between those crystals, which are occupied by a world of rapidly moving atoms, constitute the weak links in the material. Hence, if that spar fails, it will likely crack along the boundaries between those crystals. While engineers know approximately how long that aluminum spar will last, what they don't know is where and how it will crack. So, to prevent any untoward surprises, they over-design that spar, making it heavier than it would have to be if the reasons behind a potential failure were more fully understood.

And that's where Dieter Wolf and his colleagues at Argonne National Laboratory (ANL) come in. Wolf is a senior scientist and leader of the Interfacial Materials Group, which is responsible for developing computer simulations to learn what is going on in those spaces between the crystals.

"Material's microstructure is what controls the properties of the material," explains Wolf. "The nature of bonding of the atoms in individual interfaces and the entire spatial network of these grain interfaces is what controls the properties, and that's what we try to understand by simulation."

And because Wolf's group uses simulation methodology that is not material specific, their work can apply to virtually everything from aircraft wings and jet engine turbine blades down to the submicron-sized features of modern electronic circuitry.

### Predicting and Preventing

Two primary goals drive this work. One is to develop an understanding of the characteristics of materials so that engineers have a better understanding of the tolerances of the materials used in their designs. The other is to provide information that can be used to improve the characteristics of various materials for given applications during the manufacturing process.

An engine's turbine blades, for instance, are exposed to hundreds of degrees of temperature as they spin at thousands of revolutions per minute. As a result, small changes are constantly taking place in the blades' polycrystalline microstructure, but these changes cannot be empirically observed as the turbine blades turn.

"So, it is extremely important to develop predictive models that allow you to simulate the performance

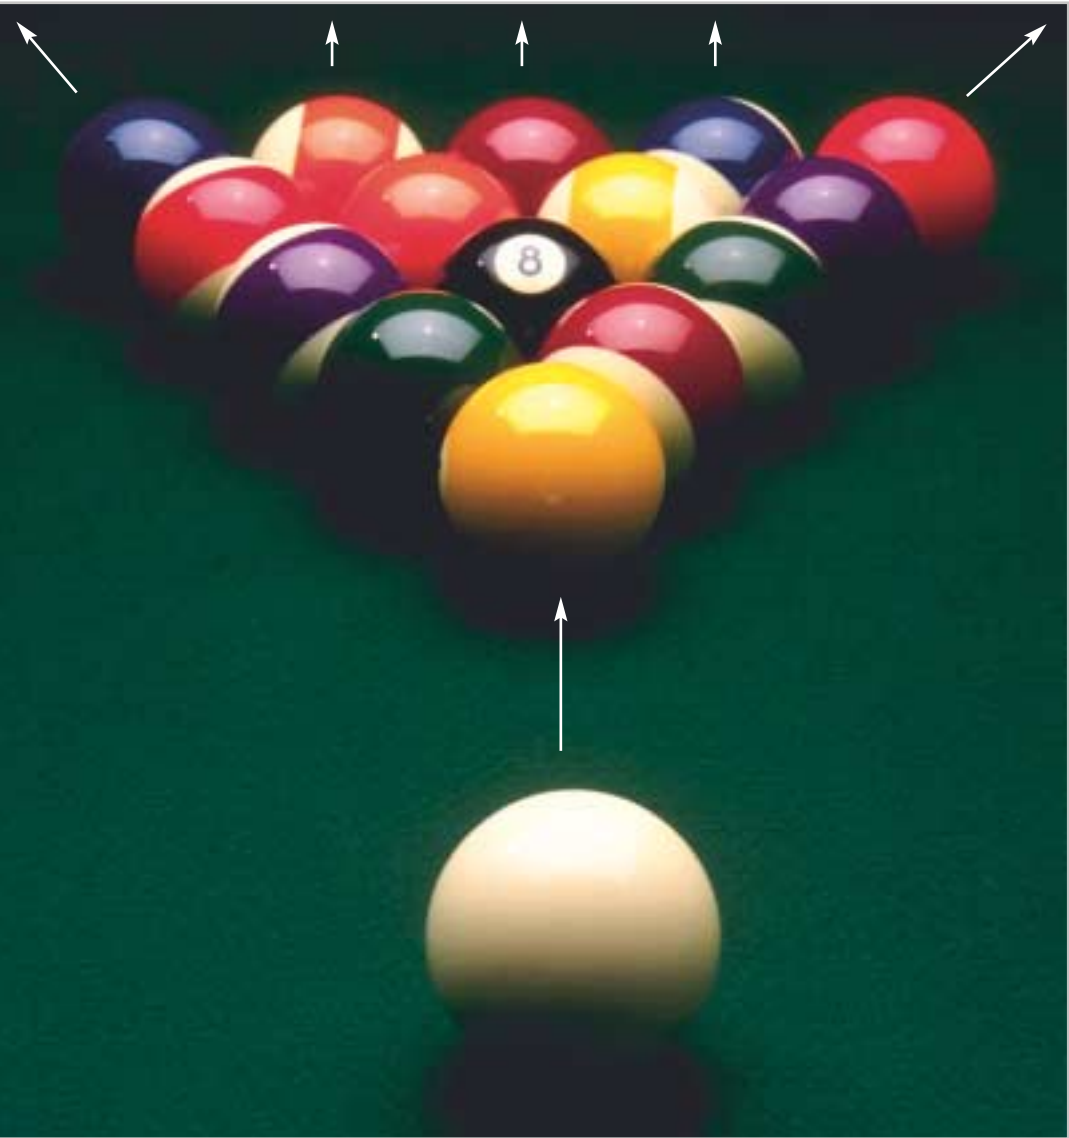
of this material under these high temperatures and stresses so you can predict how long it will last...The goal is to develop a predictive modeling tool such that the engineers one day down the road have a better way of predicting how the material will behave over time," says Wolf.

And, he adds, "if you want to reduce the weight of a jet engine, it would help enormously if you understood the underlying processes, so the goal of materials simulation is to put a lot more of the physics of the microstructure into [engineering] models," says Wolf. "If you could do that, you could tailor the microstructure of the materials. The big goal is to increase the power-to-weight ratio by a factor of 50 to 100 percent over the next 20 years."

### A "Catch-22"

The fact that most solids have crystalline structures has been known for about 100 years, but only since the mid-1990s — as a result of the development of powerful computers — have researchers had the ability to use multi-scale materials simulation to probe the secrets of the microstructures within these polycrystalline materials.

Yet the microscopic boundaries between crystals do not readily reveal their secrets. One of the



quandaries they pose is the "Catch-22" of materials science: to view the internal boundaries of a material would logically require cutting into the material, but once that is done, the interior surfaces become the exterior surfaces and their very nature is changed. It was to overcome such empirical limitations that researchers turned to computer simulation.

And that brought them right up against yet another seemingly insurmountable blockade. Specifically, to create computer simulations requires at least some

empirical evidence of what is going on in the crystals' boundaries, or else the models would be pure guesswork. To skirt this problem the researchers devised an ingenious two-pronged approach.

First, they applied a basic understanding of physics, such as Newton's laws of motion. These same laws of motion apply to atoms as much to a large ball. And while understanding this is

"The nature of bonding of the atoms in individual interfaces and the entire spatial network of these grain interfaces is what controls the properties, and that's what we try to understand by simulation."

## DIETER WOLF

Dieter Wolf received his PhD from the Max-Planck-Institute for Metals Research in Stuttgart, Germany. He has been the Senior Physicist and Group Leader of the Interfacial Materials Group, Materials Science Division at Argonne for over 15 years. Dieter won the Basic Energy Sciences Award for Sustained Outstanding Research in Metallurgy and Ceramics for work on Computer Simulation of Interfacial Materials and Phenomena in 1997.

### Further Reading

*Grain-boundary diffusion creep in nanocrystalline palladium by molecular-dynamics simulation*, V. Yamakov, D. Wolf, S.R. Phillpot and H.Gleiter, Acta Mater. 50, 61-73 (2002).

*Exact Method for the Simulation of Coulombic Systems by Spherically Truncated Pairwise 1/r Summation*, D. Wolf, P. Keblinski, S.R. Phillpot and J. Eggerbracht, J. Chem. Phys. 110, 8254-82 (1999).

*Materials Interfaces: Atomic-Level Structure and Properties*, A book containing 27 invited chapters; edited by D. Wolf and S. Yip, Chapman and Hall, London, 1992, 716 pages.

### Contact:

Dieter Wolf  
wolf@anl.gov

important to the development of materials simulations, it is also important for the scientists to realize that at the microstructural level — one of the three length scales that multi-scale materials simulation examines — the physics is different. Specifically, as polycrystalline materials are subjected to stresses, the boundaries between their crystals move. As they do so, they follow a viscous law, in that the forces exerted upon them cause them to move at a constant speed — much like a glass bead sinks slowly and steadily in a glass of water. The result of this motion is heat rather than increasing velocity. This distinction between Newton's law and the viscous law of motion is important when it comes to modeling the activity of materials.

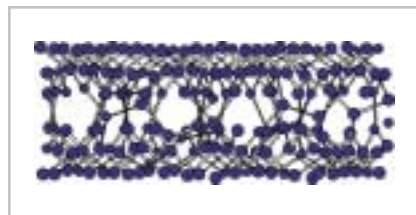
Wolf, himself a physicist, and his group apply this knowledge to building models that look at both the activity within individual grain boundaries and the interaction of all the boundaries within a material. They take this approach, rather than examining what happens to all of the boundaries at the atomic level, simply because the number of atoms involved make such a task onerous. There are simply too many of them. Hence, says Wolf, "You throw away thinking about the atoms, but you incorporate what the atoms do in the grain boundaries."

### A Physicist's Best Friend

In addition to applying these basic physics concepts to the development of materials simulations, "in some very specialized cases there are experimental techniques available to statically probe the atomic structure of the interface," says Wolf. But even then, he adds, "it's extremely difficult to obtain that information and you can only get such information on certain types of internal interfaces."

Nanocrystalline diamonds provide a case in point, since their innards can be probed nondestructively by light in a technique known as Raman scattering. Nanocrystalline diamonds are synthetic diamonds being developed at Argonne, which have grain sizes in the nanometer range. This is considerably smaller than the micron-sized grains found in conventional polycrystalline materials — a nanometer being one-billionth of a meter, or about 10 times the size of a single atom. And, in fact, natural diamonds, such as those used in jewelry, are single crystals and therefore have no grain boundaries.

The non-invasive Raman light scattering applied to nanocrystalline diamonds is possible because, though different chemically and structurally, diamond and graphite are two different phases of the same material: carbon. Diamond is the hardest material known, while graphite is soft enough to be used as pencil lead. And, in fact, industrial diamonds are created by compressing graphite.



*Above: Atomic structure of a representative grain boundary in nanocrystalline diamond, revealing that approximately half of the atoms are graphite-like. The top and bottom perfect-crystal grains are terminated by diamond surfaces which are rotated relative to each other, thus forming the grain boundary.*

By probing a nanocrystalline diamond with light and observing how the light is scattered, the extent to which graphite is present can be determined, and from that, other characteristics of the material can be extrapolated. "The Raman scattering gives a peek inside" the diamond, says Wolf, but "the simulations teach you a lot more than the experiments."

These experiments have shown that two to three percent of all the atoms in a nanocrystalline diamond are graphite-like. And since it is known that the crystals in a diamond are pure and contain no graphite, whatever graphite shows up in the light-scattering experiment should be in the boundaries between the crystals. It is also known from experimentation and simulation that about five percent of all the material in a nanocrystalline diamond is in the grain boundaries.

# diamond



very good insulators, by turning them into conductors of electricity. These ultra-small crystal diamonds are created using vapor deposition, a process similar to that used in the manufacture of semiconductor chips. "The idea is to choose deposition conditions such that as the vapor gets deposited on an amorphous substrate, little crystallites form," says Wolf. The result is diamond crystals with grain sizes in the 5 to 10 nanometer range, which gives the nanocrystalline material unique properties.

"You have steep gradients in the structure, which gives an opportunity to develop entirely new materials with properties that one never expects for these materials, and to understand those we use computer simulation," he says. It was just such simulation that showed that these diamonds are conductive because of the graphite found in their grain boundaries.

This new type of diamond material is also being tested for use as a virtually friction-free coating that holds promise for extending the life of machines in which severe operating conditions create significant wear.

### Purpose of the Work

"The whole purpose of the work here is to develop a fundamental understanding of materials microstructure, which is relevant to many different technologies," explains Wolf.

"We develop the fundamental knowledge that allows the application-oriented programs to develop more sophisticated approaches to the

processing of their materials to understand why they fail and how they fail. The ultimate goal is to develop an industrial type modeling tool which connects with the fundamental materials physics in a much more truthful way than has ever been possible."

He adds, "To have physically faithful models that properly incorporate all of the atomic-level physics into the long time-scale and spatially large-scale problem of predicting the lifetime of engineering structures is a real challenge."

## A SCANDINAVIAN WARRIOR

**>>> Wolf and his crew have proven to be as resourceful in fashioning the computer hardware on which they run their highly complex simulations as they are in developing the materials modeling software itself.**

Their computer configuration is known as a *Beowulf* cluster, after a Scandinavian warrior whose name is the title of an Old English epic. In this case, the warrior's super powers are those of a supercomputer composed of off-the-shelf personal computer components. The idea of tying a cluster of less powerful computers together to make one superpowerful computer originated during the mid-90's and has since become popular with researchers and academics.

Wolf and his colleagues opted for the *Beowulf* approach for a very simple reason: money.

True supercomputers cost millions of dollars, and demand for time on them is great. On the other hand, "We have a cluster supercomputer that you can buy at Circuit City for \$50,000," says Wolf. And, they have access to their homemade supercomputer 24/7. Actually, Wolf's Interfacial Materials Group splurged and spent \$65,000 on their 100 node cluster, but that includes the graphics terminals and the switching equipment that allows the processors to be connected in parallel so they can communicate with each other while working on the same problem.

**"In some very specialized cases there are experimental techniques available to statically probe the atomic structure of the interface."**

# Ken Comer

## Procter & Gamble

**WHILE COMPLETING HIS PHD** in computational multi-phase flow analysis at North Carolina State University, Ken Comer had a lot of ideas as to the companies or government labs where he might find a job. A diaper maker wasn't one of them.

**“W**hen they called me up and asked me if I wanted to interview for a job at Procter and Gamble, I thought, ‘Procter and Gamble? They’re making paper towels and diapers. Why would they be doing modeling?’” says Comer, now a computational fluid flow analyst at the company’s Corporate Engineering Technology Laboratory in West Chester, Ohio.

The multinational consumer product maker, with about 100,000 employees worldwide, produces nearly 300 brands, many of which are household names, including Tide, Swiffer, Pampers, Clairol, Tampax and Pringles. As Comer now makes clear, Procter and Gamble (P&G) is highly dependent on engineering and research and development savvy in order to prosper in the highly competitive consumer products business. Last year, P&G spent about \$1.7 billion on R&D.

“In the last three years, modeling and simulation has really taken off here, basically because modeling has come of age,” says Comer, who was hired by P&G in 1998, prior to completing his PhD “Our goal is to make our products faster and cheaper. Modeling helps deliver that, especially quicker development and cheaper and faster processing.”

As a member of a corporate group of a half-dozen modelers tackling problems across the P&G empire, Comer says he’s faced with an exciting variety of fluid mechanics problems, many of them outside his particular expertise. His job isn’t to always personally know the answer to a problem, but to know how to find it.

This means acting as a bridge to the larger modeling community and maintaining a broad network of government lab and university researchers, including key contacts made during the DOE CSGF program. His role is also to scour the research horizon for emerging technologies that can provide a competitive leg up, all the while keeping P&G’s trade secrets under wraps.

And, when it comes to finding answers, speed is of the essence.

“One of the first things I noticed is the speed with which you have to turn around projects,” Comer says.

“In a lot of our projects, somebody needs results within a short period of time. Instead of years, we’re talking months or weeks. In industry [as compared to graduate school research] you have to make a conscious decision about the level of accuracy you can afford in order to deliver the results. Because even if you have the perfect results, if it’s after the time when the decision had to be made, then you’ve delivered useless results.” Comer says he thrives on the mix of near-term and longer-term projects that he must juggle. The more immediate projects, though not always the most technically challenging, provide what he calls ‘brain candy’—that quick, positive boost from seeing the tangible benefits of your work.

One such modeling project involved his group’s input on a new system to improve the performance of a paper machine that was being slowed by dust control problems. Keeping these machines, which cost up to a half-billion dollars each, running 24/7 is crucial to the bottom line. To try to minimize the machine’s downtime, the modelers worked with the engineering consultants to simulate their designs and provide suggestions for improvements.

“When they put the system in, it worked perfectly. Our modeling had a direct impact on the business and you could see it,” says Comer.

**“In the last three years, modeling and simulation has really taken off here, basically because modeling has come of age.”**



# John Dolbow

## Duke University

**MOST DOCTORAL** students worry about getting a job, but for John Dolbow the worrying — albeit short-lived — started on the job. In 1999 he was recruited as an assistant professor of Civil and Environmental Engineering at Duke University, while still completing his PhD in theoretical and applied mechanics at Northwestern University.

**“I**t was pretty intimidating at first,” says Dolbow of his new academic credentials. “You’re given more than a few responsibilities right off the bat that you weren’t necessarily exposed to as a graduate student.”

There is the professor’s perennial need to write research proposals for further funding, advising graduate students (he currently has two), and then developing and teaching undergraduate courses.

“You have to be able to juggle all of these things simultaneously and still continue with your own research. As a PhD student you’re used to being able to devote weeks and months to a single problem and to thinking about it on a daily basis. That’s a luxury that you don’t really appreciate at the time,” he says.

Now thriving in this multitasking role, Dolbow credits his rapid entry into academia in part to his successful graduate work with advisor Ted Belytschko, a pioneer in the engineering application of the finite element method (FEM). Belytschko’s algorithms for explicit dynamics are widely used throughout industry, in particular within the car crash simulation codes of the Big Three automakers.

One of the problems with the standard FEM, however, is that “it can’t represent an arbitrary fracture path,” says Dolbow. “My PhD research was concerned with the development of mesh-free methods, which are really a way of circumventing some of the inherent limitations the FEM has, principally for modeling the fracture and failure of brittle materials.”

His proposal writing efforts and research are coming together in a new National Science Foundation-sponsored project to computationally study the high-speed machining of aluminum alloys. The project is an industry-university collaboration involving researchers from Duke and Notre Dame Universities and the aluminum company Alcoa.

Thin sheets of aluminum are created by using a knife-like cutting tool to skim a layer from the top of quickly translating bulk stock. The process sometimes produces unwanted segmented chips, or fractures in the metal.

“At the outset, we’re trying to understand the fundamental failure process that occurs in the aluminum alloys at the grain scale. And to do that we’re modeling the high-speed machining process using the computational facilities and numerical

algorithms developed here at Duke,” says Dolbow, who’s very familiar with the supercomputer used on the project. He built the 2-node Beowulf Cluster as one of his first tasks as a new professor.

Three years later, the initial stage fright is a distant memory. In fact, Dolbow is creating innovative ways to make undergraduate students feel more comfortable.

In conjunction with colleague Dr. Henri Gavin, he has created WEAVE: the Web-based Educational Framework for Analysis, Visualization, and Experimentation. The on-line project, supported by a \$340,000 NSF grant, allows engineering students to conduct lab experiments on-line when, and from where, it suits them. (see [www.weave.duke.edu](http://www.weave.duke.edu))



**“At the outset, we’re trying to understand the fundamental failure process that occurs in the aluminum alloys at the grain scale.”**

# Marcus Martin

*Sandia National Laboratories*

**MARCUS MARTIN IS GIVING AWAY** his research. The Sandia National Labs molecular modeler wants computational and lab-based chemists to use his Monte Carlo for Complex Chemical Systems (MCCCS) program. Please.

**T**he goal of the project is to try and get people in industry interested in using this code, and to try to get this particular code adapted by some of the academic groups. And that tends to happen a lot better if they can go in and modify the code themselves without worrying about endless licensing issues,” says Martin, a staff scientist in Sandia’s Computational Materials and Molecular Biology division.

Being a national lab researcher and code promoter means the former DOE CSGF fellow works in a realm intermediary between academic and industry researchers. He’s interested in pursuing the “big code projects” more associated with academia, but also in seeing his work applied.

The MCCCS simulation is primarily designed to understand the behavior

of molecules involved in a two-phase equilibrium, an important issue in the chemical industry. The program simulates two wall-less boxes, with the several hundred molecules involved being able to move between the boxes and therefore separate into vapor and liquid phases. The program was initially developed by Martin and his then-supervisor at the University of Minnesota, J. Ilja Siepmann, beginning in 1995. “The chemical industry is at least interested enough that they’re trying it out,” Martin says. “The prediction of vapor-liquid co-existence curves is a big deal in the chemical industry. The most common way to separate things is through distillation columns, and you need a good knowledge of the vapor-liquid co-existence to design good distillation columns.”

He’s also working with a chemical company that’s interested in applying the program to study the solubility of small molecules in polymers.

The current version — the MCCCS Towhee configurational-bias Monte Carlo program code — is available for download under a GNU public license, the same as the one used by LINUX. The project is funded in part by the DOE’s Office of Industrial Technologies to support its distribution to the chemical community.

While promoting this initial application of his code, Martin is also exploring new possibilities, including its application in drug discovery and in the development of bio-sensors. The program is particularly useful in relation to exploring the drug-docking problem — how a small molecule joins to a protein to interrupt its function.

“Part of what you want to know is how does the molecule partition between water and the active site in the protein,” Martin says. To be effective, a drug needs to preferentially absorb in the active site of the protein, yet it must also be water soluble so that it can be absorbed into the blood stream.

Martin says that extending the MCCCS Towhee program to other applications is facilitated by being part of Sandia’s critical mass of modelers with similar computational expertise that is different enough to allow for synergistic collaborations. He’s currently working with a colleague to combine the Towhee code with the Tramonto molecular theory code in order to do more efficient simulations in water.

It’s collaboration that has Martin knocking on new doors — those of biologists this time — and looking for ways to make his innovative program useful to a new community of scientists.

To download the Towhee program see: <http://www.cs.sandia.gov/projects/towhee/index.html>

Being a national lab researcher and code promoter means the former DOE CSGF fellow works in a realm intermediary between academic and industry researchers.

# Alumni Directory

## A

**Asohan Amarasingham**

*Brown University*  
Cognitive Science  
Fellowship Years: 1998-2002  
Current Status: Brown University

## B

**Edward Barragy**

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

**William Barry**

*Carnegie Mellon University*  
Computational Mechanics  
Fellowship Years: 1994-1998  
Current Status: University of Technology, Bangkok

**Martin Bazant**

*Harvard University*  
Theoretical Physics  
Fellowship Years: 1992-1996  
Current Status: Massachusetts Institute of Technology

**Edwin Blosch**

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

**Dean Brederon**

*University of Utah*  
Computer Science  
Fellowship Years: 1996-1998  
Current Status: University of Utah

**Paul Bunch**

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

**Jeffery Butera**

*North Carolina State University*  
Applied Mathematics  
Fellowship Years: 1993-1997  
Current Status: Hampshire College

## C

**Brandoch Calef**

*University of California - Berkeley*  
Applied Mathematics  
Fellowship Years: 1996-2000  
Current Status: University of California — Berkeley

**Patrick Canupp**

*Stanford University*  
Aerospace Engineering  
Fellowship Years: 1991-1995  
Current Status: Petty Enterprises

**Kent Carlson**

*Florida State University*  
Mechanical Engineering  
Fellowship Years: 1991-1995  
Current Status: 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: Systems Engineering — GE Medical Systems

**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  
Current Status: 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: 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*  
Physics  
Fellowship Years: 1991-1995  
Current Status: University of Massachusetts

**Robert Cruise**

*Indiana University*  
Physics  
Fellowship Years: 1997-2001  
Current Status: 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*  
Applied Physics  
Fellowship Years: 1992-1994

**John Dolbow**

*Northwestern University*  
Theoretical and Applied Mechanics  
Fellowship Years: 1997-1999  
Current Status: Duke University

**Brian Dumont**  
*University of Michigan  
Aerospace Engineering*  
Fellowship Years: 1994

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

E

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

**Susanne Essig**  
*Massachusetts Institute of Technology  
Aeronautics/Astronautics*  
Fellowship Years: 1997-2002  
Current Status: Massachusetts Institute of  
Technology

F

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

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

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

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

**Robert Fischer**  
*Harvard University  
Computer Science*  
Fellowship Years: 1994-1998

**Gregory Ford**  
*University of Illinois  
Chemical Engineering*  
Fellowship Years: 1995  
Current Status: Degree in  
Theological Studies

**Oliver Fringer**  
*Stanford University  
Environmental Fluid Mechanics*  
Fellowship Years: 1997-2001  
Current Status: U. of Western Cape,  
South Africa/Stanford

G

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

**Charles Gerlach**  
*Northwestern University  
Mechanical Engineering*  
Fellowship Years: 1995-1999  
Current Status: Aerosim Technologies

**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  
Current Status: Boston University

**Kevin Glass**  
*University of Oregon  
Computer Science*  
FellowshipYears: 1996-2000  
Current Status: 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: Scholars Advocate/NC  
A&T State University

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

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

**John Guidi**  
*University of Maryland  
Computer Science*  
Fellowship Years: 1994-1997

H

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

**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: University of Wisconsin

**Eric Held**  
*University of Wisconsin - Madison  
Plasma and Fusion Physics*  
Fellowship Years: 1995-1999  
Current Status: Utah State University

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

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

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

**E. McKay Hyde**  
*California Institute of Technology  
Applied 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: University of  
California — Berkeley

J

**Nickolas Jovanovic**  
*Yale University  
Applied Mathematics*  
Fellowship Years: 1992-1994  
Current Status: University of Arkansas

K

**Jeremy Kepner**  
*Princeton University  
Astrophysics*  
Fellowship Years: 1993-1996  
Current Status: 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

**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: Apex Nano Technologies

**Christie Lundy**  
*University of Missouri – Rolla  
Physics*  
Fellowship Years: 1991-1994

M

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

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

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

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

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

**Erik Monsen**  
*Stanford University  
Aerospace and Astronautical Engineering*  
Fellowship Years: 1991-1994  
Current Status: 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: Universal Instruments

**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*  
Fellowship Years: 1993-1996  
Current Status: Lawrence Livermore  
National Laboratory

**Debra E. Nielsen**  
*Colorado State University  
Applied Mathematics*  
Fellowship Years: 1992-1996

**P**

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

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

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

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

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

**Robert (Chris) Penland**  
North Carolina State University  
Nuclear Engineering  
Fellowship Years: 1993-1997  
Current Status: Physiome Science, Inc.

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

**Q**

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

**R**

**Nathan Rau**  
University of Illinois  
Civil Engineering  
Fellowship Years: 2000-2001

**Clifton Richardson**  
Cornell University  
Physics  
Fellowship Years: 1991-1995

**John Rittner**  
Northwestern University  
Materials Science  
Fellowship Years: 1991-1995

**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: Scripps Research Institute

**S**

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

**James Strzelec**  
Stanford University  
Scientific Computing  
Fellowship Years: 1992-1994  
Current Status: Autodesk

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

**T**

**Shilpa Talwar**  
Stanford University  
Scientific Computing/  
Computational Mathematics  
Fellowship Years: 1992-1994  
Current Status: Iospan Wireless

**Mayya Tokman**  
California Institute of Technology  
Applied Mathematics  
Fellowship Years: 1996-2000  
Current Status: University of  
California — Berkeley

**Mario Trujillo**  
University of Illinois  
Mechanical Engineering  
Fellowship Years: 1997-2000  
Current Status: Los Alamos  
National Laboratory

**V**

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

**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: Princeton University

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

**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: Courant Institute

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

**C. Eric Williford**  
Florida State University  
Meteorology  
Fellowship Years: 1992-1996  
Current Status: Florida State University

**Lee Worden**  
Princeton University  
Applied Mathematics  
Fellowship Years: 1998-2002  
Current Status: Princeton University

**Peter Wyckoff**  
Massachusetts Institute of Technology  
Chemical Engineering  
Fellowship Years: 1992-1995  
Current Status: Sandia National  
Laboratories — California

**Z**

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

**Scott Zoldi**  
Duke University  
Physics  
Fellowship Years: 1996-1998  
Current Status: HNC Software

For more information on  
DOE CSGF Alumni visit  
[www.krellinst.org/csgf/alumni](http://www.krellinst.org/csgf/alumni)



**Allison Baker**  
*University of Colorado  
Applied Mathematics*

**Advisor:**  
Elizabeth Jessup  
**Practicum:**  
Sandia National Laboratories –  
New Mexico  
**Contact:**  
allison.baker@colorado.edu  
**Research Synopsis:**  
My current research focuses on reducing memory access costs when solving the linear system  $Ax=b$  (A large, sparse and nonsymmetric). I am focusing on variants of the generalized minimum residual (GMRES) algorithm. I would like to gain efficiency by reducing memory access costs (primarily by increasing re-use of A) through algorithmic changes in GMRES.



**Jarrod Chapman**  
*University of California-Berkeley  
Computational Biology*

**Advisor:**  
Daniel Rokhsar  
**Practicum:**  
Stanford Synchrotron Research Laboratory  
**Contact:**  
jchapman@uclink.berkeley.edu  
**Research Synopsis:**  
My current research involves the simulation of the unfolding pathways of small protein molecules in aqueous solution using all-atom molecular dynamics (MD) methods. In addition to using standard MD packages (NAMD2) to study the behavior of experimentally well-characterized proteins, I have worked on original software for the analysis and visualization of large MD data sets as well as creating a “reduced protein” model which can be used to study the long timescale dynamics of small protein systems.



**Lewis Jonathan Dursi**  
*University of Chicago  
Astrophysics*

**Advisor:**  
Robert Rossner  
**Practicum:**  
Sandia National Laboratories –  
Livermore  
**Contact:**  
ljdursi@flash.uchicago.edu  
**Research Synopsis:**  
My current research involves working on and using our center’s ‘FLASH code’ to understand turbulent (thermonuclear) combustion in compressible fluids. Another interesting part of my work is working on modules or algorithms which may be known in other communities, but not necessarily the astrophysical community. I’m working on a ‘Level Set’ module for the FLASH code to do flame tracking or general interface tracking; this is a very useful technique, but one which is still fairly novel in these sorts of simulations.



**Glenn Hammond**  
*University of Illinois  
Environmental Engineering & Science*

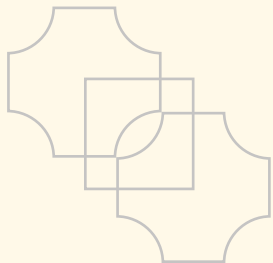
**Advisor:**  
Albert Valocchi  
**Practicum:**  
Los Alamos National Laboratory  
**Contact:**  
ghammond@uiuc.edu  
**Research Synopsis:**  
I am looking at developing a method of preconditioning which better preserves

parallel efficiency. Recently, researchers at Los Alamos National Laboratory developed a physics-based (operator-split) preconditioning for their radiation transport codes. I am applying this technique to the more sophisticated (more degrees of freedom per spatial node and more complex chemistry) multicomponent transport utilizing a Jacobian-free, Newton-Krylov as a solver (PGMRES). This parallel transport model is built upon the PETSc library developed at Argonne National Laboratory.



**Judith Hill**  
*Carnegie Mellon University  
Mechanics, Algorithms & Computing*

**Advisor:**  
Omar Ghattas  
**Practicum:**  
Sandia National Laboratories –  
Livermore  
**Contact:**  
jhill@cyrus.andrew.cmu.edu  
**Research Synopsis:**  
In the past, cohesive zone models have successfully modeled different forms of elasticity, crack nucleation, and crack propagation. In my current work, a finite element model (FEM) with a Xu and Needleman cohesive surface potential implemented into the framework has been used. The Xu and Needleman potential, derived from molecular dynamics and applicable for modeling fracture surfaces, is rate-independent and phenomenal when applied to large length scales. Current efforts are directed toward validating this model with experimental measurements from a brittle system.



**Charles Hindman**  
*University of Colorado  
Aerospace Engineering*

**Advisor:**  
Mark Balas  
**Practicum:**  
Lawrence Livermore National Laboratory  
**Contact:**  
hindmanc@colorado.edu  
**Research Synopsis:**  
The course of research undertaken under this fellowship involves applying ideas from control theory to the problem of large-scale aeroelastic system simulation. Specifically, the goal is to develop algorithms and techniques for integrating the effects of active control surfaces on an existing 200,000+ degrees of freedom aeroelastic system simulation.



**Jason Hunt**  
*University of Michigan  
Aerospace Engineering &  
Scientific Computing*

**Advisor:**  
Kenneth Powell  
**Practicum:**  
Lawrence Livermore National Laboratory  
**Contact:**  
jhunt@engin.umich.edu  
**Research Synopsis:**  
The goal of my research is to extend current Computational Fluid Dynamics (CFD) techniques for solving moving geometry problems to three dimensions and to incorporate current developments in Cartesian mesh generation for component-based geometry and finite-volume schemes for the Euler equations. Both turbo-machinery and helicopters, two common motivational examples, have rotating components that are difficult to model with conventional techniques. This research will produce methods to obtain more representative simulations of

flow through compressors, rotor/stator configurations (e.g., stator blades remain motionless while the rotor blades sweep through the domain), and even flow about helicopters.



**Eric Lee**  
*Rutgers University  
Mechanical Engineering*

**Advisor:**  
Contantinos Mavroidis  
**Practicum:**  
Sandia National Laboratories-  
New Mexico  
**Contact:**  
chingkui@eden.rutgers.edu  
**Research Synopsis:**  
One of the most important spatial, task oriented robotic system design problems is the Rigid Body Guidance Problem. This is the calculation of the geometric parameters of a mechanical system so that it guides a rigid body in a number of specified spatial locations. This project will focus on the development of computational design algorithms for the spatial rigid body guidance problem with serial manipulators.



**Diem-Phuong Nguyen**  
*University of Utah  
Chemical Engineering*

**Advisor:**  
Philip Smith  
**Practicum:**  
Lawrence Livermore National Laboratory  
**Contact:**  
diem@crsim.utah.edu  
**Research Synopsis:**  
In the area of numerical simulation of accidental fires and explosions, there is no computational fluid dynamics (CFD) code available that calculates both turbulent reacting flows and complex chemical kinetics. Due to limitations of

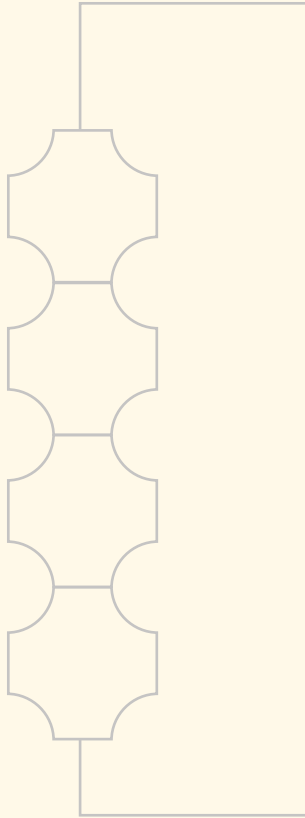
computational time and space, existing CFD codes available are codes which model only laminar flows and complex chemistry or turbulent flows and no chemistry or simple one-step chemistry. The most innovative aspect of my research then is that through introduction of a subgrid scale (SGS) reaction model, practical reacting flows may be accurately simulated. SGS models perform calculations which allow for simulation of practical reacting flow situations by bridging microscopic details to macroscopic domain.



**Christopher Oehmen**  
*University of Memphis  
Biomedical Engineering*

**Advisor:**  
Semahat Demir  
**Practicum:**  
Lawrence Livermore National Laboratory  
**Contact:**  
coehmen@memphis.edu  
**Research Synopsis:**  
My research deals with the modeling of highly nonlinear biological systems. Specifically, my aim is to create 2D and 3D models of pacing cardiac cell networks which take into account electrical activity of the cell membranes, as well as the electrical effects from diffusion and electric fields acting on ions in the tissue. My hypotheses are that 1) diffusion, which is conventionally considered trivial, is a crucial element of entrainment and propagation in the sinoatrial node (SAN), or the origin of pacing, and 2) incorporation of geometrical information into the diffusive and electric field calculations for electrical activity in the SAN will improve the ability of tissue level models to explain the experimentally observed effects of drugs and pathophysiological conditions on propagation in the SAN.

**Notable:**  
Invited to and attended the 2001 Noble Laureate Conference in Landau, Germany



## THIRD YEAR FELLOWS

**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

**Devin Balkcom**

*Carnegie Mellon University  
Robotics*

**Advisor:**  
Matt Mason

**Practicum:**  
Sandia National Laboratories  
– New Mexico

**Contact:**  
devin@ri.cmu.edu

**Gavin Conant**

*University of New Mexico  
Biology*

**Advisor:**  
Andreas Wagner

**Practicum:**  
Sandia National Laboratories  
– New Mexico

**Contact:**  
gconant@unm.edu

**Ryan Elliott**

*University of Michigan  
Aerospace Engineering*

**Advisor:**  
Nicolas Triantafyllidis

**Practicum:**  
Los Alamos National  
Laboratory

**Contact:**  
elliott@umich.edu

**Matt Fago**

*California Institute  
of Technology  
Aeronautical Engineering*

**Advisor:**  
Michael Ortiz

**Practicum:**  
Lawrence Livermore  
National Laboratory

**Contact:**  
fago@caltech.edu

**Catherine Grasso**

*Cornell University  
Bioinformatics*

**Advisor:**  
Golan Yona

**Practicum:**  
Lawrence Berkeley National  
Laboratory

**Contact:**  
cgrasso@cam.cornell.edu

**Boyce Griffith**

*New York University –  
Courant Institute  
Applied Mathematics*

**Advisor:**  
Charles Peskin

**Practicum:**  
Lawrence Livermore  
National Laboratory

**Contact:**  
griffith@cims.nyu.edu

**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

**Ahmed Ismail**

*Massachusetts Institute  
of Technology  
Chemical Engineering*

**Advisor:**  
Gregory Rutledge

**Practicum:**  
Oak Ridge National  
Laboratory

**Contact:**  
aei@mit.edu

**Benjamin Keen**

*University of Michigan  
Mathematics*

**Advisor:**  
Smadar Karni

**Practicum:**  
Lawrence Berkeley National  
Laboratory

**Contact:**  
bkeen@umich.edu

**Justin Koo**

*University of Michigan  
Aerospace Engineering*

**Advisor:**  
Iain Boyd

**Practicum:**  
Lawrence Livermore  
National Laboratory

**Contact:**  
kooj@engin.umich.edu

**Michael Kowalok**

*University of Wisconsin  
Medical Physics*

**Advisor:**  
Douglass Henderson

**Practicum:**  
Oak Ridge National  
Laboratory

**Contact:**  
mkowalok@students.wisc.edu

**Heather Netzloff**

*Iowa State University  
Physical Chemistry*

**Advisor:**  
Mark Gordon

**Practicum:**  
Lawrence Berkeley National  
Laboratory

**Contact:**  
netzloff@iastate.edu

**Catherine Norman**

*Northwestern University  
Applied Mathematics*

**Advisor:**  
Michael Miksis

**Practicum:**  
Lawrence Berkeley National  
Laboratory

**Contact:**  
c-norman@nwu.edu

**Robert Sedgewick**

*University of California –  
Santa Barbara  
Physics*

**Advisor:**  
Robert Sugar

**Practicum:**  
Lawrence Livermore  
National Laboratory

**Contact:**  
rds@physics.ucsb.edu

**Collin Wick**

*University of Minnesota  
Computational Chemistry*

**Advisor:**  
Joern Ilja Siepmann

**Practicum:**  
Oak Ridge National  
Laboratory

**Contact:**  
wick@chem.umn.edu

**Notable:**  
Awarded first prize in the  
2001 IBM Graduate Student  
Awards in Computational  
Chemistry by the American  
Chemical Society Division of  
Physical Chemistry.

## SECOND 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

**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

**Nouvelle Gebhart**

*University of New Mexico  
Chemistry*

**Advisor:**  
Martin Kirk

**Practicum:**  
Los Alamos National  
Laboratory

**Contact:**  
nouvelle@unm.edu

**Notable:**  
Invited to and attended  
the 2002 Noble Laureate  
Conference in Lindau,  
Germany

**Sommer Gentry**

*Massachusetts Institute  
of Technology  
Optimization/Control Theory*

**Advisor:**  
Eric Feron

**Contact:**  
sommerng@mit.edu

**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

**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.

**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 2001-2002.

**Richard Katz**

*Columbia University  
Geodynamics*

**Advisor:**  
Marc Spiegelman

**Contact:**  
katz@ldeo.columbia.edu

**Benjamin Kirk**

*University of Texas  
Applied Mathematics*

**Advisor:**  
Graham Carey

**Practicum:**  
Sandia National Laboratories  
– New Mexico

**Contact:**  
benkirk@mail.utexas.edu

**Seung Lee**

*University of Illinois, Chicago  
Mechanical Engineering*

**Advisor:**  
Francis Loth

**Contact:**  
slee48@uic.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  
– Livermore

**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

**Contact:**  
rtm@cs.wm.edu

**Julian Mintseris**

*Boston University  
Bioinformatics*

**Advisor:**  
Zhiping Weng

**Contact:**  
julianm@bu.edu

**Elijah Newren**

*University of Utah  
Mathematics*

**Advisor:**  
Aaron Fogelson

**Practicum:**  
Lawrence Livermore  
National Laboratory

**Contact:**  
newren@math.utah.edu

**Joyce Noah**

*Stanford University  
Theoretical Chemistry*

**Advisor:**  
Hans Andersen

**Practicum:**  
Los Alamos National  
Laboratory

**Contact:**  
jnoah@stanford.edu

**Christopher Rinderspacher**

*University of Georgia  
Chemistry*

**Advisor:**  
Peter Schreiner

**Practicum:**  
Sandia National Laboratories  
– Livermore

**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

**Contact:**  
maw@duke.edu

## FIRST 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

#### Contact:

mfbarad@ucdavis.edu

### Jaydeep Bardhan

Massachusetts Institute  
of Technology  
Electrical Engineering

#### Advisor:

Jacob White

#### Contact:

jbardhan@mit.edu

### Mary Biddy

University of Wisconsin  
Engineering

#### Advisor:

Juan de Pablo

#### Contact:

mary\_biddy@hotmail.com

### Nawaf Bou-Rabee

California Institute  
of Technology  
Applied and Computational  
Mathematics

#### Advisor:

Thomas Hou

#### 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  
Structures

#### Advisor:

Keith Hjelmstad

#### Contact:

kbergero@uiuc.edu

### Gregory Davidson

Oregon State University  
Nuclear Engineering

#### Advisor:

Todd Palmer

#### Contact:

davidsg@engr.orst.edu

### Michael Driscoll

Boston University  
Bioinformatics

#### Advisor:

James Collins

#### Contact:

driscoll@ncbi.nlm.nih.gov

### Mary Dunlop

California Institute  
of Technology  
Mechanical Engineering

#### Advisor:

Tim Colonius

#### Contact:

mjdunlop@alumni.  
princeton.edu

### Michael Greminger

University of Minnesota  
Mechanical Engineering

#### Advisor:

Bradley Nelson

#### 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

#### Contact:

yan+fellowship@cs.cmu.edu

### Benjamin Lewis

University of California –  
Berkeley  
Computational Biology

#### Advisor:

Steven Brenner

#### Contact:

ben\_lewis@post.harvard.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

#### Contact:

n\_r\_morgan@yahoo.com

### Gregory Novak

University of California -  
Santa Cruz  
Theoretical Astrophysics

#### Advisor:

Douglas Lin

#### Contact:

novak@ucolick.org

### Courtney Roby

University of Colorado  
Optics

#### Advisor:

Rafael Piestun

#### Contact:

robby@colorado.edu

### David Schmidt

University of Illinois  
Communications and  
Signal Processing

#### Advisor:

Bruce Hajek

#### 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

#### Contact:

jjw36@cornell.edu

### Michael Wu

University of California –  
Berkeley  
Computational Neuroscience

#### Advisor:

Jack Gallant

#### Contact:

wafting@uclink4.berkeley.edu

