2005 - 2006

DEIXIS

THE DOE CSGF ANNUAL

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

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

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

DOE CSGF

2005 - 2006
2005 - 2006
DEIXIS – THE DOE CSGF ANNUAL
DEIXIS – THE DOE CSGF ANNUAL

THE DOE CSGF ANNUAL

DEIXIS

TABLE OF CONTENTS

Department of Energy Computational Science Graduate Fellowship

2005 - 2006

Editor
Shelly Olsan

Copy Editor
Ron Winther

Design
Juls Design, Inc.

Contributors
Jacob Berkowitz
Alan S. Brown
Victor D. Chase
Karen Hede
Michael Szpir

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

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

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

DOE CSGF

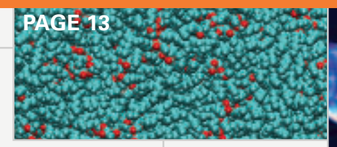

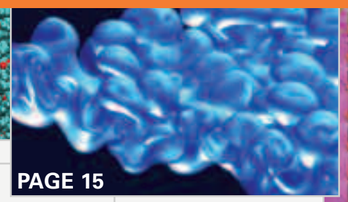

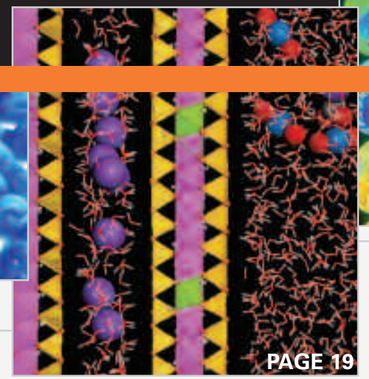

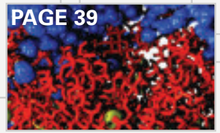

DEIXIS (ΔΕΙΞΙΣ) *transliterated from classical Greek into the Roman alphabet, (pronounced dāksis) means a display, mode or process of proof; the process of showing, proving or demonstrating.* DEIXIS *can also refer to the workings of an individual's keen intellect, or to the means by which such individuals,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

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

The Researcher's Roadtrip

THE DEPARTMENT OF ENERGY Computational Science Graduate Fellowship (DOE CSGF) program supports the brightest science and engineering students as they prepare for a future in their respective fields. Nearly 15 years old, the DOE CSGF has grown to over sixty active fellows and over 200 alumni.

Left to right: Teresa Bailey, Mary Biddy, Sam Stechmann and Bree Aldridge

TO A COLLEGE STUDENT, summer typically means warm beaches, lazy weekends, and no classes, but the DOE CSGF fellows live out a different kind of dream for the warmest three months of the year. During the summer, fellows take a practicum in one of the Department of Energy's laboratories — an experience that may impact the future of their research.

The opportunity for hands-on experience with world-class researchers and the latest technology is more than just a chance to hone analytical skills; it is the quintessential road trip of one's career (and maybe even self) discovery. This trip is an opportunity to see what you have become, and, perhaps, decide where you are going.

Measuring the Stroke of a Butterfly's Wing

BREE ALDRIDGE

Massachusetts Institute of Technology | Pacific Northwest National Laboratory | Story by Karyn Hede

More than 30 years ago, meteorologist Edward Lorenz poetically described how small events can have large consequences when he posed the question: can a butterfly flapping its wings in Brazil cause a tornado in Texas two weeks later? Lorenz' musing was no mere philosophical question. He proceeded to show mathematically how small changes in starting conditions can cause large effects in complex, interdependent systems.

Today biologists are teaming with mathematicians to apply quantitative approaches to complex biological systems that have much in common with Lorenz' meteorological conundrum. Department of Energy Computational Science Graduate Fellowship (DOE CSGF) Fellow Bree Aldridge has set her sights on measuring the molecular equivalent of Lorenz' butterfly effect in the complex milieu of the living cell.

Aldridge and her graduate advisors, Douglas Lauffenburger and Peter Sorger of Massachusetts Institute of Technology, Cambridge, Massachusetts, have teamed with Steve Wiley of Pacific Northwest National Laboratory (PNNL), Richland, Washington, to craft a large-scale mathematical representation of interacting cell signaling pathways.

"The general goal of studying these signaling networks is so that we can understand how cells process so many signals and always make the correct decision," says Aldridge. "Or another way of looking at this is for cancer therapy. Today, drug companies see protein A is allowing the cancer cell to survive when it should die, so they try to inhibit protein A. But maybe in some cells protein A doesn't matter or the cell can alter its pathway a bit to go around protein A. It may be that protein A or protein C or protein F is important. So inhibiting each of these pathways a little bit might be more effective than completely shutting down just one of them. That's the sort of information we hope to get someday out of these quantitative models."

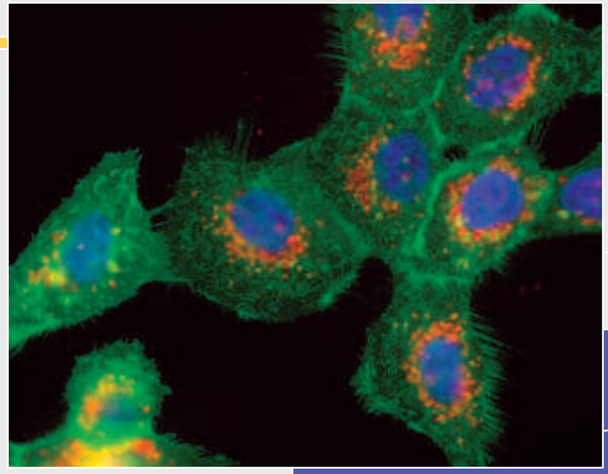

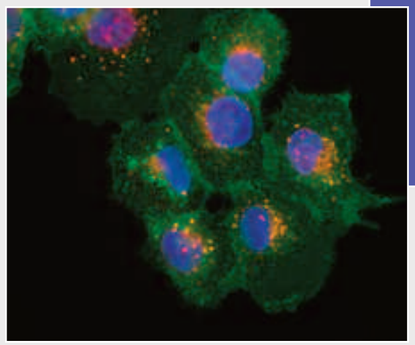

New imaging technologies can reveal the location of oncogenic proteins in living cells. Shown are human epithelial cells expressing the oncogene HER2 that emit green fluorescence. Lysosomes are stained red and the cell nucleus is stained blue for reference.

The idea is to eavesdrop on the cell's communication channels and discern, from seemingly indecipherable "chatter," patterns of signaling that cells use to make crucial decisions, such as when to divide or when it's time to die.

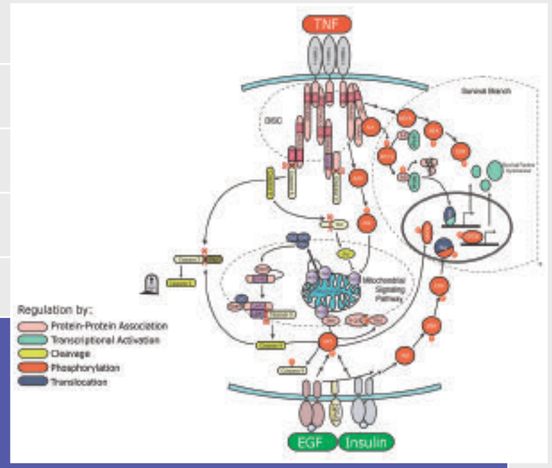

A schematic of the signaling networks induced by cytokine activation. Cells use this complicated and intertwined protein pathway to process extracellular cues into a life-or-death decision. (Drawn by Peter Sorger and Douglas Lauffenburger, MIT).

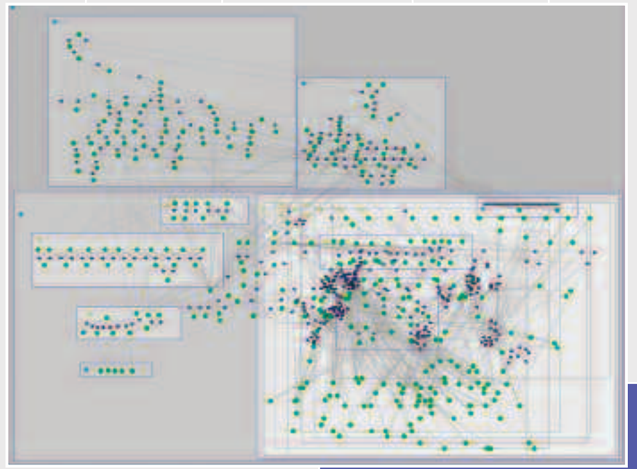

Maintaining larger ODE-based as a list of equations is increasingly difficult as the network models are expanded. In order to begin merging model together, we needed to represent each model in the same format (SBML) and annotate it using a graphical modeling tool (Teranode VLX Design Suite). This screen shot shows how complicated just a part of the model is! (graphic: Bree Aldridge)

The idea is to eavesdrop on the cell's communication channels and discern, from seemingly indecipherable "chatter," patterns of signaling that cells use to make crucial decisions, such as when to divide or when it's time to die.

"Every cell in your body sees information and makes a decision to do something, and its decision depends on its context, its state and its condition," says Wiley. "The cells make decisions based on information they derive from their extracellular environment. One of the things we've tried to figure out is how do cells know where they are. For example, cancer cells don't know where they are. Or they do know where they are and they just respond inappropriately. Everything a cancer cell does is perfectly normal to a cell in a different environment, but it's inappropriate to a cancer cell in its environment. If you understand why the cell is doing the wrong thing at the wrong time you have an approach to try to convince it to do the right thing, which for a cancer cell is 'stop growing; drop dead.'"

Wiley's laboratory studies the epidermal growth factor (EGF) receptor pathway, which is implicated in many types of cancer. Lauffenburger's group studies tumor necrosis factor alpha (TNFalpha), another pathway that often gets activated in cancers. EGF and TNFalpha are signaling molecules called cytokines that communicate "danger" signals to the cell, which typically causes it to cease growing and often to die.

"When we hit the cells with TNF, some of the cells die and some of them don't," says Aldridge. "We want to know, why is that? But these systems are so complex that you have to use modeling to try and understand what's happening."

Each research group had built its own model of how signals are transduced through the cell. Haluk Resat built the EGF model at PNNL, and Birgit Schoeberl and Suzanne Gaudet built the TNF model in the Lauffenburger and Sorger groups. Aldridge's task was nothing less than to combine the two models into a single model that incorporated both the EGF and TNF pathways, with all their various signaling molecules, networks and complex interactions. "We are using ODEs (ordinary differential equations), which describe mass action kinetics, to understand how cells make decisions," says Aldridge.

The project involved both computational and biological issues. For one thing, the models were built by use of different software, even different computer languages. The PNNL group programmed in Fortran, while Lauffenberger's lab used Matlab.

We simulate models of the signaling pathways to gain insight into the complex signaling resulting from extracellular cues (graphic: Suzanne Gaudet, adapted by Bree Aldridge)

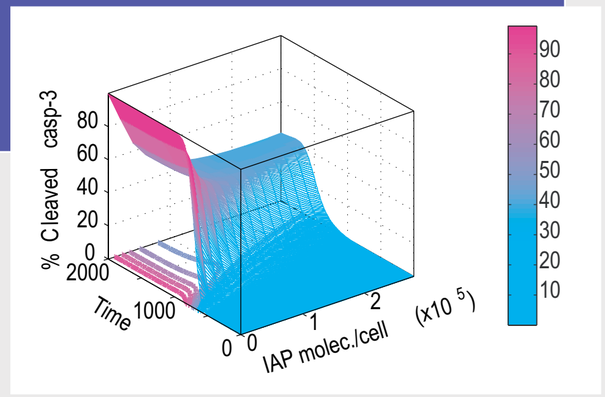

"There is no consistent language in systems biology," says Aldridge. "This makes sharing models difficult. One of my goals was to merge our two models together so we can continue and expand our collaboration."

Another problem was that the models were built based on data from two different kinds of cells. The Lauffenberger group studies a colon cancer cell line, while Wiley's group studies mammary cancer cells. The models the two labs developed were different, but complementary. And Aldridge wanted to make them work together.

To do that, she used two approaches. First, she transferred the two models to a single software package, a program called Teranode, first developed at the University of Washington, Seattle, and now available commercially.

"The size of the model that she is building is much greater than anything else we've ever seen," says Wiley.

But Aldridge soon discovered that her system was too complex to run on the existing software package, and she shifted her attention to reducing the number of interacting molecules required to run the model.

"We are using gene expression and proteomic data to filter the diagram and cull out what are the real connections and interactions," says Wiley. "One of the things Bree was involved in was taking a look at the wiring diagram and looking at our list of genes that we had identified as actually being involved and then, experimentally, looking to see if we knock out a particular piece, does it have the predicted effect on the wiring diagram?"

When the practicum was over, Aldridge had not completed merging the two models, but she had developed a better understanding of what it will take to do so and how the two labs can continue to refine the existing models by use of experimental results. For example, she had discovered that signals can get transduced through the system in different ways.

"I learned that we need to pay more attention to what's happening outside the cell," she said. "Instead of just saying, 'How do these pathways talk to one another inside the cell?' it may be as important to understand extracellular methods of cross talk between the pathways."

The butterfly's wing flutters once more.

PROGRAM REQUIREMENTS

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

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

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

Prior to the third year of the fellowship, fellows must complete a practicum assignment at a Department of Energy laboratory. Currently, approximately 15% of fellows who graduate from the DOE CSGF program work or have worked in a Department of Energy laboratory.

DISCIPLINES PURSUED

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

Modeling the Ocean's Invisible Waves

SAM STECHMANN

New York University - Courant Institute | Los Alamos National Laboratory | Story by Alan S. Brown

Stand on a beach and watch the waves. Driven by tides and sculpted by the islands and atolls that break and reform them, no two are alike. No wonder mathematicians find it difficult to represent the changing sea.

Such challenges — and the sea itself — were the furthest thing from Sam Stechmann's mind while he was growing up in Red Wing, Minnesota, a small town about an hour's drive southeast of Minneapolis/St. Paul.

As a boy, Stechmann, a 2004 Department of Energy Computational Science Graduate Fellowship (DOE CSGF) Fellow, was entirely devoted to sports, especially hockey. "Until fourth grade, I only read sports books," he recalls. "Every school project was about sports. My school invention was a hockey stick that curved both ways so you could shoot better backhand."

Finally, when he was in middle school, his mother made him enter a math contest to broaden his horizons. He surprised himself by winning a regional competition and continued to place throughout high school.

Although he entered the University of St. Thomas in St. Paul as a pre-med major, Stechmann had no idea what he really wanted to do. That changed in freshman chemistry, where he made friends with several physics majors who took him to see their laboratory projects. By the end of his sophomore year, Stechmann had switched his major to a combination of physics, chemistry, and math.

Stechmann's interest in combining physics and mathematics led him to the Courant Institute at New York University (NYU) to study fluid dynamics and computational science.

He also wondered how he would fare so far from home. Until he graduated from college, he had never been further than the Iowa border. His first real road trip, however, did not take him to New York. Instead, he received a summer internship at Los Alamos National Laboratory (LANL) in New Mexico. There he met Darryl Holm and learned about waves.

Submarines

The US Navy had asked Holm, a Laboratory Fellow and member of LANL's Computer and Computational Science Division, to investigate the ocean's internal waves.

On the surface, internal waves are all but imperceptible. Below, in the thermocline — the layer of water where temperatures drop rapidly — they are massive. Powered by tides, ranks of waves 90 meters high, 1 kilometer thick, and stretching hundreds of kilometers long race through the ocean for thousands of kilometers. From the Space Shuttle, they look like great lines on the sea.

When those waves reach a submerged submarine, the greater density of the cold wave changes the buoyancy of the sub. "There's a big jerk as the sub accelerates," says Holm. "You'd better be wearing a helmet if it's happening." Internal waves also reflect active sonar, the submarine's "eyes".

The Navy asked Holm to find a way for submarines to predict when internal waves would strike. "There's no way we could do a numerical simulation of every tide," he explains. "To do it, we needed some sort of approximate equations."

Solitons

To develop appropriate equations, Holm turned to a phenomenon first described by famed nautical engineer John Scott Russell in 1834. When a boat drawn through a canal suddenly stopped, Russell saw the water it had put in motion accumulate around the vessel's prow and then shoot forward as a solitary wave.

Classical theory held that such waves should quickly dissipate. This one did not. Russell galloped after it on horseback for more than a mile. It took 60 years before mathematicians explained how Russell's "waves of translation" could propagate over such long distances in shallow water.

Interest subsided until the 1960s, when two researchers discovered another property of the waves: They could pass through one another, transferring momentum but still retaining their shape. Because they acted more like solitary particles than waves, the researchers renamed them "solitons."

Since then, mathematicians and physicists have probed soliton properties. When Stechmann arrived after graduating from St. Thomas, Holm, an authority on solitons, had begun to apply these techniques to internal waves.

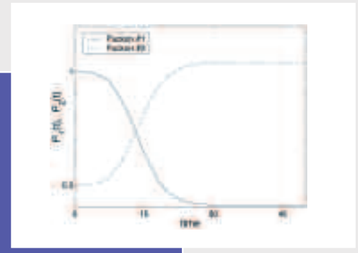

Numerical solution of the ODE for the puckon's height (velocity) as a function of time. Again, you can compare it with the PDE solution to see that the two agree.

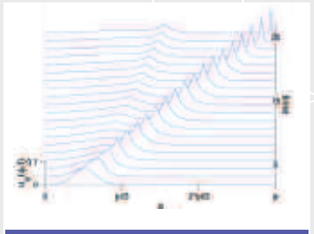

This picture shows how an initial bump breaks up into puckons. No matter what type of initial condition you pick, it will break up into puckons. The picture shows what shape the bump has as it evolves in time. The bottom is the smooth bump at the initial time. After awhile it sheds a puckon, which has a peak.

Planes

Holm asked Stechmann to find a way to visualize solitons on a two-dimensional plane. Past researchers had found that for any single point on a soliton, velocity is proportional to height. That made it possible to describe the velocity and height of a single point on a soliton as a function of one space variable, say "x."

Representing the arc of soliton waves moving across a plane is far more complex, since it requires two space variables, say "x" and "y." Yet Holm and Vakhtang Putkaradzehad of the University of New Mexico found a way to simplify the problem. They imagined the waves as concentric circles. While movement of these circles across a plane could be defined by their "x" and "y" position on a grid, the entire circle could also be described as a function of its distance from the center, i.e., its radius. By looking at only this one variable, radius, they vastly reduced the computer power needed to describe a solution.

Stechmann's goal was to make the model work. He started by playing with solutions that described solitons as a series of waves of different frequencies.

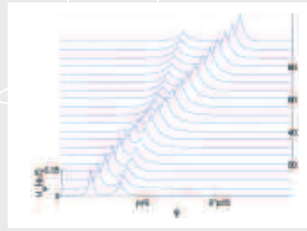

This is a rear-end collision of two puckons. A puckon's velocity is proportional to its height. The taller puckon catches the shorter one, they collide, and then they move apart – in the same way two billiard balls or Cadillacs would collide.

This generated waves that looked as if someone had dropped a hula-hoop in a shallow basin.

"We had difficulties when the radius of the waves generated inside the hula-hoop approached zero," says Stechmann. "The model didn't know what to do. It became very unstable and the code blew up." Stechmann believes he could have eventually puzzled out a solution. To get something done by the end of summer, he tried a different approach.

Stechmann turned to computational fluid dynamics (CFD), a technique that breaks a plane into grid points. To describe a soliton, he subtracted the momentum of one point from that of another. Since a soliton's velocity and height are related to its momentum, the resulting cascade of numbers approximated a wave. "It took longer but it worked," he laughs.

Spheres

After finishing his first year of graduate school, Stechmann returned to Los Alamos for his DOE CSGF practicum to work with Holm and his collaborator, Jonathan Munn of Imperial College, London, on solitons on a spherical surface. To most people, a sphere looks like a three-dimensional object. To Stechmann, it looked like a curved plane. He figured he could use the same approach he had tried the summer before.

"It didn't work," he explains. "When we created circles on the sphere, they move toward the pole and got smaller and smaller. Eventually, their height and momentum fell to zero at the pole, but not fast enough to work. Instead of the circles dying off, funny things happened at the poles and then propagated everywhere. They soon screwed up everything."

It took the summer for Stechmann to resolve the issue. He continued to use a grid approach, but he laid his grid out on a sphere, and created boundaries around both poles. Once a soliton passed beyond the boundary, it ceased to matter in the simulation.

The resulting model worked about as well as one from the previous summer. Visualizing soliton waves is only part of the larger attempt to predict the ocean's internal waves, but vital nonetheless. Stechmann's contribution has given researchers a way to visualize their equations.

Meanwhile, Stechmann has returned to NYU, where he continues to study applied mathematics. He also continues to play hockey, and last year helped his team place second in the national club hockey championships.

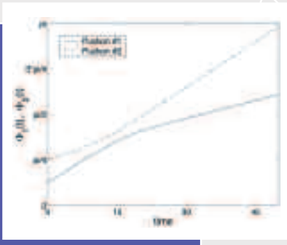

A numerical solution of the ODE for the puckon's latitude as a function of time. Solving the ODEs numerically is much faster and easier than solving the PDEs numerically. You can see that the two results agree.

SCOPE OF PROGRAM

Since its inception, the DOE CSGF program has supported nearly 225 students in more than 50 universities all over the U.S. Currently it supports 64 students in 22 states.

For nearly 15 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.

Every school project was about sports. My school invention was a hockey stick that curved both ways so you could shoot better backhand.

Faster Math Ignites Fusion Experiments

TERESA BAILEY

Texas A&M University | Lawrence Livermore National Laboratory | Story by Alan S. Brown

Teresa Bailey likes a challenge. A graduate student in nuclear engineering at Texas A&M University, College Station, Texas, and a DOE CSGF Fellow, she got one when she spent the summer of 2004 at Lawrence Livermore National Laboratory (LLNL), home of the National Ignition Facility (NIF), 50 miles east of San Francisco.

The multibillion dollar facility was built to create a fusion reaction that generates more energy than it consumes. And it has challenges aplenty for anyone in the field.

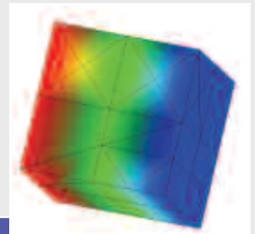

Example of a polyhedral mesh.

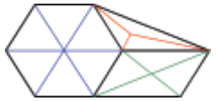

A 3 cell polygonal mesh.

Ignition and Lasers

NIF plans to fuse hydrogen isotopes in BB-sized pellets, using 192 high-powered lasers. To achieve successful ignition, though, researchers first need to optimize their equipment, materials, and fusion conditions.

Because NIF's hardware is far too expensive to build and rebuild for each test, researchers are using supercomputers to model what will happen when they turn on all 192 lasers.

These are not simple calculations. NIF's approach to fusion starts with fuel pellets of two hydrogen isotopes, tritium and deuterium. The pellets are surrounded by a crust of frozen isotopes encased in a plastic skin. Each pellet is then hung in a hohlraum, a chamber that resembles a very small soda can plated with gold on the inside and open at both ends.

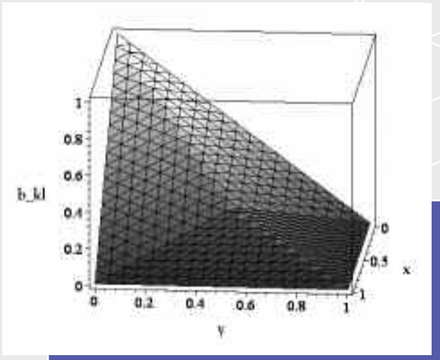

A piecewise linear (PWL) basis function in 2D.

When the laser beams strike the inner surface of the hohlraum, they instantly vaporize the gold into boiling plasma. The plasma surrounds the pellet, emitting high-energy X-rays. This heats the pellet so fast that it explodes inward, squeezing the deuterium and tritium together until they fuse.

Models

The model is as complex as the process. It must account for variables ranging from laser pulse, power, and energy fluctuations to the deuterium:tritium ratio of the pellet. It must also calculate the interactions of the laser beam moving through the plasma, the plasma generating the X-rays, the shockwave caused by the X-rays, and the thermonuclear reactions among particles. Everything happens within nanoseconds at extremely high temperatures and pressures.

To simulate what happens inside the hohlraum, LLNL is building massively parallel multiphysics computer codes as part of the Department of Energy's Advanced Simulation and Computing (ASC) Program.

The multiphysics part refers to the model's ability to simulate a vast range of physical properties simultaneously. The model is run on a massively parallel supercomputer, an IBM SP2 capable of training 8,000 powerful processors on a single problem.

Even with all that power, says Mike Zika, who heads a multiphysics model development effort, some simulations have lasted as long as four months. Clearly, the lab would welcome any approach that saves calculation time.

Finding Physics

Enter Bailey. As a high school student in Anchorage, Alaska, she thought she might want to study English or political science in college. She even interned for Alaska Senator Ted Stevens the summer after she graduated.

When it came time to pick a college, though, Bailey wanted more educational options than Alaska offered. She opted for nuclear engineering at Oregon State University. "I really liked what I had learned about physics in high school, and my sisters were already there," she explains.

These problems are really a challenge to solve, but you can use the results in everything from astrophysics and nuclear reactor design to shooting a beam of photons at a tumor.

The program was small, and that proved to be a good thing for Bailey. "The faculty cared about the students," she recalls. "There was a lot of interaction. They challenged us and they trained us well. I liked it right from the start."

In her senior year, she became fascinated with radiation transport, how radiation interacts with a medium as it moves through it. "These problems are really a challenge to solve, but you can use the results in everything from astrophysics and nuclear reactor design to shooting a beam of photons at a tumor," she explains.

She went on to study radiation transport in graduate school at Texas A&M. There, she worked closely with Marvin Adams, who developed a faster method to calculate radiation diffusion during fusion. In the summer of 2004, Bailey joined Zika, another Adams protégé, to see if this new method would work.

The Task

To understand Bailey's task, consider radiation diffusion in the hohlraum after ignition. The plasma inside the hohlraum gives off X- rays. These energize the electrons of surrounding atoms, scattering some and causing others to emit more X-rays, which then interact with more electrons.

One set of equations describes how the radiation diffuses in time and space. A second set reveals how it changes the energy of the surrounding matter. Unfortunately, both sets of equations are partial differential equations (PDEs). "Computers don't know how to solve PDEs," says Zika. "They only know how to deal with discrete operations like addition, subtraction, multiplication, division, and exponents. To solve PDEs, we have to translate them into discrete operations. For a good simulation, you need lots of points, and calculating all those points eats up computer time. There are many ways to make these approximations, and they all have advantages and disadvantages."

Communication Calculations

The disadvantage of the existing diffusion model is that it is equation-intensive. In order to create enough points for an accurate simulation, the model divides space into hexahedrals, the corners of which form the points. This creates a matrix with millions or even tens of millions of points.

Each point communicates only with its nearest neighbors to describe the interplay of energy and matter. Typically, each point is connected to 27 other points by means of one or more equations that define how each pair of points communicate with one another.

The existing diffusion model is asymmetric: The equation that describes how Point A talks to Point B is different from the equation that describes how Point B talks back to Point A. This means that each point on the grid needs separate equations for each of the 27 other points with which it communicates. Multiply that by millions of points on the grid and it is no wonder the calculations tax even the largest supercomputer.

Tackling the Code

Adams, says Bailey, found a way to make the problem symmetric, so that one equation communicates between two points in both directions. This halved the amount of memory needed to run a simulation. "In really large systems, this is a big deal," says Bailey. "We can either run simulations faster or increase their accuracy in the same amount of time."

Bailey's job was to program the new approach into the radiation diffusion code. "Teresa had never coded in C++ or used massively parallel supercomputers," says Zika. "This went well beyond her graduate school experience, where she and two or three others might work on a few thousand lines of code. Here, she was working with 20 of us and hundreds of thousands of lines of code."

Bailey not only completed the code, but began testing it as well. "Both methods work really well," she says. "It's extremely hard to find areas where they get bad answers."

For Bailey, it was a chance to tackle a challenging problem and learn new skills. "I think it broadened her horizons and gave her a real flavor of what's out there for her when she graduates," Zika concludes.

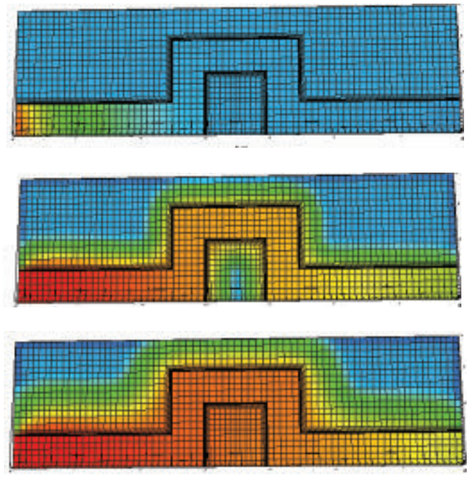

Top Hat Plots

DOE CSGF HIGHLIGHTS

- ***Payment of tuition and required fees***
- ***Yearly stipend of $31,200***
- ***A $1,000 yearly academic allowance***
- ***Matching funds of up to $2,500 for a computer workstation purchase***
- ***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

Salad Days for Plant-based Lubricants

MARY BIDDY

University of Wisconsin | Sandia National Laboratory | Story by Karyn Hede

Should you go with canola or corn, safflower or soybean? It sounds like a question about salad dressing or perhaps frying oil, but in Mary Biddy's research the ratio of vegetable oils is key to whether they will perform in car engines. Yes, that was *car engines*.

Biddy and her graduate advisor, Juan de Pablo, Professor of Chemical and Biological Engineering, University of Wisconsin, Madison, Wisconsin, are trying to perfect the formula for a vegetable oil that will perform as well as, or perhaps even better than, motor oil in internal combustion engines and other industrial applications, without polluting the environment. Their research could help the United States to wean itself from dependency on petroleum products. But to compete with petroleum-based lubricants, often called mineral oils, vegetable oils must prove themselves the equal of mineral oils over a wide range of temperatures and must withstand the friction of moving engine parts.

It turns out that vegetable oils acquit themselves quite nicely in most of the head-to-head comparisons with conventional motor oils, except for one key feature. Vegetable oils, which are composed of triglycerides, comprised of glycerol linked to fatty acids, turn to sludge when they get too cold.

This property, called the oil's pour point, is crucial to making vegetable oils practical for use in cars.

"It's a very important property because someone in Wisconsin, for instance, doesn't want the lubricant in the car to have a pour point temperature that is warmer than the temperature outside," says Biddy. "One of the main disadvantages of these vegetable oil-based lubricants is this pour point property. Petroleum and synthetic oil have pour points of –40ºC (-40ºF) while vegetable oil's pour point is –18ºC (0ºF). This is one property that has to be changed if vegetable oil is ever to be used in cars. Our idea is that this pour point temperature occurs because there is ordering that occurs on a molecular level in these triglycerides. The idea is to understand what is happening on a molecular level, so we can engineer a lubricant that has a lower pour point temperature."

Biddy and de Pablo use molecular dynamics simulation techniques to understand the potential for localized ordering of liquid triglycerides. They had run simulations that suggested local ordering of triglyceride molecules in the liquid phase, but it was hard to make a case for ordering with the small number of molecules they had been able to simulate using the computational resources at Wisconsin.

That's why Biddy's summer practicum at Sandia National Laboratory (SNL) in Albuquerque, New Mexico, was so valuable to her research. Using a computer code called LAMMPS – Large-scale Atomic/Molecular Massively Parallel Simulator — developed by Steven J. Plimpton at Sandia, Biddy was able to simulate her system on a much larger scale, both in terms of number of molecules and time frame of the simulation.

"It allowed us to see that on a larger scale these systems do tend to have some localized ordering and that's something that we are currently investigating," she says. "We want to understand what effect the local ordering has on the pour point transition."

Working with SNL scientist John Curro, Biddy developed a new technique for approximating the molecular forces the triglyceride molecules exert on one another. "We developed a coarse-grained model, which means we developed a method for lumping the atoms into "super-atoms" for the purposes of modeling. When we can lump many atoms into one "bead" we can model things on longer time scales. Using an atomistic model, we could only model 40 molecules, but with the coarse-grained model we can begin to model larger systems with shorter computation times to predict viscosity, which is an important property of lubricants."

For Curro and his SNL colleagues, the collaboration with Biddy helped validate their approach to simulating complex polymers.

"We are interested in simulating the behavior of high molecular weight polymer mixtures and blends for basic research and defense applications," says Curro. "Simulating triglycerides, which are low molecular weight, helped us compare PRISM theory (a statistical mechanics theory) with Mary's calculations. It helped validate the theory. It's clear both of us benefited."

In another project, Biddy worked with SNL's Laura Frink and Amalie Frischknecht to calculate the behavior of large numbers of triglycerides by use of classical density functional theory,

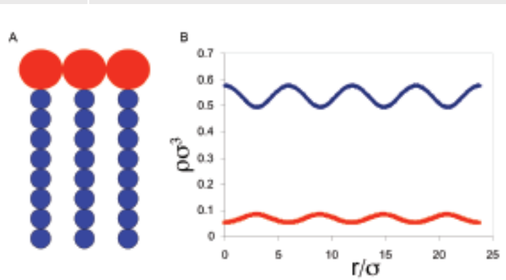

DENSITY FUNCTIONAL THEORY SIMULATIONS
(a) A schematic of the coarse grained triglyceride molecule that has been studied
(b) Density profiles of a triglyceride bilayer near the order-disorder transition

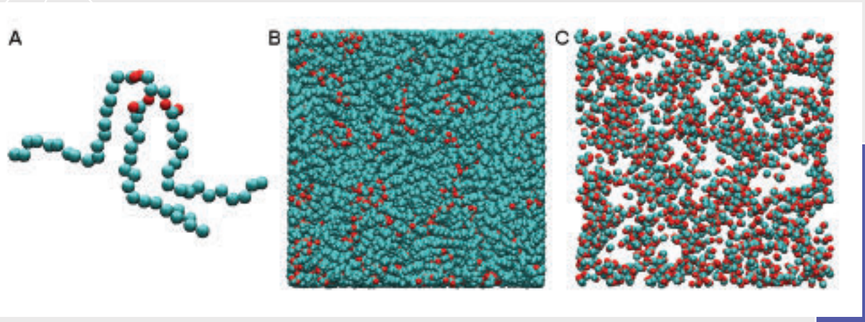

UNDERSTANDING THE POUR POINT TEMPERATURE OF VEGETABLE OILS
(a) The molecular structure of an individual triglyceride. The snapshots of the simulation box of a canola oil below its pour point temperature display (b) all the atom sites and (c) only the carbonyl groups of the triglyceride molecules. It is believed that the clustering of the carbonyl groups leads to the pour point in vegetable oils.

a method developed at SNL to model lipid bilayers, highly-ordered collections of fatty molecules.

"We wanted to see if triglycerides can form bilayers," says Biddy. "There is no experimental evidence that they do, but it is a difficult parameter to measure — you can't see it in experiments. The only way to understand the structuring is by modeling. We modified Laura Frink's model to see if we could understand how electrostatic interactions, partial charges on the molecule, affect ordering. If we include partial charges or leave it off in the model, does it change the ordering?"

That project is still ongoing, and results are not in yet. Part of the reason, says Biddy, is that she took on four projects during her practicum. "In retrospect, it probably wasn't realistic to expect to finish them all," she admits.

But one problem that did have an immediate payoff was her participation with SNL scientist Marcus Martin (a former DOE CSGF fellow) to predict the vapor pressures and heats of vaporization for acetone and butyramide under specific conditions that mimic those of industrial applications. The project was part of an international contest called the Industrial Fluid Properties Simulation Challenge.

"I was lucky to be able to work on another project that fits with my thesis work in that it looks at the industrial relevance for modeling simulations," says Biddy.

The contest involved groups of scientists in research laboratories around the world who were asked to develop simulation methods to predict physical properties of defined chemical agents. Working separately, scientists at the National Institute of Standards and Technology (NIST) and Dow Chemical, the contest's sponsors, made experimental measurements of the properties, which were then used to judge the contestants' predictions.

"Mary instigated the project and quickly came up to speed using a new (to her) simulation code called Towhee, a publicly available software program that uses Monte Carlo simulation to predict complex chemical systems," says Martin. "I added a few new features to enable us to pursue this research. We then divided up the simulation work as I did one of the molecules while Mary did the other. We managed to take second prize. Not bad for six weeks of work."

Biddy and Martin received their award and a cash prize at a special ceremony during the American Institute of Chemical Engineers annual meeting in 2004.

Martin and Biddy published their results in a paper titled "Monte Carlo Molecular Simulation Predictions for the Heat of Vaporization of Acetone and Butyramide," in the journal *Fluid Phase Equilibria.*

"It was challenging to work on a truly difficult real-world problem and to be exposed to the whole broader world of molecular modeling," says Biddy. "It helped me to realize that there are scientists who think about problems in a really different way than I have learned in my graduate training. It gave me a lot to think about when I start applying for jobs."

To compete with petroleum-based lubricants, often called mineral oils, vegetable oils must prove themselves the equal of mineral oils over a wide range of temperatures and must withstand the friction of moving engine parts.

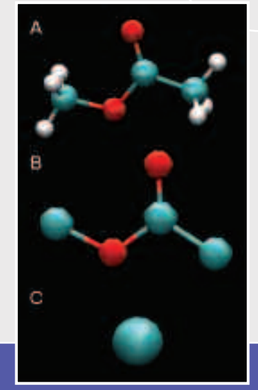

MODEL REPRESENTATIONS OF METHYL ACETATE
Methyl acetate at different levels of molecular detail: (a) all-atom where all atom sites are modeled, (b) united-atom where the methyl groups are represented as a single group, and (c) coarse-grained where all atoms are represented as a single super-atom.

By Michael Szpir

Inside a Flame

SOME ANTHROPOLOGISTS BELIEVE that our ancestors were tinkering with fire in a controlled way nearly 1.5 million years ago. It's been called our oldest technology, and we rely on it dearly. About 85 percent of the energy in the United States comes from burning fossil fuels. What's amazing is that we still don't fully understand the process.

"One of the things that's not well appreciated by most people is the complexity of a flame," says John Bell, group leader of the Center for Computational Sciences and Engineering at Lawrence Berkeley National Laboratory (LBNL) in Berkeley, California. Over the past 10 years, Bell and his team have put together a set of numerical algorithms that allows them to simulate a laboratory-scale flame to see what's going on inside. It's not what you might imagine.

When Bell looks at a flame, he sees a "reacting flow" — a fluid consisting of a number of chemical species that react as they move through space. Consider a simple methane flame, the subject of Bell's simulations. In high school you may have learned that burning methane with oxygen gives you heat, carbon dioxide and water. In reality, the process of burning methane involves about 50 chemical species and more than 300 chemical reactions.

"There are a lot of chemicals with short lives that are created and destroyed inside the flame," says Bell, "and the details of that chemistry make all the difference in how the flame propagates." Knowing the details is crucial if you want to design a combustion-based system — say, a turbine or an internal-combustion engine — that burns cleanly and efficiently. "You can't accurately describe either the detailed motion of the flame or the emissions from an engine without a detailed understanding of what's happening inside the flame," says Bell.

Unfortunately, experimental scientists who study real flames in the laboratory can provide only part of the picture. "Flames are hot. They're dynamic. It's hard to take measurements," Bell says. "You can't stick a meter into a flame and read off all the different chemicals in the flame at that point." Numerical simulations can tell you all about the chemistry in the flow, but these kinds of simulations wouldn't be possible without a whole suite of new mathematical tools.

"The traditional way to simulate a flame would be based on the equations of compressible fluid mechanics, as well as equations for all the chemical reactions, and transport, which represents how the chemical species diffuse inside the flame," explains Bell. "You can send all that to a supercomputer and let it grind away, but with the existing computer hardware it would take years to get a solution."

Bell and his colleagues recognized a way to sidestep the equations associated with compressible fluid mechanics, which eat up a lot of computing time. Those equations track the sound waves in a fluid very accurately, but it turns out that's not necessary. The fluid velocity in a laboratory-scale flame is only about 10 meters per second (and the flame moves even more slowly), whereas a sound wave travels about 300 meters per second. "You can show that in many cases the propagation of sound waves doesn't affect the system," Bell says. "There's a separation of scales, so you don't need compressible fluid mechanics."

Instead, Bell's team has devised a numerical code based on a "low Mach number formulation," so called because it describes fluids that are moving much more slowly than the speed of sound. (The Mach number is the ratio of the fluid speed to the speed of sound.) The low Mach number model is more complicated mathematically than the equations for compressible flow.

Over the past 10 years, Bell and his team have put together a set of numerical algorithms that allows them to simulate a laboratory-scale flame to see what's going on inside. It's not what you might imagine.

"The simplest low Mach number model is the incompressible Navier-Stokes equations, which describe how a fluid moves when it can't expand or contract," explains Bell. "The low Mach number model that we use for combustion is a generalization of those equations. But what all low Mach number models have in common is that, because they don't resolve the sound waves, they can advance the system in time with a much larger time step than the equations for a fully compressible flow."

This approach saves dramatically on computing time, running about a hundred times faster than standard codes. It's a crucial savings that begins to make detailed simulations possible, but it's not enough to accurately model a real flame. Fortunately, there's another way to save even more computing time. This method exploits a feature of the flame that is not obvious at first glance.

Flaming Illusion

When you look at a flame, it appears to have a very smooth, tapered shape. "That's just an illusion, an average of the flame's real appearance," says Bell. If the eye and the brain could parse time into nanoseconds we'd see a much more interesting object, one with kinks and wrinkles and corrugations.

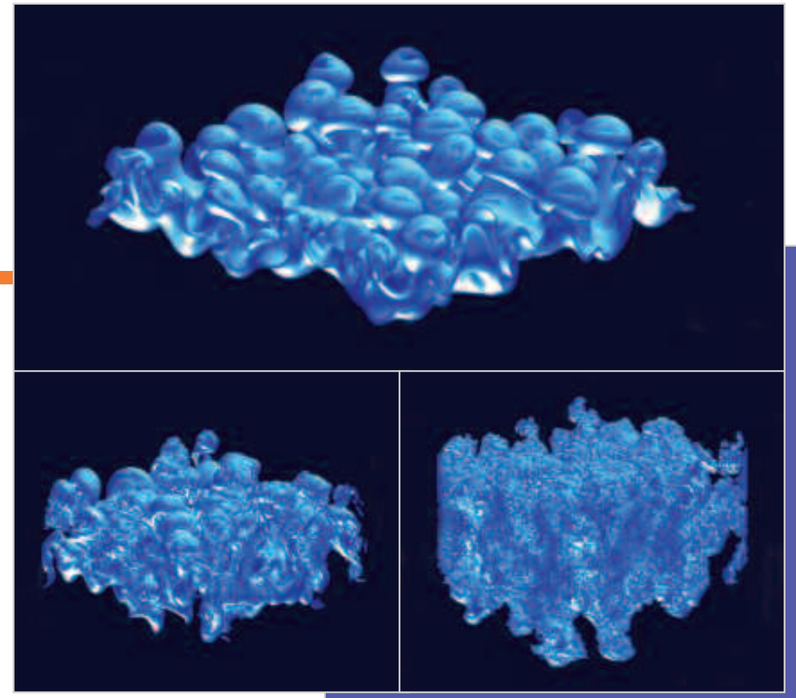

Sequence of images of a Rayleigh-Taylor unstable nuclear flame in a Type Ia supernova. In the simulation the flow transitions from laminar to turbulent. As the flow changes, the mode of combustion changes from the "flamelet" regime to "distributed" combustion. Simulations of this type play a key role in validating subgrid models required for full-star simulations.

That wrinkly object is also oscillating very quickly, giving the impression that a flame is much thicker than it really is. A flame that looks like it's about 10 millimeters across is really a very thin sheet, roughly one millimeter thick, which is flopping all over the place. The thinness of a flame has consequences for how the scientists use their precious computing time. "You want to focus your computational power where the flame is," says Bell, "not centimeters away where not much is going on."

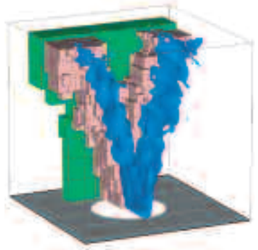

Image of simulated V-flame showing the location of the grid patches used by the adaptive mesh refinement algorithm. Each box represents a number of computational cells aggregated into a grid to discretize a local region of space. The green boxes indicate one level of refinement of the base grid and the pink boxes indicate an additional level of refinement. (Some boxes are removed to show the flame surface.)

JOHN BELL

John Bell received his B.S. (1975) degree from the Massachusetts Institute of Technology and his M.S. (1977) and Ph.D. (1979) degrees from Cornell University, all in mathematics. He is currently a Senior Staff Scientist and Group Leader for the Center for Computational Sciences and Engineering at Lawrence Berkeley National Laboratory. Prior to joining LBNL, he held positions at the Lawrence Livermore National Laboratory, Exxon Production Research and the Naval Surface Weapons Center.

Bell's research focuses on the development and analysis of numerical methods for partial differential equations arising in science and engineering. He has made contributions in the areas of finite difference methods, numerical methods for low Mach number flows, adaptive mesh refinement, interface tracking and parallel computing. He has also worked on the application of these numerical methods to problems from a broad range of fields including combustion, shock physics, seismology, flow in porous media and astrophysics. His group's web page is at http://seesar.lbl.gov/ccse/index.html.

Dr. Bell was the recipient of the 2005 Sidney Fernbach Award from the IEEE Computer Society which was presented at Supercomputing 2005 in Seattle, WA.

Further Reading:

J. B. Bell, M. S. Day, I. G. Shepherd, M. Johnson, R. K. Cheng, J. F. Grcar, V. E. Beckner, M. J. Lijewski *Numerical Simulation of a Laboratory-Scale Turbulent V-flame*, LBNL Report LBNL-54198-Journal, Proc. Natl. Acad. Sci. USA, 102(29) 10006-10011, 2005.

M. Zingale, S. E. Woosley, C. A. Rendleman, M. S. Day, and J. B. Bell *Three-dimensional Numerical Simulations of Rayleigh-Taylor Unstable Flames in Type Ia Supernovae*, LBNL Report LBNL-56966, accepted, Astrophysical Journal, Vol. 632, page 1021, October 2005.

M. S. Day, J. B. Bell, J. F. Grcar, V. E. Beckner *Numerical Control of 3D Turbulent Premixed Flame Simulations*, LBNL Report LBNL-56882 Ext. Abs., Proceedings of 20th International Colloquium on the Dynamics of Explosions and Reactive Systems, July 31-August 5, 2005.

Contact:

John Bell
JBBell@lbl.gov

PRACTICUM COORDINATOR

Joseph Grcar ..... JFGrcar@lbl.gov

That is important when you consider how a reacting flow is modeled in a computer. Computational scientists divide the region, or "domain," they want to simulate into millions of small boxes, or "grid cells." Each grid cell carries information about the fluid velocity, its temperature and density, as well as the relative amounts of each of the chemical species — adding up to about 25 variables in a typical methane-flame calculation. Each of these quantities must be updated in every cell at every time step, adding up to billions of calculations. The smaller each grid cell, the more accurate the calculation. But the more grid cells there are, the more work the computer must do.

Using a technique called "adaptive mesh refinement," the simulations performed by Bell's team use different sized grid cells in different regions of the flame. The middle of the flame contains small grid cells that allow for accurate resolution of the detailed chemistry there, but less resolution is needed outside the flame. By not using tiny grid cells in the whole domain, the researchers can reduce the total number of grid cells that the computer must track.

"The adaptive mesh algorithm automatically adds and removes grid cells based on what it thinks it needs to be accurate," says Bell. "At its peak, the calculation has about 40 to 50 million active grid cells." That's a huge number, but it turns out to be just a tenth of what the simulation would need without adaptive mesh refinement.

The flame that Bell and his colleagues conjure in the computer lasts for only a few tenths of a second, but the group's simulation is considered to be the best of its kind. "Ours is the only simulation of a real three-dimensional turbulent, pre-mixed laboratory-scale flame," says Bell. "It was possible because we developed code technology that's about three or four orders of magnitude faster than traditional simulation approaches."

A Burning Love

Bell's work on reacting flows has already been used in a number of applications, including the measurement of nitrogen-oxide emissions, simulations of vortex-flame interaction and, perhaps surprisingly, the study of supernovae (see sidebar: *Star Fire*). As a mathematician, Bell gets a tremendous amount of pleasure from seeing his work applied to these problems. "When you study partial differential equations in traditional mathematics, you might prove there is a solution, or come up with some complex, infinite series solution, but it's not tangible," he says. "In computing you use all sorts of sophisticated mathematical ideas to create a new algorithm, but you end up with something that is more substantial. That's neat."

The success of the reacting flow simulations has spawned a close collaboration with experimental scientists who study real flames in the laboratory. "We've gotten to the point where there's actually a lot of synergy," says Bell. "The experimentalists are starting to say, 'How can we provide the right experimental data to advance the computational study of these things?' It's a very exciting time." And he adds, "But you have to be careful what you say, because they're likely to go to the lab and do an experiment!"

Bell has worked in several different fields since his career began in the late 1970s, including aeronautics, seismology, and petroleum reservoir simulation, but his work on combustion at LBNL has provided a unique level of collaboration. "This is the first time in my career that there was such a good partnership between the people who do experiments and the people who do computation," he says.

Formerly at Lawrence Livermore National Laboratory, Bell started working at the national laboratories because they were among the few places that had the computer hardware to tackle complex problems in computation, but he ended up staying because of the people. "When I first began working in the field, the application of numerical methods for differential equations was in its infancy. A lone investigator could program a state-of-the-art method in a few days," he recalls. "Now we have very sophisticated nonlinear algorithms like the low Mach number model, and we do calculations on hundreds of processors on a fairly complex parallel architecture. It's more than one person can do. At LBNL I can assemble a group of people who have all the different kinds of expertise one needs to do cutting-edge research."

There are seven people in Bell's group, as well as two students from the DOE CSGF program for the summer of 2005: Jasmine Foo from Brown University, Providence, Rhode Island, and Amber Sallerson from the University of North Carolina at Chapel Hill, Chapel Hill, North Carolina. Both are working on problems related to fluid mechanics, although neither was involved in the flame simulation.

If you were to visit LBNL you might find Bell walking around the halls or standing at a whiteboard, talking over ideas with his colleagues, but often he just sits at his desk and works on problems. "I enjoy doing the computation and getting involved in the coding," he says. "To an outside observer it may look like I just scribble on a whiteboard or a yellow pad, punch something into my workstation and drink a lot of coffee." He pauses, then adds, "But I love it."

All the scribbling has resulted in some important contributions to applied mathematics. In July 2003, Bell was recognized by the Society for Industrial and Applied Mathematics and the Association of Computing Machinery with the first SIAM/ACM Prize in Computational Science and Engineering for his "outstanding contributions to the development and use of mathematical and computational tools."

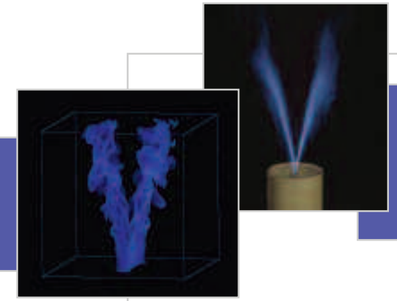

Image of simulated V-flame showing the instantaneous location of the flame.

Photographic image of a laboratory V-flame, courtesy of Robert Cheng, EETD/LBNL. The image is taken with a shutter speed of 1/30 sec.

STAR FIRE

>> ***A chance meeting at a scientific conference in 2002 has united two seemingly unrelated endeavors: terrestrial combustion research and the study of supernovae. John Bell had been standing by his presentation during a poster session in a large conference room when, during a lull, he casually asked the people at the poster next to his about their work. They happened to be two astrophysicists, Stan Woosley and Mike Zingale, from the University of California, Santa Cruz.***

As Bell recalls, "They talked about their research on supernovae and explained how hard it was because the flames inside the stars move so slowly. So I said, 'Why don't you do this with a low Mach number model?' They asked, 'What's that?' And I said, 'We do these calculations of flames and we use an algorithm that doesn't track acoustic waves. Don't people do that in astrophysics?'" Well, it turns out they hadn't, but now they do.

On the face of it, it's hard to imagine there's any similarity between a modest laboratory flame and a type Ia supernova — an annihilating explosion of a white-dwarf star that releases about 10^{44} joules of energy and ejects matter at speeds of up to 10,000 kilometers per second. As it happens, both processes involve reacting flows that travel well below the speed of sound.

Just before a white dwarf obliterates itself, a thermonuclear flame ignites deep inside the star. The flame begins at subsonic speeds, but it accelerates rapidly, ultimately propagating throughout the star and causing it to explode. Although the reactions in the white dwarf involve nuclear chemistry — the fusion of nuclei — the reacting flows are similar mathematically to the laboratory flame.

In collaboration with Woosley and Zingale, Bell and his group have been able to extend their low Mach number code to describe events inside a white dwarf. The joint work has provided both a better understanding of how the flame accelerates and a new tool for astrophysicists who study these types of stellar flows.

By Victor D. Chase

Radioactive Retention

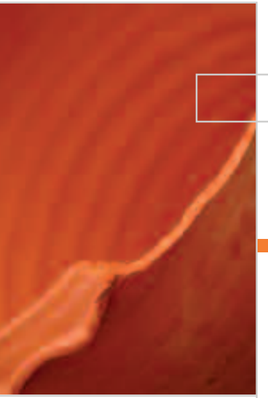

NUCLEAR POWER, be it for weaponry, electricity generation, or medical applications, is a messy business. First comes the mining of uranium ore, a process that produces radioactive tailings that can be dangerous if not dealt with properly. Then there is the processing of the uranium into nuclear fuel or weapons grade material which creates radioactive residue. And finally there is the spent product, which after serving its intended purpose must be disposed of in a safe and secure manner.

Overseeing the life of this nuclear material from mine to burial is the responsibility of the U.S. Nuclear Regulatory Commission (NRC), which is charged with preventing radioactive material from entering the environment.

One of the NRC's major concerns is the potential for radioactive wastes to leach into ground water, thereby affecting plants, animals and humans. Such leaching has already taken place at some abandoned uranium mine sites, primarily in the Western United States.

The process by which this radioactive waste has entered the environment is described in a U.S. Department of Energy (DOE) document entitled, "Uranium Mill Tailings Remedial Action" (web.em.doe.gov/bemr96/umtra.html). It states, "During the 1950s and 60s, private firms processed most uranium ore mined in the United States for the Atomic Energy Commission, a predecessor of the Department of Energy. The processing plants were shut down, and the tailings piles from mill operations were abandoned. These sites presented a potential long-term health hazard because they contained low-level radioactive and other hazardous substances that migrated to surrounding soil, ground water, and surface water. Furthermore, the piles often emitted radon gas. The tailings and other contaminated material were also used as fill dirt or incorporated into various construction materials at thousands of offsite locations."

Responding to public concern about these sites, Congress passed the Uranium Mill Tailings Radiation Control Act in 1978. It called on the DOE to "stabilize, dispose of, and control uranium mill tailings and other contaminated material at 24 uranium mill processing sites and approximately 5,200 associated vicinity properties." Responsibility for regulating and general oversight of the cleanup fell to the NRC.

An Exceptional Sponge

And that is where Randall Cygan and Jeffery Greathouse, of Sandia National Laboratories' Geochemistry Department, enter the picture. Cygan, a distinguished member of the technical staff, heads a computational sciences project aimed at determining how well minerals in the environment, especially in clay, can adsorb radioactive material. His group has focused their molecular simulations on clay minerals because of the exceptional ability of these materials to "sponge up" radioactive substances that have dissolved in water, and because clay minerals are commonly found in soils around abandoned mines. Their particular charge was to help determine how to deal with radioactive waste from one of the 24 sites covered by the Congressional Radiation Control Act — this one a mining site near Naturita, a small town in the San Juan Mountains of southwestern Colorado.

[Cygan's] group has focused their molecular simulations on clay minerals because of the exceptional ability of these materials to "sponge up" radioactive substances that have dissolved in water, and because clay minerals are commonly found in soils around abandoned mines.

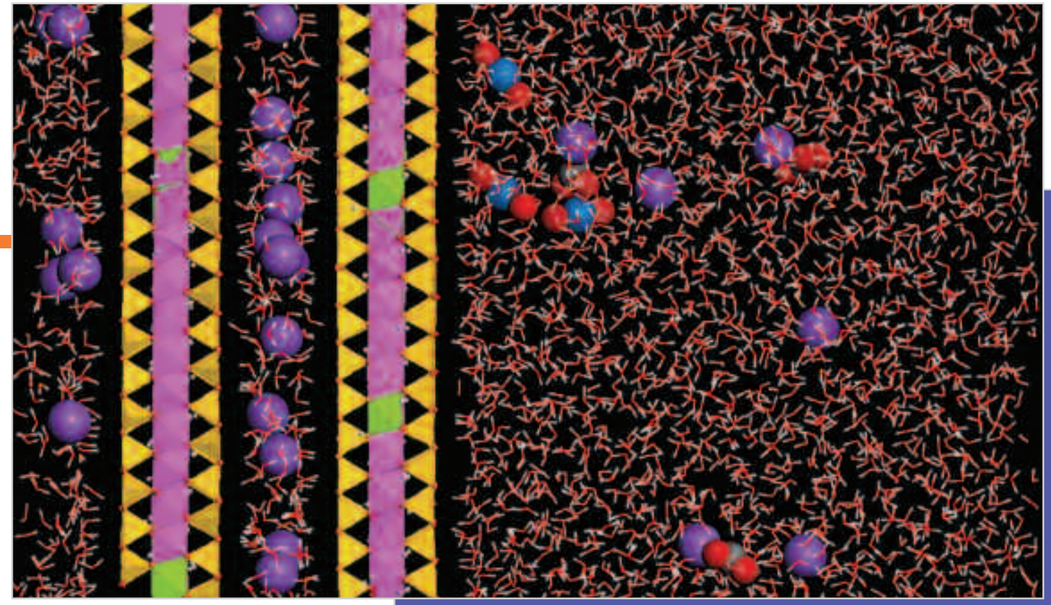

Snapshot from a uranyl-clay adsorption simulation. SiO_4 and AlO_6 groups in the two clay sheets are shown as tetrahedral and octahedral, respectively. Negative charge sites within the clay sheets are shown as green MgO_6 octahedra. The aqueous regions are composed of water molecules (red/white), sodium ions (purple), and uranyl ions (red/blue).

"A lot of waste was just dumped into waste hills at the Naturita mill plant that was abandoned in the '60s," says Cygan. And though some subsequent cleanup efforts did take place, some of the uranium leached into the ground water. Cygan characterizes the radioactive concentrations as "not huge, but enough to be of concern," especially locally, where the concentrations are quite high and people have been instructed not to use well water.

Cygan began the computer modeling work on the problem in the mid-1990s, with Greathouse joining the team in 2004. They are approaching the completion of their assigned task.

And while their models have focused on Naturita, their work is also applicable to any other site where nuclear waste is either a real or a potential problem. As a result, the work is part of a much larger project conducted by the U.S. Geological Survey and the NRC to determine how to limit the spread of radioactive waste in groundwater not only at abandoned mill sites, but at nuclear waste storage facilities as well, should there ever be a breach in containment.

10,000 Years

Current U.S. policy calls for the disposal of spent nuclear wastes at one of two sites. Since 1999, defense wastes have been deposited at the Waste Isolation Pilot Plant (WIPP) in southern New Mexico, near Carlsbad. The second site, known as the Yucca Mountain site, located in southern Nevada and intended primarily to accept commercial waste from nuclear power plants, is still undergoing excavation and has been the source of considerable controversy.

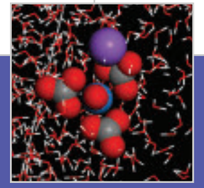

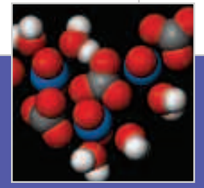

Carbonate-containing uranyl species such as $[Na(UO_2)(CO_3)_3]^{3-}$ (left) and $[(UO_2)_4(CO_3)_3(H_2O)_3]^{2}$ (right) prevent uranyl ions from adsorbing onto clays. These carbonate complexes are prevalent at high pH conditions.

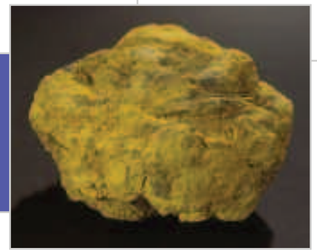

Hand specimen of carnotite uranium ore, a hydrated potassium uranyl vandadate mineral.

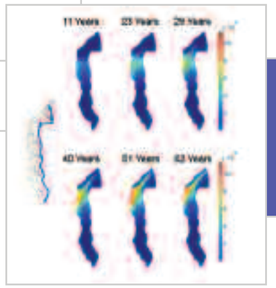

Map of groundwater wells, elevation contours, and hydrologic prediction of uranium concentrations in the alluvial aquifer at Naturita, Colorado. From Curtis (2005); U.S. Nuclear Regulatory Commission.

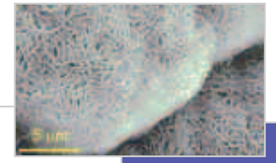

Electron microscopic image of a natural clay sample (montmorillonite) showing the very fine grain size (less than a micron).

Guidelines call for these sites to be designed to safely contain radioactive wastes for a period of 10,000 years without a breach. To ensure that these waste depositories measure up to that stringent requirement, NRC scientists run detailed analyses known as performance assessments. As part of these assessments, the scientists try to determine what would happen should their best-laid plans fail and a radioactive leak occur.

"Part of the performance assessment process is to run computer models of ground water flow, expecting the worst case, a breach at a waste site where the uranium is leached out into the ground," says Cygan. "The NRC wants to know if we can determine for a 10,000-year period how far that uranium will be transported. Will it rapidly move into the ground water, or will it remain localized and not present a problem?"

These are complex questions, the answers to which are the focus of the modeling being conducted by Cygan's group. And although the answers vary from one waste site to another, there are certain similarities in results from different sites, such as how certain clays will react when submerged in uranium-contaminated water. It is these similarities that make the Sandia models universally applicable.

A Price To Pay

The specific purpose of the computer modeling is to determine how various kinds of clays will sponge up a solution contaminated with uranium, in the form of uranyl cation (UO_2^{2+}). In addition to the fact that clay occurs naturally at Naturita, as well as at other waste sites, its efficiency at soaking up radioactive polluted water makes it suitable for use as a secondary barrier to prevent uranium that has leaked from its primary confinement vessels from reaching ground water.

"Clay is far and away the best natural material for uranium sorption," says Cygan. One reason for this, he explains, is that it is a natural nanomaterial, in that it is very fine-grained, with large total surface area and very small spaces between and within the grains. As a result, very large total surface area is available to be exposed to passing fluid, which can be soaked up, while the extremely small spaces between the grains allow clay to act as a very fine sieve.

But there is one problem with clay as a uranium adsorbing material, in that it is extremely difficult, if not impossible, to predict with certainty precisely how a given type of clay will adsorb radioactive material. In other words, clay is unpredictable. As a result, when scientists run bench-scale tests to determine how well specific clays do soak up radioactive water, the results can differ greatly from those obtained from natural materials found in the field. Compounding the problem is the fact that it is not known why the results vary widely.

Eliminating Variables

To overcome this hurdle, the Sandia team precisely controls and constrains the input data for their computer models of clay, thereby eliminating the little-understood variables that can influence laboratory testing. This enables the Sandia researchers to provide NRC scientists with consistent information about how various clays will react to contaminated ground water exposure. "Clay can adsorb uranium by a couple of different mechanisms, and our job is to understand how the chemical process of adsorption occurs and what types of clays can better adsorb than other clays," explains Cygan.

In computational chemistry terms, what the Sandia group provides to the NRC is known as the distribution coefficient, denoted K_d, which is a quantitative measurement of how much sorption occurs. The NRC then uses this information to develop accurate performance assessments.

"When we do computer modeling with molecular simulation we know exactly the type of clay we're dealing with. We know exactly what the surfaces are like. We know exactly what the composition of the clay is and, therefore, we can go through a matrix of simulations that provide all of the variables. We control the variables, then look at the resulting K_ds," explains Cygan.

To make sure they cover the gamut of all of the variables that may occur in nature, the Sandia Geochemistry team provides what Jeffery Greathouse describes as "end member cases," which are best- and worst-case scenarios that enable the NRC hydrologists to consider the full range of what might happen in the real world. As Greathouse puts it, "We are trying to bracket what can be expected in the natural world because of the complex nature of these soils.

"Because these K_ds are not always known as accurately as they need to be known for the performance assessment models, we are providing a theory to at least constrain in terms of bracketing what would be expected in a natural clay. Our clays in the computer experiment are ideal, but we can also zero in on a particular type of clay that might best represent what might be in the field and use that as a basis for our computer simulations. We've done work on maybe a half-dozen different types of clays," says Greathouse.

Modeling Clay with LAMMPS

The program the Geochemistry team is using to run their K_d modeling is known as LAMMPS, which is freely available on Sandia Laboratory's web site (www.cs.sandia.gov/~sjplimp/lammps.html). It is written specifically to be run on massively parallel computers. The clay simulations are run on the Geochemistry Department's own 34-processor cluster.

In all, there are four scientists on the clay modeling team who interact with other groups within Sandia, as well as with other DOE laboratories and researchers at several universities. One of the Sandia scientists who is not in the group but who serves as what Cygan calls "a great resource" is Marcus Martin, who was a DOE Computational Science Graduate Fellow from 1997 to 1999 and is now a Sandia staff member.

Additionally, Cygan is finding that his group is receiving considerable attention from researchers at other institutions whose interest is in materials science rather than environmental chemistry. This is because clay is increasingly being incorporated with polymers to create new materials, and the Sandia team's work provides substantial predictive capability for those materials as well. As a result, says team manager Cygan, "People have been contacting us in terms of how to use our methods."

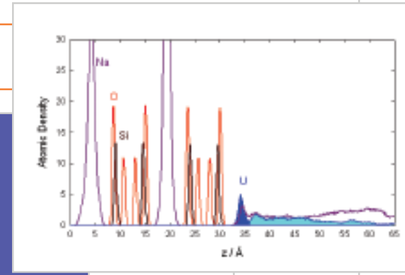

Atomic density profiles from molecular dynamics simulation. Silicon and oxygen atoms (red and black) in the two clay layers are shown along with sodium (purple) and uranyl (blue). The two shaded regions under the uranyl graph represent adsorbed (dark blue) and diffuse (light blue) ions.

RANDALL T. CYGAN

Randall Cygan received his Ph.D. degree in geochemistry and mineralogy in 1983 from Pennsylvania State University where his research emphasized chemical diffusion and dielectric behavior of silicate minerals. In late 1983, he joined the Geochemistry Department of Sandia National Laboratories in Albuquerque, New Mexico, where presently he is a distinguished member of the technical staff. He also spent two years as an assistant professor in the Geology Department at the University of Illinois in Urbana, Illinois (1987-1988).

His research interests are varied, including investigations of mineral equilibria, chemical kinetics, surface chemistry of minerals, sorption and dissolution of minerals, shock metamorphism, and atomistic modeling of minerals and geochemical processes. Randall Cygan has published extensively in various geological, chemical, and materials science journals, and has received numerous honors including membership in Phi Eta Sigma and Phi Kappa Phi honor societies. He is a Centennial Fellow of the College of Earth and Mineral Science at Pennsylvania State University and a Fellow of the Mineralogical Society of America. He is also a member of the American Chemical Society, the American Geophysical Union, the Geochemical Society, the Materials Research Society, and the Mineralogical Society of America.

Further Reading:

R. T. Cygan and J. D. Kubicki *Molecular Modeling Theory: Applications in the Geosciences.* Reviews in Mineralogy and Geochemistry, Mineralogical Society of America, Washington, D.C., 531, 2001.

R. T. Cygan, J. J. Liang, and A. G. Kalinichev *Molecular models of hydroxide, oxyhydroxide, and clay phases and the development of a general force field.* Journal of Physical Chemistry B, 108(4), 1255-1266, 2004.

J. A. Greathouse and R. T. Cygan *Molecular dynamics simulation of uranyl(VI) adsorption equilibria onto an external montmorillonite surface.* Physical Chemistry Chemical Physics, 7(20), 3580-3586, 2005.

Contact:
Randall Cygan
rtcygan@sandia.gov

PRACTICUM COORDINATOR

Marcus Martin. . . . . marmart@sandia.gov

By Victor D. Chase

Facilitating
Warp Speed Computing

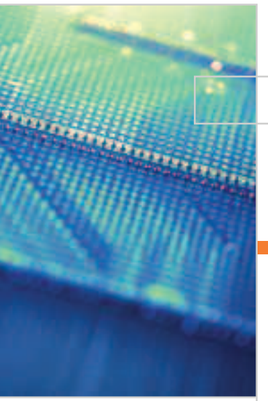

JEFFREY VETTER AND THE EIGHT OTHER MEMBERS of the Future Technologies Group he leads at the Department of Energy's (DOE) Oak Ridge National Laboratory (ORNL) could be considered the Batman and Robins of computational science, in that it is their job to seek out and assist others who can benefit from their help in preparing for the next generation of high-end computing.

As doers of good computational deeds, Vetter's group surveys applications software being used by those within DOE's Office of Science, at ORNL and at other DOE laboratories, to determine what they may need in terms of computing performance enhancement. They then help those groups adapt their applications so they can take advantage of the world's most powerful computers when they come online.

The Office of Science, to which Vetter's group is attached, is one of two main branches that make up the DOE, the second being the National Nuclear Security Administration. The primary mission of the latter branch is to maintain and support the nation's nuclear stockpile. The Office of Science includes scientists involved in a wide range of other computational science activities, including areas such as climate modeling, fusion simulation, materials modeling, and astrophysics simulations. This presents a wide range of challenges to the Future Technologies computer scientists.

In addition to working with DOE researchers to optimize their applications software, the Future Technologies Group is charged with ensuring that the DOE always has the world's most advanced computing hardware at its disposal. Toward that end, Vetter's group is also involved in evaluating advanced computing technologies that are currently on the cusp of creation, such as optical processing and reconfigurable computing, for potential use by DOE laboratories. As part of that process, the Future Technologies Group works closely with the companies developing this new technology to get a jump-start for DOE when their products do become available.

Folding Proteins Faster

Though it was formed less than two years ago, Future Technologies has already succeeded in helping the Computational Biology team at ORNL gear up so that they can make use of two new supercomputers — the Cray XT3 and IBM's BlueGene/L.

The DOE is interested in investigating biological systems for a number of reasons, among them being the development of a greater understanding of the role microbes play in energy production, and also to learn how they can be used to help clean the environment. Some of the microbes being investigated, for example, have the capacity to absorb carbon dioxide from the atmosphere. Excess carbon dioxide in the atmosphere from fossil fuel burning is generally considered to be at least partially responsible for global warming.

A key part of the Computational Biology team's endeavors has been aimed at studying a phenomenon called protein folding, a self-assembly process that proteins undergo so that they can perform their functions. Little is understood about protein folding, and such understanding is essential to learning more about the function of various proteins. Gaining such knowledge is important because proteins perform most of the functions of all living cells — they are the machines that drive life. Without proteins, which are the molecules that make up enzymes, antibodies, and the structural elements of our bodies, including bones and muscles, there would be no life.

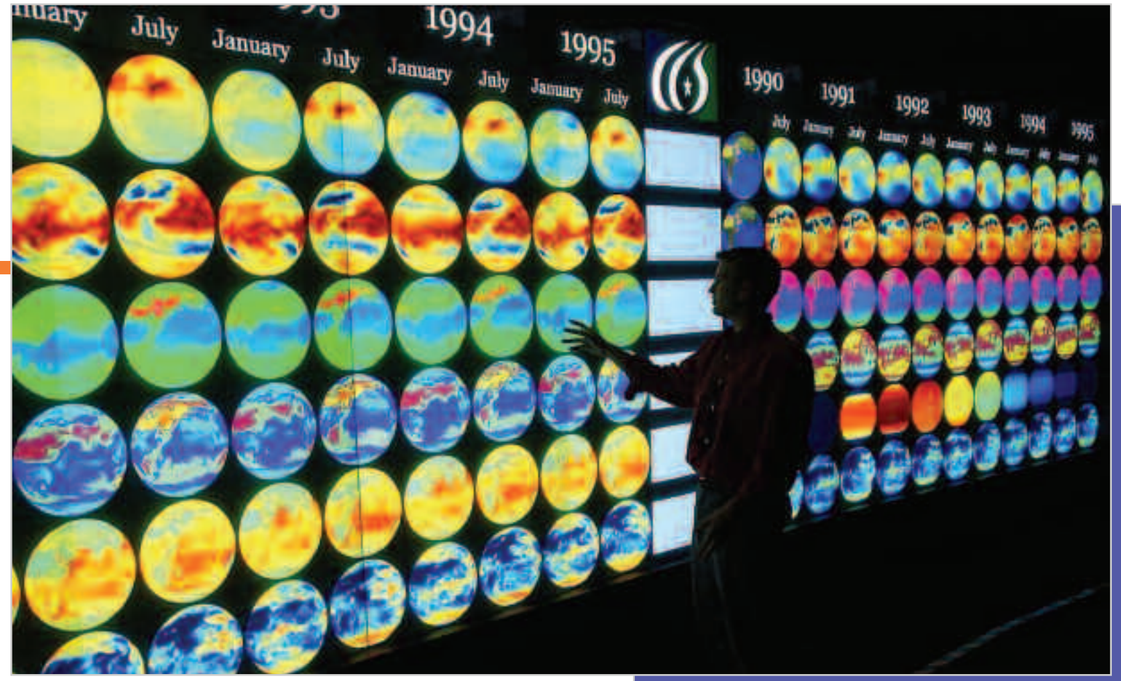

Jamison Daniel studies a climate modeling technology in the EVEREST visualization laboratory at Oak Ridge National Laboratory's Center for Computational Sciences. ORNL's Visualization Task Group seeks out projects that might benefit from applying innovative visual data understanding techniques to scientific data.

But studying the folding process is an extremely difficult task, and that is where computer modeling comes in. As Pratul Agarwal, an associate staff scientist in the Computational Biology team, puts it, "Protein folding is one of the hardest problems to crack, so we thought, 'Let's try starting from the hardest case to understand what we need to do to solve some of these problems.'" Agarwal and his colleagues opted to use computer modeling to try to solve the mysteries of protein folding because studying the folding of actual proteins is at best difficult and costly; it is frequently impossible, since the conditions realizable in the laboratory are considerably different from those found in nature.

As doers of good computational deeds, Vetter's group surveys applications software being used by those within DOE's Office of Science, at ORNL and at other DOE laboratories, to determine what they may need in terms of computing performance enhancement.

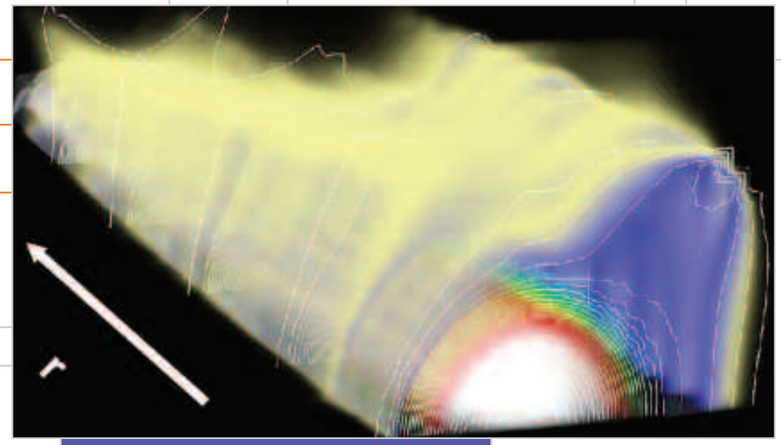

Particle distribution as a function of radius (CQL3D).

JEFFREY VETTER

Jeffrey Vetter is a computer scientist in the Computer Science and Mathematics Division (CSM) of Oak Ridge National Laboratory (ORNL), where he leads the Future Technologies Group. His research interests lay largely in the areas of experimental software systems and architectures for high-end computing. Jeff earned his Ph.D. in Computer Science from the Georgia Institute of Technology; he joined CSM in 2003. Jeff is also an Adjunct Professor in the College of Computing at Georgia Tech.

Further Reading:

T. H. Dunigan, Jr., J. S. Vetter, J. B. White III, P. H. Worley *Performance Evaluation of the Cray X1 Distributed Shared Memory Architecture.* IEEE Micro 25(1):30-40, 2005.

M. Smith, J. S. Vetter, and X. Liang *Accelerating Scientific Applications with the SRC-6E Reconfigurable Computer: Methodologies and Analysis. Proc. 12th Reconfigurable Architectures Workshop* (RAW). Denver, 2005.

J. S. Vetter, B. R. de Supinski, J. May, L. Kissel, and S. Vaidya *Evaluating High Performance Computers.* Concurrency and Computation: Practice and Experience, 2005.

Contact:
Jeff Vetter
vetterjs@ornl.gov

PRACTICUM COORDINATOR

Jeff Nichols .....nicholsja@ornl.gov

Running AMBER

To do this modeling the biology group has been using a software application, known as AMBER (Assisted Model Building with Energy Refinement). AMBER takes a basic sequence of letter codes and, from these codes, develops a three-dimensional structure of a protein. The computer power needed to accomplish this monumental task is tremendous. Even using supercomputers with between 128 and 256 processors, Agarwal's group required a full day of computation time to simulate approximately two billionths of a second of protein folding, according to Vetter.

To increase their computing capability, the computational biologists recently gained access to the National Leadership Computing Facility, which enabled them to use next-generation supercomputers, such as the Cray XT3, which has some 5,212 processors, and IBM's BlueGene/L computer at Argonne National Laboratory, with its 2,048 processors. An even larger BlueGene/L supercomputer at IBM's Thomas J. Watson Research Center has on the order of 40,960 processors, according to Vetter. Being able to use supercomputers with thousands, rather than hundreds, of processors presented great potential, but along with that came problems as well.

Specifically, AMBER, the protein folding program, had to be adapted to enable it to run on the new supercomputers. This is where Vetter's Future Technologies Group entered the picture.

"When you write up a simulation on parallel processors, the two things you need to worry about are computation and communications," says Vetter, referring to the fact that the many processors of a supercomputer run in parallel and need to communicate with each other. "So we had to go in and modify their code, use a different communication mechanism so it would scale up." The end result was that the speed at which protein folding can be modeled was increased by almost three times, says Vetter.

As a result, "We've been able to make significant progress in understanding the protein function," says Agarwal. In addition to being able to model protein folding faster, the increased computing power has allowed the biological group to model more complex systems as well. However, notes Agarwal, the work done with AMBER and the new supercomputers is not yet being used on an everyday basis. It was a learning experience, allowing his group to prepare for the future. "We are in the very initial stage, being able to have this technology ready to investigate systems," he explains.

More Power to the Desktop

In addition to working with various groups to prepare their software for next-generation computing, Vetter's group is also very much involved in evaluating and developing computational hardware of the future.

One of the promising areas they are investigating is reconfigurable computing, which, says Vetter, "involves the use of devices that you can actually reprogram the hardware to accelerate computations." In addition, he continues, "We are also looking at an optical processor that actually uses light to compute rather than electrical signals." The group is also evaluating several types of accelerators that can be plugged into conventional desktop systems to potentially make them ten to twenty times faster.

Describing the reconfigurable computing concept, Vetter explains that conventional microprocessors, such as those built into garden variety desktops or laptops as well as super-computers, are designed to respond to a specific set of instructions. Any application intended to run on a specific type of computer, regardless of whether it is a word processing, graphic design, or voice recognition program, must conform to those instructions. With reconfigurable computing, those limitations do not apply. Essentially, the hardware can be reconfigured to adapt to the program, rather than the other way around, without any parts having to be replaced. This is accomplished using a combination of special hardware and software.

Field Programmable Gate Arrays

The types of devices that allow reconfigurable computing are known as field programmable gate arrays (FPGAs). Though such devices have been in existence for over twenty years, their capability is just now reaching the point that developers can consider them for use in scientific computing. As Vetter puts it, the growth of their capability is analogous to Moore's Law, which states that the number of transistors that can be built onto a single integrated circuit will double every eighteen months. In the case of FPGAs, the applicable elements are cells rather than transistors. "If you have a specific algorithm you might generate a configuration of these cells that uses them entirely differently than another application that comes along," says Vetter. Furthermore, these cells can be reprogrammed every ten milliseconds.

Expanding on his Moore's Law analogy, Vetter says, "One of the motivations over the past 15 years of computing has been Moore's Law. The same theory applies to field programmable gate arrays. Our hope is that over the next three to four years we are going to see the same kind of phenomenon in field programmable gate arrays as we did in microprocessors ten years ago."

The advantage to reconfigurable computing is that "all the hardware is working explicitly for one application," says Vetter. The bottom line, then, is that it is more efficient than the conventional configuration. And by using the hardware more efficiently, the speed of the computations being run increases significantly.

The Speed of Light

When it comes to optical computing the idea is to use the physics of light, which moves a lot faster than electrons, to perform computations. "Some clever folks at a company called Lenslet figured out how to actually shine lasers onto a spatial light modulator that allows one to multiply a matrix and a vector, for example," says Vetter. "There are no transistors involved in the actual operation. Instead, there is a cavity that these lasers are projected into, and using analog-to-digital technology they can sample the result of those interactions and multiply a matrix and a vector together." The result is lightning-fast computing.

Lenslet, an Israeli company, recently introduced what it claims to be "the world's first commercial optical digital signal processor," which it says can run "1,000 times faster than any known DSP (digital signal processor)." Vetter hopes to obtain one of these processors for his group to evaluate for possible DOE use. They will also evaluate the considerable ramifications of having to adapt or completely redo highly intricate software applications to run on this completely different type of computer, carefully weighing the pros and cons.

Though the Future Technologies Group might not have a Batmobile, one can rest assured that if help with a new computing technology or hardware configuration is needed, they will be the first on the scene.

Pratul Agarwal (left) and Stewart Dickson (right) discuss the impact of protein vibrations in the enzyme Cyclophilin A (background).

Simulating an
Urban Disaster

By Karyn Hede

THE TOXIC PLUME MOVES SILENTLY, emanating from a point near Madison Square Garden and spreading to Penn Station before heading down 33rd Street toward the Empire State Building. Busy commuters go about their business, totally unaware that a deadly gas is swirling around them.

It's a nightmare scenario that emergency planners hope they never have to face. But emergency managers know that planning for the worst is necessary to ensure that first responders are ready to handle even this worst-case scenario. In such situations, time is of the essence, and emergency managers need to know where to deploy disaster teams first.

For them, a resource housed at Lawrence Livermore National Laboratory (LLNL), Livermore, California, offers vital information. The National Atmospheric Release Advisory Center, NARAC, is a 24/7 national resource that provides real-time predictions of the spread of hazardous materials released into the atmosphere. NARAC supports the new DHS-led Interagency Modeling and Atmospheric Assessment Center (IMAAC). The center was developed by the Department of Energy primarily to respond to radiological releases, but today's disaster planning has taken a new direction with the potential threat of a deliberate release of chemical or biological agents in an urban environment. If the release occurs in or near a city, many millions of people may be affected and emergency personnel need to know who is at greatest risk.

"Urban releases pose new challenges in predicting flow and spread of hazardous substances," says Andy Wissink, a computer scientist in the Applied Math group at the LLNL's Center for Applied Scientific Computing (CASC). "The fast plume models used operationally at NARAC do not adequately account for the presence of buildings. Our goal is to develop an operational urban model to allow real-time simulations in an urban environment."

Wissink is leading a collaborative effort between CASC and NARAC to develop a fast, accurate simulation capability of atmospheric dispersion in urban areas. The overall goal of the project, called the Adaptive Urban Dispersion Integrated Model (AUDIM), is nothing less than to be able to predict within minutes the path of a toxic plume released in any city in America. The project is funded by the Department of Homeland Security (DHS).

The team chose to model Manhattan to support a larger Department of Homeland Security sponsored project called the Urban Dispersion Program (UDP), a four-year project to measure the potential movement of contaminants in and around New York City. Scientists involved in the project are conducting releases of inert gases and using those results to, among other things, validate AUDIM as well as other computer models. Emergency management, law enforcement, and intelligence personnel will use the resulting simulations to plan for, train for and respond to disaster scenarios, including a potential terrorist attack or an accident involving release of toxic industrial chemicals.

"The computational challenge was to be able to resolve a complex urban environment while accurately representing the building structures," says Wissink.

In order to describe the location of buildings, they are placed in a Cartesian grid, explains Wissink. "The previously-used approach used a conforming mesh, meaning the mesh had to wrap around the buildings. It was a very tedious and time-consuming effort to generate the grid, and because of this the simulations were limited to on the order of 100 buildings. Beyond that it just became too difficult to do," he says.

To solve the gridding problem, Wissink and his team applied a technology called embedded boundary grids, which allowed them to resolve very complex geometry quickly and accurately. The key technology, explains Wissink, is to automatically enhance the fidelity of the grids where buildings are added.

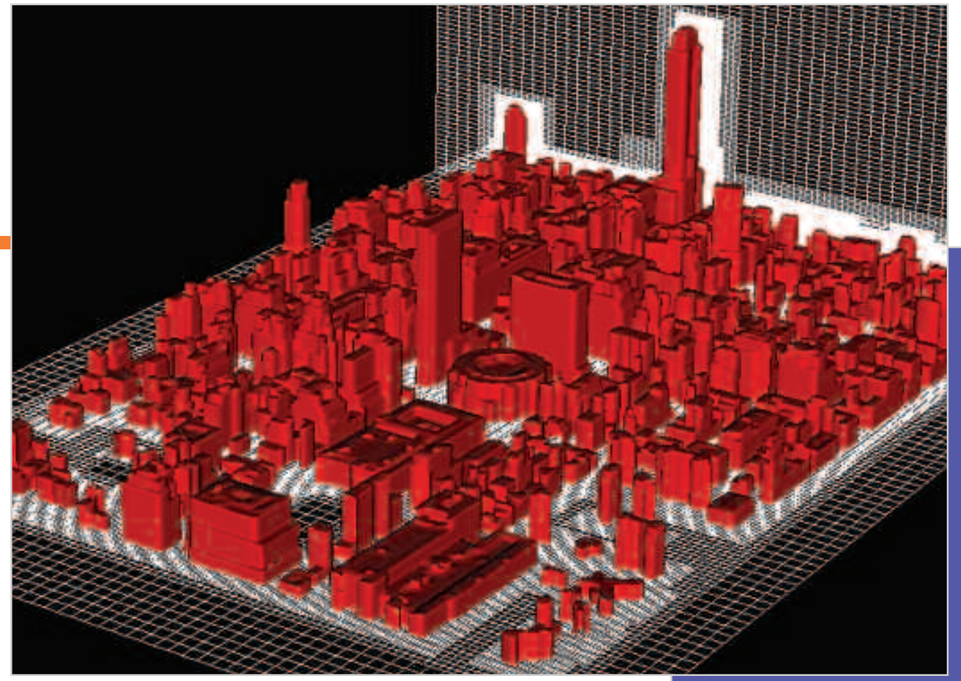

An adaptive Cartesian grid around the Madison Square Garden area in midtown Manhattan.

"The buildings cut through the grid like a cookie cutter," says Wissink. "With the Manhattan simulation, we were able to avoid all those tedious gridding steps, so it enables us to model on the order of 1,000 buildings. The grid generation is done automatically in only a couple of minutes."

The AUDIM programming team, which includes programmers Kyle Chand, Anders Petersson, Craig Kapfer, and former DOE CSGF fellow Brian Gunney, as well as atmospheric scientists Branko Kosovic, Tina Chow, and Stevens Chan, all from LLNL, are applying these new gridding techniques to a computational fluid dynamics code FEM3MP developed over the past decade at NARAC. The code solves the Navier-Stokes (NS) equations to resolve the atmospheric physics and has both Reynolds Averaged NS (RANS) and Large Eddy Simulation (LES) turbulence models. They are also combining this program with two other programs: Overture, which has tools for rapid geometry-to-mesh conversion, and SAMRAI, which supports parallel computing applications that use adaptive mesh refinement.

Each program contained elements the programmers needed to create their next-generation simulation capability, but none was sufficient on its own.

"We needed to both develop new capabilities within the programs and get them to work together," says Wissink.

Anyone who has ever visited New York City can attest to the gusting winds that can whip down narrow city streets. So in

Urban releases pose new challenges in predicting flow and spread of hazardous substances.

ANDREW WISSINK

Andy Wissink is a computer scientist/math programmer in the Applied Math group at the Center for Applied Scientific Computing (CASC). His research interests are in computational fluid dynamics, parallel computing, adaptive mesh refinement, and atmospheric modeling.

Andy earned his Ph.D. in Aerospace Engineering from the University of Minnesota in 1997, and joined Lawrence Livermore National Laboratory in 1999. He is a member of the Association for Computing Machinery (ACM) and the Society for Industrial and Applied Mathematics (SIAM).

In the past, he has worked with overlapping grids for aerodynamics calculations and with iterative methods for the solution of PDEs. Currently, he is working with adaptive mesh refinement on the SAMRAI project and is the project lead for the Adaptive Urban Dispersion Integrated Model (AUDIM) project.

Further Reading:

A. Wissink *Urban Environments Using Adaptive Grids*, Society for Industrial and Applied Mathematics (SIAM) Annual Meeting, New Orleans, LA, July 11-15, 2005. UCRL-PRES-213439.

A. Wissink, B. Kosovic, K. Chand, T. Chow, B. Gunney, C. Kapfer, S. Chan, M. Berger *High-resolution Prediction of Airborne Dispersion Around Buildings: The Adaptive Urban Dispersion Integrated Model (AUDIM)*, in Proc. of R&D Partnerships in Homeland Security, Boston, MA, April 27-28, 2005. UCRL-POST-211537.

A. Wissink, D. Hysom, and R. Hornung *Enhancing Scalability of Parallel Structured AMR Calculations*, in Proc. 17th ACM International Conference on Supercomputing (ICS03), San Francisco, CA, June 23-26, 2003, pp. 336-347. UCRL-JC-151791.

Contact:
Andy Wissink
wissink2@llnl.gov

PRACTICUM COORDINATOR

Tom Epperly . . . . . tepperly@llnl.gov

addition to the computing challenges, the team had to get the program to accurately represent the air turbulence found in an urban environment.

"We are working with the atmospheric physics group here to get the adaptive solver to represent air turbulence," says Wissink. "Actually, that part is an ongoing challenge and we are still working on it."

The group also collaborated with Marsha Berger, professor of computer science, New York University, who pioneered the development of embedded boundary gridding methods for flow over geometrically complex aerospace vehicles. The team was able to utilize software called Cart3D, which was written by Prof. Berger and her collaborator Dr. Michael Aftosmis at NASA Ames Research Center.

The project took 18 months to develop and has now been tested in two-dimensional simulations using the adaptive mesh method and in three-dimensional simulations without the adaptive mesh. The final test, combining adaptive mesh with three-dimensional simulation, is planned for Fall 2005.

The simulations run on LLNL's Linux cluster supercomputer, named "Thunder," a 19.94 teraflop massively parallel computer that uses 1,024 California Digital 6440 servers running on 4,096 Intel Itanium 2 processors, and, as of mid-2005, was the third most powerful computer in North America.

By combining efficiencies in computing scalability and time savings using the adaptive mesh framework, the AUDIM research team has moved in a mere two years from simulations involving 100 buildings and taking days to weeks to set up and solve, to the current simulations involving 1,000 buildings and taking minutes to set up and hours to solve.

In addition, the technique allows users to program higher-resolution simulations in areas more likely to be affected. "We can fine-tune the simulation," says Wissink. "The adaptivity provides the way to dial in the resolution you want. If you need a quick solution in 30 minutes we can do that, but we can also do a higher-resolution simulation for a targeted area in hours, with the code running on large-scale parallel computer systems. And we can quickly change the resolution in the field. You might not be interested in a section of a city not likely to be affected, but you might want a high-resolution simulation on an area likely to be in the path of a plume. You could refine the area around a plume while leaving the rest of the large computational area with coarser resolution, requiring significantly fewer computational resources than refining the grid globally."

All the team's preparation has been in anticipation of a crucial test of their simulation software. The AUDIM grid generation tools, coupled with the FEM3MP flow solver, were used to predict the dispersion of a tracer gas around Madison Square Garden in midtown Manhattan in support of the March 2005 DHS UDP field experiments. They were also used for predictions of UDP experiments around Rockerfeller Plaza in August 2005. The experimental gas release was measured by portable, battery-powered samplers attached to light poles and on top of surrounding buildings, in addition to ground-level gas samplers situated throughout the test area. Wind speed and direction were measured with temporary weather stations in the study area.

The AUDIM team would like to be able to include weather data such as wind speed and direction, temperature, and air viscosity as "boundary conditions" that add to the simulation accuracy. "Eventually, we would like to add chemical reaction terms and data from chemical sensors that would further improve our predictions," says Wissink.

They likely won't know how well their simulations performed until all the data are analyzed in late 2006. But that won't stop the group from continuing to refine their model.

"In the longer term we want to try to couple image processing technology to our simulations," says Wissink.

To do that, the AUDIM team is pursuing a collaboration that would add the final piece of the puzzle to allow them to take satellite data of an atmospheric release and automatically convert that data into its gridding software. Such technology is available through a technique called LIDAR, which is a laser remote-sensing technique that can measure both building outlines and clouds and aerosols from satellite and earth-based stations.

"Currently, our simulations are limited to a few cities where we have building shape files generated," says Wissink. "LIDAR data is available on almost any U.S. city."

If that piece is put in place, the AUDIM project would be close to its goal of delivering real-time urban release simulations that can help provide better estimates of where contaminants may travel and to enhance emergency response when time is of the essence.

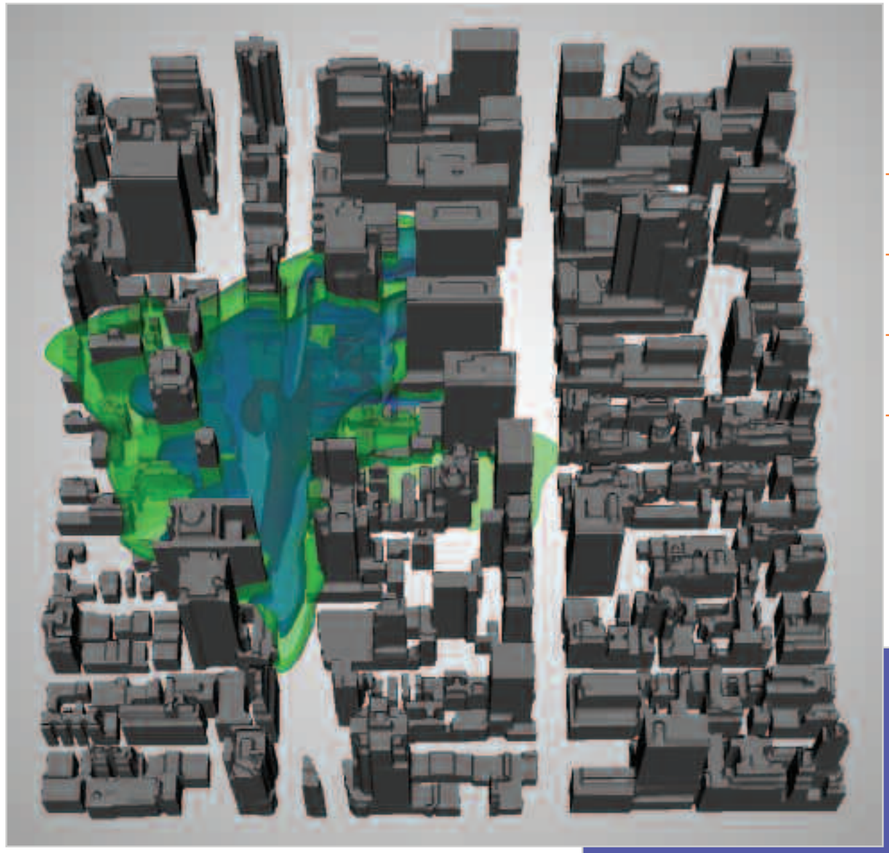

Concentration contours of a Computational Fluid Dynamics (CFD)-based dispersion simulation around the Times Square area of Manhattan.

Emergency management, law enforcement, and intelligence personnel will use the resulting simulations to plan for, train for and respond to disaster scenarios, including a potential terrorist attack or an accident involving release of toxic industrial chemicals.

By Jacob Berkowitz

Fusion and Ice

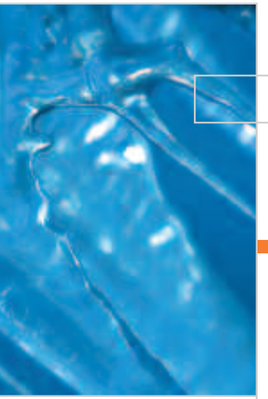

WHAT HAPPENS IF you throw an ice cube into a fire? It's the kind of question that could easily inspire grade school students to experimentation. Roman Samulyak and Paul Parks are working on a high-tech variation of this question. In their case, the "ice cube" is a tiny, frozen deuterium fuel pellet — colder than any natural thing on earth — that's launched into a roiling, 100 million degree Celsius radioactive mix.

And the results of this fire-and-ice research will help guide scientists and engineers in a decades-old quest — to bring the power of the stars to Earth.

Samulyak, a scientist in Brookhaven National Laboratory's and an expert on mathematical modeling and scientific computing, Computational Science Center, and Parks, an internationally recognized pioneer in the field of fusion fueling at San Diego-based General Atomics, are bringing their combined expertise to bear on one of today's greatest engineering and physics challenges, fusion power. They're part of the broad U.S. team that is designing and building the fueling mechanism for the International Tokamak Experimental Reactor (ITER).

ITER, a Department of Energy (DOE) top priority, is the next step in the decades-long dream of making nuclear fusion energy a source of electricity for homes and industries (see sidebar: *Going for Q*). The fusion group at DOE's Oak Ridge National Laboratory is fabricating the ITER fueling device. And as one part of this international effort, Samulyak and Park's computational simulations are part of a collaborative effort involving experimentation, theory and advanced computation that is pushing fusion technology from fantasy to fact.

Fusion Fundamentals

"Nowhere in today's science is computational simulation more important than in fusion physics," says Parks. "Theory and computation are driving the technology for fusion fueling."

This is because actually experimenting with thermonuclear fusion is phenomenally difficult. The heart of fusion energy is the burning plasma, a very hot mix of the hydrogen isotopes deuterium and tritium, which are the raw materials for fusion energy. Plasma, the fourth form of matter and the most common form in the Universe, is a super-heated gas that is so energized that an atom's electrons and nuclei are separated to create a soup of high-speed positively and negatively charged particles.

A star's plasma is contained by its massive gravitational pull. On Earth, the challenge is to heat the plasma to temperatures exceeding that of the Sun's core, and safely and effectively contain it by using massive magnetic forces inside donut-shaped experimental reactors called tokamaks, like the D-IIID tokamak at General Atomics.

The combination of heating and containment causes deuterium and tritium nuclei to collide and fuse. This atomic marriage releases an alpha particle, which further heats the plasma, and a high-energy neutron whose energy heats the containment vessel and thus can be used to produce steam to make electricity.

A Pellet's Fiery Fate

The Sun has enough hydrogen to burn as fuel for the next six billion years. By comparison, ITER's plasma will be a fusion spark. Once a fusion reaction starts, the fuel will burn out in a matter of seconds. This is what makes refueling so critical.

… Samulyak and Park's computational simulations are part of a collaborative effort involving experimentation, theory and advanced computation that is pushing fusion technology from fantasy to fact.

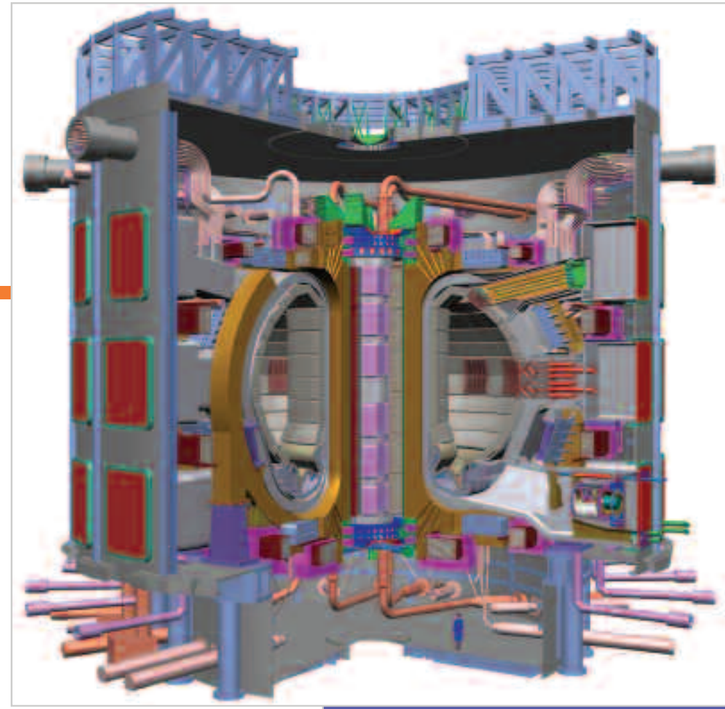

International Tokamak Experimental Reactor (ITER) cutaway.

As with some existing tokamaks, ITER will be fueled with solid, frozen, cylindrical pellets of deuterium-tritium, three to five millimeters (the latter is about the width of a standard paper clip) in diameter and length. They'll be cooled to only 10 degrees above absolute zero. The current plan is to fire about three fuel pellets per second into the plasma. The pellets will be fired like bullets through one of two semi-circular metal guide tubes connected to ITER's inner plasma chamber wall.

The big unknown is whether this fueling system will effectively deliver fuel deep enough into the burning plasma core where the fusion reaction is taking place. ITER's volume is 40 times larger — and the proposed pellet firing speed slower — than in current tokamaks with pellet injection systems.

"We know how far the icy pellet goes. What we need to know now is how far the pellet's vapor goes," says Parks, who in the late 1970s developed the theory describing the pellet's penetration into the plasma, a theory that's well supported by experimental evidence.

The computational simulations under development by Samulyak and Parks are central to characterizing the initial conditions — including temperature, density and width — that will ultimately determine the fate of this deuterium-tritium vapor cloud.

While Parks brings world-class fusion theory to the problem, Samulyak's work numerically captures a wild frontier of physics that's ideal for predicting the vapor cloud's behavior in the plasma. In determining how deep a fuel pellet will penetrate ITER's plasma, Samulyak says it's necessary to understand and model several key components: How quickly does the frozen pellet ablate, changing from solid directly to a gas? How does the growing neutral (i.e., not-yet-ionized) ablation cloud affect the bits of still-frozen pellet? And how will the ablation cloud be heated and thus ionized by the plasma, eventually becoming an integral part of it? Yes, and all of this is taking place in several milliseconds.

The questions draw on Samulyak's expertise in mathematically modeling complex free flows of material, especially ones involving phase transitions, such as the pellet's ablation. Some computational modeling involves solid objects, such as a car chassis, whose shape doesn't change. But a free flow model represents a material that is constantly morphing, such as the pellet ablation cloud. Added to this already difficult computational problem are the phase transitions. They bring so-called strong discontinuities, boundary areas, such as that between

COLLABORATORS

Roman V. Samulyak is a scientist at Brookhaven National Laboratory's Computational Science Center (CSC). His research interests are in the area of fluid dynamics and magnetohydrodynamics, large scale scientific computing, computational accelerator physics, and mathematical physics. Currently, he is leading a computational MHD project at the Brookhaven National Laboratory and working on the development of mathematical models, numerical algorithms, and parallel software for the simulation of free surface magnetohydrodynamic flows of conducting liquids and weakly ionized plasmas in the presence of phase transitions and interaction with intense particle or laser beams. The main feature of the numerical approach is the method of front tracking for material interfaces, and adaptive grid method for elliptic problems in geometrically complex domains, and accurate physics models for phase transitions. The developed code is being used for the study of hydro and MHD processes in liquid mercury targets for future accelerators such as the Neutrino Factory/Muon Collider and Spallation Neutron Source, and in tokamak refueling devices.

Paul D. Parks is a highly versatile plasma physicist who has made significant contributions to a number of diverse theoretical and applied areas connected with fusion and non-fusion research. He is internationally recognized for his pioneering work on pellet ablation in tokamak plasmas. He is actively involved with researchers at the National Institute for Fusion Science in Japan on a 2D numerical simulation of pellet ablation. Dr. Parks has applied his strong analytical skills to the areas of rf current drive, fast wave antenna/plasma coupling, rf ponderomotive stabilization of MHD instabilities, plasma-surface interactions, impurity transport and control, advanced propulsion concepts, and MHD generators. He helped build the pellet-alpha diagnostic program, which yielded the first measurements of energetic alpha particles inside TFTR. Since 1996 his more recent contributions have been in the areas of disruption mitigation using liquid jets/killer pellets. He has authored and co-authored over 75 papers in the refereed Physical Journals and holds one U.S. patent.

Further Reading:

R. Samulyak, Y. Prykarpatskyy, *Richtmyer-Meshkov instability in liquid metal flows: influence of cavitation and magnetic fields*, Mathematics and Computers in Simulations, 65 (2004), 431-446.

R. Samulyak, T. Lu, Y. Prykarpatskyy, *Direct and homogeneous numerical approaches to multiphase flows*, Lecture Notes in Comp. Sci., 3039 (2004), 653-660. Springer-Verlag, Berlin-Heidelberg, (2004).

B. Fix, J, Glimm, X. Li, Y. Li, X. Liu, R. Samulyak, Z. Xu, *A TSTT integrated FronTier code and its applications in computational fluid physics*, Journal of Physics: Conf. Series, 16 (2005), 471-475.

P. B. Parks and L. R. Baylor *Effect of Parallel Flows and Toroidicity on Cross-Field Transport of Pellet Ablation Matter in Tokamak Plasmas*, Phys. Rev. Letters 94 (2005) 125002.

R. Ishizaki and P. B. Parks et al., *Two-Dimensional Simulation of Pellet Ablation with Atomic Processes*, Physics of Plasmas 11 (2004) 4064.

P. B. Parks, T. E. Cowan, R.B. Stephens, and E.M. Campbell, *Model of Neutron Production Rates from Femtosecond Laser-Cluster Interaction*, Phys. Rev. A 63 (2001) 063203.

Contact:

Roman Samulyak
rosamu@bnl.gov

Paul Parks
parks@fusion.gat.com

PRACTICUM COORDINATOR

James Davenport. . . . . daven@bnl.gov

the ablation cloud and surrounding plasma, where critical properties like density, temperature and conductivity change dramatically.

"A fusion plasma code alone couldn't model the entire pellet ablation process," says Samulyak, now in his sixth year at Brookhaven, after graduating from the New Jersey Institute of Technology/ Rutgers University joint program in applied mathematics in 1999. "A single code usually works within some mathematical approximations. Fusion plasma codes assume conductivity is very high. But the conductivity of the initial ablation cloud is actually very low."

So Samulyak is developing the mathematical models, numerical algorithms and computer codes necessary to create the first 3D models of pellet ablation in a tokamak.

To do this, he's building on the FronTier code — a very large code used for modeling complex physics processes, from astrophysics problems to modeling particle accelerators — in order to model the magnetohydrodynamics of the pellet ablation cloud in ITER. The code also draws on high performance software for meshing and discretization of geometrically complex domains developed by the Terascale Simulation Tools and Technologies (TSTT) Center, and high performance solvers of the Terascale Optimal PDE Solvers (TOPS) program as part of the DOE SciDAC program.

"Our computational technique is based on tracking the boundaries of geometrically complex and evolving areas. It allows the use of various physics models, mathematical approximations, and numerical solvers in different domains occupied by the solid pellet, the ablation cloud, and the ambient plasma," says Samulyak. "Running these 3D simulations will require substantial computational time." He's primarily using Brookhaven's 150-CPU Linux-based clusters for the task.

Simulating the Fusion Future

"We hadn't really applied supercomputing to this problem until relatively recently," notes Parks. "We don't yet have a rigorous 2D numerical model of the ablation cloud."

In 2004, he and Japanese computational scientist Ryuchi Ishizaki published a paper describing a 2D simulation of pellet ablation. Although this 2D model is a step in the right direction, Parks says that it didn't account for the tokamak's intense magnetic fields, which are included in Samulyak's models.

And this, he says, is why Samulyak's 2D and 3D simulations are important. There are significant doubts as to whether the presently proposed ITER fueling system has enough fire power to propel the pellet deep enough into the plasma. The central problem is that ITER's inner diameter is too small to contain a straight guide tube. Yet, because of the direction in which the ionized and ablated pellet material drifts inside the plasma, this is where the pellet must be launched from. This necessitates the creation of a curved guide tube that carries the pellet on a slower roller coaster ride from the outer to the inner diameter of the tokamak.

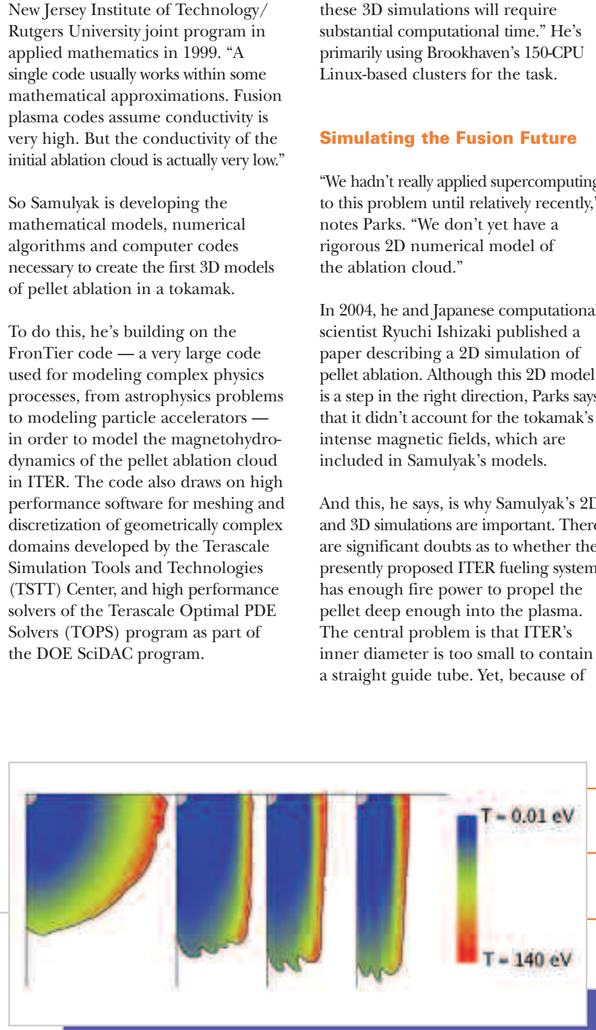

Formation of pellet ablation clouds in magnetic fields ranging from 0 (left image) to 5 Tesla (right image). Temperature distribution in ablation clouds is shown at 2 microseconds after the start of ablation.

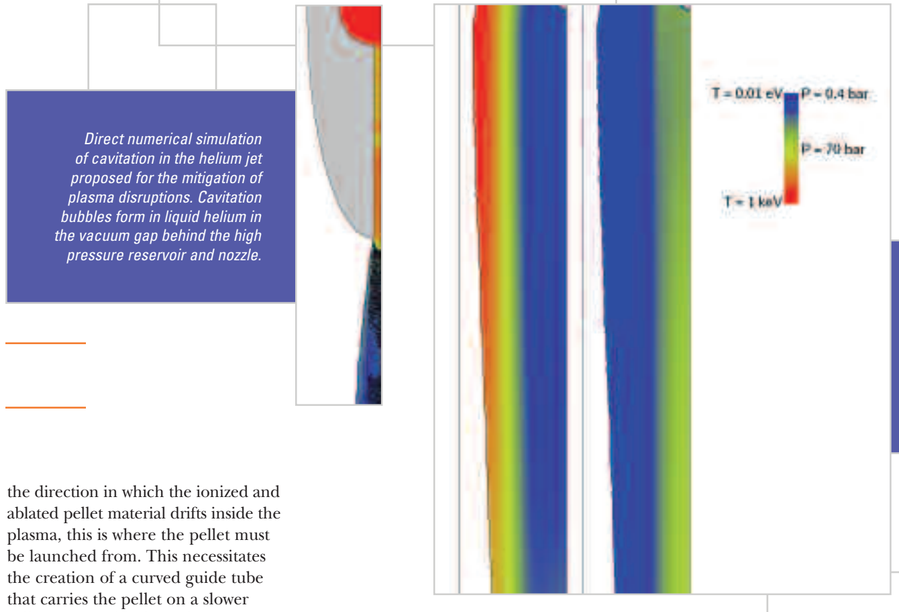

Direct numerical simulation of cavitation in the helium jet proposed for the mitigation of plasma disruptions. Cavitation bubbles form in liquid helium in the vacuum gap behind the high pressure reservoir and nozzle.

Quasi steady state pellet ablation channel in a 5 Tesla tokamak magnetic field. Left image: temperature distribution, right image: pressure distribution.

But Parks believes that the new simulations might significantly clarify the effectiveness of this approach. It's known that the magnetic field forces make the ablation cloud form a long tube parallel to the magnetic field lines. This extended ablation tube acts as a pellet shield, since electrons (also flowing along the magnetic field lines) have to pass through the length of the cloud. This shielding could provide the pellet with a few milliseconds of grace, enough for it to drift into the heart of the plasma and feed the fusion "fire".

"What we want to know is how good a shield this tubular cloud is, and our simulations could provide this information," says Parks.

And time is now of the essence. Although the final ITER design is still in flux, it's possible that the first cutting of steel and pouring of concrete for this fusion machine could start in 2006.

While developing the 3D simulations, Samulyak also collaborates with Brookhaven postdoctoral research associate Zhiliang Xu on modeling of liquid jets of lithium or helium that might be used to mitigate plasma disruption in ITER. The two have already determined that these jets would cavitate while crossing the vacuum barrier between the plasma and the inner wall of the tokamak and might not be able to effectively penetrate the plasma.

In addition, although a newcomer to plasma problems, like a long line of researchers before him, Samulyak is already looking to the big picture — tempted by the allure of helping harness a source of nearly infinite power. He's begun collaborating with researchers of Princeton Plasma Physics Laboratory to integrate the pellet-level simulations with large-scale plasma models for tokamaks.

"For the future, the challenge will be to combine the pellet ablation numerical techniques with tokamak codes for high temperature plasmas," he says. "Because we're talking about adding to the simulation of the entire device, not just the separate components."

GOING FOR *Q*

>> ***Since the earliest thermonuclear fusion research in the 1950s, two generations of physicists have tried to achieve Q. Dubbed the "magic parameter," Q is the amount of fusion power produced divided by the amount of power put into the system to heat the plasma and kick-start a fusion reaction. To this day, Q has always been less than one. In other words, fusion experiments have required more energy than they've produced.***

This is why the International Tokamak Experimental Reactor (ITER), to be built at Cadarache, France, is the focus of such intense international effort. ITER's ultimate objective is to finally achieve a *Q* of 10 or more.

"It's hard to believe that we still haven't gotten more fusion power out than auxiliary power put in," says General Atomics theoretical physicist Paul Parks. For example, Princeton's tokamak TFTR, which completed experiments in 1997, reached a *Q* of 0.3 — for about a quarter of a second.

ITER's initial goal is to achieve ignition, or a *Q* of greater than one. This is the point at which the fusion reaction products, in the form of energetic alpha particles, can transfer enough energy to the plasma to sustain the fusion reaction without the input of any additional energy. To get to ignition, scientists use microwaves or neutral beams to pump megawatts of energy into the plasma and heat it to the point where the fusion reactions take over.

If ITER can achieve a *Q* near10, it would pave the way to the next major fusion milestone — a prototype commercial reactor that efficiently converts fusion energy into electricity.

Of Tsunamis, Thermonuclear

Explosions, & Asteroids

By Victor D. Chase

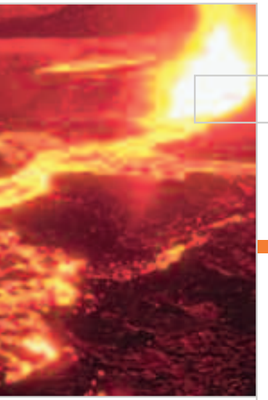

AS THE VOLCANO ERUPTS and one side of the mountain gives way, huge quantities of rock, tumbling faster than a waterfall, drop into the ocean, creating in the surface of the sea a crater that quickly closes. As the pouring rock continues through the depths to the ocean floor, it creates an underwater landslide that comes to rest several miles away.

The mighty force of the collapsing mountain also generates extremely powerful waves, otherwise known as a tsunami, that radiate in all directions, with the shores of the Canary Islands, Africa, and Europe in their paths. But the coastline of the western Atlantic Ocean, including North and South America, is spared.

Such is the scenario should the Cumbre Vieja volcano on the island of La Palma in the Canary Islands erupt with a resulting collapse of its western flank into the North Atlantic. In fact, part of Cumbre Vieja did drop away during a 1949 eruption, but came to a halt before falling into the sea.

The projection of what would happen should the side of the volcano hit the water during its next eruption was developed at the Department of Energy's (DOE's) Los Alamos National Laboratory (LANL) as part of a computer modeling project being conducted by members of the Thermonuclear Applications Group, which goes by the James Bond-like designation X-2.

During the Cold War era, LANL scientists, working in conjunction with Science Applications International Corporation (SAIC), a DOE contractor, attempted to determine what would happen to the shores of America were an enemy to set off an underwater nuclear explosion. Working in the safety of land-locked New Mexico, the team developed a sophisticated computer code designed to model such explosions. Dubbed SAGE, the code models the movement of air, water and rock, and the interactions between them. Doing so takes massive computing power, which is provided by thousands of linked Linux computers at LANL. To make the most efficient use of the computers, SAGE utilizes continuous adaptive mesh refinement, a process that automatically focuses the power of the computers on those regions in the computational process where most of the action is taking place. Without this capability the modeling would be a much slower, more laborious process than it is.

Bombs to Boulders

When Galen Gisler joined X-2 in 2001, the LANL researchers had turned their attention from the study of underwater nuclear explosions to examining what would happen if an asteroid were to hit the ocean. The similarities in the dramatic impact and devastation that could be caused by either event made SAGE the perfect code with which to examine the asteroid question. Because Gisler is a University of Cambridge trained astrophysicist, this project was of great interest to him.

About a year after joining the group, Gisler and his colleagues Bob Weaver, also of LANL, and Mike Gittings of SAIC presented a paper about their asteroid work at a tsunami symposium in Hawaii. The fact that they were dealing with a phenomenon as exotic as an asteroid strike, coupled with the obvious sophistication of their computing ability, sent a wave of excitement through their audience and generated a good deal of interest in the tsunami community generally. "As a result, I started getting interested in the general phenomenon of tsunamis," said Gisler.

In hindsight, his timing could not have been better. He began considering eruptions of volcanoes, both underwater and on land but near a shore, and underwater earthquakes, such as the one that caused the devastating December 26, 2004, tsunami in the Indian Ocean, although his group did not model that particular event. As with nuclear explosions and asteroid strikes, erupting volcanoes and underwater earthquakes unleash tremendous energy and have the potential for great devastation, making them appropriate fodder for SAGE.

Gisler and his colleagues turned much of their attention to possible occurrences in the Atlantic Ocean, where tsunamis do occur, although much less frequently than in the Pacific and Indian Oceans. One such event took place on November 1, 1755. It was a Sunday morning when the underwater Lisbon earthquake hit. Because seismology did not exist at the time, the precise location of the quake remains unknown, but it is called the Lisbon quake because that is where the most damage occurred. Many people attending church were killed as walls fell on them. The tsunami that followed washed over coastal areas of Portugal, Spain and North Africa. And, much like the 2004 Indian Ocean tsunami, the Lisbon tsunami was felt on the opposite side of the Atlantic as water surged onto land in the Eastern Caribbean and along the Florida coast.

The fact that they were dealing with a phenomenon as exotic as an asteroid strike, coupled with the obvious sophistication of their computing ability, sent a wave of excitement through their audience and generated a good deal of interest in the tsunami community generally.

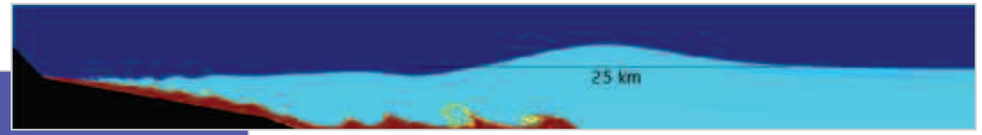

Density raster plot in a 2-dimensional run for a schematic Cumbre Vieja geometry at a time of 180 seconds after the start of the slide. The reflective region representing the unchanging basement of La Palma is at left in black, the basalt fluid slide material is red, the water is light blue, and the air is dark blue. Intermediate shades represent the mixing of fluids, in particular the turbidity currents mixing water and basalt are readily apparent. The water wave leads the bull-nose of the slide material by a small amount; the forward-rushing slide material (with a velocity of 190 meters/second almost matching the wave velocity) continues to pump energy into the wave. The wave height at this time is 1,500 meters, and the wavelength is roughly 50 km. The computational domain extends to the right to a distance of 120 km.

GALEN GISLER

Galen Gisler is a technical staff member in the Thermonuclear Applications group X-2 at Los Alamos National Laboratory (LANL) a founding member of the ROTSE and RAPTOR collaboration teams.

He served as Principal Investigator for the Earthwatch expeditions of 1997-1999, and as Co-PI for the expeditions of 2000-2003. These expeditions were Earthwatch Student Challenge Award Program Campaigns in Transient Astrophysics at LANL. Dr. Gisler received his Ph.D. in Astrophysics from Cambridge University, Cambridge, UK, in 1976.

Further Reading:

G. Gisler, R. Weaver, C. Mader, M. Gittings *Two- and Three-Dimensional Simulations of Asteroid Ocean Impacts*, Science of Tsunami Hazards, 21, 119, 2003.

G. Gisler, R. Weaver, C. Mader, M. Gittings *Two-Dimensional Simulations of Explosive Eruptions of Kick-'em Jenny and other Submarine Volcanos*, Proceedings of the Caribbean Tsunami workshop, in press, 2005.

Contact:
Galen Gisler
grg@lanl.gov

PRACTICUM COORDINATOR
Aric Hagberg .....aric@lanl.gov

"It Really Hit Me"

One of the X-2 group's early undertakings was to study the potential hazards of an eruption of the underwater volcano aptly named Kick 'Em Jenny, located in the Eastern Caribbean near Grenada. It was while they were working on modeling the underwater landslide that might result if Jenny actually did decide to "kick 'em" that the Sumatra tsunami occurred in the Indian Ocean.

"Up to that time the whole exercise was intellectual and academic. Then, suddenly over the Christmas holidays it really hit me how vitally important this work is to people, to societies," said Gisler.

As soon as he returned from Christmas vacation, he called the International Tsunami Information Center, near Honolulu, and offered the services of his group. Thinking that they would be asked to model the Indian Ocean event, Gisler was surprised when he was instead asked to model a potential eruption and collapse of Cumbre Vieja. This volcano had become important because the Sumatra tsunami caused a good deal of concern about its threat to the east coast of the U.S. X-2's modeling showed that the fears were unfounded.

Short and Long Waves

Since beginning the tsunami work, Gisler has become convinced that the potential for another event such as the Lisbon quake poses a greater threat to the shores of the Atlantic than does possible volcanic action. This is not to say that a volcano could not cause tsunami havoc, however.

When the Krakatau volcano blew in August of 1883 in the Sunda Strait between Java and Sumatra — not far from the region of the December 2004 earthquake — it produced one of the largest recorded explosions on Earth.

"If it was just an eruption it would not have cost the 36,000 lives it did," said Gisler. "But the eruption produced a collapse of the mountain, which caused big waves that wiped out whole villages." He also noted that if Mount St. Helens had been near water it too would have caused a huge tsunami when it erupted on May 18, 1980, "because the whole northern side of the mountain just went away."

But as massive as the potential devastation caused by volcanic tsunamis, such as the one produced by Krakatau, the destruction is usually not as widespread as is that caused by large underwater earthquakes.

At right, a density raster plot in a horizontal slice through a 3D calculation for La Palma, at 500 meters altitude at a time of 209 seconds, and at left, three vertical slices through the same calculation at the positions indicated. The reflective region representing the island of La Palma is in black, and the original slide position has been vacated. Deep blue represents the density of air, orange the density of water, and shades in between indicate mixtures of air and water. The peak of the large-amplitude wave is directed towards the SSW, and a following crest is following in the same direction. A line extended through the peak strikes the northeast coast of South America. The vertical slices at left (exaggerated 5 times in the vertical direction) show that the highest wave is at the southern end of the slide. Note that there is considerable wave energy at this 500 meter level directed towards the other islands in the Canaries archipelago and toward the African and Spanish mainland, indicating the possibility of locally dangerous waves impinging on those shores.

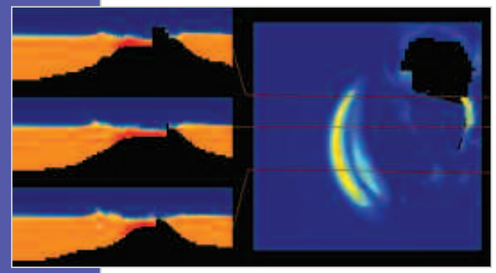

Comparison of an inviscid-slide calculation (top frames) with a plastic-flow slide calculation at 270 seconds after the start of the slide. The plastic-flow slide is significantly retarded in velocity with respect to the inviscid case, and the ripples in the basalt slide (which would represent turbidite layerings) are further apart and suppressed in amplitude. The maximum horizontal velocity of the slide material is 150 m/s in the inviscid case, 130 m/s in the plastic-flow case. The position and velocity of the water wave are very similar in both. Note the strong velocity gradient across the water-air interface at and behind the wave crest in both cases. The Kelvin-Helmholtz instability will very likely cause the wave to degenerate further into high-frequency components.

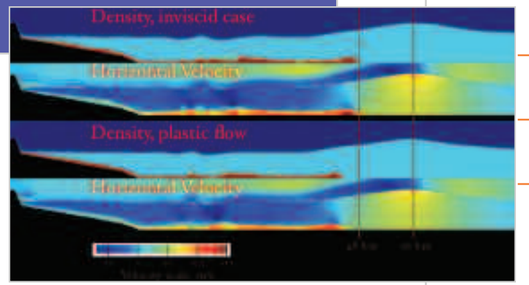

The reason for this has to do with the different nature of the waves they produce. Waves resulting from volcanic eruptions tend to be high, short, and choppy, much like the sea whipped up by a strong wind. They radiate in all directions, and are subject to having their energy sapped by other forces, such as wind. As a result, they spend themselves relatively quickly. Earthquake-generated waves, on the other hand, are long — perhaps 60 miles long — and slow, flow in only one direction, and are not readily dissipated. People on boats at sea may not even notice such a wave, other than to feel the boat rise somewhat and then fall, as was reported by fishermen on the Indian Ocean as the December 2004 tsunami headed for Sumatra and beyond. And as the world has seen, the power of these waves is great, and does not easily dissipate.

As Gisler explains it, "A landslide from a volcano is a point source. It will disperse laterally and the waves fall off with distance faster because they are propagating out in all directions, whereas a line source pushing in one direction doesn't have that much dispersion and becomes more dangerous. That's what we saw in the Indian Ocean, and that's what you can see potentially from a large earthquake along the same lines of what happened in 1755."

Are there fault lines under the Atlantic, in addition to the one that caused the Lisbon quake, that could result in a devastating tsunami? Gisler points to the Caribbean plate, which is slowly sliding under the Atlantic plate in the Eastern Caribbean as one potential site. However, he notes, it is not as active as those in the Pacific and Indian Ocean region, known as the "Ring of Fire."

As for what would happen to the Eastern United States should such an underwater earthquake occur, that remains for X-2 to model. In the meantime, Gisler cautions against excessive concern.

"As scientists, we're interested in hazards. It's important for us to keep things in proper perspective and do real analyses of the risks and dangers and not try to scare people to get publicity," he said.

THE WORK OF MODELING DISASTERS

>> ***As dramatic as tsunamis are, whether caused by underwater nuclear explosions, asteroids landing in the ocean, or underwater volcanoes and earthquakes, modeling them is rather routine business.***

As Galen Gisler, spokesperson for the X-2 group that is doing just such computer modeling at Los Alamos National Laboratory, describes it, "The day-to-day work of science is actually pretty tedious. You type out files that are illegible to any but the trained eye, you stick them on the server, invoke an executable, and it goes off and runs the thing. A few days later you come back and look at the results and realize you have done something wrong, so you go back and do it again."

For Gisler, the exciting part of the work comes with the intellectual exercise of understanding the physics of what the computer is telling him, and validating that information.

"The difficulty is trying to convince yourself that the numbers that are coming out of the machine represent physical reality," says Gisler. "So a tremendous amount of work goes on in the background, trying to make sure that things are valid; that you are not violating any fundamental physical assumptions; that the energies that you get out are reasonable and consistent with one another; and that the code produces results that are consonant with observations of particular events.

"That is particularly difficult when you are talking about something that hasn't happened and are trying to gauge what might happen under circumstances that are extremely uncertain."

One thing that is certain, however, is that Gisler and his group will continue to examine the uncertain. Their plans for the future include fine-tuning their modeling of a Cumbre Vieja volcano eruption on the island of La Palma in the Canary Islands based on more precise data about the sea floor in the area, which they hope to obtain from the Spanish government.

They then plan to turn their attention to learning more about the science of underwater earthquakes, such as the devastating Lisbon quake of 1755 and the Sumatra tsunami of 2004.

By Michael Szpir

Heavy Metal Blues

>> Combustion
>> Geochemistry
>> Computational Performance
>> Computational Performance
>> Advanced Scientific Computing
>> Fusion Energy
>> Thermonuclear Applications
>> Computational Biosciences
>> Computational Tool Kit

THERE ARE TENS OF THOUSANDS of hazardous waste sites in the United States, and the U.S. Environmental Protection Agency has designated more than 1,200 of these as Superfund sites — places so hazardous that the Federal government has dedicated billions of dollars for their remediation. Notable among these is the Hanford Site on the Columbia River in south-central Washington, which is engaged in the world's largest environmental cleanup.

Established during World War II as the Hanford Engineering Works, the site manufactured the plutonium needed for the first atomic bomb. Hanford continued to produce weapons-grade plutonium and other radionuclides well into the 1960s, but it also produced more radioactive waste and highly toxic chemicals than anyone knew what to do with. Much of the waste is now buried and sealed in lined pits or large tanks or even encased in glass…but not all of it.

A few miles down the road from Hanford, T.P. Straatsma, Laboratory Fellow and Associate Division Director of Computational Biology and Bioinformatics at the Pacific Northwest National Laboratory (PNNL) in Richland, Washington, considers the problem of cleaning up this nasty legacy of the Cold War. "A lot of the heavy metals and radioactive isotopes that were produced at Hanford and other Department of Energy (DOE) lands have made their way into the environment, into the subsurface," says Straatsma. Uranium and technetium are among the most notorious offenders, and they tend to be positively charged ions. "These metals are soluble, which means that they can make their way into groundwater or surface waters," Straatsma says. "If we can reduce some of these metals (by adding electrons to them), they become insoluble, which immobilizes them and prevents transport to the groundwater or the rivers."

As it happens, certain bacteria are more than happy to contribute their electrons to the cause. Straatsma leads a team of computational scientists who are simulating how the bacterial membrane responds when it comes in contact with a metal ion or the surface of a mineral. "We're particularly interested in microbes that are able to reduce metals by donating electrons to their environment," he says. The scientific questions Straatsma's team are asking are so fundamental and the potential payoff so enormous that the transfer of electrons between microbes and the environment is the central theme of DOE's Grand Challenge in Biogeochemistry.

Atom by Atom

There are a number of microbial candidates for the simulation, but Straatsma and his colleagues chose *Pseudomonas aeruginosa*, a bacterium that is infamous for being an opportunistic human pathogen (see sidebar: *The Bad Side of Pseudomonas*). "It's a ubiquitous microbe that's commonly found in sediments and the subsurface, so it is environmentally relevant," says Straatsma. "And because of its medical importance, there has been a lot of research on this bug, including the structure of the molecules in its membrane. This is the kind of information that we need to do the simulations."

Pseudomonas has an external "skin" made of two membranes — an outer and an inner — that are separated from each other by a periplasmic space. The outer membrane is a heterogeneous bilayer consisting of lipopolysaccharides on the outside and lipids on the inside. The transfer of electrons to the environment is believed to be mediated by the lipopolysaccharide molecule, and it's the key to the simulation developed by Straatsma and his group.

Over the course of the past six years, Straatsma's team developed a molecular model for the lipopolysaccharide and set up a system that resembles the outer membrane. "It's a very small piece of the membrane, only a few square nanometers of its surface, but it's an enormous computational challenge," he says.

That's because every atom in the membrane is described explicitly in the computer code. Each simulation has hundreds of thousands of atoms and every atom has a force field that interacts with neighboring atoms. The forces impart a velocity to each atom, and the scientists need to calculate the effects on each one of them to understand how the membrane responds to ions or mineral surfaces that approach it from the environment.

"In principle this is an N-squared problem, where N is the number of atoms," explains Straatsma, "...but the approximate methods that we use make this nearly linear. Even then, we have a thousand atoms in our system, we have about a million interactions to calculate. Of course, we have many times that number, which should give you some idea of the calculations."

To make matters worse, those calculations must themselves be repeated millions of times. That's because the simulation evolves in tiny "time steps" as the forces on each atom give it a slight nudge that changes its position and velocity. The forces, velocities and positions must be recalculated for each time step.

"Our simulation time is on the order of nanoseconds (or billionths of a second), with a half-million time steps per nanosecond," says Straatsma.

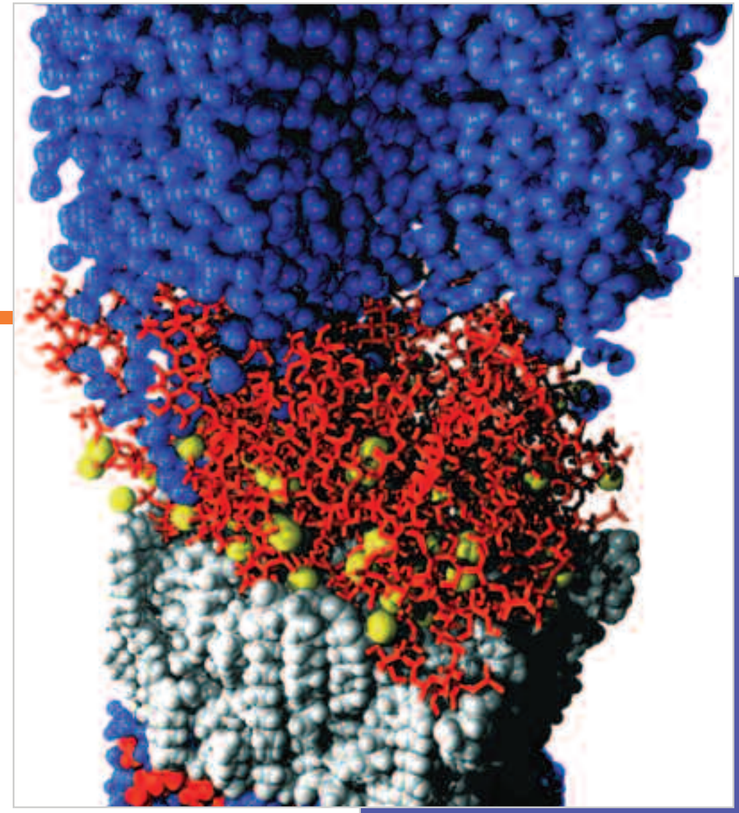

The exterior of the bacterial membrane was assembled by replication of a single lipopolysacchride (LPS) molecule, and the interior by replication of a single phosphatidylethanolamine molecule. Molecular dynamics (MD) simulations of the rough LPS membrane of P. aeruginosa were carried out under periodic boundary conditions, such that the membrane consists of a periodic double layer of LPS/phospholipid molecules externally exposed to aqueous environments, as illustrated above.

The scientific questions Straatsma's team are asking are so fundamental and the potential payoff so enormous that the transfer of electrons between microbes and the environment is the central theme of DOE's Grand Challenge in Biogeochemistry.

T. P. STRAATSMA

T. P. Straatsma is a Laboratory Fellow and Associate Division Director for Computational Biology and Bioinformatics at Pacific Northwest National Laboratory (PNNL). Dr. Straatsma is an internationally recognized scientist with more than 25 years of experience in the development, efficient implementation, and application of advanced modeling and simulation methods as key scientific tools in the computational study of chemical and biological systems. With an extensive background in quantum mechanics and classical mechanics approaches he has enabled and contributed to technical programs focused on the understanding of the relationship between structure, dynamics and function of molecular and biomolecular systems.

The unique combination of expertise in electronic structure calculations, molecular dynamics simulations, and the evaluation of thermodynamic properties allowed Dr. Straatsma to establish an international scientific reputation for the development of new molecular simulation methodologies as well as the application of computational modeling and simulation methods to chemical and biological molecular systems. His research interests include the development, efficient implementation and application of molecular dynamics simulation as a key scientific tool in the study of chemical and biochemical systems, complementing analytical theories and experimental studies. His contributions include the development of computational techniques that provide unique and detailed atomic level information that is difficult or impossible to obtain by other methods and that enable the understanding of the properties and function of these systems. He is known for the design of efficient implementations of these methods on modern, complex computer architectures, including vector processing and massively parallel computer systems. Dr. Straatsma's current application research focuses on the use of simulation methods to study microbial membrane mediated geochemical processes at mineral surfaces, protein mediated transport processes across microbial membranes, and the development and application of novel methods to study complex enzymatic reaction mechanisms that involve electron or proton transfer steps.

Dr. Straatsma's expertise in computational chemistry and computational biology has led him to establish and provide the technical leadership of the Computational Biology and Bioinformatics research group at the Laboratory. He is serving as principal investigator for multi-laboratory and multi-disciplinary research programs, and provides a leadership role in the definition of hardware and software requirements for data intensive computing applications for complex biological systems.

Dr. Straatsma earned a Ph.D. in Mathematics and Natural Sciences from the University of Groningen.

Further Reading:

R. D. Lins and T. P. Straatsma, *Computer simulations of the lipopolysaccharide membrane of Pseudomonas aeruginosa*, Biophysical Journal, 81, 1037-1046 (2001).

T. A. Soares, J. H. Miller, and T. P. Straatsma, *Revisiting the Structural Flexibility of the Complex p21ras-GTP: The Catalytic Conformation of the Molecular Switch II*, Proteins: Structure, Function and Genetics, 45, 297-312 (2001).

R. M. Shroll, and T. P. Straatsma, *Molecular Dynamics Simulations of the Goethite-Water Interface*. Molecular Simulation, 29, 1-11 (2003).

Contact:
Tjerk Straatsma
tps@pnl.gov

PRACTICUM COORDINATOR

Stephen Elbert ..... steve.elbert@pnl.gov

"That is why we do these simulations on massively parallel computers, using hundreds, up to maybe a thousand processors simultaneously. It can take several months to run and analyze a membrane simulation, so we are using new methods to improve the implementation in a massively parallel fashion." The implementation has been highly successful. Indeed, the massively parallel code that Straatsma and his colleagues developed, called NWChem, was recognized by *R&D Magazine* with an R&D 100 award in 1999.

Why Bacteria Love Dirt and Heavy Metal

The membrane simulations have produced some critical discoveries. "One of our biggest findings is that there is an intrinsic electrostatic difference, a dipole, across the membrane so that it's always negatively charged on the outside," says Straatsma. "It explains why these microbes adsorb to positively charged mineral surfaces." Knowing how these bacteria stick to dirt is important if you want to use them for bioremediation. As Straatsma puts it, "We need to know the strength of the interaction between the microbes and mineral surfaces because we may not want them clinging to the first thing they see. On the other hand, if the interaction is too weak, the microbes may wash out of the site immediately."

"A second and related finding is that the negative charge on the surface of the membrane is extremely dynamic," says Straatsma. "There are regions of negative charge that readily adapt to anything that comes close. So as a positive ion approaches, the membrane immediately pushes negative charge toward it to grab the ion," he says. Remarkably, the membrane can change its electrostatic signature to accommodate an ion within picoseconds, or a few trillionths of a second.

Straatsma's team has also determined the amount of work it takes to draw the ions into the membrane by calculating the free energy involved. "It turns out that if you do that calculation for positive ions, it doesn't require work," says Straatsma. "In fact, you get energy back!" These results provide an atomistic explanation for experimental observations, which show that positive ions are pulled into the membrane. "That's the value of these computational studies," notes Straatsma. "You can confirm experiments, or even help to interpret them."

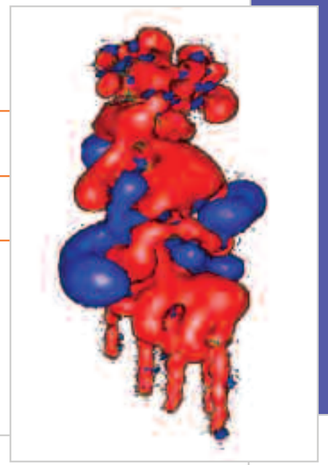

For the appropriate description of the intermolecular interactions within the lipopolysaccharide membrane it is important to have an accurate electrostatic 'signature' of the molecular components. The electrostatic potential is shown here of the lipopolysaccharide molecule of Pseudomonas aeruginosa *that was used to fit the partial atomic charges used in the molecular dynamics simulations of the complete membrane.*

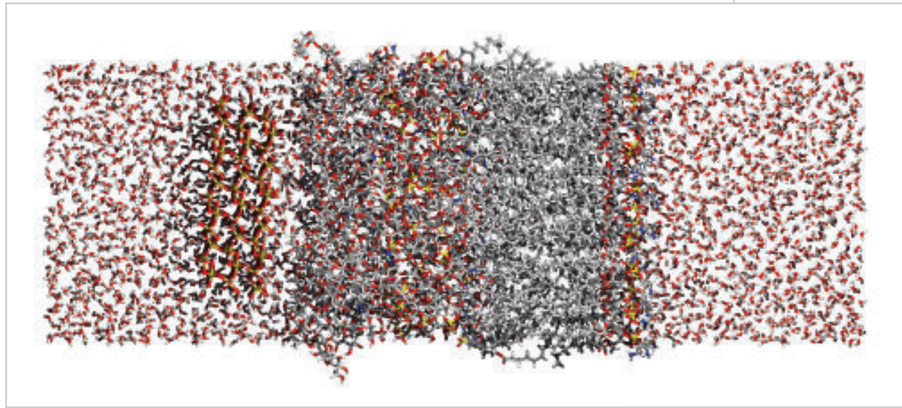

To study the interaction between microbes and mineral surfaces, Straatsma and his team are carrying out simulations of the outer lipopolysaccharide membrane in contact with mineral fragments. The illustration is of a snapshot out of one of the molecular dynamics simulations used to study the specific atomic interactions that determine the adhesion of Pseudomonas aeruginosa *with the mineral goethite.*

A More Perfect Membrane

Straatsma eventually hopes to simulate even larger portions of the bacterial membrane. That will require not only more computational power, but also more information about the atomic structures of other membrane components, such as proteins that play fundamental roles in the microbe's interaction with its environment. Such proteins include cytochromes, which are involved in shuttling electrons; porins, which let water and small ions pass through; and transporters such as the iron-citrate receptor, which transports iron across the membrane.

If the work isn't done at PNNL, it probably won't be done anywhere else. "It turns out that nobody else is simulating membranes that involve lipopolysaccharides, because these molecules are so large that the calculations become computationally very expensive. But we have this big computer sitting here and so we are able to simulate large molecules," says Straatsma.

Adding large proteins into the membrane models will make the simulations more realistic, but it means that there's still an enormous amount of research that must be done. Straatsma's team works very closely with other groups at PNNL who are determining the structures of the molecules. "The wonderful thing about working here at PNNL is the availability of resources — that includes not only the hardware but other scientists," says Straatsma. "More so than anywhere else I've worked, there is the sense of accomplishing things in a team. You are working with a large group of people, not sitting alone in your office, and it makes you feel that what you're doing is really worthwhile because it's part of a larger mission. It's a great place to work," he says.

Several groups at PNNL are now studying various microorganisms that may someday join the work force that's cleaning up Hanford and other DOE sites. Biologists and chemists take on the experimental work, while environmental scientists and engineers are sorting out how these microbes might be applied to the job. Some may find it ironic that these simple life forms will be cleaning up the mess created by some of the greatest scientific minds in history. But the bacteria have been adapting to changes on (and in) the Earth for nearly 3.5 billion years. They have an ancient relationship with dirt, and it may be this timeless wisdom we need to acquire if we are to maintain our own relationship with the planet.

THE BAD SIDE OF *PSEUDOMONAS*

>> ***Although bacteria may ultimately come to our rescue, the microbe Pseudomonas aeruginosa has caused more than its share of grief to the human species.***

The bacterium rarely infects uncompromised tissues, but it will readily attack almost any part of the body that has been weakened by disease or injury. A *Pseudomonas* infection has a frighteningly high mortality rate for patients who are hospitalized with cystic fibrosis or burn wounds. Unfortunately, the microbe is naturally resistant to many antibiotics because the lipopolysaccharide molecules in its outer membrane serve as a barrier that prevents the entry of these drugs.

Now that Straatsma and his team have developed a model of the *Pseudomonas* membrane, they are also studying ways to combat the microbe's defense mechanisms. "We are looking at how these bugs excrete molecules that are harmful to them and how the membrane interacts with antibiotics," he says.

The means to achieving this goal may be a simple but important structural component of the *Pseudomonas* outer membrane. Even though there is an electrostatic gradient across the membrane, the overall structure has a neutral charge. That's because counter ions, such as calcium (Ca^{+2}), act to balance out the negative charges that are concentrated on the outer surface of the membrane. The counter ions may prove to be the key.

"It turns out that positively charged antibiotics are able to replace some of the counter ions within the outer membrane. They have a particular molecular shape that disrupts the integrity of the membrane. If the membrane breaks, then the microbe dies," says Straatsma. The electrostatic properties of the bacterial membrane that scientists may exploit for bioremediation might also prove to be the microbe's Achilles' heel.

By Karyn Hede

PETSc Engines

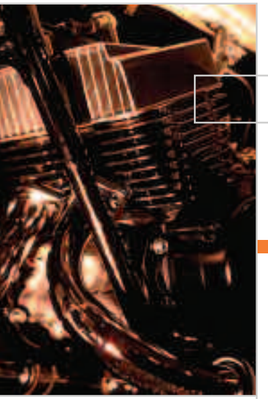

IF EVERYONE WHO WANTED TO OWN A CAR had to build their own internal combustion engine, imagine how different the world would be. Luckily for us, we had Ford to build the first automobile engine and chassis, while today we just enjoy the ride.

In the world of scientific computing, the U.S. national laboratories have built the engines and chassis that drive many of today's high-performance computational simulations. For the scientists and engineers, the availability of rigorously-tested dependable computational toolkits means spending less time writing software code and more time doing science. One such computational "engine," unassumingly called Portable, Extensible Toolkit for Scientific Computation or PETSc (pronounced "pet c") provides a library of software for solving partial differential equations, the most common mathematical tool for modeling relationships between physical processes.

"If you are working on a particular engineering problem, you might have to combine many different numerical algorithms and that can get complicated," says Barry Smith, one of PETSc's early developers at Argonne National Laboratory in Illinois. "By encapsulating the commonly used algorithms in libraries, each user doesn't have to understand all the details of all the algorithms they are using in order to use them. Trying to do so would just be too overwhelming."

PETSc has become the go-to source for what PETSc user Mark Adams calls "software that does all the heavy lifting."

Adams, a research scientist in applied math at Columbia University, uses PETSc as the engine for a number of high-level programs he developed to solve solid-state mechanics problems that use finite element analysis.

One such problem, a collaboration with Tony Keaveny and Panayiotis Papadopoulos of the University of California - Berkeley's Orthopaedic Biomechanics Laboratory, required modeling the strength of the spinal vertebra to high resolution.

"People working in the area of bone modeling have been using what are called matrix-free, element-by-element preconditioners, which are mathematically simple linear solvers that do not scale very well, but their advantage is they use very little memory," Adams says. "To be able to accurately model bone requires very large mesh models and people have been pushing the limits of these simple solvers. Several groups have these problems that they literally run for weeks and weeks on end on small parallel computers like 32-processor SGIs."

Adams worked with UC-Berkeley mechanical engineering graduate student Harun Bayraktar to convert the problem from a non-scalable linear solver to a non-linear, highly scalable solver.

"PETSc provided a substantial amount of code, on the order of 40,000 lines of code, that is really very general, very unstructured, but at the same time provides fully functional and highly optimized matrix, vector and other parallel linear algebra objects," says Adams. "It lets people like me concentrate on developing good algorithms to build on top of it."

PETSc has become the go-to source for what PETSc user Mark Adams calls "software that does all the heavy lifting."

Adams used PETSc as a foundation to build a solver that would take the large finite element meshes and perform large-scale, non-linear whole bone modeling on high-end parallel computers.

"Using this method, we were able to speed up these calculations by literally two orders of magnitude in CPU time and about three or four orders of magnitude in wall clock time," says Adams. "We went from seven weeks to about two minutes to solve some of these problems."

Adams and his colleagues presented the work at the SC2004 supercomputing conference and received a Gordon Bell Award, one of high-performance computing's most prestigious honors, for the work.

The researchers are hoping these simulations will give them a better understanding of the behavior of bone tissue that is thinning as a result of bone loss. Ultimately the goal of the research is to be able to take a CT scan of bone and calculate the risk of fracture.

"The thesis of the research is that bone density alone is not a good measure of the risk of fracture," says Adams. "By understanding how the microstructure of bone degrades, they are hoping to come up with more accurate prediction of who is at risk for fracture."

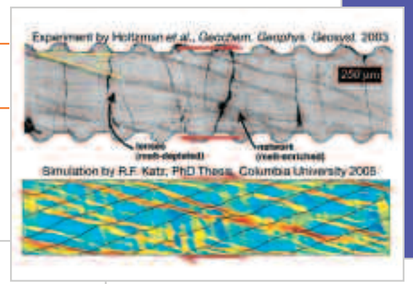

Top panel: a cross section through an experiment from Holtzman et al. (Geochem. Geophys. Geosys. 2003) in which a partially molten aggregate of rock representing the earth's mantle has been subjected to a shear strain of 370%. After a strain of about 100% the initially uniform distribution of melt and shear strain have concentrated into bands at a low angle to the shear plane. Black subvertical cracks occur on quench of experiment and are not modeled.

Bottom panel: one timestep from a PETSc-based simulation by Richard Katz (Ph.D. thesis, Columbia University, 2005) of non-Newtonian two-phase creeping flow. Colors show the vorticity perturbation induced by a localization instability. Black lines are passive strain markers that were initially straight and vertical.

BARRY SMITH

Barry Smith is a Computer Scientist in the Mathematics and Computer Science (MCS) Division at Argonne National Laboratory. His research interests include the scalable solution of algebraic equations and their use in PDE simulations.

Barry obtained his B.S. in mathematics from Yale University and received his Ph.D. from New York University. His interest in scientific computing was sparked in an undergraduate numerical analysis class taught by Bill Gropp, who is currently the Associate Division Director and Senior Computer Scientist in the MCS Division at Argonne.

Barry's current main interest is in the efficient translation of results in academic numerical analysis to the scientific and engineering end users.

Further Reading:

B. Smith, P. Bjorstad, and W. Gropp *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations* Cambridge University Press.

T. Bittner, M. Donnelly, B. Smith *Endurants and Perdurants in Directly Depicting Ontologies* in AI Communications, Vol. 17, Issue 4, October 2004.

A. Kumar, B. Smith, and D. Novotny *Biomedical Informatics and Granularity: Conference Papers* in Comparative and Functional Genomics, Vol. 5, Issue 5-6, August 2004.

Contact:

Barry Smith
bsmith@mcs.anl.gov

PRACTICUM COORDINATOR

Raymond Bair . . . . . bair@mcs.anl.gov

Adams' application of PETSc is just what Smith had in mind when he and Argonne co-developer William Gropp set out in the early 1990s to develop a set of numerical tools that would make the most out of the tremendous gains that had been made in processing speed in high-end scientific computing. The project, which was sponsored by DOE, has had nine developers over the years and currently keeps four developers busy full-time.

"What we saw was a lack of flexible software for high-performance computing," says Smith. "We set out to create software that was designed for solving partial differential equations, which can be used to represent the relationships between physical quantities such as pressure, velocity, and force. These are continuous processes, but in order to solve them in a computer you always have to represent them in a discrete way. The process of discretization results in algebraic equations which you can then solve using PETSc."

Its first users were scientists at Argonne who were modeling fluid dynamics, but in the ten years since its release its reputation has spread far and wide in the computer science community. Users now include geochemists, biologists, physicists and environmental engineers, among others. Since the software works fairly seamlessly with any hardware, except vector machines such as the CRAY X1, it has a large fan base among a substantial contingent of scientists.

"As an open source project, all of the developers and many users have contributed valuable design ideas as well as the actual source code. It has evolved a great deal to match the needs of our users," Smith notes.

Former DOE CSGF fellow Allison Baker, now a post-doctoral fellow at Lawrence Livermore National Laboratory, spent a summer at Argonne in 1999 learning PETSc and implementing an additional linear solver routine into PETSc.

"Spending that summer at Argonne really helped me see how important it is to have good scientific code publicly available," says Baker. "It really makes life a whole lot easier."

Her summer learning experience set the stage for Baker's thesis work at the University of Colorado where she developed a new linear solver algorithm that reduces the movement of data through machine memory. Baker used PETSc as the backbone of her solver, allowing her to move more quickly into the heart of the research question.

"My experience with learning PETSc showed me that a good implementation can be ten times or a hundred times faster than a naïve one," says Baker. "I just really learned a lot about how important it is to think carefully about implementation details and what a big difference it can make. These are just things you don't learn about in class or in textbooks."

For Richard Katz, also a DOE CSGF alum, the attraction of PETSc was its menu of tools to solve a host of different problems and its "inherent scalability."

"One of the great things about PETSc is that you can write your application and test it on a laptop. Then without much thinking you can scale it to a supercomputer. For me this was just fantastic."

Katz spent his practicum at Argonne in 2003 learning PETSc and applying it to solve a set of equations that describe the flow of the earth's mantle over time in subduction zones, regions deep within the earth where the edge of one of the earth's tectonic plates is forced under an adjacent plate.

"The code that comprises PETSc is highly structured and very clearly written," says Katz. "It has a very disciplined interface and so it teaches you to write good code. I learned a lot about writing software just by using PETSc."

Armed with the basic knowledge of applying PETSc, Katz set out to do something no one previously had been able to accomplish: modeling what happens when the earth's mantle, which is solid, interacts with molten magma and both substances flow. (See Katz's profile in the 2004-2005 issue of *DEIXIS*.)

"The behavior of this two-phase system is described by partial differential equations that are highly non-linear. In the past this has been very difficult to solve," he said. "PETSc allowed me to simulate a system that no one had been able to model."

In addition to using PETSc, Katz' research has led to an extension that will be added to the PETSc numerical library. The method, called a semi-Lagrangian advection scheme, is a means of tracking the advective movement of particles in a flow. "Another great feature of PETSc is its extensibility. You can build your own component if you don't find what you need."

For the novice just thinking about trying to use PETSc, Katz recommends checking out the library's extensive collection of example codes.

"The typical way for someone to get started is to find an appropriate example code and modify it to their needs," says Katz. In addition, he says the PETSc developers are very helpful and responsive to questions about the library.

It's just like driving a car, where you have to learn how to operate it before starting the ignition, and it certainly helps if you can check the fluid levels and change a tire. Likewise, Katz says it is important to have some background in software and be able to do a little bit of coding before trying to use PETSc. But, he adds, once you get started the library has a lot of documentation to help solve that next big problem.

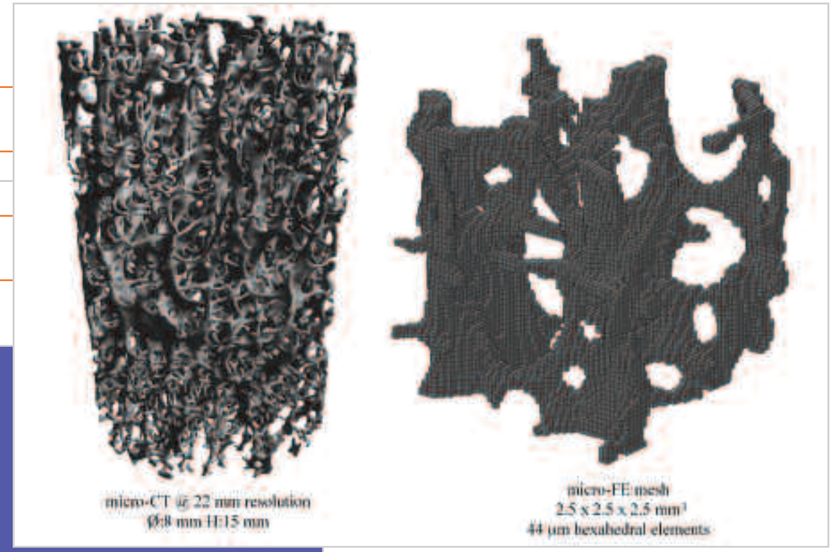

Rendering of a human vertebral bone cylindrical specimen and detail showing voxel finite elements.

Chris Oehmen

In Edgar Allen Poe's FAMOUS SHORT STORY the *Tell-Tale Heart*, the main character is tormented by the relentless, increasingly loud beating of his murder victim's heart. Chris Oehmen knows as much as anybody about seemingly dead hearts that keep beating.

His doctoral research provided evidence for a previously unknown mechanism that drives the heart's pacemaker, a remarkable bio-clock that will keep on ticking even when a heart is removed and suspended in a blood-like substance. But these days, the heart this DOE CSGF alumnus hears beating in his ears is his own. Just over the two-year mark as a scientist at DOE's Pacific Northwest National Laboratory (PNNL) in Richland, Washington, Oehmen is captivated by his role in helping push the boundaries of biology, in silico.

"It's almost intoxicating. I'm still adjusting to it," says Oehmen, a member of PNNL's 25-person Bioinformatics and Computational Biology Group. "In grad school you get really focused on solving your problem. Here I've been exposed to so many problems that have huge possibility for impact."

Oehmen's graduate research demonstrated the role that cellular-level electrical resistance plays in regulating the pacemaker's, or sinoatrial node's, electrical wave that excites the heart muscle to contract. It also showed Oehmen, a student steeped in math and physics, two crucial things about biology: it involves an enormous level of complexity; and computational science is indispensable to understanding this complexity.

"What we showed about the sinoatrial node could only be done with modeling," says Oehmen of his Ph.D. research in the Joint Graduate Program in Biomedical Engineering program at the University of Memphis and the University of Tennessee Health Science Center.

Now Oehmen is parlaying this experience into helping biologists maximize their use of computation.

Biologists are presently faced with a data conundrum: they have too much of it. The genome sequencing of the past 20 years on everything from worms to humans has produced gigabytes of sequence databases. Together, these databases enable crucial pair-wise alignment studies — the comparison of newly discovered gene or protein sequences with known ones to help identify the role of the gene or protein. But these databases are now so large that they're a bottleneck in the process.

"The target databases are increasing like crazy," observes Oehmen. "They're growing so fast they soon won't fit on a single computer and this is creating huge performance issues."

Oehmen has created two data-intensive computing techniques that hold great promise in breaking this data logjam. First, he's applied the Global Array code, developed by a team led by PNNL's Jarek Nieplocha, to share the genome database on a high-performance parallel platform.

Second, he's adapted BLAST, one of the codes most widely used by scientists around the world for sequence comparisons, for use on a parallel multiprocessor architecture.

In initial trials using real data from PNNL colleague Heidi Sofia, Oehmen's ScalaBLAST version showed it's the Indy 500 version of the original.

"Overnight I was able to turn around 16 full genomes with about 64,000 genes, whereas that would have taken her a month and held down her machine so she couldn't do anything else with it," says Oehmen.

While he revels in helping supercharge the genomics and proteomics revolution, Oehmen doesn't stray far from things of the heart. He's usually home by 5:00 p.m. to help feed, bathe and read to his six- and two-year-old children.

"I don't want to be a dad that only shows up at home periodically and works 80-hour weeks," he says, noting that the Battelle Corporation which manages PNNL is an excellent employer when it comes to employees balancing their work and family roles.

Says Oehmen, "If you do things the right way, I believe you can find a way to work smarter not harder."

Chris Oehmen knows as much as anybody about seemingly dead hearts that keep beating.

Joel Parriott

Joel Parriott's Ph.D. is in COMPUTATIONAL ASTROPHYSICS. But today he's working to understand and help others grasp the forces at play not within the Milky Way but rather within the Washington D.C. Beltway. And he's a force himself.

The annual revolution of the federal budget cycle sees Parriott — Program Examiner in the Office of Management and Budget (OMB) responsible for the DOE's Office of Science — scrambling each autumn to provide advice as that year's budget request makes its way through the Executive Office of the President and ultimately to Congress.

To those physicists, chemists and other scientists he meets on visits to FermiLab, the Stanford Linear Accelerator Center and other DOE facilities, Parriott is tantalizingly close to the purse strings for their beloved projects. It's Parriott's job to make recommendations on a portfolio of $3.5 billion. But during the presentations he gives about his role and that of OMB in the Executive Office, his first PowerPoint slide shows him on the organization's bottom rung — five levels from the ear of President George W. Bush.

"In my presentations I'm trying to educate folks about how decisions are made and what they might do to improve the case they're trying to make for more resources," says Parriott from his office across the street from the White House.

A key part of this decision making, says Parriott, is the reality of hard trade-offs. And this, he says, is a difficult reality for many scientists to face and consequently a stumbling block when they come to make the funding pitch for their projects.

"When you've committed your life to an academic pursuit, you have a hard time believing that somebody else wouldn't share your excitement and your conviction that this is the best way to spend money," he notes. "So people say 'We could do my research if you could build just one less tank.' But that's just not the way decisions are made here."

It's Parriott's job to provide the informed technical recommendations for the budget decisions that are made, ensuring that they reflect the President's policy priorities, increase management efficiency and ultimately result in the biggest research bang for U.S. citizens' tax dollars.

This means he asks the tough questions he didn't dream of positing while doing his doctoral dissertation as a DOE CSGF Fellow at the University of Michigan. What's the real value of new investment in supercomputers? Are there other, more cost-effective ways of addressing a scientific problem?

"This job has totally changed my thinking about science," says Parriott, who nonetheless sees himself as a card-carrying member of the scientific community. "I know the limits of computation. I'm trying to push scientific computing to areas where there's predictive capability. We need to be able to run simulations and get results that drive what we're doing experimentally."

Parriott notes that while he's now an odd duck in the world of science graduates, he's also an outsider of sorts in the OMB — one of the few program examiners with a science Ph.D. among offices of MBAs and public policy graduates. After receiving his doctorate in 1998, Parriot honed his policy smarts as a program officer on the National Research Council's Board on Physics and Astronomy. It was here that he developed a keen sense for the push and pull of science politics through committees, reports and lobbying.

When the opening at OMB came three years ago, Parriott says he jumped at the chance to put his scientific and policy know-how to work in a career position that emphasized technical advice on complex and varied scientific issues. Today these issues range from the U.S. role in the international fusion energy project ITER and international competitiveness in high-performance computing, to DOE's role in biological research, and the future of Big Science projects such as particle accelerators.

Says Parriott, "The thing I like about my job is that I have to understand the politics, but I'm explicitly forbidden, thankfully, from playing any of the politics."

Parriott notes that while he's now an odd duck in the world of science graduates, he's also an outsider of sorts in the OMB — one of the few program examiners with a science Ph.D. among offices of MBAs.

Mayya Tokman

Mayya Tokman STUDIES THE SUN, but she's not blinded to questions here on Earth. While her research is helping discover the cause of solar storms, she's also helping shape U.S. foreign policy through science and education exchanges and is about to embark on a journey to build a unique interdisciplinary applied mathematics program.

"I've always been interested in the broad impact of science because sometimes you publish a paper but you don't see the big picture of science and society," says Tokman, who for the past three years has been a visiting assistant professor in the mathematics department at the University of California, Berkeley.

Tokman learned early and personally how small, incremental changes barely perceptible on the surface of things can lead to enormous changes in the big picture. In 1991, the 19-year old Tokman arrived in the U.S. with her parents as a Jewish refugee from her native Azerbaijan, just two weeks before the fall of the Berlin Wall and the collapse of the Soviet Union. Her first academic appointment was as a financial aid worker at Santa Monica Community College, where she "learned English very fast" and helped other Russian speakers.

Her zeal for math soon took her to UCLA to complete her undergraduate degree and then to Caltech where, as a DOE CSGF Fellow, she earned a Ph.D. in applied and computational mathematics.

Through initial collaborations with Paul Dellan, a Caltech experimental physicist interested in replicating astrophysical phenomena on Earth, Tokman (who also has a deep fascination with all things astrophysical) found her niche in the mathematical modeling of the behavior of space and laboratory plasmas. Plasma, a super-hot soup of charged particles, is the fourth state of matter, along with solids, liquids and gases.

Her current research is focused on modeling the behavior of plasma in the solar atmosphere that results in coronal mass ejections (CMEs). Also known as solar storms, these are the massive eruptions of billions of tons of plasma into space. This stream of charged particles creates the beautiful dancing Northern and Southern Lights in the Earth's upper atmosphere. But it can also wreak havoc on satellites, communication systems and power grids. For this reason, the ability to predict these eruptions would be a major scientific coup.

CMEs start as coronal loops — massive arches up to 300,000 miles high of magnetically confined streams of plasma in the sun's atmosphere, or corona, that occur in groups called coronal active regions. Tokman's doctoral research used computational simulations to demonstrate that the structure of the plasma eruptions is in part defined by the rotation of the two footpoints (or sunspots) where the coronal loop joins the sun's surface.

Most recently, she's extended this work to propose a new theory to explain the trigger that turns a staid coronal loop into a cataclysmic solar flare or larger CME.

Recently returned from the American Geophysical Union Conference, Tokman says observations are providing evidence that supports her model.

"Most of the solar flaring that's being observed is coming from boundary zones within the coronal active regions. And this fits exactly with what my model predicts and has not been understood before," says Tokman.

But this year, Tokman's solar modeling is on the back burner. She's turning her attention to creating models of good scientific relations here on Earth as an American Association for the Advancement of Science Diplomacy Fellow in the State Department. Based in Washington, D.C., she's helping improve U.S. relations with other countries through programs in science and math education.

And in September 2006, she'll begin a tenure track appointment as a founding member of the mathematics department at the new Merced campus of the University of California. Tokman says it's her dream job — to be part of an energetic, visionary faculty at a campus founded on an interdisciplinary model.

She knows it will take time, but there's also the knowledge that with incremental steps there's room for big changes.

Tokman learned early and personally how small, incremental changes barely perceptible on the surface of things can lead to enormous changes in the big picture.

Howes Scholars

THE FREDERICK A. HOWES SCHOLAR

in Computational Science award was established in 2001 to honor the late Frederick Anthony Howes who was a champion for computational science education.

Before his death, Dr. Howes was the program manager for the Applied Mathematical Sciences (AMS) Program in the Department of Energy's Mathematical, Information and Computational Sciences (MICS) Division. He had held that position for eight years. Dr. Howes was highly respected by his peers and admired for his energy, dedication and personal integrity.

One of Howes' responsibilities was to oversee the Department of Energy's Computational Science Graduate Fellowship (DOE CSGF) program. Without his support and dedication, the fellowship program may have not survived. Currently, the program is considered one of the most prestigious computational science fellowships in the country, and it is a tribute to his energy and passion.

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

A DOE CSGF fellow is eligible to be named the Howes Scholar if he or she has completed all the requirements for his or her Ph.D. program while being supported by the fellowship program or having been supported by the program for the maximum number of allotted years.

2005 Scholars

An exceptional list of nominees prompted the selection of two awards this year. The winners were Dr. Judith Hill of Carnegie Mellon University and Dr. Ryan Elliott of the University of Minnesota. Dr. Hill, a DOE CSGF fellow from 1999-2003, graduated from Carnegie Mellon University in 2004 with a Ph.D. in Computational Science and Engineering. Dr. Elliott was in the Fellowship program from 2000-2004 and received his Ph.D. in Aero Engineering and Scientific Computing from the University of Michigan in 2004.

Dr. Hill is currently a Post-Doctoral Research Associate at Carnegie Mellon University, and Dr. Elliott is working as an Assistant Professor in the Department of Aerospace Engineering and Mechanics at the University of Minnesota.

Both award recipients were on hand at the DOE CSGF annual fellows' conference where they presented their research and received their awards. David Brown from Lawrence Livermore National Laboratory presented the awards.

For more information on this program, contact the Krell Institute at 515.956.3696 or email Rachel Huisman at rachel@krellinst.org.

David Brown of Lawrence Livermore National Laboratory presents awards to the two 2005 Howes Scholars, Ryan Elliott and Judith Hill.

Judith Hill discusses her current research projects with the attendees at the annual DOE CSGF fellows conference.

Ryan Elliott presents his research to the attendees at the annual DOE CSGF fellows conference.

Alumni Directory

A
B
C

A

Marcelo Alvarez
University of Texas
Astrophysics
Fellowship Years: 2001-2005
Current Status: Student, University of Texas

Asohan Amarasingham
Brown University
Cognitive Science
Fellowship Years: 1998-2002
Current Status: Staff, University of Jaffna, Sri Lanka

Kristopher Andersen
University of California – Davis
Physics
Fellowship Years: 2001-2005
Current Status: Naval Research Laboratory

Matthew Anderson
University of Texas
Physics
Fellowship Years: 2000-2004
Current Status: Staff, Louisiana State University

B

Allison Baker
University of Colorado
Applied Mathematics
Fellowship Years: 1999-2003
Current Status: Lawrence Livermore National Laboratory

Devin Balkcom
Carnegie Mellon University
Robotics
Fellowship Years: 2000-2004
Current Status: Faculty, Dartmouth College

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

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

Martin Bazant
Harvard University
Physics
Fellowship Years: 1992-1996
Current Status: Faculty, MIT

Bonnie Carpenter Beyer
University of Illinois
Mechanical Engineering
Fellowship Years: 1991-1995
Current Status: Rockwell Collins

Edwin Blosch
University of Florida
Aerospace Engineering
Fellowship Years: 1991-1994
Current Status: ESI-CFD Inc.

Dean Brederson
University of Utah
Computer Science
Fellowship Years: 1996-1998

Paul Bunch
Purdue University
Chemical Engineering
Fellowship Years: 1994-1997

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

C

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

Patrick Canupp
Stanford University
Aerospace Engineering
Fellowship Years: 1991-1995
Current Status: Joe Gibbs Racing

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

Nathan Carstens
Massachusetts Institute of Technology
Nuclear Engineering
Fellowship Years: 2001-2004
Current Status: Student, MIT

Edward Chao
Princeton University
Plasma Physics
Fellowship Years: 1992-1995
Current Status: TomoTherapy

Jarrod Chapman
University of California – Berkeley
Computational Biology
Fellowship Years: 1999-2003
Current Status: DOE Joint Genome Institute

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

Kevin Chu
Massachusetts Institute of Technology
Applied Mathematics
Fellowship Years: 2002-2005
Current Status: Staff, Princeton University

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

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

Gavin Conant
University of New Mexico
Biology
Fellowship Years: 2000-2004
Current Status: Trinity College, Dublin Ireland

John Costello
University of Arizona
Applied Mathematics
Fellowship Years: 1998-2002

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

Stephen Cronen-Townsend
Cornell University
Computational Materials Physics
Fellowship Years: 1991-1995
Current Status: Esko-Graphics

Robert Cruise
Indiana University
Physics
Fellowship Years: 1997-2001

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: Faculty, University of Iowa

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

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

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

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

Lewis Jonathan Dursi
University of Chicago
Astrophysics
Fellowship Years: 1999-2003
Current Status: Canadian Institute for Theoretical Astrophysics

E

Ryan Elliott
University of Michigan
Aerospace Engineering
Fellowship Years: 2000-2004
Current Status: Faculty, University of Minnesota

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

Annette Evangelisti
University of New Mexico
Computational Molecular Biology
Fellowship Years: 2001-2005
Current Status: Student, University of New Mexico

F

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

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

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

Michael Feldmann
California Institute of Technology
Computational Chemistry
Fellowship Years: 1999-2002
Current Status: Walleye Trading Advisors LLC

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

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

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

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

G

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

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

Sommer Gentry
Massachusetts Institute of Technology
Optimization/Control Theory
Fellowship Years: 2001-2005
Current Status: Faculty, United States Naval Academy

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

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: Adaptive Optics Associates

Ahna Girshick
University of California – Berkeley
Vision Science
Fellowship Years: 2001-2005
Current Status: Student, University of California – Berkeley

C
D
E
F
G

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

Larisa Goldmints
Carnegie Mellon University
Structural Mechanics
Fellowship Years: 1997-2001
Current Status: General Electric & Rensselaer Polytechnic Institute

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

Catherine Grasso
Cornell University
Bioinformatics
Fellowship Years: 2000-2004
Current Status: Soar Technology, Inc.

Kristen Grauman
Massachusetts Institute of Technology
Computer Science
Fellowship Years: 2001-2005
Current Status: Student, MIT

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

Michael Greminger
University of Minnesota
Mechanical Engineering
Fellowship Years: 2002-2005
Current Status: Seagate Technologies

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

Boyce Griffith
New York University – Courant Institute
Applied Mathematics
Fellowship Years: 2000-2004
Current Status: Faculty, New York University

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

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

H

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

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

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

Heath Hanshaw
University of Michigan
Nuclear Engineering
Fellowship Years: 2001 – 2005
Current Status: Sandia National Laboratories – New Mexico

Rellen Hardtke
University of Wisconsin – Madison
Physics
Fellowship Years: 1998-2002
Current Status: Faculty, University of Wisconsin, River Falls

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

Judith Hill
Carnegie Mellon University
Mechanics, Algorithms & Computing
Fellowship Years: 1999-2003
Current Status: Sandia National Laboratories – New Mexico

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

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

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

Daniel Horner
University of California – Berkeley
Chemistry
Fellowship Years: 2000-2004
Current Status: Lawrence Berkeley National Laboratory

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

Jason Hunt
University of Michigan
Aerospace Engineering & Scientific Computing
Fellowship Years: 1999-2003
Current Status: General Dynamics – Advanced Information Systems

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

I

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

Ahmed Ismail
Massachusetts Institute of Technology
Chemical Engineering
Fellowship Years: 2000-2004
Current Status: Sandia National Laboratories – New Mexico

J

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

K

Richard Katz
Columbia University
Geodynamics
Fellowship Years: 2001-2005
Fellowship Years: Staff, Columbia University

Benjamin Keen
University of Michigan
Mathematics
Fellowship Years: 2000-2004
Current Status: IDA Center for Computing Sciences

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

Sven Khatri
California Institute of Technology
Electrical Engineering
Fellowship Years: 1993-1996
Current Status: Honeywell, Inc.

Benjamin Kirk
University of Texas
Aerospace Engineering
Fellowship Years: 2001-2004
Current Status: NASA Johnson Space Center

Justin Koo
University of Michigan
Aerospace Engineering
Fellowship Years: 2000-2004
Current Status: Advatech Pacific, Inc.

Michael Kowalok
University of Wisconsin
Medical Physics
Fellowship Years: 2000-2004
Current Status: Staff, Virginia Commonwealth University

Yury Krongauz
Northwestern University
Theoretical & Applied Mechanics
Fellowship Years: 1993-1996
Current Status: Black Rock

L

Eric Lee
Rutgers University
Mechanical Engineering
Fellowship Years: 1999-2003
Current Status: Northrup Grumman Corp.

Seung Lee
Massachusetts Institute of Technology
Mechanical Engineering
Fellowship Years: 2001-2005
Current Status: Student, MIT

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

Mary Ann Leung
University of Washington
Theoretical Physical Chemistry
Fellowship Years: 2001-2005
Current Status: Student, University of Washington

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

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

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

M

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

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

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

Randall McDermott
University of Utah
Chemical Engineering
Fellowship Years: 2001-2005
Current Status: Staff, Cornell University

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

Matthew McNenly
University of Michigan
Aerospace Engineering
Fellowship Years: 2001-2005
Current Status: Student, University of Michigan

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

Richard Mills
College of William and Mary
Computer Science
Fellowship Years: 2001-2004
Current Status: Oak Ridge National Laboratory

Julian Mintseris
Boston University
Bioinformatics
Fellowship Years: 2001-2005

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

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

Nathaniel Morgan
Georgia Institute of Technology
Mechanical Engineering
Fellowship Years: 2002-2005
Current Status: Los Alamos National Laboratory

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

Sarah Moussa
University of California – Berkeley
Machine Learning
Fellowship Years: 2003-2005
Current Status: Google

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

N

Heather Netzloff
Iowa State University
Physical Chemistry
Fellowship Years: 2000-2004
Current Status: Staff, Australian National University

Elijah Newren
University of Utah
Mathematics
Fellowship Years: 2001-2005
Current Status: Student, University of Utah

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

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

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

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

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

Catherine Norman
Northwestern University
Applied Mathematics
Fellowship Years: 2000-2004

O

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

P

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

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

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

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

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

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

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

Q

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

R

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

Clifton Richardson
Cornell University
Physics
Fellowship Years: 1991-1995

Christopher Rinderspacher
University of Georgia
Chemistry
Fellowship Years: 2001-2005

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

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

David Ropp
University of Arizona
Applied Mathematics
Fellowship Years: 1992-1995

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

S

Samuel Schofield
University of Arizona
Applied Mathematics
Fellowship Years: 2001-2005
Current Status: Student, University of Arizona

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

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

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

Jason Sese
Stanford University
Computational Materials Science
Fellowship Years: 2003-2005
Current Status: Tutor

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: Invista

Eric Sorin
Stanford University
Chemical Physics
Fellowship Years: 2002-2004
Current: Student, Stanford University

Scott Stanley
University of California – San Diego
Mechanical Engineering
Fellowship Years: 1994-1998
Current Status: Hewlett Packard Company

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

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

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

T

Shilpa Talwar
Stanford University
Scientific Computing
Fellowship Years: 1992-1994
Current Status: Intel

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

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

V

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

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

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

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

W

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

Adam Weller
Princeton University
Chemical Engineering
Fellowship Years: 2001-2002

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

Collin Wick
University of Minnesota
Computational Chemistry
Fellowship Years: 2000-2003
Current Status: Pacific Northwest National 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: Faculty, University of California – Berkeley

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

C. Eric Williford
Florida State University
Meteorology
Fellowship Years: 1993-1996
Current Status: Weather Predict

Matthew Wolinsky
Duke University
Earth Surface Dynamics
Fellowship Years: 2001-2005
Current Status: Staff, University of Minnesota

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

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

Z

Charles Zeeb
Colorado State University
Mechanical Engineering
Fellowship Years: 1993-1997
Current Status: Los Alamos National Laboratory

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

Bree Aldridge
Massachusetts Institute of Technology
Computational Biology

Advisor:
Douglas Lauffenburger
Practicum:
Pacific Northwest National Laboratory
Contact:
breea@mit.edu
Research Synopsis:
My research goals focus on understanding the mechanisms which make cellular signaling networks robust and precise through modeling and mathematical analysis. Cells use large, highly integrated signaling networks to transduce and process multiple input signals. We use computational methods to study the underlying operation of the network and systematic quantitative experimental measurements of several key network components in order to validate the computational model used for analysis. By using mathematical models based on ordinary differential equations representing the elementary chemical reactions (mass-action kinetics), we can make predictions about the response of a cell to a given set of stimuli. Further, the model can be analyzed by studying its nonlinear dynamics, sensitivity, and topology to elucidate the important processing nodes and computing mechanisms embedded in the system.

Teresa Bailey
Texas A&M University
Engineering

Advisor:
Marvin Adams
Practicum:
Lawrence Livermore National Laboratory & Oak Ridge National Laboratory
Contact:
baileyte@tamu.edu
Research Synopsis:
Computational science and reactor physics are related fields of study. Reactor physics, the study of neutron diffusion and transport, uses computational science to model reactor behavior. Diffusion and transport problems are too complex to solve analytically because they are dependent on three spatial dimensions, energy, angle, and time. For this reason, much of the focus of reactor physics is on developing computational methods to find efficient solutions for the system of equations that governs a physical model. Computers are the essential tools to perform the large amount of calculations required during the evaluation of these models.

Michael Barad
University of California – Davis
Environmental Modeling

Advisor:
Geoffrey Schladow
Practicum:
Lawrence Berkeley National Laboratory
Contact:
mfbarad@ucdavis.edu
Research Synopsis:
My field of interest is the numerical study of environmental transport in fluid systems. My case study is the San Francisco Bay for which I have already collected a large validation dataset. My model is based on the variable density incompressible Navier-Stokes equations in 3D, including air/water and fluid/solid interfaces and the transport of passive constituents. My numerical methodology is based on a second-order accurate projection method with high-order accurate Godunov finite differencing including slope limiters and a stable differencing of the nonlinear convection terms. This is a proven methodology for hyperbolic problems that yields accurate transport with low phase error while minimizing the numerical diffusion at steep gradients typically found in 'classical' high order finite difference methods. For the fast and robust solution of the Poisson equations in my model I am using a geometric multigrid method. This methodology is placed in EBChombo, a three-dimension parallel adaptive mesh refinement (AMR) framework, developed at the Applied Numerical Algorithms Group (ANAG) at Lawrence Berkeley National Laboratory. The application of this methodology together with high-order accurate volume-of-fluid and level set interface tracking methods for modeling the air/water interface to environmental problems is new. Most of this numerical methodology was developed at DOE laboratories or through DOE Office of Science funding.

Jaydeep Bardhan
Massachusetts Institute of Technology
Electrical Engineering

Advisor:
Jacob White
Practicum:
Argonne National Laboratory
Contact:
jbardhan@mit.edu
Research Synopsis:
My research focuses on the application of boundary element methods (BEM) to electrostatics calculations in molecular biochemistry. Since electrostatic forces act over a long range, and they can have both favorable and unfavorable effects on protein-protein interactions, much rational drug design research has focused on methods to optimize the charge distribution in a ligand to make a ligand-receptor interaction as favorable as possible, in terms of both affinity and specificity.

The boundary element method can be accelerated using the precorrected FFT algorithm, resulting in a "matrix-free" calculation of the potential on the ligand surface; because even this fast method is computationally expensive, we would like to minimize the number of potential calculations required in order to find the optimal charge distribution. With a new BEM formulation in hand, I have most recently been working to develop matrix-free optimization methods for this problem. Although presently only linear inequality constraints are treated, I would like to extend the present solution method to handle nonlinear constraints as well.

My long-range research goals include the incorporation of other forces — such as hydrophobic and van der Waals forces — into the optimization procedure, and investigating BEM formulations that permit the variation of the charge locations as well as the charge values.

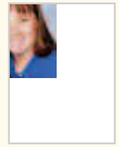

Mary Biddy
University of Wisconsin
Engineering

Advisor:
Juan de Pablo
Practicum:
Sandia National Laboratories – New Mexico
Contact:
mbiddy@students.wisc.edu
Research Synopsis:
Concern over the availability of petroleum-based products has generated considerable interest in finding alternative renewable sources for these products. In the case of industrial lubricants, vegetable oils are a promising alternative to comparable petroleum-based products. Unlike petroleum oils, vegetable oils offer the advantages of being environmentally friendly, renewable, and biodegradable resources.

Vegetable oils are essentially mixtures of triglycerides. Although triglycerides are naturally abundant, little is understood about the effect of molecular structure on their physical properties. To formulate vegetable oils with optimal lubrication characteristics, however, the influence of the individual molecular triglycerides on the overall oil properties must be understood. Computer simulations can directly relate molecular composition to physical properties and offer a valuable tool for improving our knowledge of this important class of materials.

By employing molecular dynamics simulations, we have obtained viscosities and densities that agree with experimentally observed values for both vegetable oils and triglycerides. The effect of molecular structure on physical properties is and continues to be explored for triglycerides that are difficult to isolate in vegetable oils and that are chemically modified. The low-temperature properties of vegetable oils, which are a major limitation in their use as lubricants, have also been characterized using molecular modeling methods. The ultimate goal of this work is to use the knowledge gained from molecular-level modeling to engineer a vegetable oil-based lubricant.

Nawaf Bou-Rabee
California Institute of Technology
Applied and Computational Mathematics

Advisor:
Jerrold Marsden
Practicum:
Los Alamos National Laboratory
Contact:
nawaf@acm.caltech.edu
Research Synopsis:
Currently, my research covers the following topics in geometric mechanics: geometric integration on manifolds and linear spaces, dissipation-induced instabilities, nonlinear stability theory, and Hamiltonian chaos. I am also thinking about optimal parallel computing environments for computational dynamical systems applications.

Kristine Cochran

University of Illinois – Urbana – Champaign
Structures

Advisor:
Keith Hjelmstad
Practicum:
Sandia National Laboratories – New Mexico
Contact:
kbergero@uiuc.edu
Research Synopsis:
The focus of my research is the development and implementation of continuum mechanics based constitutive models to represent cyclic metal plasticity. The intended application is the modeling of low-cycle fatigue crack growth in metals, within a three-dimensional finite element code. Two nonlinear cyclic plasticity models have been added to a research oriented finite element code. The nonlinear kinematic hardening Frederick-Armstrong model provides improved cyclic response over the commonly used linear hardening model. More recently, the generalized plasticity model, developed by Lubliner and improved by Auricchio and Taylor, was implemented into the code. This model also features nonlinear response, but solves some of the theoretical and numerical issues inherent in the Frederick-Armstrong model as applied to low-cycle fatigue modeling. Both models are implemented with substepping error control and make use of the algorithmically consistent tangent.

Gregory Davidson

University of Michigan
Nuclear Engineering

Advisor:
Ed Larsen
Practicum:
Bettis Atomic Power Laboratory
Contact:
davidsgr@umich.edu
Research Synopsis:
My field of interest is particle transport. Particle transport involves the study of radiation flowing through and interacting with a background medium. The Boltzmann transport equation which mathematically describes this process is far too complex to solve analytically, except in the most elementary and idealized of problems. Computational methods are therefore essential to finding solutions to any realistic transport problem.

The transport equation is an integro-differential equation, which typically requires an iterative method to solve. Often, the transport equation is spatially discretized using a finite-element method. Advances in computer power have made deterministic transport algorithms discretized on unstructured, arbitrary polygonal and polyhedral meshes viable. Unfortunately, the development of robust basis functions for arbitrary zone shapes has lagged behind other progress in the field. In 1975, Eugene Wachspress published a book unveiling a new class of rational basis functions that may be applied to any convex polygon or polyhedron, as well as certain polycons (zones with curved sides). We know of no one having implemented these basis functions in a finite element radiation transport algorithm.

My research involves investigating the properties of Wachspress rational basis functions, and evaluating their performance in solving radiation flow problems, particularly in the diffusive limit.

Michael Driscoll

Boston University
Bioinformatics & Systems Biology

Advisor:
James Collins
Practicum:
Lawrence Berkeley National Laboratory
Contact:
mdriscol@bu.edu
Research Synopsis:
The discovery that the human genome contains far fewer genes than had been widely believed underscores a point: the richness of the human organism derives not from genes per se, but from interactions of genes. We are fast moving away from a time when we can say "this is a gene for X", toward recognizing that biological function — and its counterpart, disease — are products of systems.

I am interested in studying the networks that represent these interactions of genes. The ability of a cell to change during development, division, or in response to its environment hinges on a balance of competing and noisy factors. Even the smallest unit of change, the turning on or off of one protein's production, is governed by several other proteins, which in turn are also regulated. These genetic regulatory networks form the architecture for a cell's genomic program, and I believe understanding their dynamics is a step towards understanding a cell's higher functions and, in the case of disease, malfunctions.

While traditional experimental work will be critical to this understanding, I intend to focus on developing numerical models of gene regulatory networks. These models rely on existing experimental knowledge, as well as principles of nonlinear systems, to predict the behavior of a given network of genes. Their ultimate validation is found by creating an engineered cellular system, where a desired genetic program — an inducible switch or a steady oscillator, for example — is achieved through the design of genes and their regulatory elements.

Mary Dunlop

California Institute of Technology
Mechanical Engineering

Advisor:
Richard Murray
Practicum:
Los Alamos National Laboratory
Contact:
mjdunlop@caltech.edu
Research Synopsis:
My research is in control and dynamical systems with applications to biology.

The field of synthetic biology has recently surged as researchers discover how to manipulate genetic circuits to accomplish simple tasks. Toy problems such as blinking bacteria, cells that use digital logic to play tic-tac-toe, and bacteria that perform basic Boolean computations are all developments that have occurred within the past 5 years. Such experimental advances leave biologists and engineers alike asking the question: "Can we redesign control systems in microbiological organisms?"

But as this fascinating field of engineering biological systems unfolds it becomes apparent that biologists need better tools to help them understand and characterize what is actually going on within the cell. Current synthetic gene network design is driven by expert biologists whose understanding of very specific pathways enables them to develop ad hoc designs, without adding to the understanding of the field of gene circuit design as a whole. Biological systems are an enticing area to work in because so many genetic pathways have been characterized experimentally without a systematic approach towards understanding overarching principles.

Owen Hehmeyer

Princeton University
Chemical Engineering

Advisor:
Athanassios Z. Panagiotopoulos
Practicum:
Sandia National Laboratories – New Mexico
Contact:
hehmeyer@princeton.edu
Research Synopsis:
For the first half of my studies, my work focused on the molecular simulation of polymers in confined geometries. Going forward, I will work towards more complex models of tethered polymers, and study their interactions with charged colloid (protein-like) particles. In my first year of research, I developed understanding of the needed Monte Carlo and other algorithms, and also created a new implementation to compute electrostatic forces in non-periodic geometries. Over the next six months, I completed work on the thermodynamics and structure of lattice homopolymers in quasi 2-D and 2-D geometries. Next, the structure of uncharged tethered homopolymers was studied theoretically. During the course of my practicum at Sandia and afterwards, I studied electrically charged, tethered polymers using molecular dynamics simulations. For the remainder of my studies, I will continue to focus on polymers in confined geometries, but will seek also to understand not only their structure, but also their interactions with other, larger particles. Such study could reveal how polymers affect the adsorption of proteins on surfaces.

Yan Karklin

Carnegie Mellon University
Computational Neuroscience

Advisor:
Michael Lewicki
Practicum:
Lawrence Berkeley National Laboratory
Contact:
yan+fellowship@cs.cmu.edu
Research Synopsis:
My current research focuses on the development of computational models of biological visual systems. Recent work, aided by probabilistic modeling techniques in computer science and information theoretical techniques in signal processing, has yielded some insights into the basic principles underlying computation in the human visual system. Using only the notion that the brain's task is to extract useful information from a noisy scene and represent it efficiently, it has been possible to deduce theoretically optimal representations for low level visual features without making assumptions about the kinds of visual structure we expect to discover.

The resulting learned representations have been shown to relate closely to the responses of cells in early visual processing areas of the brain. The appeal of this approach lies in the fact that it is not limited by our incomplete understanding of the physiology of the visual system, as it attempts to lay out the basic theoretical principles that apply to any sensory processing system.

Benjamin Lewis
Massachusetts Institute of Technology
Computational Biology

Advisor:
Chris Burge
Practicum:
Lawrence Berkeley National Laboratory
Contact:
benlewis@mit.edu
Research Synopsis:
In recent years, the rapid accumulation of biological data has given rise to a new discipline: bioinformatics. Generally, bioinformatics involves the application of techniques from statistics, computer science, and mathematics, to organize, interpret, and integrate large quantities of data describing macromolecules of biological interest. Much of the current focus of bioinformatics is on three major categories of information: DNA and protein sequences, macromolecular structures, and functional genomics data (e.g. gene expression profiles, two-hybrid interaction data). Unifying elements of biology and computer science, bioinformatics involves studying the informational content of biological molecules and the use of computational tools to perform comparisons, to identify relationships and trends, to build models of interactions and networks, etc.

Alex Lindblad
University of Washington
Structural Engineering

Advisor:
George Turkiyyah
Practicum:
Sandia National Laboratories – California
Contact:
alind@u.washington.edu
Research Synopsis:
My interest is in the simulation and design of high-performance structures — systems that can withstand extreme loading conditions with predictable performance. Earthquakes, fires, explosions, winds/hurricanes, etc., place severe demands on structural systems. Balancing economic considerations with predictable performance in these systems is a challenge. Constructing accurate usable models of these systems requires not only a firm understanding of the mechanics behind the problem, but also a considerable knowledge base in the field of mathematics with an emphasis in computation. The availability of high performance computers has made it possible to build high-fidelity numerical simulations that can predict behavior accurately, and allow us to mechanistically rather than empirically make design tradeoffs and compromises.

Gregory Novak
University of California – Santa Cruz
Theoretical Astrophysics

Advisor:
Sandra Faber
Practicum:
Lawrence Berkeley National Laboratory
Contact:
novak@ucolick.org
Research Synopsis:
My research is designed primarily to answer two questions: 1.) Considering nearby galaxies as a population, can their detailed internal dynamics (as probed by Integral Field Unit, or IFU, observations) be explained by state-of-the-art simulated galaxy merger remnants? If not, what are the physical processes missing from the models that are causing the discrepancy? 2.) What can we learn about the merger history of an individual galaxy from its detailed internal dynamics?

These questions are interesting, relevant, and ripe today because of the recent maturation of IFUs with wide fields of view and spectral resolution sufficient to study internal dynamics of galaxies (e.g. Emsellem et al. 04). Both the quantity and quality of the data are increasing dramatically.

Comparing the simulations to the observations on equal footing is not trivial because of the limited number of particles in the simulations. The brute force solution is to use more particles, but that may not be feasible in terms of computer time. Bendo and Barnes (00) and Jesseit, Naab and Burkert (04) have published computationally cheaper methods of "smoothing" the simulated observations.

David Schmidt
University of Illinois – Urbana – Champaign
Electrical Engineering

Advisor:
Richard Blahut
Practicum:
Sandia National Laboratories – New Mexico
Contact:
dschmidt@uiuc.edu
Research Synopsis:
My field of interest within electrical engineering is signal processing, and one of the main reasons I was drawn to signal processing more than other areas of electrical engineering is the amount of mathematics that is found in the field. I really like how useful the concepts I learned in real and complex analysis are to information theory, and especially how topics in algebra show up once again in understanding and evaluating error-correcting codes. I plan to do research in error-correcting codes because they bring together my favorite aspect of electrical engineering and my preferred area of mathematics. I knew it was what I wanted to study as soon as I picked up a book on coding theory and saw Galois theory used outside a math classroom. I am also greatly interested in multi-user detection and space-time processing after my introduction to them in the class I am taking on wireless communication from Professor H. Vincent Poor. Given the explosion of cellular technology and the huge gains to be enjoyed by incorporating them into one's wireless system, they seem like important areas of research, beyond simply being extremely interesting to study in their own right. Digital communications and error-correcting codes are very strongly related to computational science because of their importance to communications in general and a computer's inherent use of digital signals. Furthermore, because many signal processing models require processor intensive simulations, high performance computing is integral in the advancement of the field.

Amoolya Singh
University of California – Berkeley
Computational Biology

Advisor:
Richard Karp
Practicum:
Pacific Northwest National Laboratory
Contact:
agni@cs.berkeley.edu
Research Synopsis:
For my doctoral work I am interested in developing mathematical models and computational simulations of cellular regulatory networks.

Since this is a huge and intricate field, I would like to initially simulate a simple cellular system to gain computational and biological insights. Under analysis in the Arkin lab are several candidate problems: chemotaxis and type-1 pili phase variation of *E. coli*, sporulation and chemotaxis of *B. subtilis*, lysis/lysogeny pathways in Lambda phage, G-protein coupled signal transduction pathways in cardiomyocytes, etc. I intend to generalize the insights obtained from these specific simulations to construct high performance computational tools. I will validate these tools by showing that they can be applied to other computational biology problems such as climate change modeling, population genetics, or protein folding.

Obioma Uche
Princeton University
Materials/Statistical Mechanics

Advisor:
Salvatore Torquato
Practicum:
Sandia National Laboratories – New Mexico
Contact:
ouuche@yahoo.com
Research Synopsis:
The determination of the macroscopic and microscopic properties of materials via molecular simulation methods is of particular interest to me. The above topic is one that has its foundation in theoretical chemistry and physics. Molecular simulations methods, e.g. Monte Carlo simulation, rely heavily on high-performance computation for a tractable solution to the problem. In addition, new computer algorithms are always being written to solve problems that arise in the application of statistical mechanics to a variety of cases.

In general, my research involves the use of optimization techniques to study optimal geometric packing arrangements of particles in various geometries.

Joshua Waterfall
Cornell University
Biophysics

Advisor:
James Sethna
Practicum:
Lawrence Berkeley National Laboratory
Contact:
jjw36@cornell.edu
Research Synopsis:
Borrowing methods from statistical physics, I am studying how certain bacteria combine cell-cell signaling with gene regulation. Currently I am investigating the network which controls conjugal gene transfer in Agrobacterium tumefaciens and the pathogenic regulatory network of Pseudomonas syringae. Both of these networks alter gene expression in response to many chemical signals. In particular, these two networks include a quorum sensing module, whereby the bacterium changes its gene expression in response to the density of its neighbors. Such responses are present in many bacteria and allow not only for intraspecies communication but for cross talk between bacteria of different species and lead to large scale, collective behavior.

Michael Wu
University of California – Berkeley
Computational Neuroscience

Advisor:
Jack Gallant
Practicum:
Los Alamos National Laboratory
Contact:
wafting@berkeley.edu
Research Synopsis:
My research will focus on the computational analysis and modeling of neurophysiological data collected from experiments that are designed to unravel the visual processing in the primary visual cortex (V1). If one treats a neuron as a computing unit that transforms input (stimuli) into action potentials, then its transfer function, or kernel, can be estimated using mathematical techniques. The kernel might reveal the computations mediating extraction of visual features to which the cell is tuned. More sophisticated nonlinear kernel estimation methods could enable us to understand processing in extra-striate visual areas, such as V2 and V4.

Paul Bauman
University of Texas
Computational and Applied Mathematics
Advisor:
J. Tinsley Oden
Practicum:
Sandia National Laboratories – New Mexico
Contact:
pbauman@ices.utexas.edu

William Conley
Purdue University
Nanoscale Mechanics
Advisor:
Arvind Raman
Practicum:
Sandia National Laboratories – New Mexico
Contact:
wconley@ecn.purdue.edu

Aron Cummings
Arizona State University
Electrical Engineering
Advisor:
David Ferry
Practicum:
Sandia National Laboratories – California
Contact:
Aron@turbonet.com

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

Jasmine Foo
Brown University
Applied Mathematics
Advisor:
George Karniadakis
Practicum:
Lawrence Berkeley National Laboratory
Contact:
jfoo@dam.brown.edu

Tod Pascal
California Institute of Technology
Physical Chemistry
Advisor:
William Goddard III
Practicum:
Sandia National Laboratories – New Mexico
Contact:
tpascal@wag.caltech.edu

Christina Payne
Vanderbilt University
Chemical Engineering
Advisor:
Peter Cummings
Practicum:
Sandia National Laboratories – New Mexico
Contact:
Christina.payne@vanderbilt.edu

Emma Rainey
California Institute of Technology
Planetary Sciences
Advisor:
David Stevenson
Practicum:
Argonne National Laboratory
Contact:
emma@gps.caltech.edu

Mark Rudner
Massachusetts Institute of Technology
Physics
Advisor:
Leonid Levitov
Practicum:
Brookhaven National Laboratory
Contact:
rudner@mit.edu

Samuel Stechmann
New York University
Applied Mathematics
Advisor:
Andre Majda
Practicum:
Los Alamos National Laboratory
Contact:
stechman@cims.nyu.edu

Brian Taylor
University of Illinois – Urbana – Champaign
Engineering Mechanics
Advisor:
Scott Stewart
Practicum:
Lawrence Livermore National Laboratory
Contact:
bdtaylo1@uiuc.edu

William Triffo
Rice University
Bioengineering
Advisor:
Robert Raphael
Practicum:
Lawrence Berkeley National Laboratory
Contact:
triffo@rice.edu

Michael Wolf
University of Illinois – Urbana – Champaign
Computer Science
Advisor:
Michael Heath
Practicum:
Lawrence Berkeley National Laboratory
Contact:
mmwolf@uiuc.edu

Brandon Wood
Massachusetts Institute of Technology
Computational Materials Science
Advisor:
Nicola Marzari
Practicum:
Lawrence Berkeley National Laboratory
Contact:
brandonw@mit.edu

Erik Allen
Massachusetts Institute of Technology
Chemical Engineering
Advisor:
Kenneth Beers
Practicum:
Sandia National Laboratories – New Mexico
Contact:
ecallen@mit.edu

Michael Bybee
University of Illinois
Chemical Engineering
Advisor:
Jonathan Higdon
Practicum:
Lawrence Livermore National Laboratory
Contact:
bybee@uiuc.edu

Jimena Davis
North Carolina State University
Applied Mathematics
Advisor:
H.T. Banks
Contact:
jldavis9@unity.ncsu.edu

Jeffrey Drocco
Princeton University
Computational Condensed Matter Physics
Advisor:
Shivaji Sondhi
Contact:
jdrocco@princeton.edu

Peter Kekenes-Huskey
California Institute of Technology
Computational Chemistry/Biology
Advisor:
William Goddard
Practicum:
Sandia National Laboratories – New Mexico
Contact:
huskeypm@wag.caltech.edu

SECOND YEAR FELLOWS

Bonnie Kirkpatrick
University of California – Berkeley
Computer Science
Advisor:
Richard Karp
Practicum:
Lawrence Livermore National Laboratory
Contact:
bbkirk@eecs.berkeley.edu

Matthew McGrath
University of Minnesota
Physical Chemistry
Advisor:
Ilja Siepmann
Practicum:
Lawrence Livermore National Laboratory
Contact:
mcgrath@chem.umn.edu

Ian Parrish
Princeton University
Computational Plasma Physics
Advisor:
James Stone
Practicum:
Sandia National Laboratories – New Mexico
Contact:
iparrish@princeton.edu

David Potere
Princeton University
Remote Sensing/GIS
Advisor:
Burt Singer
Practicum:
Oak Ridge National Laboratory
Contact:
dpotere@princeton.edu

Mala Radhakrishnan
Massachusetts Institute of Technology
Physical Chemistry
Advisor:
Bruce Tidor
Practicum:
Lawrence Berkeley National Laboratory
Contact:
mradhakr@mit.edu

Amber Sallerson
University of North Carolina – Chapel Hill
Applied Mathematics
Advisor:
Roberto Camassa
Practicum:
Lawrence Berkeley National Laboratory
Contact:
asalle1@email.unc.edu

Michael Veilleux
Cornell University
Structural Fracture Mechanics
Advisor:
Anthony Ingraffea
Practicum:
Sandia National Laboratories – New Mexico
Contact:
mgv5@cornell.edu

Allan Wollaber
University of Michigan
Nuclear Engineering
Advisor:
Edward Larsen
Practicum:
Los Alamos National Laboratory
Contact:
wollaber@umich.edu

Etay Ziv
Columbia University
Computational Biology
Advisor:
Chris Wiggins
Contact:
ez87@columbia.edu

John ZuHone
University of Chicago
Astrophysics
Advisor:
Donald Lamb
Contact:
zuhone@uchicago.edu

FIRST YEAR FELLOWS

Joshua Adelman
University of California – Berkeley
Biophysics
Advisor:
George Oster
Contact:
jadelman@ocf.berkeley.edu

Zlatan Aksamija
University of Illinois – Urbana – Champaign
Electrical Engineering
Advisor:
Umberto Ravaioli
Contact:
aksamija@uiuc.edu

Jordan Atlas
Cornell University
Chemical Engineering
Advisor:
Michael Shuler
Contact:
jca33@cornell.edu

Christopher Carey
University of Wisconsin
Plasma Physics
Advisor:
Carl Sovinec
Contact:
cscarey@wisc.edu

Ethan Coon
Columbia University
Applied Mathematics
Advisor:
Marc Spiegelman
Contact:
etc2103@columbia.edu

Jeff Hammond
University of Chicago
Theoretical Chemistry
Advisor:
David Mazziotti
Contact:
jhammond@uchicago.edu

Asegun Henry
Massachusetts Institute of Technology
Mechanical Engineering
Advisor:
Gang Chen
Contact:
ase@mit.edu

Kevin Kohlstedt
Northwestern University
Bio-Polymer/Soft Matter Computation
Advisor:
Monica Olvera de la Cruz
Contact:
kohlstedt@northwestern.edu

Miler Lee
University of Pennsylvania
Genomics and Computational Biology
Advisor:
Junhyong Kim
Contact:
miler@mail.med.upenn.edu

Jeremy Lewi
Georgia Institute of Technology
Neuroengineering
Advisor:
Robert Butera
Contact:
gtg120z@mail.gatech.edu

David Markowitz
Princeton University
Computational Neurobiology
Advisor:
David Tank
Contact:
dave@princeton.edu

Peter Norgaard
Princeton University
Computational Plasma Dynamics
Advisor:
Edgar Choueiri
Contact:
norgaard@princeton.edu

Natalie Ostroff
University of California – San Diego
Bioengineering
Advisor:
Jeff Hasty
Contact:
nostroff@ucsd.edu

Christopher Schroeder
University of California – San Diego
Physics
Advisor:
Julius Kuti
Contact:
crs@physics.ucsd.edu

Stefan Wild
Cornell University
Operations Research
Advisor:
Christine A. Shoemaker
Contact:
stefan@orie.cornell.edu