2006 - 2007

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

2006 - 2007
2006 - 2007
DEIXIS – THE DOE CSGF ANNUAL
DEIXIS – THE DOE CSGF ANNUAL

THE DOE CSGF ANNUAL

DEIXIS

Department of Energy Computational Science Graduate Fellowship

TABLE OF CONTENTS

2006-2007

Editor
Shelly Olsan

Copy Editor
Ron Winther

Design
Juls Design, Inc.

Contributors
Jacob Berkowitz
Alan S. Brown
Victor D. Chase
Karyn Hede
Thomas R. O'Donnell

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

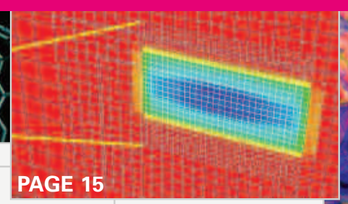

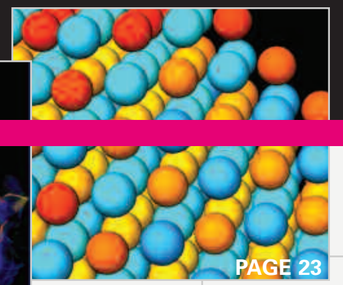

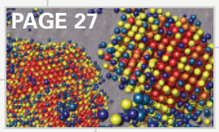

DEIXIS (ΔΕΙΞΙΣ) *transliterated from classical Greek into the Roman alphabet, (pronounced dāksis) means a display, mode or process of proof; the process of showing, proving or demonstrating.* DEIXIS *can also refer to the workings of an individual's keen intellect, or to the means by which such individuals,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

A Job That's a Journey

THE DEPARTMENT OF ENERGY Computational Science Graduate Fellowship (DOE CSGF) has a 15-year track record in training the nation's next generation of groundbreaking researchers. These young scholars enter their fields equipped to apply powerful computing resources to the most complex problems in science and engineering.

Left to right: David Potere, Brandon Wood, Bonnie Kirkpatrick and Aron Cummings.

THE PRACTICUM EXPERIENCE is one of the major aspects that distinguishes the DOE CSGF. For three months, usually in the summer, students join a research team at one of the Department of Energy's national laboratories. It's more than just a summer job — fellows take on projects that are distinct from their doctoral research. It has changed some career paths.

These stories describing four students' practica demonstrate the impact of the experience. Each faced new challenges, ventured into new areas, and learned the topography of the national laboratory system. They may be headed to a lab near you.

Tracking Wal-Marts to Make Better Maps

DAVID POTERE

Princeton University | Oak Ridge National Laboratory | Story by Thomas R. O'Donnell

David Potere has lived in 18 homes in 29 years and visited five continents in six years, first as the son of an Air Force pilot and then as a Navy lieutenant. Despite that, "I'm always lost," he jokes.

So it's no wonder Potere's path led to geography. Since entering graduate school, he's become part of a growing community of researchers who use high-performance computers and satellite imagery to make better maps — especially ones showing where people live. In his doctoral degree research at Princeton University and his Department of Energy Computational Science Graduate Fellowship (DOE CSGF) practicum, Potere (pronounced poe-TEER) is refining techniques to make global maps of human population more accurate.

With satellites transmitting complete images of the Earth about every 48 hours, "We have so much information coming out of space right now — more than a terabyte (one trillion bytes) of data per day — we need high-performance computers to sort through this," says Potere. "There's a lot of speculation about what answers might be buried in the data, and people are just figuring out what questions to ask."

For Potere, the questions have revolved around defining urban areas – and tracking Wal-Mart stores.

Potere's work at Boston University and Princeton got the attention of Budhendra Bhaduri, leader of the Geographic Information Science and Technology (GIST) Group at DOE's Oak Ridge National Laboratory (ORNL). "An announcement was circulated about DOE CSGF," Bhaduri says. "I browsed that list and I was shocked to find a geography major in there." Bhaduri invited Potere to do his practicum with GIST in summer 2005.

GIST's LandScan Global Population Project uses high-performance computers to calculate worldwide

FIGURE 2A. Durham, NC, March 1993. (Pixxures/Terraserver)

FIGURE 2B. Durham, NC. Wal-Mart 2005. (DigiGlobe/GoogleEarth)

FIGURE 2C. Durham, NC, Wal-Mart and the LEDAPS (Landsat Ecosystem Disturbance Adaptive Processing System) Disturbance Index (DI) product. DI Map representing change in the Disturbance Index from 1987-2002. Red and orange are areas of forest disturbance, yellow is stable forest, and green is forest regrowth. DI map is based on Landsat GeoCover imagery from 14 October 1987 and 24 May 2002.

With satellites transmitting complete images of the Earth about every 48 hours, "We have so much information coming out of space right now — more than a terabyte (one trillion bytes) of data per day — we need high-performance computers to sort through this," says Potere.

FIGURE 3. Pre- and post-construction NDVI time series at the Wal-Mart Distribution Center in Apple Valley, CA.

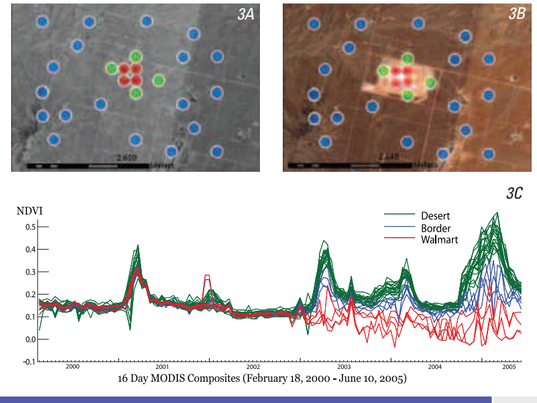

FIGURES 3A-3C. Figures 3a-3b are pre- and post-construction satellite images of an 11 hectare Wal-Mart Distribution Center in Apple Valley, CA (from GoogleEarth, DigiGlobe, and TerraServer). The circles are 250 m diameter sample points of the desert background vegetation (blue), the building footprint of the Wal-Mart distribution center (red), and several points located on the border between the desert and the Wal-Mart (green). Figure 3c is a twice-monthly vegetation index time series for all of the red, green, and blue points for the years 2000-2005. The red Wal-Mart pixels deviate from the blue and green desert and border pixels during the winter of 2002/2003 — roughly one year prior to the opening of the Apple Valley facility, in March 2003.

population distribution based on census data, topography, transportation and other data. Computer programs estimate population even where no specific census figures are available.

Just as importantly, the project estimates population at a given time, such as on a busy highway through a desert area. A census would show the desert has little or no population, although the highway may have thousands of people on it at one time. Such estimates are important to calculate risks and plan for natural disasters, terrorist attacks and industrial accidents.

LandScan is "the community standard for the whole world in terms of disaster response and disaster management," Bhaduri says. "LandScan is the finest population data that has ever been produced," estimating population in each of hundreds of millions of cells of about one square kilometer. The project produces a new map every year, but there is continuous research into improving it.

Maps are only as good as the information they're based on, Potere says. "If you want to build the kind of map we want — to talk about space and time, where has there been change and when has there been change…you need to have some sort of ground-truth" — accurate data on where and when a change occurred. Potere and his fellow researchers needed a set of large, documented geographic and demographic changes to validate their mapping tools. They settled on Wal-Mart stores because of their large size and precise opening dates.

With Neal Feierabend, a student research assistant, and Edward Bright, a remote sensing and Geographic Information System (GIS) specialist who heads the LandScan Global Project, Potere translated the street addresses of more than 3,000 stores into map coordinates. Then the researchers tapped high-resolution satellite images to precisely locate the coordinates of each store.

The researchers next looked at "signatures" for selected stores. They examined data from NASA's MODIS (Moderate Resolution Imaging Spectroradiometer) project — two satellites circling the Earth over the poles to produce a complete image of the planet every one to two days. A "greenness index" of the images indicated how much was in vegetation.

The index rises and falls with the seasons, as plants grow greener in the spring and summer and recede in the fall and winter, Potere says. "What we were looking for is a depression in that signal, which would indicate a change in use, and we found it," he says. "The greenness index still responds to the seasons, but it does so far less dramatically" over areas where a Wal-Mart store has been built.

Such data can "train" high-performance computers to recognize land-use changes in the masses of satellite images, Potere says. "If I'm interested in new facilities in the forests of North Korea, I can't go there, but I know what a forest looks like in Maine and I know what a Wal-Mart looks like in Maine," he adds. Computers could scan MODIS images from anywhere on the planet and spot similar changes. Bhaduri says computers also could spot urbanization, illegal logging, crop disease and other land-use changes.

Computer automation ties into another of Potere's practicum projects. With Bhaduri and other GIST researchers including Anil Cheriyadat, he helped refine an algorithm that precisely delineates urban boundaries in high-resolution satellite images. Good estimates of a city's boundaries are important because they affect how researchers distribute census counts.

The program Potere and his fellow researchers refined correlates gray levels in the photos with edges of geographic features. Urban areas typically have more edge features. The result is an image with urban areas that are precisely — and quickly — defined. The algorithm ran on ORNL's eXtreme Tennessee Oak Ridge Cluster, comprised of 64 2-gigahertz Pentium IV nodes and two 1.7-gigahertz Pentium IV head nodes. The parallel computing approach drastically cut execution time.

Despite his globe-trotting tendencies, Potere's formal geography studies began relatively recently. He earned a bachelor's degree in American History at Harvard and did his Navy hitch before starting in the geography program at Boston University.

Potere learned of the DOE CSGF from Michael Driscoll, a Harvard classmate who earned a fellowship to study computational biology. "He couldn't say enough good things about the program," Potere adds.

Potere's Boston University research focused on forest loss along the Appalachian Trail — a project that capitalized on documentation gathered during a four-month trek of the famed route. But after completing his master's thesis, Potere found himself more interested in the human aspects of geography. In the fall of 2005, he moved to Princeton's Office of Population Research to study demography.

His dissertation is rooted in a third ORNL-GIST project. Potere, ORNL student research assistant Karen McNeany, and Annemarie Schneider, assistant geography professor at the University of California, Santa Barbara, compared six independently produced global maps of urban land cover. Although they're based on common sources, "There are order of magnitude differences between these products," Potere says. His dissertation will use these global maps for things like modeling disease spread and predicting epidemics. Potere's doctoral advisor, Princeton Demography and Public Affairs Professor Burton Singer, says remote sensing technology like satellite images has "the potential to really change the game" in predicting the spread of diseases, particularly those borne by insects and parasites. Potere, he says, came to his program "with a much more sophisticated understanding of the remote sensing technology than any other student I've had."

Potere and his Oak Ridge colleagues plan to publish papers on the Wal-Mart project and the urban boundary algorithm. Researchers presented data on the Wal-Mart project at the annual meeting of the Association of American Geographers. Potere and Schneider also presented a paper on their global urban mapping project at the meeting.

Potere expects to earn his doctorate in spring 2008.

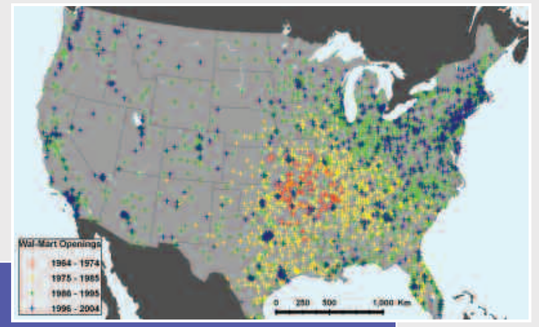

Spatial-temporal distribution of U.S. Wal-Marts 1962-2004.

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, over 20% of program alumni work or have worked at 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.

Model Puts a Charge into Nanotube Loops

ARON CUMMINGS

Arizona State University | Sandia National Laboratories – California |
Story by Thomas R. O'Donnell

Aron Cummings loves the kicks, punches and blocks of the martial arts. The physical forces governing atoms and particles also fascinate him. And he sees similarities between the two.

Cummings, a U.S. Department of Energy Computational Science Graduate Fellow (DOE CSGF) studying at Arizona State University, investigates the physics that govern the atomic-scale future of computing. He's also a second-degree black belt and Arizona State University's head instructor in Kokondo martial arts, a combination of karate and jujitsu.

"When you're approaching your black belt level, you think, 'Hey, I'm getting pretty good,'" says Cummings, a Washington state native. "Then you realize there's this guy who has his black belt, and he's much better than you are," and you realize there is much more to learn.

Physics is the same way, Cummings says: "You can understand things at a high level, but when you dig down, there's more to learn."

In both physics and the martial arts, "When you go even deeper, things actually become simpler," Cummings says. In Kokondo, "We have hundreds of techniques, but they're all really based on just a few principles. On the physics side, at the basic level there are only a few laws that govern how everything works."

Cummings' career has been driven by the desire to dig deeper. He started as a computer science student, with plans to become a programmer. "Something about it always bothered me," he says. When he wrote a program, he wondered what was really going on in the guts of the computer.

He switched to computer engineering, but "It really seemed to me the same as computer science," Cummings says. He understood how programs govern circuits and transistors, but "Those are still a level of abstraction as to what's going on below." He embarked on a master's degree at Washington State to understand the physics behind the circuits.

That led Cummings to research on computing and electronics at the nanoscale — nearly the atomic level — in his quest for an electrical engineering doctorate and in a DOE CSGF practicum at Sandia National Laboratory at Livermore, California. His unusual combination of programming, engineering and physics is perfect for simulating the properties of structures like nanotubes — one of the most promising atomic-scale structures.

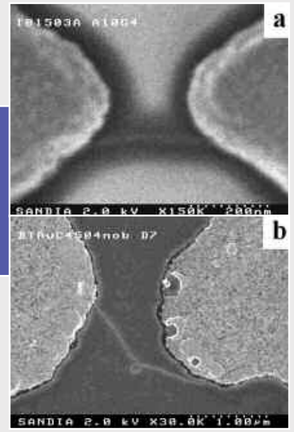

Experimental example of a nanotube loop. Applied Physics Letters, 86, 093112 (2005).

Nanotubes are sheets of carbon atoms rolled into tubes about 10,000 times thinner than a human hair. Though incredibly thin, they can be at least a millimeter long and are incredibly strong. More importantly, "Depending on the structure it can be either a semiconductor or it can be a metal. That's pretty exciting because there is no other material that acts this way," says François Leonard of the Nanoscale Science and Technology Department, Sandia National Lab at Livermore. He supervised Cummings' summer 2005 practicum.

"Because of these properties and their smallness, people have tried to use these carbon nanotubes in various applications, especially as transistors," Leonard says. "That's why people like me and Aron are interested in modeling the properties of these systems."

But "Fabricating transistors with these carbon nanotubes is not a simple process," Leonard says. Sometimes they have bends or loops, crossing over themselves between the electrodes. "The question is, is that going to be a deal killer?" he adds. "Will that render the transistor useless or degrade its properties so it won't be as useful anymore?"

Cummings set out to help answer that question during his practicum. He adapted the computer code Leonard had assembled to simulate straight nanotubes so it simulated looped nanotubes. Such simulations — and the high-performance computers to run them — are important to learning the properties of nanostructures, Cummings says.

"Building nanostructures is very expensive and, depending on your facility, is not always reliable either," he adds. While it's impossible to build a totally accurate computer model, "The upside is some modeling lets you do almost whatever you want." In a real-world experiment, a result could have many causes a scientist will have to tease out. In a computer simulation, those causes can be isolated to learn about specific effects.

Nanotubes can also be viewed as stacks of carbon atom rings, and Leonard's simulation looked at the tubes' charge and electrostatic potential as a series of parallel rings. Cummings first changed the code to geometrically simulate a three-dimensional looped nanotube and calculated the electrostatic potential and charge throughout the system.

Cummings also refined the code so it used less computer memory and ran faster. The original approach "was a pretty mathematically intense way to figure out the charge," Cummings says. "My approach was to use a more simple formula based on a more fundamental solid-state physics idea."

Cummings was able to make only a preliminary test before his practicum ended in the fall. It showed that looped carbon nanotubes couldn't conduct electricity — a result that contradicts earlier experiments. That might be because his simulation doesn't account for everything, or that the loop he simulated was small, creating problems longer tubes don't have, he says.

Leonard says he hopes to collaborate with Cummings on further simulations. "He's very motivated to continue. We just need to find the time," Leonard adds. "He was able to make some really strong improvements to some of the modeling tools we were using because of his computational background," including cutting memory demand and modifying codes so the simulation can model longer, bent nanotubes.

The practicum research capitalized on Cummings' master's thesis, which focused on heat conduction in nanotubes. Straight nanotubes have the highest thermal conductivity of any known material, Cummings says, but his master's research found the structural defects in branched nanotubes — those with a Y shape — interrupt heat flow.

Cummings' master's advisor at Washington State, Electrical Engineering and Computer Science Professor Mohamed Osman, told him about the DOE CSGF and encouraged him to apply. Leonard invited Cummings to do his practicum at Sandia after Cummings described his master's thesis to a Sandia representative at the annual fellows' conference in Washington, D.C.

Cummings' dissertation research goes beyond even nanotubes to the sub-atomic scale. He's using computer simulations to study electron spin for quantum computing, which hopes to exploit quantum mechanics for nanocomputers.

The theory is that electron spin could be manipulated for use as transistors or logic elements, says Electrical Engineering Professor David Ferry, Cummings' doctoral advisor at Arizona State. An electron with spin up could represent a one; an electron with spin down could represent a zero. Cummings' simulations show it may be possible to isolate them according to spin direction.

Ferry is impressed with Cummings' drive. "He's taken some pretty intense classes and it just doesn't seem to faze him," Ferry says. "He's kind of a laid-back guy and he just kind of absorbs it like a sponge. He's willing to work hard and he digs right in."

Leonard says Cummings' practicum experience bore out that assessment. "Aron was really stellar and I wish he could have stayed on," he adds.

For his part, Cummings says the practicum made returning to a Department of Energy lab a distinct possibility. "I really enjoyed coming into work each day," he says. "Going into the future, that wouldn't be a bad group to work for."

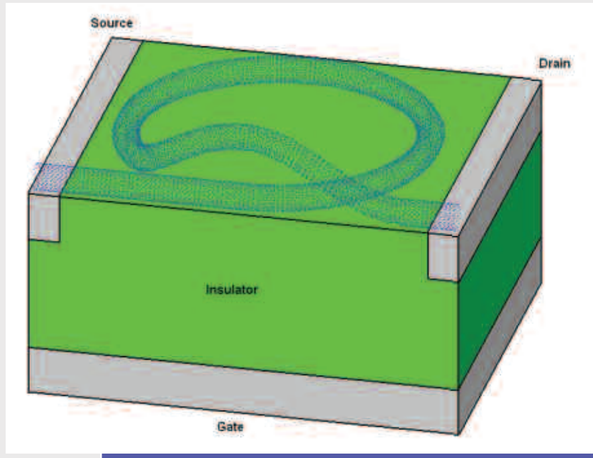

A field-effect transistor with a looped carbon nanotube.

SCOPE OF PROGRAM

Since its inception, the DOE CSGF program has supported over 250 students studying at more than 50 universities throughout the U.S. Currently it supports 62 students in 20 states.

For over 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.

In both physics and the martial arts, "When you go even deeper, things actually become simpler," Cummings says.

Project Puts Proteins on Fast Track to Classification

BONNIE KIRKPATRICK

University of California – Berkeley | Lawrence Livermore National Laboratory | Story by Thomas R. O'Donnell

Imagine constructing a family tree for people named Smith based on how much each Smith resembles other Smiths, and you'll get an idea of Bonnie Kirkpatrick's task.

Kirkpatrick first took on the challenge in summer 2005 during her Computational Science Graduate Fellowship practicum at Lawrence Livermore National Laboratory. Working with Senior Scientist Adam Zemla, Kirkpatrick helped refine computer algorithms to automatically compare and classify the structures of thousands of proteins.

"When the protein structures are very similar to each other, we have groupings," says Kirkpatrick, a doctoral student in electrical engineering and computer science at the University of California – Berkeley. "We would like to have tree structures that connect these groups like the branches of a family tree."

Proteins are metabolic workhorses. Assembled by cells from DNA blueprints, they act as hormones, enzymes and other key agents. Understanding protein structure is key to biological research, including chemical pathways and disease. If scientists know which proteins are related, they can better predict their functions. The Protein Data Bank, a repository for protein structures researchers have elucidated, has more than 35,000 entries, with about 500 added each month.

The most notable effort to organize Protein Data Bank entries by structure and evolutionary origin is Structural Classification of Proteins, or SCOP. SCOP curators, however, manually compare structures to classify newly discovered proteins — a time-consuming process. "Experimentally, we're able to solve more and more protein structures, so there are more and more protein structures available" for entry in SCOP, says Kirkpatrick, a Montana native. "SCOP can't keep up. It's a bit behind." Zemla notes that the last release, in July 2005, was based on Protein Data Bank entries from October 2004, meaning SCOP actually is out of date by about two years.

Nonetheless, Kirkpatrick says, SCOP is "the canonical standard. Everybody uses it. If there's any way to automate the classification process or provide a new classification that may not be the same but produces similar patterns, that would be a big deal."

Creating that automated system will take several steps, Zemla says. "First we would like to detect what features in protein structures are the most useful for automated classification," he adds. Protein structures are described with coordinates for the thousands of atoms comprising them. Zemla, a computer scientist, created an algorithm called LGA that automatically compares this "fingerprint" from a target protein with ones from a cluster of proteins to find features they share.

LGA produces a summary of the structural alignments. Kirkpatrick took the summaries, which are more concise than the alignments themselves, and clustered similar ones. The algorithm she and Zemla created, called STRALCP (STRucture ALignment-based Clustering of Proteins), uses statistical modeling to do the job. It creates detailed information about overall and specific similarities between pairs of protein structures, identifies fragments they have in common, and uses that information to classify the proteins.

SCOP classifies proteins according to a hierarchy, ranging from a class (describing the most distant similarities) down to a family, in which proteins are closely related. In tests, STRALCP detected relationships between proteins with 88 percent accuracy at the superfamily level, the one just above family.

"I think it's pretty good, but those are just initial results," Kirkpatrick says. "It's not as good as it could be. There's some more tuning that needs to be done on the algorithms."

One of the next steps is to examine a protein's structure and predict its expected family and superfamily. The question is how to place a newly characterized protein in its proper cluster. "In the worst case, you can do it like you would identify it — create a new alignment to all the other proteins you already have," Kirkpatrick says. "But that is slow. You'd have to create thousands of alignments."

Zemla and Kirkpatrick are exploring shortcuts, including comparing the new protein with representative fingerprints derived from each SCOP family. Early results show it predicts membership in a SCOP family with almost complete accuracy.

The process of detecting structural similarities, Kirkpatrick says, is a bit like tuning a TV. "When you see static only, it's completely noise," she says. "When you see the picture and no static, there is no noise. The hard part is picking out the feature or the summary of the alignment so you can find the signal to do the clustering. That way you're not clustering on noise."

The algorithm should improve as more structural features are defined, Zemla says. He and Kirkpatrick are verifying their current algorithm based on one type of structural feature and identifying other features for use in the algorithm or new algorithms. They're also preparing a paper on STRALCP. Though she has returned to Berkeley, "I'm continuing to work with Adam, and I find the work very exciting," Kirkpatrick says. "One of the reasons I picked computational biology is the work in that area is fairly high-impact. I have a chance to produce tools and things that help basic biology research."

Zemla agrees: "It's a pleasure to collaborate with Bonnie and share ideas."

Kirkpatrick's first experience with computational biology was as an undergraduate at Montana State University, where she earned a bachelor's degree in computer science. Gwen Jacobs, head of the Department of Cell Biology and Neuroscience and a researcher in the Center for Computational Biology, introduced Kirkpatrick to problems in the field.

That experience was one factor that led Richard Karp, an electrical engineering and computer science professor with deep interests in computational biology, to sponsor Kirkpatrick's entry to the doctoral program at UC – Berkeley. "She seemed nice and she seemed bright, so I felt it would be fun to work with her," he says. Working with Eran Halperin, a researcher at the International Computer Science Institute, a nonprofit organization affiliated with UC – Berkeley, Kirkpatrick is developing better and faster computational tools to analyze genetic information.

"Bonnie is basically a full partner in that project, primarily between her and Halperin," Karp says. He also is participating in the research, which aims to make analysis more efficient by pooling DNA samples. Pooling saves laboratory time, but computational tools are needed to dissect the information.

Kirkpatrick has bridged the computer and biology worlds, even though most of her knowledge of the latter comes from books and journals. The biggest obstacles to understanding were language, because computer scientists and biologists often use the same terms for different things, and biology's unpredictable nature. "Computers are so clean from a logical standpoint," Kirkpatrick adds. "They are very formalized and predictable. But biology is completely different. There's always the exception to the rule."

Kirkpatrick's future seems similarly unpredictable. Although she's considering academia, her national lab experience has made that atmosphere more attractive. For the present, she'll continue her studies and zero in on a dissertation topic.

Zemla hopes to have some input. "I think she enjoyed her work and being involved with the topics and projects I am currently working on, so I hope we will be collaborating for a long time," he adds.

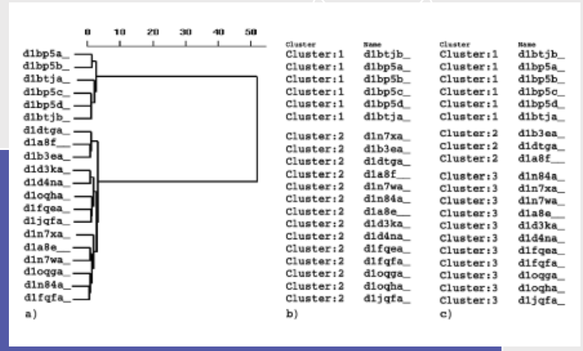

COMPARING AUTOMATED CLASSIFICATION METHODS. This figure illustrates that the results of our method, with two different sets of parameters (b) and (c), are competitive with the hierarchical, tree-like, method of clustering (a). The proteins clustered here are from the SCOP family c.94.1.2.

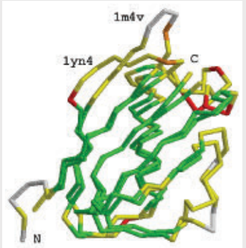

STRUCTURAL SIMILARITIES. These are the structures of the two proteins 1yn4_A and 1m4v_A (SCOP domain: d1m4va2) drawn in a superposition. Since these proteins have low sequence similarity (~14%), their similarity is not detected by traditional comparison methods. The colors highlight the 75% structural similarity detected by our method. Regions of the structure that are indicated in green are very similar. Yellow regions are slightly less similar, while red regions are dissimilar.

"One of the reasons I picked computational biology is the work in that area is fairly high-impact. I have a chance to produce tools and things that help basic biology research."

SEQUENCE REGIONS WITH STRUCTURAL SIMILARITY. These are results from an analysis of structure similarities between EAP domains from Staphylococcus aureus *and proteins from the SCOP's Superfamily of Superantigen toxins. The top green bar in the figure represents the sequence of the protein EapH2 (1yn4_A) to which all the subsequent proteins are compared. Each bar represents the sequence of a single protein with green representing structural motifs that are very similar with the corresponding region of EapH2. Yellow indicates slightly less similar structural elements, and red indicates a lack of similarity.*

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

Simulation "Bumps" Nanotubes

BRANDON WOOD

Massachusetts Institute of Technology
Lawrence Berkeley National Laboratory
Story by Thomas R. O'Donnell

Sometimes, things just need a nudge — including the minuscule carbon straws called nanotubes.

Brandon Wood, a DOE CSGF fellow and materials science doctoral student at Massachusetts Institute of Technology, did the prodding in a computer simulation he helped develop during his summer 2005 practicum at Lawrence Berkeley National Laboratory (LBNL). The tests explored how a type of impurity affects nanotubes' heat conductivity. Nanotubes, about 10,000 times thinner than a human hair, are carbon atoms joined into a sheet of hexagons that's rolled into a tube. Keeping computer chips cool could be one use for them, says LBNL researcher Joel Moore. "Chips have to dissipate a lot of heat in a small space," says Moore, who supervised Wood's practicum. "You want to put something there that will really pull the heat away. Nanotubes are good for that because they have the highest thermoconductivity we know of — especially if they're pure nanotubes" made entirely of one carbon isotope, like carbon-12.

Nanotubes made of one isotope are more expensive to make — and nanotubes aren't cheap to start with. High-purity ones sell for more than $200 per gram, compared to about $25 per gram for gold. It's more common for the tubes to contain a mix of isotopes with different weights, like carbon-13 or carbon-14.

Wood's job at LBNL was to help create computer simulations of how these random isotopic mass defects scatter phonons, the primary mechanism for conducting thermal energy, or heat. The research has broad applications, Wood says, because the program he helped devise can simulate an entire class of defects — not just isotopic ones. High-performance computers, like those at the National Energy Research Scientific Computing Center at LBNL, are necessary for such simulations. "You can do the pristine nanotube case just by using equations, but if you want to do randomized order on a potentially infinite system, you need to use computers," Wood says.

Wood collaborated with Padraig Murphy, a physics graduate student at the University of California – Berkeley, to create the simulation. It uses vibration as a model for thermal conductivity, because heat at the atomic scale is vibration. The higher the temperature, the faster atoms move.

"We artificially create an input wave — bump [the nanotube] from one side and watch how this propagates through the system" of atoms, Wood says. Moore's group was running such simulations on two-dimensional models — a flat sheet of carbon atoms. Wood and Murphy turned that into a three-dimensional tube.

They based the simulation on a classic physics "ball and spring" model — connecting the atoms with a network of massless springs. Like ripples in a pond, the movement of one atom affects the movement of its neighbors, and its neighbors' neighbors, on through the tube.

Physicists use "normal" vibration modes to describe how such atomic systems oscillate. "Any type of vibration can be characterized as a combination of these very basic types of vibration," Wood says. "We were looking at inputting the types of vibrations to see how much is transmitted through the length of the tube." Thermal conduction depends on these properties.

"It turns out nanotubes, because they're tubes, have some peculiar normal modes," Wood says. In some cases, a compression wave travels the length of the tube. In others, the nanotube may pinch in certain spots, or breathe outward, expanding along its radius. And in one basic vibration, the ends are fixed and the middle of the tube moves up and down, like a guitar string.

In the simulation of imperfect nanotubes, heavier isotopes randomly replace carbon-12 atoms. The key question is how those imperfections will change the normal modes. "The jury is still out," Wood says. During his summer at LBNL, he just had time to help finish the simulation for pristine nanotubes. The program now can simulate impure nanotubes, but there hasn't been time to run it, Moore says. "The next step would be to do a bunch of nanotube sizes and impurity concentrations," he adds. He hopes to eventually publish a paper on the research.

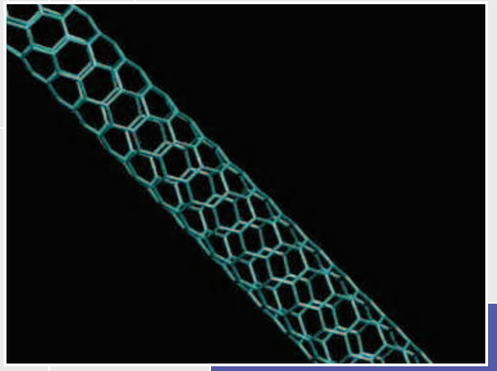

The "torsional" mode of a (4,4) armchair nanotube. This mode represents a twisting motion around the primary axis of the tube.

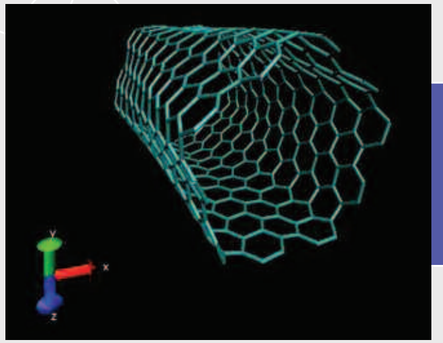

An example of a mixed radial and longitudinal distortion mode of a (10,10) armchair nanotube, viewed along the primary tube axis.

Wood's biggest contribution, Moore says, was creating visualizations — short computer animations — of the nanotube vibrations. "An important step to understand what's going on is to take that data and put it in visual form," Moore says, and Wood brought expertise in that area.

Visualization also is a key component in Wood's doctoral research at MIT, where he uses computers to simulate the properties and behaviors of materials. He focuses on ion transport through a class of materials called superionics. "They're solids that have liquid-like diffusive properties," he says. Ions — atoms that carry a negative or positive electrical charge as a result of gaining or losing electrons — move through these solids as if they were liquids at the atomic level. Superionic materials are important for alternative energy, particularly fuel cells and storage media for the hydrogen they use.

Scientists know superionics have these unusual properties, but they don't understand why. Wood focuses on their quantum physical properties to find out. "If we know why, we can perhaps improve the system" to make better materials, he adds. He's simulating how ions diffuse through the lattice of atoms in three superionic materials.

The work requires a different approach, says Woods' academic advisor, Associate Professor of Computational Materials Science Nicola Marzari. "To characterize this requires sort of topological ideas," he says, with "an understanding of how things are connected with each other."

"That's a whole direction Brandon has really developed on his own," Marzari adds.

For Wood, the work is an appropriate mix of basic science and applied research. He earned his bachelor's degree in physics at Stanford University. "I fell in love with it, I guess because it's as fundamental as science gets. It's pure." Eventually, however, "I kind of got bogged down in what I thought was a little too much theory and not enough practice." Wood chose materials science for his graduate degree to focus on more applied research. Computation provides common ground, he says: "It sits the fence very much between theory, because it's grounded in theory…but on the other hand I'm doing applied experiments."

Wood's other major interest is Russian culture — a result of two years of Mormon Church missionary work in Siberia. The experience led him to earn a second bachelor's degree in Russian Studies. He's returned once and wants to again. "I would love to do something where I work cooperatively with former Soviet scientists," Wood says. Beyond that, he's unsure what awaits when he graduates in 2007.

Moore thinks Wood would fit in at a national laboratory. "What he wants to do is longer-term research, but with applications," Moore adds. "It's harder to do in a place like industry. National labs are sort of picking up the slack" for industrial labs.

Marzari says the DOE CSGF has prepared Wood well. "It's a great initiative," he adds. "It's really helpful to prepare research scientists with the right background, forcing students to take classes beyond their usual field and putting them in touch with each other" at an annual conference. "It works very well."

Wood's job at Lawrence Berkeley was to help create computer simulations of how these random isotopic mass defects scatter phonons, the primary mechanism for conducting thermal energy, or heat.

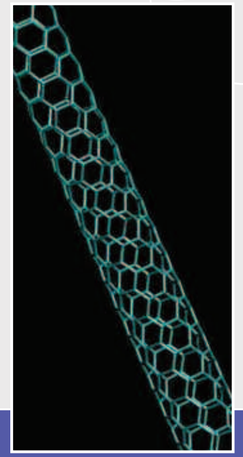

The "flexural" mode of a (4,4) armchair nanotube, characterized by a lateral pinching and bending motion transverse to the nanotube axis.

By Alan S. Brown

Adapting to the Scale of Complexity

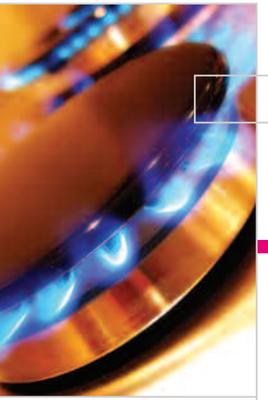

TURN ON A STOVE'S GAS BURNER and a circle of small blue tongues lick the air. Some jump out, their tips turning orange, and then recede again. The placid pulsing of the flames obscures the invisible violence of combustion as gas and air mix, react, release heat, and recombine, forming molecules that careen across the stovetop.

Such complexity is everywhere in nature, from combustion and supernovas to waves and weather. Yet capturing its turbulent glory with accurate mathematical models presents monumental challenges.

In the past, mathematicians could only approximate the complexity of the real world. They simplified models and idealized settings. Yet even these relatively crude simulations provided enough insight to rattle paradigms in fields as diverse as aerodynamics, astrophysics, biology, chemistry, fluid dynamics, materials science, and structural mechanics.

Prompted by past successes and ever more powerful supercomputers, mathematicians are now seeking to go beyond such simplified simulations. They have begun to develop new, more powerful modeling techniques.

Some focus supercomputing power on the most critical regions of an event, such as interaction of shock waves during compression or the thin layer in which combustion transforms fuel into energy. Others allow them to model events using different theoretical approaches — particles and waves for radiation, for example — and combine results.

Early examples of such models are far more realistic than past simulations. Ultimately, they may help us understand such complex physical systems as cell division, combustion, the formation of stars, and even the weather. They may eventually transform our understanding of nature in ways far more profound than the simpler models of the past.

Yet fulfilling those promises will not come easily, warns Phillip Colella, head of Lawrence Berkeley National Laboratory's Applied Numerical Algorithms Group and a leader in the push for more realistic simulations. The new modeling techniques are so different, he says, mathematicians will have to revisit many of the field's fundamental assumptions. And they will have to build a new set of tools to make these techniques practical.

Discrete

To understand what makes the new models different, consider a wave. It looks smooth and continuous. Yet its curve, like the lines in our first geometry lessons, contains an infinite number of points (or unknowns). Unlike a line, which is straight, the geometry of a curve may change along its length. A crest of a wave, for example, may curl upon itself and break into foam and spray.

"A computer representation of a continuous wave slices that curve into a large number of discrete straight lines," Colella explains. "Since each section is straight, the computer knows the exact location of any point along its length. It has transformed an infinite number of unknowns into a finite number of unknowns."

Connecting the lines creates a three-dimensional grid of locations, or cells, that contain information about an object. The cells defining a wave, for example, might contain information about composition (water or air), mass, temperature, salinity, velocity, and direction. As the model moves through time, the cells interact. Velocity carries water from one adjacent cell to another; air from another cell rushes in to take its place.

This approach works for many problems, but it has a serious limitation: What if some critical events are much smaller than the overall scale of the problem?

The new modeling techniques are so different, [Colella] says, mathematicians will have to revisit many of the field's fundamental assumptions.

Imagine, for example, a stovetop burner heating up a tea kettle. The entire system is several inches on a side. Yet the flame front — the space where combustion takes place — is only a few hundreds of microns thick.

A model of the burner-kettle system would have to be sized to capture the mixing of natural gas with air, the movement of air and fuel in the flame front, and the diffusion of heat. Yet they would be far too large to capture the heat-releasing chemical reactions that power the entire system. "Larger cells miss what's happening in the flame," says Colella. "If you have enough resolution to accurately represent the flame, then you have way more than you need for other parts of the problem."

In fact, modeling those molecular-scale details would create so many cells, the model would slow even the most powerful supercomputer to a crawl. Yet only a model that captures both large-scale and small-scale events can explain why a blue flame suddenly flicks orange. And nature — from shock waves rebounding off a wall to the mixing of hot and cold layers of air in the atmosphere — abounds with examples of small-scale phenomena influencing much larger systems.

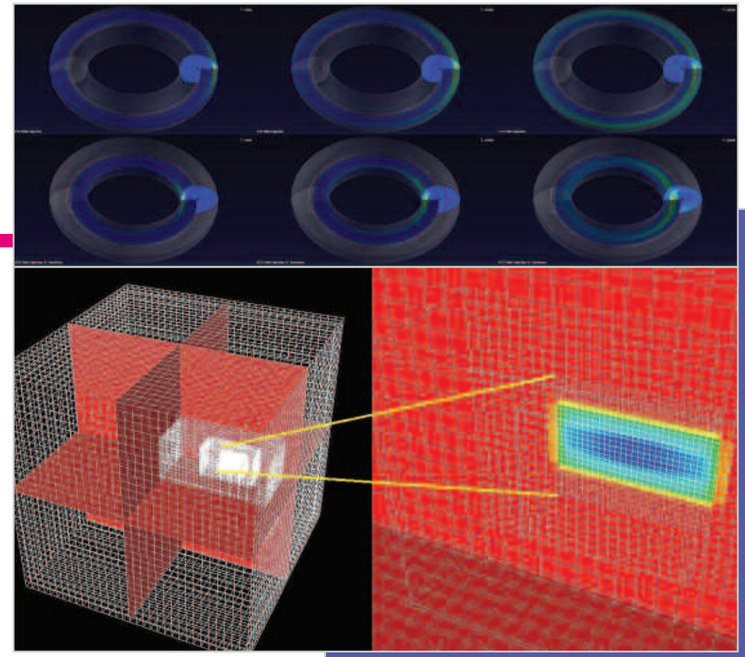

INJECTING FROZEN DEUTERIUM
One way to replenish fuel in a tokamak fusion reactor is to inject frozen pellets of the hydrogen isotope deuterium into its high-temperature plasma. The goal is to move as much hydrogen as possible along magnetic lines and into the center of the plasma. The visualization on the left shows the plasma in red. The mesh adapts to provide greater detail for the blue area, only 0.015 percent of the total volume, where critical pellet-plasma interactions occur. The picture on the right shows how hydrogen (light colors) disperses when injected from the outside (top) and inside (bottom) of the plasma torus.

New Approaches

Starting in the late 1980s, Colella, building on the work of Marsha Berger and Joseph Oliger, began developing practical ways to bridge the scale gap. It was called localized adaptive mesh refinement.

Essentially, says Colella, it focuses computing power on the most critical aspects of the action. The algorithm starts with large cells, but automatically searches for cells that require more attention. It then refines them by dividing them into much smaller cells capable of resolving small but critical details.

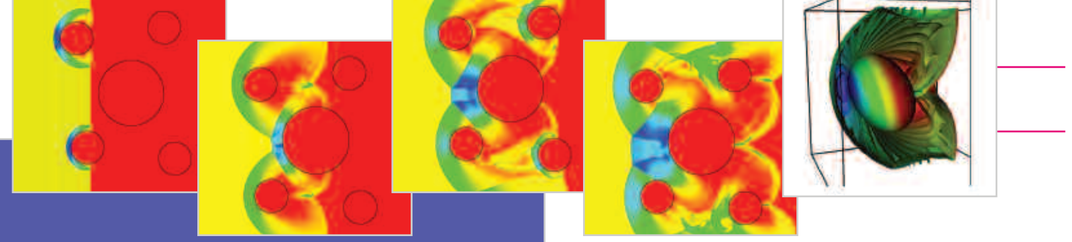

SHOCK WAVES
A shock wave of dense air (in yellow) propagates from left to right across five cylinders set amid lower density air (in red). As energy reflected from the cylinders impinges on the original wave, it creates a very complicated interaction. Adaptive meshes make it easy to generate a grid that captures both the massive onrush of force and how the small-scale interactions around the cylinders affect it. The final image shows a planar shock propagating over a three-dimensional ellipsoid.

PHILLIP COLELLA

Dr. Phillip Colella received his A.B., M.A. and Ph.D. degrees from the University of California – Berkeley, all in applied mathematics. He is currently a Senior Staff Scientist and Group Leader for the Applied Numerical Algorithms Group in the Computing Sciences Directorate at the Lawrence Berkeley National Laboratory. His research has been in the area of high-resolution and adaptive methods for partial differential equations. He has also applied numerical methods in a variety of scientific and engineering fields, including shock dynamics, low-Mach number and incompressible flows, combustion, porous media flows, and astrophysical flows. Honors and awards include the IEEE Sidney Fernbach Award for high-performance computing in 1998, the SIAM/ACM prize (with John Bell) for computational science and engineering in 2003, and election to the U.S. National Academy of Sciences in 2004.

Further Reading

D. Post, D. Batchelor, R. Bramley, J. R. Cary, R. Cohen, P. Colella, and S. Jardin. *Report of the Fusion Simulation Project Steering Committee.* Journal of Fusion Energy, 23(1), March 2005. Note: Also available as Los Alamos National Laboratory TR LA-UR-04-5930.

P. McCorquodale, P. Colella, G. T. Balls, S. B. Baden, *A Scalable Parallel Poisson Solver in Three Dimensions with Infinite-Domain Boundary Conditions.* ICPP Workshops 2005: 163-172.

T. Wen, P. Colella, *Adaptive Mesh Refinement in Titanium.* IPDPS 2005.

Contact:
Phillip Colella
pcolella@lbl.gov

PRACTICUM COORDINATOR
Joseph Grcar . . . . . JFGrcar@lbl.gov

"The grid collapses around the action, and we put fewer grid points in other regions," Colella explains. "In combustion, the grid would collapse around the area where fuel transitions from unburned to burned. In shock waves, you'd want to focus on pressure and velocity changes where the waves reflect off surfaces and interact with each other. Adaptive grids enable models to pick up the details of those subtle interactions." While developing adaptive meshes, Colella's team also turned its attention to another challenge, developing computer simulations that combine models of more than one physical process to obtain high-fidelity representations of complex natural phenomena.

Take, for example, combustion. Fluid dynamics is only one of many physical processes that must be represented to model combustion. A true description of the system must include heat transport by radiation and conduction, elaborate networks of chemical reactions, and how they occur within complex geometries.

"In the past," says Colella, "researchers tried to understand the individual pieces mathematically and glue them into a coherent whole based on experience and experiments. What we're trying to do is glue those pieces together within a coherent mathematical model."

Tools

The goal, he continues, is to create mathematical models that simulate the complex reality around us rather than just selected aspects of it.

"We want to find ways to do this automatically, within the model, so that we can use it as a true design tool to create better, more efficient products," Colella says. "But there's a whole new set of difficulties when you try to include different scales or different types of physics within a single model."

"The behavior of the combined system is not just the behavior of different individual pieces because they all have to interact," Colella explains. "How does radiation react with fluid dynamics in a complex system? How do we represent that so that it is accurate and computationally efficient?"

Well-Posed

Advanced multiphysics/multiscale models also raise mathematical issues about software stability and consistency. Physical scientists often have several ways of testing scientific problems. They can experiment, work from first principles, or seek patterns and analogies.

Computational scientists are constrained by the requirements of mapping the problem onto the computer. "A computer doesn't understand anything but a mathematical model that can be mapped to a computer by breaking it into discrete parts," Colella explains. "Such models must be internally consistent as mathematical objects." For example, they must be well-posed. According to Colella, a well-posed model has three characteristics: solutions must exist, they must be unique for a given set of inputs, and small perturbations in inputs must yield small changes in output, the way a small amount of pressure on the accelerator causes a car to gain only a little speed.

"Not only do the mathematical models have to be well-posed, but when you replace the model with a discretized set of equations, the resulting discrete system must be well-posed also," he continues. "The well-posedness of the mathematical model does not necessarily imply that the discrete system is well-posed. And some problems, like the weather, are genuinely unstable, and we need to understand the limitations that real-world instability places on the predictive ability of our models."

Typically, mathematical models of multiphysics systems are too complex for mathematicians to prove rigorously that they are truly well-posed. Yet Colella and other researchers have developed a number of ways to probe models and test their soundings. Sometimes, they study model stability empirically, feeding models a wide range of input data to check whether they produce unique answers that change in proportion to variations in inputs.

Other times, they study a smaller piece of a larger model. "Instead of looking at a full combustion simulation, we can compare a simulation of chemical dye moving through the system with experimental data to validate our fluid transport and mixing subprocesses," Colella explains.

Testing models against experiments is especially useful when modeling critical parts of a system, such as the shock waves created by compressing a fluid. "This feature so dominates the properties of that class of problems, that designing good methods for it helps you design everything better," says Colella.

Testing conventional models for well-posedness is no trivial task. The task grows exponentially harder when dealing with simulations that include adaptive grids and multiphysics models.

"They're the reason we're revisiting well-posedness," he says. "We're revisiting the fundamentals and trying to understand the underlying math so we know when a model is reliable. That way, others can identify the mathematical components that they can reassemble into new models."

Implemented

Colella's work on adaptive grids began with a complex problem: the study of how shock waves propagate, interact, and reflect off surfaces. This is a phenomenon that affects events as diverse as the ultrasonic waves used to shatter kidney stones, the explosion of supernovas, and national security.

Researchers had been looking for a reliable way to predict shock wave behavior for 125 years, says Colella. Switching to an adaptive grid enabled them to concentrate computing power on the most complex area of the shock front, where waves impinged on reflecting surfaces and generated complicated multidimensional wave behavior.

"We were able to do shock physics no one could do before," he says. "By doing this one problem well, we developed techniques that let us look at other problems in fluid dynamics, such as combustion. By focusing on a sequence of specific applications problems, we developed an algorithm and software edifice that's useful for solving many different science problems."

Looking back, he can now pinpoint two key issues that made the development of practical models so difficult. First, multiscale and multiphysics models have very different properties than simulations that use uniform grids. "What they do to the stability and accuracy of the numerical methods is subtle and differs from one problem to the next," he explains.

Second, adaptive and multiphysics algorithms are extremely difficult to program and manage. Not only do they contain more lines of code, but the codes are interconnected in many more subtle ways.

By learning to overcome these hurdles, Colella's team found the key to extending their software architecture to other areas. "Some people initially thought that having mathematicians who had worked on plasma physics now work on cell biology was crazy," says Colella. "But this range of activities is natural for our computational science group because the mathematical and computational models share many commonalities."

"We're trying to understand in a deep way what the models are telling us about the math so we can find what they have in common," says Colella. The goal of Lawrence Berkeley's Applied Numerical Algorithms Group, he says, is to create a toolkit of advanced techniques that others can use to create better simulations without having to write every component of a multiphysics computer model from scratch.

In the past, he notes, simpler modeling tools created whole industries and revolutionized the way science and industry develop and test new ideas and products in fields as diverse as weather prediction, astrophysics, chemical engineering, and aerodynamics.

Tomorrow's models promise still greater fidelity and more ways to tackle difficult problems. They promise to bring us closer to the real, continuous world around us. Perhaps they will help us probe the origins of the universe or find a practical route to nuclear fusion.

Thanks to Colella and other mathematicians, many of the tools needed to build this future already exist. They are just waiting for us to adapt them so everyone else can use them too.

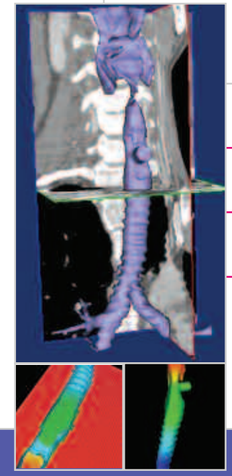

TRACHEOTOMY
Physicians must sometimes make a vertical incision in the trachea to improve airflow or remove secretions from a patient. A fast algorithm enables medical professionals to model an MRI image of a trachea (top), apply an adaptive mesh to its inner surface (bottom left) and use the information to calculate air flow and pressure within the trachea (bottom right). The technology could enable physicians to one day model vascular problems and stents or aneurysms.

Combustion Simulation

Mixes in Turbulence

By Thomas R. O'Donnell

NEXT TIME YOU'RE PRESENTED WITH A BIRTHDAY CAKE, try blowing gently on the candles before giving the big, out-they-all-go puff. If you do it right, they should burn more intensely. But blow too hard, of course, and the candles will go out.

That light puff of air is a form of turbulent mixing, a method to enhance combustion that is key to smaller, more efficient engines and other devices. Fuel and air have to mix well to get the most power out of the reaction, especially if they're initially separated, as in cars with direct injection and in aircraft jet engines. A little turbulence, like blowing gently on a candle, increases the burning intensity. Too much turbulence, though, causes problems, much like blowing on a candle too hard, says Jacqueline Chen, a researcher in the Combustion Research Facility at Sandia National Laboratory's Livermore, California, site.

"It's well established that turbulent mixing enhances combustion," Chen says. But with too much turbulence, "key chemical reaction rates can't keep up with the mixing rate, and local quenching occurs" — pockets where fuel remains unburnt. "If it's pervasive enough, total blowout may occur" — a bad thing if it happens in an airplane engine. Even local quenching cuts efficiency and leads to increased pollution as unburnt fuel goes out the exhaust.

Simulating Turbulence

Despite its important role, many key properties of turbulent mixing are poorly understood, Chen says. That's why she and postdoctoral researchers Evatt Hawkes, Ramanan Sankaran, and James Sutherland ran what is to date the largest direct numerical simulation of a turbulent, non-premixed jet flame with detailed chemistry. The simulation modeled combustion at incredibly high resolution — using up to half a billion grid points and 120,000 time steps to track 11 different kinds of molecules (called species) and 21 different reactions. It also simulated the most intense turbulence ever in a reactive flow simulation in which all of the relevant fluid and chemical scales are resolved numerically. The simulations generated some 30 terabytes of data — enough to fill the 100-gigabyte hard drives on 300 desktop computers. Engineers and researchers have few tools to understand and model combustion. "Combustion application designers are increasingly relying on CFD (computational fluid dynamics) to represent the coupling of turbulence, turbulent mixing, and reaction," Chen says, but the models CFD uses are inadequate in many ways. To run on less-powerful computers, they must use a spatial grid that is too coarse to capture the small-scale reactions and molecular mixing that governs combustion. They spatially filter or average the governing reactive Navier-Stokes and species continuity equations and rely on models to capture small-scale mixing and reaction rate effects. The models make assumptions that may or may not be correct and must be validated against experimental or simulated benchmark data.

Only high-performance computers have the power to directly simulate the many complex nonlinear equations that describe combustion in three dimensions with great detail.

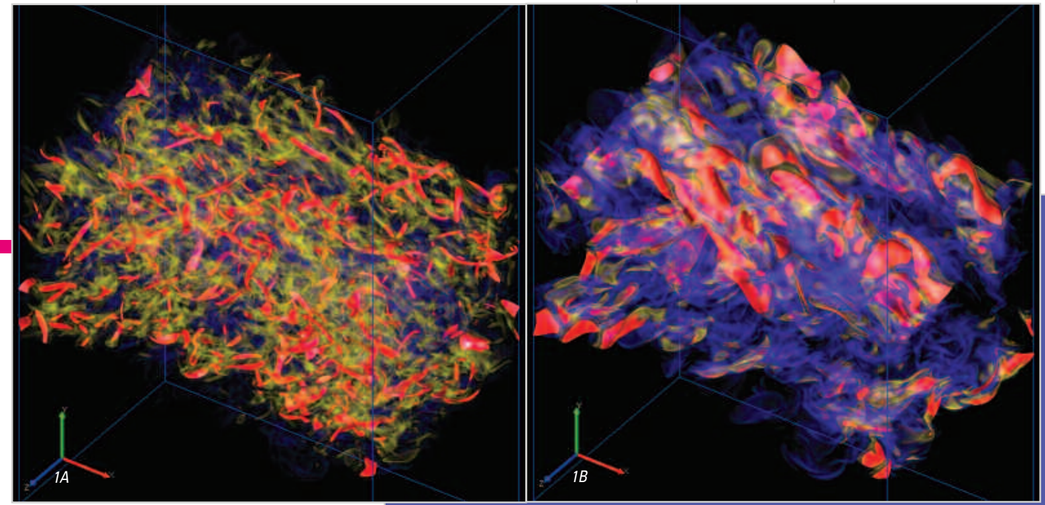

FIGURE 1A. Instantaneous vorticity magnitude in direct numerical simulation of a turbulent CO/H_2 slot jet flame undergoing extinction and reignition. The vorticity magnitude is indicative of the local turbulence structure which affects the mixing of heat and reactive species.

FIGURE 1B. A simulated planar CO/H_2 jet flame, colored by the rate of molecular mixing (scalar dissipation rate based on mixture fraction), which is critical for determining the interaction between reaction and diffusion in a flame. The image shows that high scalar dissipation regions exist in thin, highly intermittent structures aligned with principal strain directions. (Visualization created using a volume rendering application written by Hongfeng Yu, Hiroshi Akiba, and Kwan-Liu Ma, UC Davis.)

Complete Picture

It's more important than ever to have predictive computer models for combustion, Chen says. As engineers strive to make more efficient and cleaner-burning engines, "We're pushing combustion to the ragged edge. It's burning near the lean flammability limit or at temperatures that are too low to support customary flame propagation." It's at these limits that combustion is especially poorly understood and current computer models fail to be predictive.

Other researchers have used laser diagnostics to provide combustion data from experiments. Such tests give snapshots of turbulent mixing and certain molecular concentrations, which are valuable but "don't provide the complete picture," Chen says. "There's lots of quantities — intermediate species concentrations and turbulent fluctuations — that are not possible to measure, so it is often difficult to understand underlying causal relationships with incomplete information." Only high-performance computers have the power to directly simulate the many complex nonlinear equations that describe combustion in three dimensions with great detail.

Powerful Computing

The Sandia simulation tracked molecular mixing and the chemical reactions of combustion in tiny time increments and at a range of turbulence levels as measured by Reynolds number; the higher the Reynolds number, the greater the turbulence. It ran on some of the world's most powerful computers:

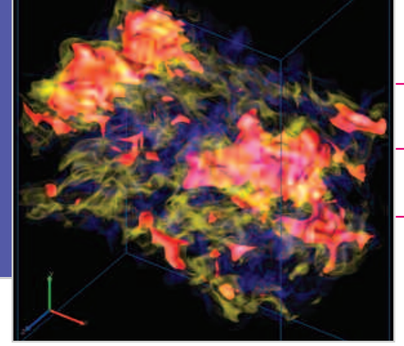

Instantaneous isocontours of the scalar dissipation rate of the hydroxyl radical showing regions of intense mixing of hydroxyl radical exist where the mixture fraction scalar dissipation rate is low. The differences in reactive species and mixture fraction scalar dissipation rate underscore the importance of considering local chemical effects which can directly affect the dissipation rate of the species by modifying their gradients directly through chemical reaction.

the National Energy Research Scientific Computing Center's (NERSC) IBM SP3 RS/6000, named Seaborg, and IBM p575 POWER 5, named Bassi, both at Lawrence Berkeley National Laboratory in California; and on the Cray X1E at Oak Ridge National Laboratory (ORNL) in Tennessee. The simulation was made possible with a grant of 2.5 million processor-hours from the Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. It was the largest award of computing time made through INCITE in 2005.

Without access to the computing resources at NERSC and ORNL, the simulation would have been virtually impossible, Sankaran says. It would have taken decades to run on less-powerful machines, and they would not have the memory needed to house the data. "Leadership computing is enabling this kind of science," he adds.

The simulation was a huge leap, Hawkes says. The group's largest previous simulation used just 4 million grid points. The INCITE grant meant scaling up to more than 400 million grid points.

"We were trying to understand a physical phenomenon — extinction and reignition — that is very sensitive to turbulence and configuration and fuel conditions," Hawkes adds. "There were a lot of variables to play with." The group's priority was to simulate a high Reynolds number, but, Sankaran says, "As the Reynolds number goes up, you need more grid points, which means you need more computing power." To complicate matters further, the size of the simulation meant the group couldn't run full-scale tests. Instead, the researchers ran dozens of two-dimensional and three-dimensional simulations with less detail. Preliminary results challenge some basic suppositions. Most models assume that all kinds of molecules (called scalars in mathematical representations) mix at the same rate, corresponding to the large-scale turbulence time scale. The Sandia simulation found that mixing rates of reactive and passive scalars may vary by as much as a factor of three because of differences in the way they diffuse and other effects. Atoms with low molecular weights, such as hydrogen, diffuse rapidly, affecting the mixing rate, Chen says. Diffusion is gradually superseded by turbulence as the Reynolds number increases, Chen says, but "Even at the highest Reynolds number we still didn't see that effect completely washed out. It was diminished but it was not negated."

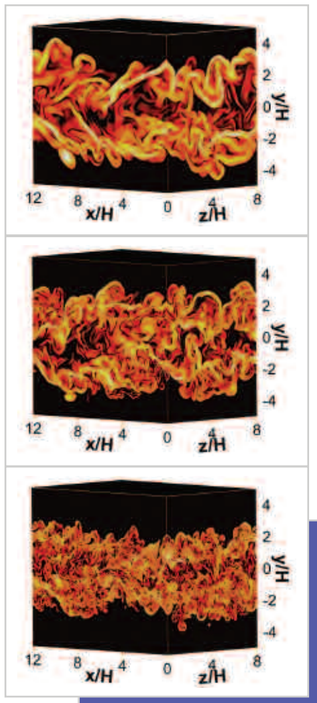

New computer models of mixing and combustion may need to take these effects into account, the researchers say. If a mixing model is inaccurate, it could predict a stable flame when the flame may actually have pockets in which fuel isn't burning.

High-performance computing power also let the researchers study intermittency — small, closely contained fluctuations in a turbulent flow. Intermittency can create areas in which combustion stops and restarts. The simulations show greater turbulence causes greater intermittency and a longer lag for reignition. It's necessary to understand how long it takes those pockets to reignite because that can affect engine performance and pollution, Chen says.

Data from the simulation could improve computer simulations of mixing and combustion, leading to better engine designs. "The biggest impact will be felt if this will translate into a model that can be used at an engineering level," Hawkes says. Research groups from several institutions already want to study the results for that purpose.

"The challenges are with the data analysis," Hawkes adds. "There are phenomenal amounts of data, and the physics are truly complex." The researchers continue to probe the results, seeking clues to local extinction and reignition. They're also focusing more on the fundamentals of the turbulence-chemistry interaction.

"We're going to be looking at this for some time," Chen says.

*Instantaneous isocontours of the mixture fraction scalar dissipation rate field for successively higher Reynolds numbers at a time when reignition following extinction in the domain is significant. The dissipation fields are organized into thin sheet-like lamellar structures, with lengths far exceeding their thickness, consistent with experimental observations in nonreactive flows. Increasingly fine-scaled structures are observed at higher Reynolds numbers. *(From E. R. Hawkes, R. Sankaran, J. C. Sutherland, and J. H. Chen, "Direct Numerical Simulation of Temporally-Evolving Plane Jet Flames with Skeletal CO/H_2 Kinetics," to appear in Proceedings of the Combustion Institute, 2006.)*

OPTIMIZING COMPUTING POWER

>> ***The combustion simulation created by Sandia National Laboratory researchers is so complex it took weeks to run on high-performance computers employing thousands of parallel processors. But it could have been even longer.***

David Skinner of the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory collaborated with the researchers to scale up their massively parallel direct numerical simulation code, called S3D, to run on thousands of processors. Similarly, Mark Fahey, a computational scientist at Oak Ridge National Laboratory (ORNL), helped Chen's group port their code to the lab's Cray X1E computer.

Because the computer scientists studied how the code runs and fine-tuned it, S3D ran twice as fast as it did before on Seaborg, NERSC's IBM SP3 computer, and more than 10 times faster than before on the Cray computer. S3D ran with 90 percent parallel efficiency on 5,000 processors on the Cray XT3 at ORNL and 80 percent parallel efficiency on 512 Cray X1E processors, says Jacqueline Chen, a combustion researcher at Sandia National Laboratory – California.

"The code pretty much scales up linearly with an increasing number of processors," she says. "These guys have been really instrumental in assisting us with scalar and vector optimization of S3D."

Now the Sandia researchers are turning to computer scientists to help sort through 30 terabytes of data from the simulation. "Conventional tools used to visualize and analyze the data simply don't work on such large datasets," Chen says, so they're collaborating with Kwan-Liu Ma, an associate professor of computer science at the University of California – Davis, to develop new tools. With Ma, they hope to create programs to automatically extract and track interesting features and to visualize results.

It's the kind of teamwork that's necessary if researchers hope to attack fundamental problems with major computer power. "No one group can wrap their brain around all these different areas," Chen says. "Performing simulations at this scale requires expertise in combustion and turbulence as well as expertise in computer science for extracting and tracking salient regions of interest in the flow and in visualization or optimization on a particular architecture."

COLLABORATORS

Jacqueline Chen received her bachelor's degree from Ohio State University in 1981, her master's degree in mechanical engineering from the University of California at Berkeley in 1982, and her doctoral degree in mechanical engineering from Stanford University in 1989. She is a distinguished member of the technical staff in the Combustion Research Facility at Sandia National Laboratories and an adjunct professor of mechanical engineering at the University of Utah. She was the co-editor of *Proceedings of the Combustion Institute*, Volumes 29 and 30, and is a member of the editorial advisory boards of *Combustion and Flame*, *SciDAC Review*, and *Computational Science and Discovery*. Her primary research interests are in direct numerical simulation of turbulent combustion, focusing on fine-grained interactions between convective and diffusive transport with detailed chemistry in the combustion of hydrogen and hydrocarbon fuels.

Evatt Hawkes has been a postdoctoral fellow at the Combustion Research Facility at Sandia National Laboratories for four years. He earned a doctorate in engineering at the University of Cambridge and bachelor's degrees in applied mathematics and mechanical engineering at the University of Western Australia. His research focuses on applying advanced high-fidelity simulation techniques — direct numerical simulation and large eddy simulation — to understanding turbulent combustion. His research has made the simulations more accurate and enabled their effective use on massively parallel computers. He has exploited them to better understand fundamental combustion physics, and used them to develop effective models for engineering-level simulations.

Ramanan Sankaran has been a postdoctoral fellow at the Combustion Research Facility at Sandia National Laboratories for two years. He has recently joined the NCCS at Oak Ridge National Laboratory. He has a master's degree in mechanical engineering from the University of Tennessee at Knoxville and a doctorate in mechanical engineering from the University of Michigan at Ann Arbor. His research interests include combustion, computational fluid dynamics, numerical methods, parallel computing and data mining.

Further Reading:

E. R. Hawkes, R. Sankaran, J. Sutherland, and J. H. Chen, *Scalar Mixing in Direct Numerical Simulations of Temporally-Evolving Plane Jet Flames with Detailed CO/H_2 Kinetics*, Proceedings of the Combustion Institute, doi:10.1016/j.proci.2006.08.079 (2006).

R. Sankaran, E. R. Hawkes, J. H. Chen, T. Lu, and C. K. Law, *Structure of a Spatially-Developing Lean Methane-Air Turbulent Bunsen Flame*, Proceedings of the Combustion Institute, 31. (2006) doi:10.1016/j.proci.2006.08.025.

Contact:

Jacqueline Chen jhchen@sandia.gov
Evatt Hawkes erhawke@sandia.gov
Ramanan Sankaran rsankar@sandia.gov

PRACTICUM COORDINATOR

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

By Thomas R. O'Donnell

Magnetic Models

Could Make Minuscule Materials into Big Capacity

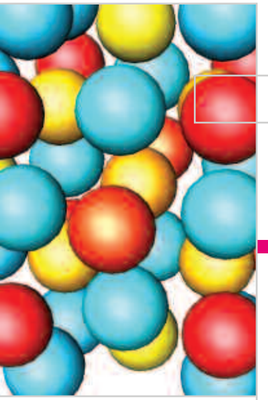

>> Algorithms
>> Combustion Modeling
>> Materials Science
>> Computational Physics
>> Biochemistry
>> Environmental Science
>> Computational Biology
>> Nuclear Engineering

IN 1987, THOMAS SCHULTHESS AND HIS BROTHER bought a 100-megabyte disk drive to attach to their Apple Mac Plus. Today, his laptop computer's disk drive has 1,000 times that capacity, but it's smaller and cheaper. Almost 20 years ago, "If you would have told me that I would have a laptop with a 100-gigabyte hard drive, I would have thought, 'What the heck would I want to do with that?' Now it's full," says Schulthess, Computer Science and Mathematics Group leader at Oak Ridge National Laboratory (ORNL).

That demand for computer storage has fueled exponential data density growth, leading to big capacity on small disk drives and other media. Now, however, scientists are encountering roadblocks to continuing that expansion.

"To get information into smaller spaces, you have to address all the scientific and technological issues — how to deal with smaller and smaller sizes of magnetic matter," says Oleg Mryasov, a researcher for disk drive-maker Seagate Technology. As more data is packed into less space, nanomagnetism — the magnetic properties of materials at the molecular and atomic level — becomes more important. Yet, those properties are poorly understood.

Mryasov, Schulthess, other ORNL researchers and academics are collaborating on nanomagnetism simulations that run on high-performance computers. The models already run at terascale speed, and are expected to take advantage of the petascale power of ORNL's upcoming generation of leadership-class computers. They'll also provide clues for scientists who will study magnetism in the lab's Spallation Neutron Source, now coming on line.

"This is a very big experimental challenge, and in order to guide experimentalists we are attempting to develop a comprehensive model of magnetic properties at the nanoscale," Mryasov adds.

The Search for Storage

Boosting magnetic storage density and reliability is critical to continue the growth of computer power, says Schulthess, who also is a researcher in ORNL's Center for Nanophase Materials Sciences. "If you introduce information technology to all the parts of life, you must be able to store and retrieve huge amounts of data. Otherwise, you're going to reach the limitation of information technology very, very quickly," he says.

To continue boosting storage capacity, scientists are looking to nanoparticles — alloy beads a few nanometers in diameter. Each acts as a tiny magnet, storing data bits as magnetic orientation, or magnetic moment, Schulthess says. Nanoparticles must maintain their magnetic moments at operating temperatures. That can be a problem, because the thermal barrier between magnetic states typically becomes smaller as particles become smaller. "That means the volume for the given material at some point will become too small and the thing will be susceptible to thermal fluctuation," Schulthess says. "That means you cannot store your information anymore."

To overcome the problem, researchers must make nanoparticles with a large magnetic anisotropy, a material constant and one of the factors governing thermal barrier size. With a large enough constant, the thermal barrier stays high enough even as the particle reaches nanometer size, Schulthess says.

"If you introduce information technology to all the parts of life, you must be able to store and retrieve huge amounts of data. Otherwise, you're going to reach the limitation of information technology very, very quickly," (Schulthess) says.

Researchers have shown it's possible to create quantities of iron-platinum particles 3 nanometers to 5 nanometers in size with sufficient magnetic anisotropy — ability to retain magnetic orientation — for use as data storage media. The particles have a high thermal barrier and are small enough to store tens of terabits per square inch, so "in principle, the storage medium part is solved for a long way out," Schulthess says. Production hard drives, in comparison, have a top density of around 150 gigabits per square inch, Mryasov says.

Addressing Recording Issues

Terabit-scale data density is still just a postulation, Mryasov stresses. "There are a lot of experimental and theoretical challenges to understand," he says. "To control the magnetic properties of these particles, particularly when you talk about not just demonstrating large anisotropy, but using it for recording… you have to go to the next level and understand how properties vary from one particle to another."

Besides, currently available technology can't write data to this nanoparticle-based medium, Mryasov says. Nanoscale write heads today can generate magnetic fields of only about 1.5 Tesla. It will take fields of 10 to 14 Tesla to selectively change the magnetic moments of iron-platinum nanoparticles. "It's about an order of magnitude short, unless new writing technology is developed," Mryasov adds.

Researchers at Seagate and other companies and institutions are attacking the write-head problem, Schulthess says. He focuses on understanding more about nanoparticles' magnetic properties. "They don't behave as simply as you would think," he says. For one thing, the smaller the particle, the larger its surface becomes in proportion to its volume. In a particle 3 nanometers in diameter, about 60 percent of the atoms are in the surface and subsurface.

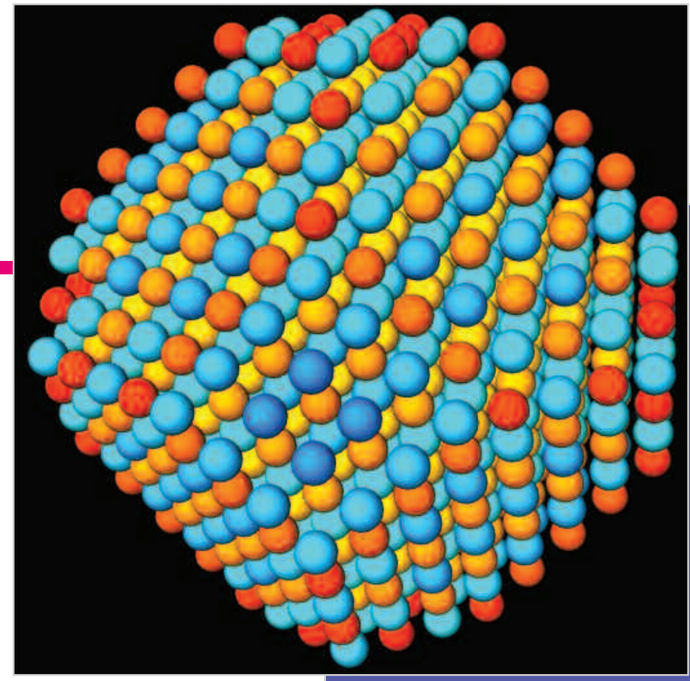

Color-coded distribution of magnetic moments computed from a first principle electronic structure calculation of an 807 atom FePt nanoparticle with full relaxation of the atoms' structure. Fe atomic moments at the surface are enhanced to about 3.2 Bohr (red) from 2.5 Bohr (yellow) in the bulk. Pt surface moments are reduced to about 0.2 Bohr (dark blue) from 0.3 Bohr (magenta) in the bulk.

COLLABORATORS

Thomas Schulthess is group leader in the Computer Science and Mathematics division and in the Center for Nanophase Materials Sciences at Oak Ridge National Laboratory, where he leads the Computational Materials Sciences and the Nanomaterials Theory groups, respectively. His research interests are largely in the area of computational nanoscience, with an emphasis on solving hard problems with practical impact on leadership-class computers. Thomas earned his doctoral degree in physics from the Swiss Federal Institute of Technology (ETH) in Zurich.

Malcolm Stocks is a corporate fellow and leader of the Materials Theory Group in the Materials Science and Technology division of Oak Ridge National Laboratory (ORNL). His research interests lie largely in the areas of electronic structure theory, alloy phase stability and magnetism. Malcolm has been involved in the development of high-performance computing on massively parallel platforms from the beginning and has pioneered the use of these platforms in materials science. He and his colleagues have twice been awarded the Gordon Bell Prize. Malcolm earned his doctoral degree in physics in 1969 and joined ORNL in 1976.

Oleg N. Mryasov has a degree in physics from Ural Technical University in Russia and a doctoral degree in physics and mathematical sciences from the Russian Academy of Sciences, Ural Branch. From 1993 to 1998 he worked in the Physics and Astronomy Department of Northwestern University in Evanston, Illinois, studying electronic structure methods and their application to understanding magnetic, mechanical and optical properties of novel materials. From 1999 through 2001 he held a joint appointment with the Materials Research Department at the University of California – Berkeley and Sandia National Laboratories, Livermore, Calif. During this period he contributed to understanding magnetic interaction in novel magnetic materials. In 2001, he joined the newly created research center of Seagate Technology. One of his current interests is the exploration of novel media design options for heat assisted magnetic recording, understanding magnetic interactions in high anisotropy materials, atomic scale simulations of the magnetic and transport properties of media and reader devices, finite size effects on magnetic properties, and chemical ordering in small iron-platinum nanoparticles.

Further Reading:

G. Brown, B. Kraczek, A. Janotti, T. Schulthess, M. Stocks, and D. Johnson, *Competition between ferro- and antiferro-magnetism in FePt*. Physical Review B 69, 052405 (2003).

C. Zhou, T. Schulthess, S. Torbrugge, and D. Landau, *Wang-Landau algorithm for continuum models and joint density of states*, Physical Review Letters 96, 120201 (2006).

Contact:

Thomas Schulthess
schulthesstc@ornl.gov

Oleg N. Mryasov
oleg.mryasov@seagate.com

Malcolm Stocks
stocksgm@ornl.gov

PRACTICUM COORDINATOR

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

That's important, because "The magnetic states in the surface region are much more complicated," Schulthess says. Those complications also are poorly understood.

"That's where this project we're doing on nanomagnetism on petascale computers is aiming," Schulthess says. Most studies of thermodynamic properties like magnetism can only consider the behavior of entire particles. The computer simulations Schulthess and his fellow researchers use probe the dynamics and structure of magnetism at the atomic level. That's important because magnetism depends on the behavior of electrons.

Adapting Algorithms

Using teraflop-scale computers, the researchers have demonstrated algorithms to discern from first principles the electronic and magnetic properties of a simple nanoparticle in the ground state — the lowest energy level. Because of the importance of temperature fluctuations in these nanometer-sized particles, magnetic properties have to be computed and studied as a function of temperature. "The next step, therefore, is to combine the teraflop-scale first-principles calculations with methods used in statistical physics, which will multiply the computational effort by orders of magnitude," Schulthess says.

The researchers have tailored the programming to the job. Typically, statistical physics methods use sampling algorithms to select information, such as magnetic configurations in a very high dimensional space. The distribution of these configurations can be summarized over a few variables and is called the density of states. It paints a simpler picture of the complex system, but Schulthess says the job can overwhelm typical sampling methods, like a Monte Carlo simulation. Researchers instead adapted an algorithm for calculating density of states based on the random walker in energy space sampling method. It's called the Wang-Landau algorithm, named for Fugao Wang and David Landau of the University of Georgia. Landau is a collaborator on the Oak Ridge magnetism simulation.

The ORNL team generalized the algorithm to compute the "joint" density of states in a certain energy range and magnetization for systems with continuous degrees of freedom as they appear in magnetic nanoparticles. The modifications save more than 90 percent of computer processor time compared to the original Wang-Landau algorithm.

Parallelizing Monte Carlo simulations is typically straightforward, since the sampling work can be distributed over many random walkers. What makes the Wang-Landau algorithm different from typical Monte Carlo algorithms, however, is that the "random walkers" doing the sampling aren't independent; it isn't "embarrassingly parallel."

"Our random walkers have to communicate the latest updates to the density of states at every Monte Carlo step," Schulthess says.

Locally Self-Consistent Multiple Scattering Method

The individual random walkers evaluate, at every sampling step, the electronic and magnetic structure, using another algorithm, called the locally self-consistent multiple scattering method or LSMS. LSMS runs at nearly 90 percent efficiency and also is parallelized to routinely run on machines capable of one trillion to 10 trillion calculations per second, or teraflops. Computing the joint density of states and free energies "requires a lot of sampling and you have to calculate as many individual states concurrently as you can," Schulthess said. Hence, the combined Wang-Landau/LSMS simulation running with 100 parallel random walkers will reach 1,000 teraflops — a petaflop.

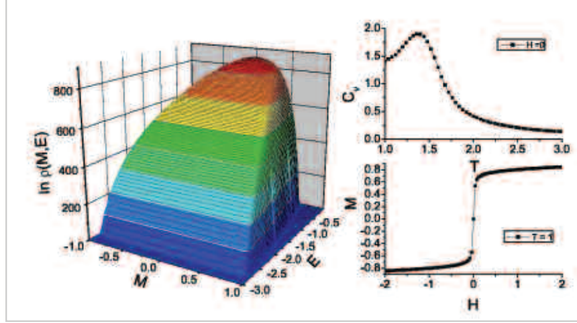

Logarithm of the density of states (entropy) computed of a microscopic Heisenberg model, plotted (left) as a function of macroscopic magnetization and energy, and from which measurable quantities such as the temperature dependent specific heat (upper right) or the field dependent magnetization (lower right) can be computed.

"Without Wang-Landau, using just regular Metropolis sampling will almost certainly not lead us anywhere on the petaflop computer," Schulthess says. "The Wang-Landau algorithm and the generalization we made is really one of the key things" leading to that scale. The high efficiency of the LSMS algorithm means the entire simulation will run at more than 800 sustained teraflops on a petaflop machine. These factors help make the simulation a top candidate to run on ORNL's petascale computer.

LSMS solves the electronic structure of large systems by approximating the electron density and density of states of individual atoms. "You want to be able to calculate the energy or the magnetic structure of some system. It could be a nanoparticle or it could be a solid," says Malcolm Stocks, one of the algorithm's creators and leader of ORNL's Materials Theory Group in the Materials Science and Technology Division. "How many atoms you have to treat, or simulate, to describe that system is critical."

When the material is an element like iron or copper, it's enough to simulate one atom because it has periodicity — each one is identical to its neighbor. But in a nanoparticle, "Atoms on the surface are going to differ from ones in the middle," Stocks says, and most are on or near the surface. "The only way to treat the system is to treat" all the atoms. "In order to do these kinds of models, you need a first-principles electronic structure method that will scale to a few thousand atoms, minimum." LSMS has that capability. It scales well partly because the electronic configuration for each atom is calculated on a single computer node. The huge number of processors in a petascale computer will allow researchers to calculate the magnetic structure of a nanoparticle made of thousands of atoms, Stocks says. Combined with the Wang-Landau algorithm, "It's an amalgamation of the need to do many, many calculations in one nanoparticle, where we can treat that with a few thousand nodes," Stocks says. "It's not just doing the same problem on more and more nodes. It's doing more problems. Without the petascale computer, you wouldn't even attempt this."

Schulthess, Mryasov and their fellow researchers will spend the next two to three years preparing the simulation to run on the petaflop-capable computer. They'll work closely with computer scientists and supercomputer maker Cray Inc. to optimize and port the needed codes.

Simulations, however, aren't enough. "Theory can predict certain dynamic properties that you can then try to measure with experiments" like those planned for Oak Ridge's Spallation Neutron Source, Schulthess says. The facility's neutron beam can probe materials at the atomic level, revealing their electronic structures and magnetic properties. The simulations, in return, will direct neutron source experimentalists to important discoveries about magnetism and help them interpret mounds of data.

The ability to combine theory with experimentation in one location makes Oak Ridge unusual, Schulthess says. Computers help test theories, he says, "but at the end of the day the physical reality is still happening in nature," and that's where truth lies. "It's the reality that we experiment in. We absolutely have to have both."

NANOMAGNETISM ISN'T JUST FOR HARD DRIVES

>> ***Magnetic nanoparticles have uses beyond computer hard drives, Thomas Schulthess says, including health care and national defense.***

One possibility is selective drug delivery. Drugs like enzymes or proteins could be packaged with magnetic nanoparticles, then injected. Doctors could use magnetic fields to direct the particles and their attached drugs to targeted locations.

Scientists at the Naval Research Laboratory also are probing the use of nanomagnetism for detecting biological warfare agents and diagnosing disease. Arrays of microscopic sensors similar to those used in computer read heads would be arrayed on a chip, with single-stranded DNA attached above them. A single chip could contain sensors and DNA strands for several biomolecules.

As a sample flows over the chip, targeted biological agents bind to the receptor DNA on the sensor sites. Next, magnetic nanoparticles coated with a second set of DNA receptors flow over the chip, binding to the target DNA already attached to the chip. A magnetic field pulls off any non-binding beads, leaving only the positive tests for the tiny sensors to detect and identify as a particular molecule.

A Naval Research Laboratory paper published in the journal *Sensors and Actuators A: Physical* says magnetic biosensing has advantages over chemical sensors, including stability and a lack of background interference. Magnetic labels also can be remotely manipulated with magnetism, the paper notes.

A Chilling Tale of

Nuclear Weaponry

By Victor D. Chase

>> Algorithms
>> Combustion Modeling
>> Materials Science
>> Computational Physics
>> Biochemistry
>> Environmental Science
>> Computational Biology
>> Nuclear Engineering

FREEZING AND MAINTAINING THE RELIABILITY of the United States' nuclear weapons stockpile would appear to be diametrically opposed concepts. Think of one and the notion of shivering in the cold comes to mind. The other conjures images of an inferno.

But in the world inhabited by Jim Glosli, a staff physicist at the Department of Energy's (DOE) Lawrence Livermore National Laboratory (LLNL), and other members of the Simulations Group, freezing and nuclear weapons mesh perfectly. Their job is to conduct computer simulations of what happens when metals freeze, or change from a liquid to a solid state.

This work is part of the Advanced Simulation and Computing Program managed by DOE's National Nuclear Security Administration. The program is designed to maintain the safety and reliability of the nation's nuclear weapons stockpile without nuclear testing. Modeling metals is particularly important to this program because, as with virtually all metal parts, the metal components of nuclear weapons are formed by converting molten metal into a solid state.

And, in something of an understatement, Glosli notes, "When liquids freeze it is not a simple process. When a material solidifies it is not necessarily homogeneous at the atomic scale. Lots of little crystallites form. The way they arrange and organize themselves affects the material properties of that metal."

Understanding those properties can provide useful information for improving the metallic materials used to build nuclear weapons. "Imagine a nuclear explosion," Glosli says. "There are violent things happening. Materials are going through lots of different states and all of those details have an effect on the performance of those weapons and their reliability. By understanding those material properties the designers can do a better job of weapon design."

A deeper knowledge of the metals that go into those weapons also can enable scientists to understand what happens to stored weapons over time. This is so because even after metal solidifies and seems stable, changes in its structure continue to take place. It is these kinds of changes at the microstructure level that lead to occurrences such as metal fatigue.

As Glosli puts it, "These structures continue to evolve over very long periods of time. The system would like to get to its lowest energy state, which would be a single crystal. But that wouldn't happen. It would just take too long — geologic sorts of time."

Pounding and Stretching

The art of heating, pounding and stretching metals to give them different characteristics, such as strength and brittleness, predates even the blacksmith's shop. Though early metalworkers didn't know it, what they were doing was manipulating the microstructures within metals. Today, it's crucial to understand metals at the atomic level. Modern metal devices are more complex and face extreme demands — especially when it comes to safety and national defense.

When metal is liquid, the arrangement of the atoms is random; but as it begins to cool, little crystallites, each containing an orderly arrangement of atoms, begin to form. At first there are a few tiny crystallite dots in a sea of liquid. As the metal continues to give off energy, crystallites continue to form and eventually merge. And though the atoms within each crystallite are lined up, the crystallites themselves do not arrange in an orderly fashion. As a result, "You have not one continuous crystal to describe the metal, but a collection of little crystals," Glosli says.

The problem for scientists studying this phenomenon is that it is extremely difficult, and frequently impossible, to empirically observe these crystals as they form. And experimentally observing what happens to metals under the temperatures and pressures that occur during a nuclear explosion is clearly out of the question. So scientists, including Glosli's group, use computer modeling to understand what happens during crystal formation, and thereby better understand the makeup of the metals.

Taking the Prize

Due to the complexity of solidification, modeling this process is also an extremely intricate undertaking. To accomplish this task the LLNL Simulations Group fired up IBM's BlueGene/L, the world's most powerful supercomputer. Their success at using this mammoth computer to create a seemingly accurate picture of the metal solidification process earned the team of six Livermore scientists — team leader Fred Streitz, Glosli, Mehul Patel, Bor Chan, Robert Yates, and Bronis de Supinski — and IBM

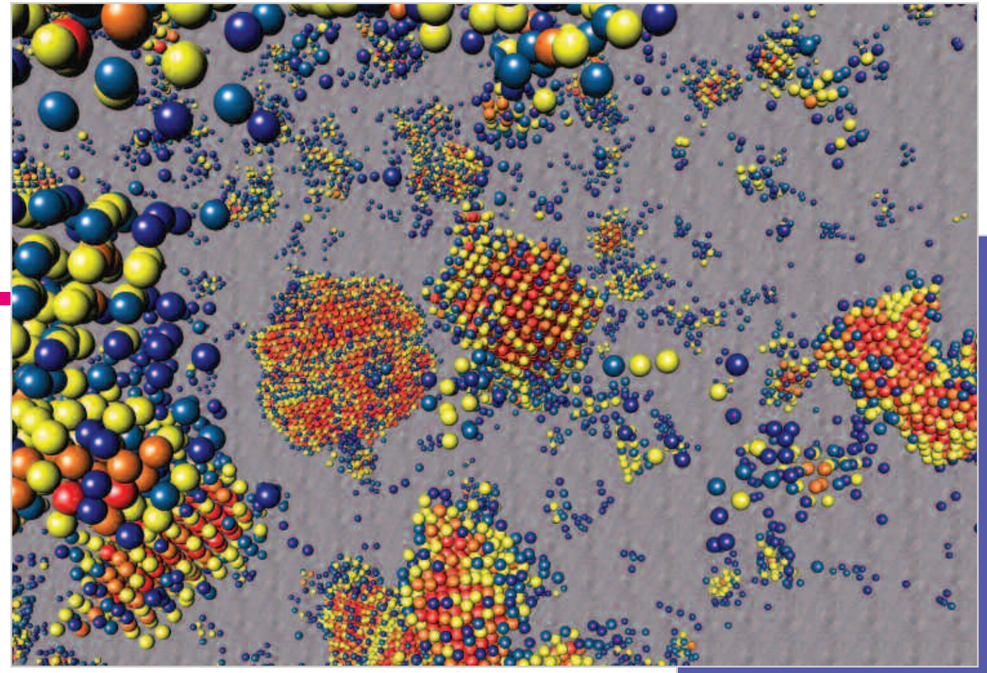

Nucleation and growth of crystallites from an atomistic simulation of molten Ta. Color shows degree of crystalline order (Blue — low, Red — high order). Most of the liquid atoms are not displayed for the purpose of visualization.

Graphics provided by Liam Krauss, a computer scientist specializing in scientific visualization for the Computation Directorate at Lawrence Livermore National Laboratory.

Today, it's crucial to understand metals at the atomic level. Modern metal devices are more complex and face extreme demands – especially when it comes to safety and national defense.

JAMES GLOSLI

Jim Glosli received a Ph.D. in physics from Simon Fraser University in Canada for his work on the statistical mechanics of energetic heavy ion collisions. In 1989 he accepted a postdoctoral position with IBM Almaden Research Center. At Almaden, using atomistic methods, he studied tribology of polymer films and structure of electrolytes at surfaces. In 1992 Glosli joined LLNL and H division. Since then he has explored nanotribology of silicon and carbon films, growth of hydrogenated carbon films, phase transformation of carbon at high pressure and temperature, and nucleation and solidification of metals. He has developed various codes including order N methods for evaluations of long range Coulomb fields (FFM and PPPM), chemical equilibrium methods for multi-component multiphase systems (CHEQ) and molecular dynamics on massively parallel architectures (ddcMD). In 2005 the ddcMD code running on BlueGene/L was the first scientific application to achieve 100 TFlops performance. Dr. Glosli and his coworkers won the 2005 Gordon Bell Prize for this work.

Dr. Glosli's research interests include: Statistical mechanics, atomistic simulation methods, fast electrostatic methods, thermo-chemical equilibrium codes, bond order potentials, algorithm optimization, massively parallel architectures and algorithms, interfacial tribology, friction and wear, interfaces in aqueous solution, carbon phase diagram, carbon films and clusters, energetic materials, and solidification of metals.

Further Reading:

Frederick H. Streitz, James N. Glosli, Mehul V. Patel, Bor Chan, Robert K. Yates, Bronis R. de Supinski, James Sexton and John A. Gunnels, *100+ TFlop Solidification Simulations on BlueGene/L*, Proceedings of Supercomputing 05. Seattle, Washington, November 2005. IEEE Gordon Bell Prize.

C. J. Wu, J. N. Glosli, G. Galli and F. H. Ree, *Liquid-liquid phase transition in elemental carbon: A first-principles investigation*, Phys. Rev. Lett. 89, 135701 (2002).

J. N. Glosli and F. H. Ree, *Liquid-liquid phase transition in carbon*, Phys. Rev. Lett. 82, 4659-4662 (1999).

Contact:
Jim Glosli
Glosli1@llnl.gov

PRACTICUM COORDINATOR

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

researchers James Sexton and John Gunnels the coveted Gordon Bell Prize in 2005. Named for one of the founding fathers of supercomputing, the prize recognizes outstanding achievements in high-performance computing.

The BlueGene/L supercomputer has 65,536 nodes — each node containing two processors — for a grand total of 131,072 processors running in parallel. Its computational speed is such that it can solve in three days a problem that would take a typical desktop computer 1,000 years. In technical terms this means that BlueGene/L ran at over 100 teraflops, a teraflop being one trillion floating point operations per second, to model the freezing of metal.

This speed allowed the materials researchers to view atom-by-atom what happens when crystallites form as liquid metal transforms into a solid. The group focused primarily on the metal tantalum because it is similar to plutonium in some ways, and it is not fraught with classification issues. And while DOE can (and does) conduct classified research, "There is great value in doing work in the open environment, with being able to discuss results with collaborators both inside and outside of the DOE complex," says Glosli.

The power of BlueGene/L allows the researchers to model what happens in tantalum over nanoseconds and at angstrom spatial levels. The distance between atoms usually is a few angstroms.

"When you do an atomic simulation we are always constrained by how large you can go. We could never do the number of atoms we would really like to do because the amount of computation resources required would be more than we could possibly imagine," says Glosli, who was primarily responsible for writing the computer code to model tantalum on BlueGene/L. Yet the work done at LLNL comes close.

Verifying Findings

But how do the LLNL physicists in the Simulations Group know the results of their computations actually mimic reality? As with the solidification process, the answer is complex. "When push comes to shove we are not simulating a real material," says Glosli. "This is only a model of a real material. But we hope that our models are close. And 10 years from now we hope that our models will be even closer. It is certainly something that will evolve. That said, we feel confident that we are at least in the class of real materials."

That confidence results from painstaking efforts to understand the real material being modeled as much as possible. That is another reason tantalum has been the metal of choice, rather than plutonium. "You try to make as many connections to experiments as you can, and tantalum is a much easier material to work with than plutonium and more readily available to a bigger audience of scientists, so there is a larger scientific database on tantalum," explains Glosli.

Knowledge gained from quantum theory also is applied to produce calculations that give confidence in the computer simulation results.

The final element used to verify results is consistency. "We do our simulations and we view them to make sure we don't see anomalies that are not consistent with what we understand of these materials," Glosli says.

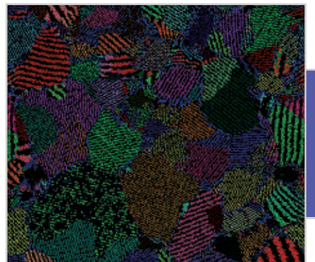

Using a vector order parameter allows extraction of crystallite orientation. Each color represents a different orientation.

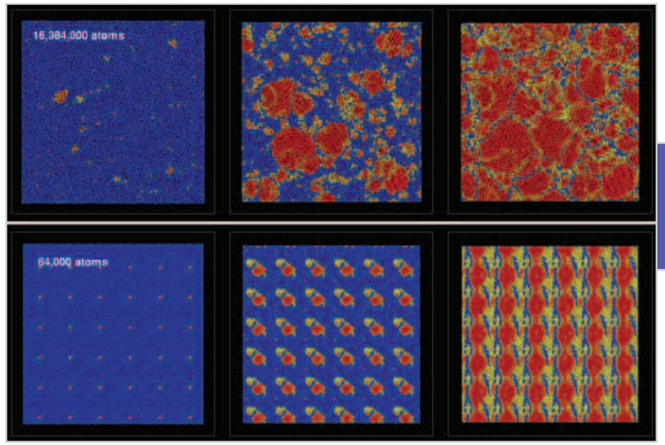

Small simulations can lead to unrealistic microstructure.

When the team started their work on smaller computers they could simulate only a 64,000-atom system. The resulting arrangement was too orderly to represent the real metal. A more realistic picture emerged when they went to two million atoms. The most recent simulations on BlueGene/L enabled them to simulate the solidification of tantalum at the 16 million-atom level, which showed a crystallization pattern much closer to what is known about the actual metal.

In the final analysis, says Glosli, "It is a convoluted path that we have to navigate. We go through this path and try to make as many connections as we can, and to the extent that we see consistency we gain confidence. If we see inconsistencies we take a step back and try to understand what caused them. Is it some fundamental physical process that we haven't thought of before? Are we discovering something new? Or is there some pathology that we have introduced into our model that is giving rise to some phenomenon that you would never see in the real system?"

The metals modeling project at LLNL began in 2000 when team leader Fred Streitz and Mehul Patel developed a single-processor code to model metal solidification under the auspices of LLNL's Metals and Alloys Group. (The metals modeling team recently spun off from that group to form the Simulations Group.) Glosli joined the group in 2002, and increased the single processor performance by a factor of 20. He then set about adapting the code to run on a parallel supercomputer, a process that took some six months. But, he notes, "That was just the beginning. To do more things you add more bells and whistles." Doing so is an ongoing process, as is the metals simulation project.

There clearly are more things to do. "From a scientific point of view there are lots of additional questions we have to ask, and there is an enormous amount of data that we have generated on BlueGene that we need to, and are going through at this point," says Glosli.

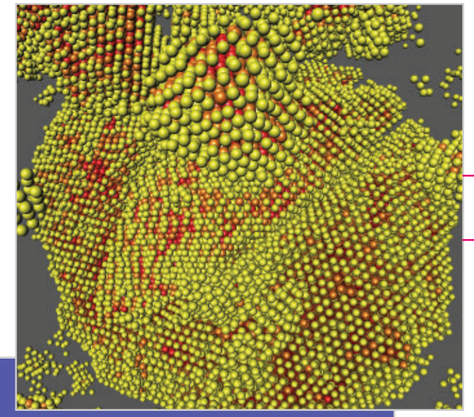

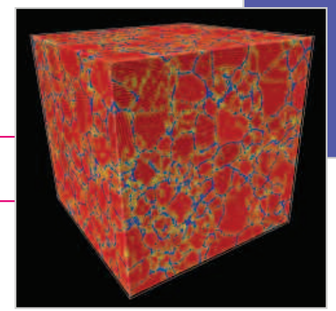

Close-up of large growing crystallites showing formation of vicinal planes.

At left: Coalescence of growing crystallites results in formations of microstructure.

Graphics provided by Liam Krauss, a computer scientist specializing in scientific visualization for the Computation Directorate at Lawrence Livermore National Laboratory.

Solving a
Mystery in Green

By Jacob Berkowitz

Botanists are turning to mathematical modeling to understand the inner lives of plants — and how we can use this knowledge to our advantage.

>> Algorithms
>> Combustion Modeling
>> Materials Science
>> Computational Physics
>> Biochemistry
>> Environmental Science
>> Computational Biology
>> Nuclear Engineering

IN HIS OFFICE at Brookhaven National Laboratory in Long Island, New York, Jörg Schwender leans back in his chair and asks a beguilingly simple question: What makes a soy seed different from a canola seed?

The answer might seem like a straightforward case of genetic differences. After all, the two seeds are obviously different, and not just in terms of the plants they produce. Almost half of a canola seed (also known as rapeseed) is oil, which is why it's used to make everything from industrial lubricants to cooking oil. In contrast, a soybean (the soy plant's seed) is primarily protein, and only about one-quarter oil — hence the origins of protein-rich tofu.

What isn't clear is why it's so difficult to bioengineer changes to these seed characteristics. Scientific attempts to create more commercially valuable seed strains by genetically altering the ratios of oil, protein and starch often come up short. The reason, Schwender says, is that understanding a plant's genes is not enough.

The image on Schwender's office computer screen provides an example. It shows a tangle of colorful lines, arrows and associated acronyms — a map of the chemical pathway a parent plant's sugar travels as it's transformed into a fat stored inside the embryonic seed. And it's in this maze of chemical reactions — the plant's metabolism — that scientists are finding a plant's inner secrets are more fully revealed.

Plant biochemists have created these metabolic maps for more than 50 years. But today, Schwender is among a group of scientists pioneering the use of a powerful combination of cutting-edge mathematical and computational techniques and experimentation to literally redraw them. He's adding hitherto unknown metabolic pathways and in the process helping lay the groundwork for a new green revolution through the bioengineering of plants for food, new medicines and crop-based fuels.

Rewriting Textbooks

"Today, most plant researchers interested in metabolism and manipulating these pathways are focused on the genetic level," Schwender says in his softly spoken, German-accented English. "They're looking at whether a particular gene is turned on or not, or whether its protein is present or not. But this isn't enough to understand metabolism. Because what you observe with these methods are only the individual parts of the living cell. And the way all of these components interact in a living cell is just too complex to be able to predict an outcome by simply knowing the level of any one component."

In essence, Schwender says, the whole of a plant's metabolism is larger — or in this case much more complex — than the sum of its parts.

The key component in the chemical chain of plant life, and the focus of Schwender's research, is the element carbon. Carbon is a plant's primary chemical currency. A plant absorbs carbon from the air as carbon dioxide (a carbon atom with two oxygen atoms). Using energy from sunlight, the plant initially packages the carbon as the atomic backbone of the sugar glucose. The plant then converts the sugar into oils, proteins, or starches. These in turn are either used to build the plant or stored in its seeds.

Plant biochemists have created these metabolic maps for more than 50 years. But today, Schwender is among a group of scientists pioneering the use of a powerful combination of cutting-edge mathematical and computational techniques and experimentation to literally redraw them.

Beginning with his Ph.D. research at the University of Karlsruhe in Germany in the mid-1990s, Schwender has become adept at breaking new ground in our understanding of carbon pathways in plants. To date he's rewritten the textbook definitions of two major chemical pathways.

"What Jörg's work has shown is that a lot of our understanding about the metabolic pathways in plants was simply wrong. It was based on faulty assumptions," says John Ohlrogge of Michigan State University, a leading plant biochemist and in whose lab Schwender spent several years as a post-doctoral researcher. "He's rewritten key aspects of our textbook understanding of how seeds work."

This means that any future genetic engineering of these seeds stands a much better chance of succeeding.

Schwender's first scientific coup was co-discovering an alternate, completely distinct metabolic pathway for the creation of terpenoids. The discovery was the equivalent of learning that cars aren't made just in automobile plants — they're also being produced in donut bakeries. Terpenoids are an important class of plant molecules that include the well-known pigments that give fruits and vegetables their yellow, orange or red color. They also create the scents of cinnamon, clove and ginger. Now pharmaceutical companies are studying them for a variety of medicinal qualities, including the ability to kill germs.

"There are many researchers interested in manipulating these terpenoid pathways," Schwender notes. "But without a proper understanding of the particular steps that lead to this product, you can't effectively bioengineer the plant. So, in part, many of the previous attempts to do this have been working in the dark."

Schwender's experimental method relies on chemically tracing the movement of carbon in living plants and their developing seeds (*see sidebar*). But the experimental results gain their full power when combined with computational mathematical models.

JÖRG SCHWENDER

Jörg Schwender is a member of the biology department at Brookhaven National Laboratory. He received his Ph.D. from the University of Karlsruhe in Germany.

Dr. Schwender's research is focused on metabolic flux analysis and pathway analysis in plants by employing labeling experiments, mathematical models and computer simulation to describe and analyze metabolism quantitatively. In particular he uses steady-state stable isotope labeling to determine flux ratios through branch points of metabolism. *Brassica* napus embryos are labeled with a variety of 13C-labeled precursors, and individual C-atoms are traced through the metabolic network by analyzing the label in metabolites and end products by GC/MS and NMR. This methodology can investigate fluxes in vivo in systems unperturbed by cell disruption, mutation or transgenes. A particular challenge in plants is the sub-cellular compartmentation of enzymes and substrates.

Further Reading:

F. D. Goffman, A. P. Alonso, J. Schwender, Y. Shachar-Hill, and J. B. Ohlrogge. *Light enables a very high efficiency of carbon storage in developing embryos of rapeseed*, Plant Physiol. 138(4), 2269-2279 (2005).

J. Schwender, F. Goffman , J. B. Ohlrogge, and Y. Shachar-Hill, *Rubisco without the Calvin cycle improves the carbon efficiency of developing green seeds*, Nature 432, 779-782 (2004).

J. Schwender, Y. Shachar-Hill, J. B. Ohlrogge, *Mitochondrial metabolism in developing embryos of* Brassica napus. J. Biol. Chem. 281: 34040-34047 (2006).

J. Schwender, J. B. Ohlrogge, Y. Shachar-Hill, *A flux model of glycolysis and the oxidative pentosephosphate pathway in developing* Brassica napus *embryos*. J. Biol. Chem. 278, 29442-29453 (2003).

Contact:
Jörg Schwender
schwend@bnl.gov

PRACTICUM COORDINATOR

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

Jörg Schwender in front of the gas chromatograph / mass spectrometer (GC/MS) which is used to analyze his carbon-13 labeled seeds.

Mathematical models of metabolism are quantitative blueprints that outline the network of metabolic reactions constituting a plant's life. Change the quantity of this input and you change the quantity of that output. These numeric biological models are already widely used in the industrial bioengineering of bacteria. Using an understanding of a bacterium's metabolism, chemical companies and others are bioengineering them to manufacture products such as citric acid and lysine, both major food additives.

"I was reading about the mathematical modeling of bacterium metabolism and the power of the approach and the insights blew me away," Ohlrogge says. "I wondered, could we do this with plants?"

His answer came with Schwender's arrival at his Michigan State lab in 2000. Ohlrogge suggested Schwender apply bacterial-type mathematical modeling to study the metabolism of oil production in canola seeds. It was an ambitious challenge. While a bacterium is a single cell, a plant cell's metabolism is not only more complex but segmented into sub-compartments, including chloroplasts, the sites of photosynthesis. Nonetheless, Schwender succeeded, producing astounding results. In a report on the work published in *Nature*, one of the world's leading science journals, Schwender, Ohlrogge and Michigan State colleague Yair Shachar-Hill and co-workers showed that embryonic canola seeds possess a previously unknown mixed metabolism. Not only are they turning simple sugars into oil; they're using a metabolic pathway associated with photosynthesis to re-use some of the carbon by-product of these reactions. Schwender showed that the canola seeds are truly green in an eco-warrior sense — they're literally recycling their own carbon.

Ramping-up the Computation

The canola seed mixed-metabolism results relied in part on the use of computational flux balance analysis. A flux balance analysis is a mathematical model based on the mass balance of all the molecules present as reactants and products in individual biochemical reactions. With complex metabolic networks it's difficult, if not practically impossible, to determine the actual pathway a carbon atom followed from the dozens, if not hundreds, of potential paths. The mathematical model is used to determine the metabolic pathway that best fits the experimental results.

As such, Schwender's research was a breakthrough not just in plant biochemistry but also in using computation to understand how seeds produce the stored products that feed us.

"There just aren't that many plant scientists who have realized the power of mathematical modeling and have the training that's required in both math and biology to really tackle this," says Ohlrogge, who's continuing to collaborate with Schwender on studying canola oil seed metabolism.

At Brookhaven, Schwender and his group want to take the journey to unravel the secrets of plant metabolism to the next logical step — the whole plant. How are the metabolic networks of leaves (producers of sugars) and seeds (sugar consumers) linked and dependent on each other?

"What ultimately determines how carbon gets partitioned into different end products? The plant's central metabolic network is what distributes the stream of carbon into the various potential end products. If you can model this network, you create models that are actually predictive," says Schwender. This central metabolism modeling is being done with the seeds of the plants *Brassica napus* (canola) and *Arabidopsis thaliana*. The latter is the lab workhorse of the botany world, the green equivalent of the lab mouse, so understanding its metabolic pathways will be an enormous boon for future botanical research.

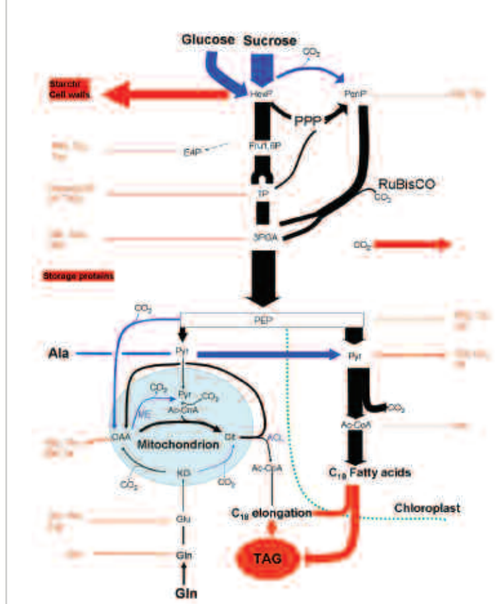

An image of a plant metabolic flux map.

Schwender also is helping take the mathematical modeling of plant metabolism to the next level. He's working with his post-doctoral colleague Björn Junker, who is leading a project to develop a kinetic computational model of central metabolism. Whereas flux balance models are akin to still photos, a kinetic model is similar to a movie, simulating the flow of matter over time. As such, a kinetic model of plant metabolism promises to be a far more powerful tool for describing the moment-to-moment change in carbon flow in response to changes in enzyme levels.

"I see no end to the important insights that Jörg's rigorous mathematical approach to studying plant metabolism will bring," says John Shanklin, a senior biochemist at Brookhaven. "With the current emphasis on renewable sources of feedstocks and biomass for liquid fuel production, his techniques may have important applications in those areas."

At present, Schwender's models are run using Matlab, a well-known mathematical modeling software, on a computer using the Linux operating system. With a growing body of metabolic pathway information, and more advanced mathematical models, he says he's ideally positioned to draw on Brookhaven's high-performance computing research expertise and hardware to one day more clearly see the difference between soy and canola seeds revealed *in silico*.

Says Schwender: "Our models are becoming increasingly predictive and this means that computers will play a bigger and bigger role in understanding the inner lives of plants."

FOLLOW THE CARBON

>> Jörg Schwender's metabolic detective work at Brookhaven National Laboratory is based on adding a rare form of carbon to a growing plant. Carbon-13 is a naturally occurring, stable carbon isotope (or non-radioactive variant) of carbon. Ninety-nine percent of Earth's carbon atoms are of the "standard" carbon-12 form, with carbon-13 constituting most of the other one percent.

Chemically, carbon-13 behaves the same way as carbon-12. What makes carbon-13 a great research tool is that while it's chemically identical, it's slightly heavier. This means scientists like Schwender can pick it out of the carbon crowd using Nuclear Magnetic Resonance (NMR) or Mass Spectrometry (MS) methods.

Schwender uses sugars containing carbon-13 and then feeds these to growing seeds. He then waits and watches to see where the labeled sugars end up — either as part of a fat, oil or starch.

Based only on where the carbon-13 is in the molecular structure of the metabolic product, for example an oil, Schwender chemically backtracks to determine the metabolic path the sugar took to get there. It's like painting a penny red and putting it into circulation at a convenience store. Later, when the penny lands in the cash register of a dry cleaner, you could trace back the transactions that brought it there.

"A large part of the time you spend with this research is the data analysis with the software," Schwender says.

What's powerful about this carbon-tracing technique is that it works in live plants and it's quantitative. It tells researchers not just where the carbon went, but also just how much of the carbon traveled a particular pathway.

How to Build a Better

Greenhouse Gas Trap

By Karyn Hede

>> Algorithms
>> Combustion Modeling
>> Materials Science
>> Computational Physics
>> Biochemistry
>> Environmental Science
>> Computational Biology
>> Nuclear Engineering

ACCORDING TO AMERICAN INDIAN LORE, the Iroquois, Algonquin, Seneca and other tribes that inhabited what is now the eastern and central United States marked the end of hostilities between tribes by burying the hatchet, a symbolic gesture that peace was at hand. Today, scientists and engineers are trying to bury a substance more ephemeral than a hatchet, but which could pose a much bigger threat: carbon dioxide.

Within a couple of years a new kind of exploratory well will be drilled in several locations throughout the United States. But this drilling will not be in search of new oil or gas reserves to tap. Instead, it will inject carbon dioxide, a byproduct of the burning of such fuels, into deep crevices where engineers hope it will remain for thousands of years. It's a kind of peace treaty with Mother Earth — one they hope will slow the rate of warming of the earth's atmosphere caused by the heat-trapping gas's greenhouse effect.

The idea of injecting supercritical liquid CO_2 deep within the earth's crust is called carbon sequestration, and it is among several technologies the U.S. Department of Energy is exploring as a way to mitigate the effects of greenhouse emissions from fossil fuel-fired power plants. In fact, the United States, in the spirit of international cooperation, is also participating in a new multinational effort to capture CO_2 in several promising locations worldwide. The "Carbon Sequestration Leadership Forum" is examining emerging technologies for permanently isolating carbon dioxide and other greenhouse gases in former oil and gas reservoirs, coal seams and under ancient lava flows. Some of the technology required for such an effort already exists. Oil and gas companies regularly use carbon dioxide injection to help recover oil in difficult-to-reach locations.

Burying carbon dioxide emissions from power plants is an attractive way to help reduce the amount of greenhouse gas that reaches the atmosphere. But before embarking on such a larger-scale endeavor, scientists and engineers want to understand more about how such gases will behave underground, how they will react physically and chemically with geologic formations, and how to ensure they don't escape.

Such a problem is tailor-made for large-scale massively parallel computing simulation, since many of the questions about the flow of such gases are similar to questions Department of Energy scientists have been asking about the flow of underground contaminants in groundwater.

"One of the things we are interested in is trying to see how rapidly this CO_2 plume would dissipate," says Peter Lichtner, an earth scientist at Los Alamos National Laboratory. "One of the ways that could happen is that it would dissolve into subsurface brine. Another thing that could happen is that it could become mineralized. You could produce calcite, magnesite or other minerals that contain CO_2. These are the type of reactions that we will evaluate."

Lichtner leads a multidisciplinary team of earth scientists, engineers, and computational scientists who are developing a massively parallel simulator that models subsurface, multiphase, multi-component reactive flow and transport of supercritical carbon dioxide, which behaves as a liquid. The software development evolved over several years, with several people contributing programming and testing its capabilities. The result, called PFLOTRAN, now consists of two separate modules, PFLOW and PTRAN, which can be run either in stand-alone or coupled modes. PTRAN, which solves multi-component reactive transport equations, was derived from an earlier code called PARTRAN, written by DOE CSGF alumnus Glenn Hammond, now at Pacific Northwest National Lab. Another DOE CSGF alumnus, Richard Mills, developed the prototype of PFLOW, which solves multiphase flow equations for oil, water, and supercritical carbon dioxide. Mills built the PFLOW prototype during his summer practicum, using the PETSc suite of nonlinear equation solvers as a platform *(see sidebar)*.

"Peter was using a Newton-Krylov method in PTRAN, using the linear solvers provided by PETSc inside a Newton solver that he had written," Mills said. "When I started writing PFLOW we adopted a similar numerical approach, but I used the nonlinear solver framework provided by PETSc, rather than writing my own Newton solver. Using PETSc's nonlinear solvers gives us a much wider array of options for solving the system of nonlinear equations."

The idea of injecting supercritical liquid CO_2 deep within the earth's crust is called carbon sequestration, and it is among several technologies the U.S. Department of Energy is exploring as a way to mitigate the effects of greenhouse emissions from fossil fuel-fired power plants.

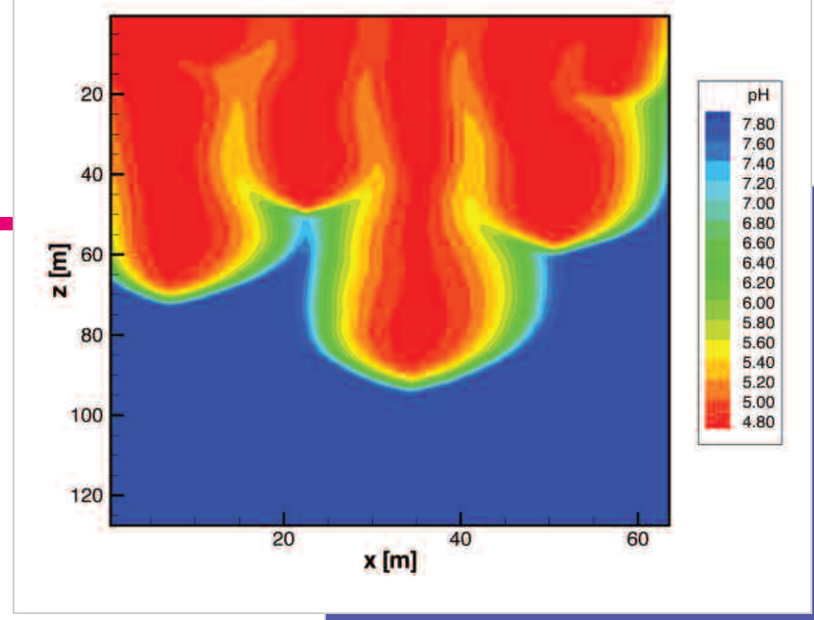

The pH profile after 200 years resulting from a density instability as CO_2 at the upper boundary diffuses downward into the domain increasing the density of the fluid in the sandstone pores. The pH ranges from the initial pH value of 8 to approximately 4.8 at the center of the high-CO_2 lobes.

Postdoctoral scientist Chuan Lu, a petroleum engineer, is working on a multi-scale, multiphase component of the code that will allow scientists to resolve some of the complexities of CO_2 behavior. He has added modules that resolve transport in liquid brine, oil and gas phases. The result is a software suite that can model what happens to CO_2 over thousands of years of underground storage in a variety of media, such as basalt in the Pacific Northwest, sandstone formations in the Midwest, and limestone formations in Texas.

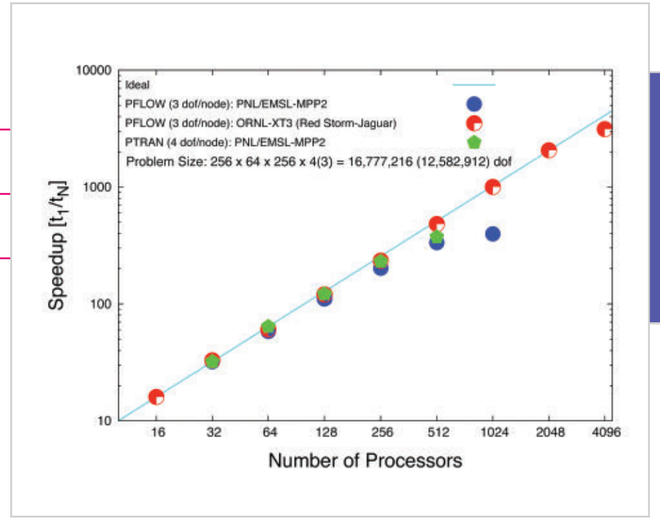

The performance of PFLOW and PTRAN running a single phase thermo-hydrologic benchmark problem on a 256 x 64 x 256 grid with three and four degrees of freedom per node, respectively (approximately 12.6 million degrees of freedom total). The benchmark problem was run on both the MPP2 cluster at PNNL/EMSL, a cluster of 1960 1.5 GHz Itanium 2 processors with Quadrics QS Net II interconnect, and Jaguar, the 5294 Opteron processor Cray XT3 at ORNL/NCCS. PFLOW scales quite well on both machines, bottoming out at around 1024 processors on MPP2, and scaling exceptionally well on Jaguar, displaying linear speedup all the way up to 2048 processors, and still displaying modest speedup when going from there to 4096 processors. PTRAN scales similarly, which is not surprising because its computational structure is nearly identical to that of PFLOW.

"The idea is to store the CO_2 close to existing power plants so you don't have to be trucking it far from the source," said Lichtner. "We are trying to understand how it is likely to behave in different situations."

For example, one of the most attractive sites for sequestration is Texas, but the state is riddled with thousands of spent oil wells and mines, and nobody knows where all of them are located. This presents a safety concern because if the CO_2 escaped it could blow like a geyser, or, more ominously, escape in a slow leak with tragic consequences. Pure carbon dioxide is denser than air and can "pool" at low elevations, causing immediate suffocation. Such a tragedy occurred in 1986 at Lake Nyos in Cameroon, Africa, when a cloud of naturally occurring carbon dioxide gas was released from the depths of the lake. More than 1,000 people died of suffocation. Therefore, the engineers and geologists want to ensure none will travel too far from its injection point and escape.

For example, the scientists are trying to find out what physical and chemical changes to expect when CO_2 is injected. This information is crucial to ensuring that carbon sequestration will not cause radical change to an underground rock formation and that CO_2 storage is secure and environmentally safe.

"It's important to understand the chemistry and flow behavior of the CO_2," says Lichtner. "At these depths, hundreds of meters, CO_2 is less dense than water and 'floats' on top of it. That's what makes spent oil wells an attractive sequestration site. They have the cap-rock that would be necessary to contain the CO_2."

Modeling such a complex system requires massively parallel computing available at only the largest national lab computing facilities. PFLOTRAN has run on the MPP2 cluster at the Environmental Molecular Sciences Laboratory at Pacific Northwest National Laboratory, a cluster of 1960 1.5 GHz Itanium 2 processors; Jaguar, the Cray XT3 at Oak Ridge National Laboratory; and a number of "big-iron" machines at Los Alamos National Laboratory.

"Once you get into three dimensions you drastically increase the number of degrees of freedom and that's where supercomputing becomes the name of the game," Lichtner says. "In these problems it's three dimensions in space and then you have to add all of these chemical components if you want to describe the interaction with the substrate, clays, minerals, and limestone in the geologic formation. That adds at least 10 degrees of freedom. You multiply the number of spatial nodes by the number of chemical components and you're looking at a massive computational challenge."

For example, the research team is studying the density-driven instabilities resulting from an increase in fluid density as CO_2 dissolves into water, a phenomenon scientists call "fingering."

"You get these fingers that protrude downward and they affect the rate at which CO_2 dissipates," Lichtner says. "This puts a constraint on the modeling we can do because we need to resolve it spatially. The size of a spatial node that you use with these big 3D simulations, even with the supercomputing, is on the order of the size of a small conference room. And there are millions of these nodes. But within one block you have these centimeter- and millimeter-scale processes going on, and so now the issue is how do you incorporate these small-scale effects?"

To solve this multi-scale problem, the scientists applied for funding through the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC), a program developed to support research into interdisciplinary problems that require massively parallel computing and the combined expertise of investigators from several scientific disciplines. Lichtner's team received a five-year, $4 million grant that will enable them to add a multi-scale sub-grid algorithm that will allow calculation of small-scale processes and many other features.

"We are also planning to add an adaptive mesh framework that will allow us to dynamically adapt and resolve fine grid effects without having to apply the finest grid over the whole domain," Lichtner said.

The research team is also working to validate the model by comparing their results to a carbon sequestration field test at a west Texas oil reservoir owned by Kinder Morgan. Lichtner is collaborating with doctoral candidate Weon Shik Han and his advisor Brian McPherson at the New Mexico Institute of Mining and Technology to compare their analysis of CO_2 movement in a planned sequestration demonstration project to PFLOTRAN simulations of the test site.

"The whole idea is to be able to take this code and apply it to real-world situations," Lichtner says.

PETER LICHTNER

Peter Lichtner is currently a Staff Scientist at Los Alamos National Laboratory in the Earth and Environmental Sciences Division. He holds a B.S. in Mathematics and Physics from the University of Wisconsin – Milwaukee, a M.S. in Physics from the University of Maryland, and a Ph.D. in Physics from the Universitat Mainz, in Mainz, Germany.

Dr. Lichtner is a member of the American Geophysical Union and was awarded the Buffon Society Special Investigator Award in 1995.

Contact:
Peter Lichtner
lichtner@lanl.gov

PRACTICUM COORDINATOR

Aric Hagberg.....aric@lanl.gov

FRUITFUL PARTNERS

>> ***Richard Tran Mills was looking for a way to combine a life-long interest in geology with his knowledge of computational science when he heard a presentation that was to alter the trajectory of his career.***

It was the annual Department of Energy Computational Science Graduate Fellowship meeting and Glenn Hammond, then a fourth year fellow, was giving his required talk. During his practicum with Peter Lichtner at Los Alamos National Laboratory, Hammond worked on a code that would simulate chemical reactions over time between species in groundwater. Mills was interested in the idea of modeling geological processes and after the presentation he talked to Hammond about his experience. Those conversations led Mills to contact Lichtner and eventually to spend his practicum writing the flow code that, when combined with Lichtner's existing transport code, became the foundation for PFLOTRAN.

"I came in and talked with Peter about what kinds of equations we wanted to solve and then I just began writing the flow code from scratch," Mills says.

As the foundation of his solver, Mills used modules from PETSc (pronounced "pet c") — Portable Extensible Toolkit for Scientific Computation — a library of software for solving partial differential equations. The code provided a modular framework that gave the scientists more flexibility when designing their groundwater simulations.

Mills eventually became so adept at using PETSc that he began working with the code's developers at Argonne National Laboratory to debug the software and make modifications to the code. That partnership eventually led to Mills' contributing to PETSc development from his current position at Oak Ridge National Laboratory, where he is a computational scientist. Mills works with the Cray XT3, dubbed "Jaguar" and now one of the top 10 fastest computers in the world, performing up to 54 trillion calculations per second.

"My work with Peter was a very fruitful partnership," Mills says. "I learned a lot from him, and I think he learned some things from me."

By Karyn Hede

Harvesting the Fruits of the Genome Revolution

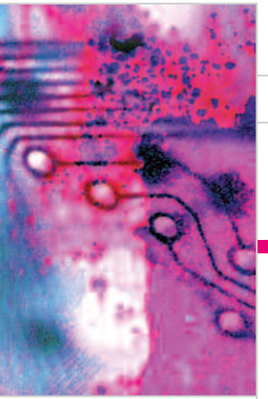

IF SOMEONE WERE TO DESCRIBE A STRUCTURE consisting of anti-parallel strands of 2-deoxyribose joined by phosphate bonds and annealed by hydrogen bonds between complementary nucleotide bases, you may not know what they're talking about. Even a person untutored in biology, however, probably will recognize a picture of a DNA strand as the iconic symbol of life. Add color to highlight certain parts of the strand, and it becomes easy to see how DNA's structure suggests its function.

That kind of biological visualization has accelerated advances that otherwise may not have been made at all in areas such as protein structure and function determination, drug design, and other fields.

Today, a similar revolution is underway in the growing field of "-omics" biology: genomics, proteomics, metabolomics, transcriptomics, and a host of subspecialties that deal with enormous amounts of data generated by high-throughput devices. At Pacific Northwest National Laboratory (PNNL) in Richland, Washington, large research teams are engaged in several proteomic projects ranging from spotting breast cancer biomarkers to identifying organisms that degrade hazardous waste or create biofuels. The sheer quantity of data generated by these new biology fields requires a computational solution. But even with computers helping to reduce data to manageable chunks, biologists must make judgments about the biological relevance of their results.

In proteomics, a field in which a single experiment can generate more than 100,000 data points on a spreadsheet, the burden of analyzing experimental results can quickly become overwhelming. At PNNL alone, the high-throughput proteomics facility produces hundreds of gigabytes of experimental data per day. Analytical techniques and data filters can help reduce that burden, but what has been lacking is a way to visually scan the results so meaningful patterns are easy to identify.

"I realized very early on that we have a data mining problem in biology," says Josh Adkins, a PNNL biochemist who is working in proteomics. "That's our biggest issue right now, even more so than experimental design. Unfortunately, there aren't a lot of tools available to help solve the problem."

In response, a team of bio-mathematicians and computer scientists at PNNL set about creating a software program that would help them quickly scan their results and zero in on potentially meaningful ones.

The result is a visualization tool called PQuad, for Peptide Permutation and Protein Prediction. It can take the raw data representing thousands of protein fragments, reassemble them, and create a visual overlay that shows where they are encoded on a chromosome. It also assigns each peptide a color that is linked to its presence or absence in the experimental sample. Finally, and perhaps most importantly for the biologists, PQuad also can compare the results of experiments run under two different conditions, such as with and without a specific nutrient, highlighting the proteins that change abundance under the two conditions. In addition, the software allows the user to "zoom in" to three levels of detail: global (whole chromosome), intermediate (partial chromosome), and detailed (short stretches of DNA sequence).

"We understood that being able to see the peptides in their genomic framework is really the only way that you can get to the information about how a protein is being expressed in the context of its neighbors," says Bobbie-Jo Webb-Robertson, a senior research scientist specializing in statistical inference models for bioinformatics, and leader of the team that developed PQuad.

An underlying principle of genetics, she explains, is that genes that are very close to one another often are controlled by a single regulatory element and are turned on or off in concert. Therefore, understanding where proteins are located in relation to one another on a chromosome can provide valuable insight into their control, or co-expression.

"PQuad is nice in that the developers have recognized that people are looking for ways to highlight the relevant information," Adkins says. "PQuad is good in that it will do comparisons between experiments and place the results over sequence information. By doing this, you can then go over it visually and then by a color map as you scan the genome you can quickly recognize areas of interest."

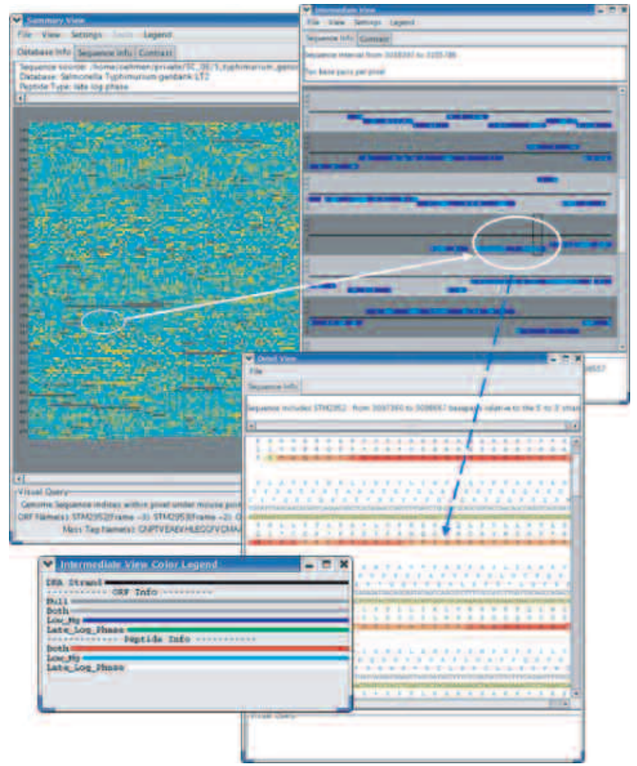

The basic visualizations allow interactive refinement of views based on browsable information representations. The Genome View (top left) provides an overview of the entire proteomics dataset at the chromosome level. The next level of resolution is the ORF View (top right). This view shows individual peptides (short, light blue bars) localized to their respective proteins or ORF's (long, dark blue bars); these peptide and protein "bars" are positioned with respect to their reading frame (indicated on left side of the window). The Sequence View (bottom right) gives genome-level information for a selected ORF, allowing the user to observe the amino acid string associated with a peptide identification.

The sheer quantity of data generated by these new biology fields requires a computational solution. But even with computers helping to reduce data to manageable chunks, biologists must make judgments about the biological relevance of their results.

COLLABORATORS

Bobbie-Jo Webb-Robertson is a Senior Research Scientist at Pacific Northwest National Laboratory's Computational Biology and Bioinformatics Division. She earned her B.A. in Mathematics from Eastern Oregon University and her M.E. in Operations Research and Statistics and Ph.D. in Decision Sciences and Engineering Systems from Rensselaer Polytechnic Institute.

Chris Oehmen is a Senior Research Scientist at Pacific Northwest National Laboratory's Computational Biology and Bioinformatics Division. He earned his B.A. in Physics and Mathematics from Saint Louis University and his M.S. and Ph.D. in Biomedical Engineering from University of Memphis & The University of Tennessee Health Science Center. Dr. Oehmen was a DOE CSGF fellow from 1999-2003.

Further Reading:

C. S. Oehmen, T. Straatsma, G. A. Anderson, G. Orr, B. M. Webb-Robertson, R. C. Taylor, R. W. Mooney, D. J. Baxter, D. R. Jones and D. A. Dixon. (2006) *New Challenges Facing Integrative Biological Science in the Post-Genomic Era.* Journal of Biological Systems, 14(2):275-293.

K. M. Waters, M. Singhal, B. M. Webb-Robertson, E. G. Stephan and J. M. Gephart. (2006) *Breaking the High-Throughput Bottleneck.* Scientific Computing, 23(5).

K. F. Ferris, B. M. Webb-Robertson and D. Jones. (2006) *Information-Based Development of New Radiation Detectors.* Mater. Res. Soc. Symp. Proc., 894: 307-312.

W. R. Cannon, K. H. Jarman, B. M. Webb-Robertson, D. J. Baxter, C. S. Oehmen, K. D. Jarman, A. Heredia-Langner, K. J. Auberry and G. A. Anderson. (2005) *Comparison of Probability and Likelihood Models for Peptide Identification from Tandem Mass Spectrometry.* Journal of Proteome Research, 4(5): 1687-1698.

B. M. Webb-Robertson, C. S. Oehmen and M. M. Matzke. (2005) *Bayesian Alignment Improves Sensitivity for Protein Family Classification from Sequence Based Support Vector Machines.* Computational Biology and Chemistry, 29(6): 440-443.

Contact:

Bobbie-Jo Webb-Robertson
bj@pnl.gov

Chris Oehmen
Christopher.oehmen@pnl.gov

PRACTICUM COORDINATOR

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

As the computational biologists worked with biologists, they realized that to make PQuad useful, they would have to work together and learn the other field to some extent.

That's where Department of Energy Computational Science Graduate Fellowship alumnus Chris Oehmen helped accelerate the process, Webb-Robertson says. Oehmen's training in modeling biological systems (during his practicum, he built heart function models) came into play as he helped with user evaluation and further development of PQuad.

"Understanding the language of biology has made it really easy for me to talk to people in biology," says Oehmen. "I know what a protein is and how it is made. I can kind of be the middle person. I can talk to the software design people, I can talk to the people who are creating the data, and I can understand both sides enough to keep things moving along."

But Does It Work?

The PQuad team has released the software for public use at http://ncrr.pnl.gov/software/software_register.asp?id=23. Now, they are eager to show what it can do. Oehmen and Webb-Robertson, along with PNNL colleagues Adkins, microbiologist Lee Ann McCue, and programmers, data analysts, mathematicians, and graphics experts, collaborated to enter the Supercomputing 2006 analytics challenge. The challenge lets researchers showcase computationally intensive applications that use high-performance computing, networking, and visualization to solve real-world, complex problems.

For their demonstration problem, the research team analyzed the biochemical pathways that the bacterium Salmonella uses to produce toxins that cause food poisoning. The project addresses the laboratory's mission to understand the fundamental biochemical pathways of microorganisms. The group used proteomic datasets from samples taken under growth conditions designed to mimic environments in which the organism produces or doesn't produce toxins. The sample of all proteins produced under the two growth conditions was first digested by enzymes into small pieces called peptides, and then aerosolized and injected into a mass spectrometer, which records the mass-to-charge ratio of each peptide fragment.

The data analysis portion of the project combines several technologies developed by the PNNL research team. First the team feeds the raw proteomic data through a program called Polygraph, which compares the experimental data to an "ideal" dataset of all potential Salmonella peptide fragments and produces a statistical measure of confidence for each fragment. For the dataset examined in the analytics challenge, the program must process 800,000 data points, which requires parallel processing to complete in near real-time. The team uses an Itanium II cluster of 512 processors, which can process each dataset in about 6 minutes.

Once the data has been processed, the researchers import it into PQuad. It's tied to an integrated Bioinformatics Resource Manager (BRM) that can access publicly available applications, including tools available through a Web interface. This multi-functionality makes it possible to identify a protein and then open another window to visualize the pathway associated with a particular protein of interest.

"The workflow enabled by BRM sped up that whole process of going from spectra to protein identification by something like 20-fold," McCue says. "This method puts the information in the hands of biologists much more quickly."

In the case of the Salmonella problem, the research team zeroed in on what are called "pathogenicity islands" — DNA segments that produce many proteins responsible for making a microbe toxic or "virulent." Using the visual tools of PQuad, McCue was able to identify proteins present in the virulent growth condition that were not present in the non-virulent growth condition, providing validation for the experimental growth conditions. "We went to those pathogenicity island regions and one of them in particular was really standing out as coming up under one condition and not the other. That gave us some confidence that this could be pretty interesting and we could look for other areas of interesting proteins," McCue says.

She then identified proteins in a region called type III secretory system, which appeared to be turned on as well.

Oehmen says it's an example of what the team is moving toward. "In this demonstration project we showed that we could identify something a biologist would find interesting that could help them direct their next set of experiments."

"One of the nice things about this project," Adkins says, "is that we are now in a position to follow-up and validate some of our observations." Adkins is an investigator in the Biodefense Proteomics Research Center, a collaboration between PNNL and the Oregon Health and Science University in Portland, that is trying to find new treatments for possible biowarfare agents. "What makes this work meaningful is that we have been able to take piles of files that contain raw experimental results and quickly convert them," Oehmen says. "That's why analytics is such an important piece of high-performance computing."

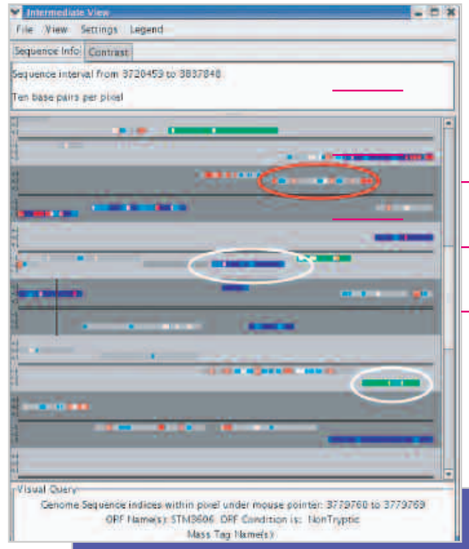

The comparative proteomics capability of PQuad allows concurrent visualization of peptides and proteins identified under two distinct conditions, such as virulent versus non-virulent. The peptides are colored based on identification conditions: virulent (light blue), non-virulent (white), and both (red). For easy distinction of proteins uniquely identified in the virulent versus non-virulent conditions the proteins are colored based on the likelihood of identification in only one of the conditions: virulent (dark blue) and non-virulent (green). Two examples of uniquely expressed proteins are circled in white; a case of a protein expressed in both conditions is circled in red.

By Jacob Berkowitz

X-treme Waste Modeling

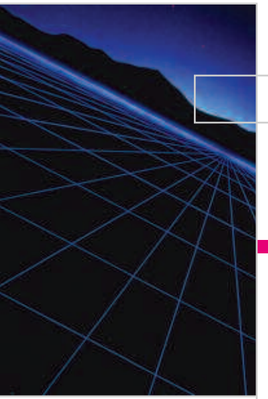

MARK PETERS HAS A DREAM. He's a half-mile underground inside a mountain. In front of him are hundreds of shiny metal containers. He watches as, atom-by-atom, microscopic changes take place in the containers. Neutrons and beta-particles from the containers' radioactive contents ricochet through the surrounding rock. Eons-old water seeps over the containers. In a scientific reverie he watches this subterranean scene unfold for generations — for tens of thousands of years.

Today it's a dream that occupies a large part of the DOE Argonne National Laboratory senior manager's waking hours — because it's a dream he wants to see realized. Peters is at the forefront of a movement to dramatically improve the computational modeling of long-term nuclear waste disposal to literally see the future of high-level waste nuclear fuel by modeling it first *in silico* on the silicon chips of the nation's most powerful computers.

It's a scientific mission driven by a political and economic one: the U.S. Global Nuclear Energy Partnership (GNEP). The GNEP is an ambitious plan to increase the nation's and the world's use of nuclear energy, while minimizing the risk of nuclear weapons proliferation. It's the framework for a new era of nuclear energy built on a 21st century vision of recycling nuclear fuels, new reactor types and new types of radioactive wastes, and long-term disposal techniques.

What Peters and others know is that just as the GNEP plans to transform nuclear power in the U.S. and beyond, achieving this goal will partially rely on a near-revolution in our ability to computationally simulate nuclear energy systems. That's especially important when it comes to the most politically and technically challenging nuclear issue of them all: long-term geologic disposal of radioactive waste.

Handling the Waste of a Nuclear Renaissance

The GNEP is a key part of the current U.S. energy policy that's helping jump-start a nuclear industry renaissance at home and abroad. Driven by domestic political concerns over greenhouse gas emissions and energy security, commercial nuclear power plants once again are on the drawing board in the U.S. after a hiatus of nearly two decades. The U.S. Nuclear Regulatory Commission has more than a dozen expressions of interest in new U.S. nuclear power plants.

But the long-term success of this nuclear resurgence still depends on what has been the industry's Achilles' heel.

"Fundamentally, with nuclear power the technical challenge is managing the wastes," says Peters, who for eight years led the science and engineering testing for the DOE's Yucca Mountain Repository Project.

Yucca Mountain, located about 100 miles northwest of Las Vegas, Nevada, within the boundaries of the Nevada Test Site, is the only proposed U.S. geological repository for high-level nuclear waste and spent nuclear fuel. These radioactive wastes, from nuclear power generation and national defense programs, now are stored at 126 sites around the nation. Congress and the president approved the Yucca Mountain site in 2002, and DOE is applying for a Nuclear Regulatory Commission license to proceed with construction of the repository.

View of the rear of the tunnel boring machine showing the laser guidance system in operation at Yucca Mountain's Exploratory Studies Facility.

"Yucca Mountain is an essential part of the nuclear fuel cycle," says Peters. "At the same time, if we're really going to establish nuclear power as a major, long-term energy source, we have to come up with a more sustainable system that involves recycling the majority of the waste."

Nuclear fuel recycling is at the core of the GNEP vision. It's based on the use of so-called fast burner reactors, a technology already in use elsewhere, including at France's Phenix nuclear reactor. Today's commercial nuclear reactors in the U.S. are thermal reactors, which "burn" enriched uranium — mined uranium in which the fissionable atoms are concentrated. After a few years, the enriched fuel no longer reacts efficiently and must be disposed of as high-level waste, eventually in the Yucca Mountain repository.

The GNEP is an ambitious plan to increase the nation's and the world's use of nuclear energy, while minimizing the risk of nuclear weapons proliferation. It's the framework for a new era of nuclear energy built on a 21st century vision of recycling nuclear fuels, new reactor types and new types of radioactive wastes and long-term disposal techniques.

MARK T. PETERS

Mark Peters received his B.S. in geology from Auburn University and his Ph.D. in geophysical sciences from the University of Chicago. He is currently the Deputy to the Associate Laboratory Director for Applied Science and Technology at Argonne National Laboratory. Technical responsibilities include the development of new program opportunities, management and integration of the energy and environment-related LDRD program, and support of the Global Nuclear Energy Partnership (GNEP) program. His duties also include responsibility for program management and development in high-level nuclear waste disposal across the Laboratory, management and technical leadership of Department of Energy (DOE) long-term science and technology work related to radionuclide source terms, and technical support to senior DOE management on the Yucca Mountain Project (YMP).

Previously, Peters was on a two-year detail to DOE Headquarters as a senior technical advisor to the Director of OCRWM. Prior to that, he was the YMP Science and Engineering Testing Project Manager. He was responsible for the technical management and integration of science and engineering testing in the laboratory and field on the YMP.

Earlier, Peters was the YMP Office Manager for Thermal Testing and Engineering Support. He was responsible for the technical integration of science, construction, and design organizations and technical leads engaged in field testing at Yucca Mountain.

Before joining the YMP in 1995, Peters had a research fellowship in geochemistry at the California Institute of Technology. He conducted research in trace-element geochemistry and presented the results at national and international meetings. He has also been published in several major scientific journals.

Activities and Achievements
Member of Geological Society of America, American Geophysical Union, Geochemical Society, and Mineralogical Society of America.

Member of Geological Society of America Committee on Geology and Public Policy

Member of Sigma Xi, Scientific and Engineering Research Honorary Society

Member of Sigma Gamma Epsilon, Earth Sciences Honorary Society

Published over 60 abstracts and papers in peer reviewed journals and proceedings.

Further Reading:
M. T. Peters and R. C. Ewing, 2006. A science-based approach to understanding waste form durability in open and closed nuclear fuel cycles: *Journal of Nuclear Materials* (in press).

J. Lian, F. X. Zhang, M. T. Peters, L. M. Wang, and R. C. Ewing, 2006. Ion beam irradiation of La- and Th-doped $Y_2Ti_2O_7$ titanates: *Journal of Nuclear Materials* (in press).

T. Kiess, R. Budnitz, D. Duncan, M. Peters, J. Wengle, and H. Williams, 2005. OCRWM's inaugural science and technology program for Yucca Mountain: *American Nuclear Society Radwaste Solutions*, May/June Issue, 41-48.

Contact:
Mark Peters
mtpeters@anl.gov

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

Fast burner reactors significantly reduce the quantities of high-level waste by recycling the spent fuel. After the initial enriched uranium fuel has been used once, it's removed and processed to create a variety of other radioactive waste products and recycled fuel products. These techniques could reduce the amount of high-level waste by as much as 100 times.

Beyond Yucca

The challenge is that, until now, the Yucca Mountain repository has been designed to handle wastes from existing reactors, not from recycling in fast burner reactors.

"The radioactive wastes produced by recycling in the fast burner reactors envisioned by the GNEP are significantly different than those currently planned for Yucca," says Peters. They're different in two key ways: the radioactive wastes from fast burners have a shorter radioactive half-life — they decay to non-radioactive forms thousands of years sooner than wastes from thermal reactors — and they produce less heat.

"We need the tools to be able to computationally model how these different wastes will behave in the rock. The problem is we don't have the best tools now to answer these questions in a comprehensive and timely manner for a wide range of geologic settings. So what I'm after goes beyond Yucca Mountain," Peters says.

It's a desired computational approach shared by colleagues. This past summer, a DOE workshop on Simulation and Modeling for Advanced Nuclear Energy Systems brought together, for the first time, leading scientists in nuclear energy, applied mathematics and computer science. Their job: Envision how to ramp-up U.S. ability to simulate advanced nuclear energy systems, from new reactor designs and fuel types to geological repositories. The workshop's final report concluded: "Current [computer codes] for advanced nuclear systems lack the predictive ability needed to achieve the goals of the GNEP."

Aerial view of Yucca Mountain.

It's this predictive ability that's key to successful nuclear waste modeling, says University of Wisconsin Engineering Physics Assistant Professor Paul Wilson, who co-chaired the repository modeling session in Washington, D.C.

"Predictive modeling of materials in extreme environments is a recurring theme throughout the simulation of advanced fuel cycles," says Wilson. "But the big difference with repository modeling is the time scale. You have to model tens of thousands of years into the future. This is the aspect that's unique to repository modeling." The fuel cycle is the nuclear fuels' entire journey from the ground, through use and recycling, and finally back into the ground in a geologic repository.

Back to Basics

Argonne's Peters is acutely aware of the challenges involved in the long-term predictive modeling of waste fuels in geologic repositories. He began at Yucca Mountain in 1995, as DOE computational scientists were building a comprehensive computational model to forecast how radioactive wastes from spent nuclear fuel placed in the mountain would behave over the millennia.

The total system performance assessment (TSPA) model was specifically designed for assessing the Yucca Mountain repository's long-term performance. The underpinnings of the TSPA are primarily empirical models — their predictive ability is based on experimental data that's then extrapolated into the future. This is similar to drawing a line through existing points on a graph and extending the line out past the graph. The model includes more than a dozen factors, from water seepage in the repository tunnels to the transport of radioactive particles through the tunnel floor.

Peters says next-generation repository models must be built on first-principle physics and chemistry. This means understanding the fundamental chemical and physical behavior of the waste, the containers and the surrounding rock, and then integrating those behaviors. These models must be able to simulate wastes not just from today's thermal reactors, but also from future fast burner reactors.

"We need a first-principle understanding of how the waste behaves," says Peters. "We want to demonstrate that we understand how any given waste will perform in a variety of geochemical environments. Then we could apply this broadly and provide feedback on developing waste types that are safest for disposal in a given geologic repository."

Peters' current focus is on assembling a research group and securing funding to develop new advanced computational models for the repository source term — the combined performance of the waste and disposal container. "This is the most important question in terms of how the repository fits in with the overall fuel cycle — it's the place where the waste meets the rock," notes Peters.

Peters' proposed research team is a rock-and-radiation scientific group. It will include university and DOE geochemists and materials scientists with experience in high-radiation and aqueous environments. The advanced computational models they create will be able to predict the heat, chemical and radiation stability of advanced materials, such as specifically tailored ceramics, composite materials, and super alloys.

The models they build will be coupled, tying together vast amounts of data from the sub-atomic scale to literally the mountain scale. Coupled models combine a wide number of different physical phenomena. For example, coupled climate models integrate the various forces in the atmosphere, in the oceans and on land that drive climate patterns. These coupled advanced nuclear waste repository models will require the power of DOE supercomputers to model the enormous number of interactions over tens of thousands of years. They'll also require the integration of existing computational codes developed for the Yucca TSPA.

"The major challenge as I see it is making use of existing legacy simulation models and coupling them with new physics-based models in an efficient way that takes advantage of high-performance computing systems," says Sandia National Laboratory scientist Dr. Robert MacKinnon, an expert in TSPA analysis of radioactive waste repositories, and co-chair of the Washington, D.C. workshop on repository modeling.

The research also will push the boundaries of uncertainty analysis and validation tools. How do you determine whether a computational simulation of nuclear waste behavior in rock over 100,000 years actually is accurate? "There's no way to prove this experimentally," says Peters. Thus, the advanced repository simulations will draw on the latest applied mathematics techniques for determining levels of uncertainty in computational models.

With the political push on to build more nuclear reactors, Peters and colleagues know time is of the essence. They need to build the computational models quickly enough so that they can inform the actual steel-and-concrete building. But Peters says computational models could provide a key technical and political ingredient that will help turn the GNEP vision into a reality: reliable facts. New reactors and repositories require regulatory approval, with particular emphasis on predictive science. The DOE's Yucca Mountain website notes that "Our license application must present a scientifically defensible position that we can construct, operate, and close a repository without unreasonable risk to people and the environment."

Says Peters, "If we can come up with a set of models for understanding how the waste form behaves in a geological repository, we demonstrate a better scientific understanding, and it could also enhance our ability to select and license repositories as part of the future global nuclear energy enterprise."

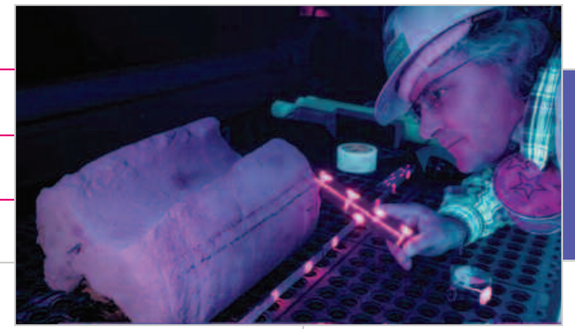

A scientist uses ultraviolet light to study how fluids move through rock.

Patrick Canupp

There's a FUNDAMENTAL DIFFERENCE between Patrick Canupp's job and what he did for more than a dozen years as a student and professor.

"In academia you're constantly trying to show the world what you've done," says Canupp, a DOE CSGF fellow from 1991 through 1995. Now, as chief aerodynamicist for an auto racing team, "You don't want anyone to know."

That's the competitive nature of NASCAR, the National Association for Stock Car Auto Racing. When cars speed by at more than 200 miles per hour, even the tiniest advantages are vital.

"If we can see gains that are on the order of, say, half a tenth of a second, those are generally considered good gains and you try to get as many as you can," says Canupp, a member of the Joe Gibbs Racing team, based in Charlotte, N.C. "You're kind of searching for needles in a haystack and doing fine-tooth combing."

Canupp is one of only two NASCAR aerodynamicists with doctoral degrees. He earned his in aeronautics and astronautics from Stanford University in 1997 with a dissertation on modeling neutral gas dynamics in high-density plasmas. He taught at Gardner-Webb University in his native North Carolina before getting into racing.

"This has turned out to be pretty fascinating, but it wasn't what I saw myself doing at all," says Canupp, who wasn't a NASCAR fan before his first racing job. "It maybe even looked a little bit strange to me, after developing my skills in computational fluid dynamics."

Canupp rarely uses computer modeling now, relying instead on wind tunnel and track testing. His goal is keep his team's cars on the track. Drivers must back off the gas as they go through the turns, and "The key is who can get back on the throttle and who can be at 100 percent speed leaving the corner" without skidding, Canupp says. "The guy who can get to the throttle quickest will have the fastest car and the greatest chance to win."

Canupp and the engineers and builders on the Gibbs team tune the aerodynamics to create downforce, pushing the car onto the track. It sounds simple, but it's difficult to do consistently. Cars may be identical in every way engineers can find and still perform differently, Canupp says. "When you start digging into why that is, it's kind of mind-boggling," he adds. "There are a lot of complex things that, even though I have a Ph.D. in aerospace engineering, I can't explain."

NASCAR drivers and teams have the kinds of fame and followers previously reserved only for rock stars. Yet, Canupp says his job is rarely glamorous. He spends most of his time in the shop, crunching test results and tweaking the car's design within the strict NASCAR limits. In a previous racing job he often rode the company plane to watch races from pit row. Now Canupp watches most of them from the comfort of his home.

"I view that as a feedback session — sometimes in excruciating detail," he adds.

"In academia you're constantly trying to show the world what you've done," says Canupp, a DOE CSGF fellow from 1991 through 1995. Now, as chief aerodynamicist for an auto racing team, "You don't want anyone to know."

Oliver Fringer

Oliver Fringer has a DIFFERENT PERSPECTIVE than most people when he visits beaches in the San Francisco Bay area, where he lives. That's because he's made his reputation creating computer models of ocean activity in coastal areas.

"I always look at the way the waves and small ripples propagate" as they break across the shore, says Fringer, a DOE CSGF fellow from 1997 to 2001.

"They do it in the same way large waves do," and "you can describe how very large nonlinear waves interact by describing how those small waves interact."

Fringer is an assistant professor of environmental fluid mechanics and hydrology at Stanford University, where he earned his master's degree in aeronautics and astronautics in 1996 and his doctoral degree in civil and environmental engineering in 2003. Fringer was serving a postdoctoral fellowship at Stanford when he developed what has become the centerpiece of his academic reputation and research: the Stanford Unstructured Nonhydrostatic Terrain-Following Adaptive Navier-Stokes Simulator, or SUNTANS.

Fringer's model is different because it simulates nonhydrostatic pressure. Most ocean simulation codes have been hydrostatic because most ocean flow is long and horizontal, Fringer says. That's fine for global and regional ocean circulation, but it's less accurate when modeling wave action along coastal areas, where nonhydrostatic pressure plays a bigger role. SUNTANS, which is supported by the Office of Naval Research, simulates the massive internal waves that begin deep in the ocean, then gain strength as they rise and break near the shore.

However, nonhydrostatic models make huge demands on computer resources. "The reason traditional models haven't gotten into nonhydrostatic pressure is you have to solve the three-dimensional elliptic equation. It's very expensive to solve elliptic problems in the ocean," Fringer says. Ten years ago, a simulation like SUNTANS would have been impossible, but researchers have since found ways to solve the 3D elliptic equation more efficiently.

SUNTANS can model flows of any coast, river or estuary, Fringer says. The graduate students he supervises use the basic hydrodynamic kernel he devised to study sediment transport, internal waves in the South China Sea, and other ocean activity.

Such models can be instrumental in understanding how pollutants move through the ocean and how weather systems form. "The ultimate goal is to predict weather and climate," Fringer says. "To do that you have to be able to accurately model the interaction of the ocean and atmosphere. To do that, you have to understand how energy moves around the ocean in the form of internal waves."

Part of Fringer's current research focuses on integrating SUNTANS with the Regional Ocean Modeling System (ROMS), a University of California – Los Angeles project, to simulate ocean activity over a larger range of space and time.

Fringer's interest in research was inspired by his boyhood remote-control plane hobby and his contact with a Massachusetts Institute of Technology aerospace engineering graduate student. Now, Fringer's profession already is influencing Avery, the daughter his wife, mechanical engineer Krista Donaldson, gave birth to in June.

"Poor Avery," Fringer laughed. "Whenever I give her a bath, I already tell her about fluid mechanics."

Fringer's interest in research was inspired by his boyhood remote-control plane hobby and his contact with a Massachusetts Institute of Technology aerospace engineering graduate student.

Aric Hagberg

Aric Hagberg's analyses of COMPLEX PATTERN formation have gained an unusual fan: a Boston artist who bases works on Hagberg's colorful demonstrations.

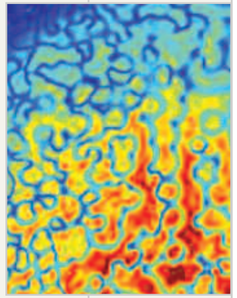

Hagberg, a DOE CSGF fellow from 1992 to 1994, has a broad portfolio as a scientist at Los Alamos National Laboratory. The group he belongs to focuses on mathematical modeling and analysis, but Hagberg has dabbled in computer networking, programming, and even assembly of an early high-performance computer cluster.

Hagberg's doctoral dissertation and subsequent research focus mainly on mathematical analysis of how chemical patterns propagate in systems like the Belousov-Zhabotinsky reaction. The mathematical models — and experiments they're compared against — generate pictures and movies of spirals, vortices, and labyrinths. (Go to http://math.lanl.gov/~hagberg/Movies/) That caught the attention of artist Brian Knep, who asked to borrow the computer codes for the simulations. "I thought, 'Should I answer some crazy artist's questions?'" Hagberg recalls. Hagberg sent Knep some information, and the artist took it from there. His computer creations interact with gallery visitors and mimic patterns found on Hagberg's Los Alamos Web site. (See Knep's work at http://www.blep.com/projects.htm)

"They're beautiful. He really captured the essence," Hagberg says. "He is a very skilled software engineer. I'll send him anything he wants now."

The art connection is the most unusual aspect of a Los Alamos career that started with a life-changing summer job there. "I was a bit lukewarm about graduate school, but after that I was energized," says Hagberg, who is the lab's DOE CSGF practicum coordinator. "The national labs are really where the rubber meets the road. They solve real-world problems."

One example is Hagberg's work on network dynamics. In one project, he's working with biologists and engineers to improve wireless sensor networks for ecological research. The idea is to have sensors analyze data and make initial judgments about their quality. That will make it easier for biologists to sift the piles of information the sensors generate. Hagberg also authored NetworkX, an open-source program to create, manipulate and analyze complex networks. People from around the world, including the U.S. military, have used it.

It's typical for mathematical modeling and analysis researchers to dive into such a variety of projects, Hagberg says. "There's 30 of us and there are probably 40 or 50 projects" they're working on, he adds. "We're like consultants. We search for interesting and hard problems."

That spirit led Hagberg to collaborate with other Los Alamos researchers on assembling what was rated as the 113th fastest machine in the world in 1998. Spending just $313,000, the team bought 140 personal computers by mail order and connected them in a cluster. "We showed people you could do it. You could definitely build a Top 500 supercomputer" for less cost with commercial technology, Hagberg says. "Now, you go into any university, and they have lots of these Linux clusters around."

The exercise illustrates the entrepreneurial spirit Hagberg enjoys at Los Alamos. He sounds like someone destined to be there a long time.

"I know a good thing when I see it," Hagberg adds. "It's a great place to work, and I was lucky to find it."

> **"The national labs are really where the rubber meets the road. They solve real-world problems."**

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.

Dr. Kevin Chu of Princeton University and Dr. Matthew Wolinsky of the University of Minnesota have been named 2006 Frederick A. Howes Scholars in Computational Science.

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

One of Howes' responsibilities was to oversee the Department of Energy's Computational Science Graduate Fellowship (DOE CSGF) program. He was extremely committed to this program that supports graduate students in computational science. This program is unique, as it requires candidates to take courses in mathematics, computer science and an applications discipline, such as physics or engineering. The DOE CSGF program currently supports over 70 graduate students and is administered by the Krell Institute.

To honor his memory and his dedication to the Department of Energy's Computational Science Graduate Fellowship program, one DOE CSGF fellow is chosen each calendar year as a Howes Scholar. But because there were so many outstanding nominees for the award this year, two winners were selected. This award will provide the recipients with a substantial cash award, a Tiffany 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 DOE CSGF program or having been supported by the DOE CSGF program for the maximum number of allotted years.

Dr. Chu was a fellow from 2002-2005. He received his Ph.D. in Applied Mathematics from Massachusetts Institute of Technology in 2005. He is currently a Postdoctoral Research Associate in the Mechanical and Aerospace Engineering Department of Princeton University.

Dr. Wolinsky was a fellow from 2001-2005, and received his Ph.D. in Earth and Ocean Sciences from Duke University, also in 2005. He is currently a Postdoctoral Associate at the National Center for Earth-Surface Dynamics 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 (center) with the two 2006 Howes Scholars, Matthew Wolinsky (left) and Kevin Chu (right).

Kevin Chu receives his award from David Brown (LLNL) at the 2006 annual DOE CSGF fellows' meeting in Washington, D.C.

Matthew Wolinksy presents his research during the 2006 annual DOE CSGF fellows' meeting in Washington, D.C.

Alumni Directory

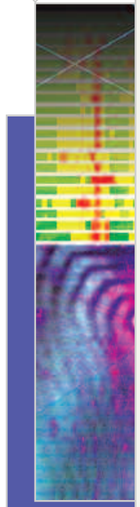

A
B
C

A

Bree Aldridge
Massachusetts Institute of Technology
Computational Biology
Fellowship Years: 2002-2006
Current Status: Student, MIT

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

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

Teresa Bailey
Texas A&M University
Nuclear Engineering
Fellowship Years: 2002-2006
Current Status: Student, Texas A&M University

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

Michael Barad
University of California – Davis
Civil & Environmental Engineering
Fellowship Years: 2002-2006
Current Status: Staff, Stanford University

Jaydeep Bardhan
Massachusetts Institute of Technology
Electrical Engineering
Fellowship Years: 2002-2006
Current Status: Argonne National Laboratory

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

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

Mary Biddy
University of Wisconsin
Chemical Engineering
Fellowship Years: 2002-2006
Current Status: British Petroleum

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

Nawaf Bou-Rabee
California Institute of Technology
Applied & Computational Mathematics
Fellowship Years: 2002-2006
Current Status: Student, California Institute of Technology

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

Kristine Cochran
University of Illinois
Civil Engineering
Fellowship Years: 2002-2006
Current Status: Student, University of Illinois at Urbana-Champaign

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: Staff, 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
Current Status: Staff, Central Michigan University Research Corporation

D

William Daughton
Massachusetts Institute of Technology
Plasma Physics
Fellowship Years: 1992-1996
Current Status: Faculty, University of Iowa

Gregory Davidson
University of Michigan
Nuclear Engineering
Fellowship Years: 2002-2006
Current Status: Student, University of Michigan

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

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

Michael Driscoll
Boston University
Bioinformatics & Systems Biology
Fellowship Years: 2002-2006
Current Status: Student, Boston 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

Mary Dunlop
California Institute of Technology
Mechanical Engineering
Fellowship Years: 2002-2006
Current Status: Student, California Institute of Technology

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

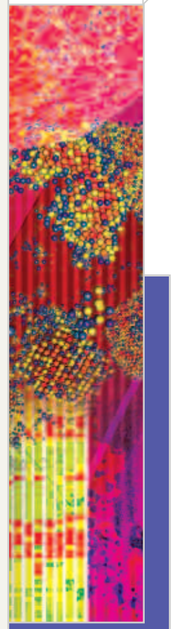

C
D
E

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
Current Status: Quant

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

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

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: Student, University of Michigan

Kristen Grauman
Massachusetts Institute of Technology
Computer Science
Fellowship Years: 2001-2005
Current Status: Faculty, University of Texas

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: Staff, 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: Pacific Northwest National Lab

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

Owen Hehmeyer
Princeton University
Chemical Engineering
Fellowship Years: 2002-2006
Current Status: ExxonMobil Upstream Research Corporation

Eric Held
University of Wisconsin
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: Los Alamos National Laboratory

William Humphrey
University of Illinois
Physics
Fellowship Years: 1992-1994
Current Status: NumeriX LLC

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: General Electric

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

Yan Karklin
Carnegie Mellon University
Computational Neuroscience
Fellowship Years: 2002-2006
Current Status: Student, Carnegie Mellon University

Richard Katz
Columbia University
Geodynamics
Fellowship Years: 2001-2005
Current Status: Staff, University of Cambridge

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.

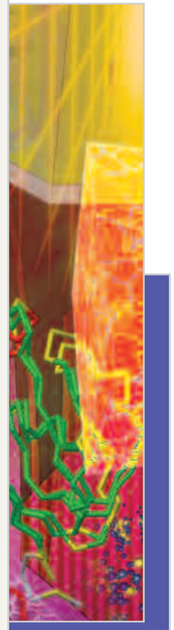

Michael Kowalok
University of Wisconsin
Medical Physics
Fellowship Years: 2000-2004
Current Status: Staff, Virginia Commonwealth University

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

Benjamin Lewis
Massachusetts Institute of Technology
Computational Biology
Fellowship Years: 2002-2006
Current Status: Student, MIT

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

Alex Lindblad
University of Washington
Civil Engineering
Fellowship Years: 2002-2006
Current Status: Sandia National Laboratories – California

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 Years: 1998-2002
Current Status: Staff, Harvard Medical School

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
Current Status: Pfizer

Erik Monsen
Stanford University
Aerospace and Astronautical Engineering
Fellowship Years: 1991-1994
Current Status: Staff, Max Planck Institute of Economics, Germany

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
Current Status: Center for Naval Analysis

Gregory Novak
University of California – Santa Cruz
Theoretical Astrophysics
Fellowship Years: 2002-2006
Current Status: Student, University of California – Santa Cruz

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

Emma Rainey
California Institute of Technology
Geological and Planetary Sciences
Fellowship Years: 2003-2006
Current Status: Arete Associates

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
Current Status: Staff, University of Georgia

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

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

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

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

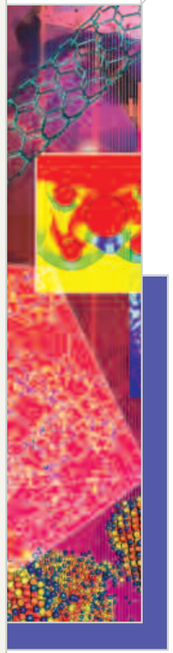

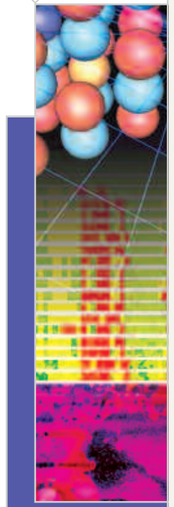

S

David Schmidt
University of Illinois
Elecctrical Engineering
Fellowship Years: 2002-2006
Current Status: Epic Systems

Samuel Schofield
University of Arizona
Applied Mathematics
Fellowship Years: 2001-2005
Current Status: Staff, Los Alamos National Laboratory

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

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

Amoolya Singh
University of California – Berkeley
Computational Biology
Fellowship Years: 2002-2006
Current Status: European Molecular Biology Lab, Heidelberg Germany

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 Status: 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

U

Obioma Uche
Princeton University
Materials/Statistical Mechanics
Fellowship Years: 2002-2006
Current Status: Sandia National Laboratories – California

V

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

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

Joshua Waterfall
Cornell University
Biophysics
Fellowship Years: 2002-2006
Current Status: Staff, Cornell University

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

Michael Wu
University of California – Berkeley
Computational Neuroscience
Fellowship Years: 2002-2006
Current Status: Student, University of California – Berkeley

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

Paul Bauman
University of Texas
Computational & Applied Mathematics

Advisor:
J. Tinsley Oden
Practicum:
Sandia National Laboratories – New Mexico
Contact:
pbauman@ices.utexas.edu
Research Synopsis:
The major thrust of our research is to develop a general approach to multi-scale modeling based on the notion of *a posteriori* estimation of modeling error and on adaptive modeling using so-called Goals algorithms. Events at different scales require, in general, different models, so multi-scale models should involve a blending or adaptation of models of one scale of events with those of another. It would seem to follow, therefore, that successful multi-scale modeling techniques should be able to compare models of different structure and to adapt features of different models to deliver desired results of an accuracy sufficient to capture essential features of the response or to make engineering decisions. This is the basis for the methodologies developed in the present work.

William Conley
Purdue University
Mechanical Engineering

Advisor:
Arvind Raman
Practicum:
Sandia National Laboratories – New Mexico
Contact:
wconley@ecn.purdue.edu
Research Synopsis:
It is estimated that a lack of understanding of tribology accounts for 420 billion dollars of wasted energy each year in the U.S. alone. By understanding the fundamental mechanisms which give rise to friction, it may be possible to modify and control macroscopic frictional properties through atomic level modifications of surfaces. Our theoretical predictions match with the majority of experimental literature. One recent paper suggests the existence of adequately small contact forces which results in no net frictional losses which Tomlinson's model predicts. Additionally, our theoretical work predicts a variety of nonlinear phenomena which few other researchers are examining. These include several different bifurcation mechanisms and the existence of separate isola due to subharmonic excitation. Future work includes a more detailed analysis of the cantilever using both continuum mechanics and molecular dynamics techniques. Additionally, lateral force experiments are planned on a Molecular Imaging atomic force microscope in the near future.

Aron Cummings
Arizona State University
Electrical Engineering

Advisor:
David Ferry
Practicum:
Sandia National Laboratories – California
Contact:
aron.cummings@gmail.com
Research Synopsis:
Over the past several decades, incredible gains have been made in the performance of computing technology. These improvements have been due primarily to the ability of manufacturers to fabricate MOSFETs at continuously smaller dimensions. However, eventually a point will be reached when it becomes physically impossible to make these transistors any smaller. Therefore, it is necessary to find an alternative method of building computational circuits. My course of study focuses on nanoscale electronics, and how one may approach quantum computing in this context. I am currently investigating various approaches to manipulating the spin of an electron, and how those manipulations could be used to do quantum computations.

Krzysztof Fidkowski
Massachusetts Institute of Technology
Computational Fluid Dynamics

Advisor:
David Darmofal
Practicum:
Argonne National Laboratory
Contact:
kfid@mit.edu
Research Synopsis:
My research deals with the development of a high-order finite element solver for Computational Fluid Dynamics. The target application is complex aerodynamic systems, although the software is developed to be general enough for other equation sets. The development project is being undertaken at the MIT Aerospace Computational Design Laboratory, in a team consisting of half a dozen graduate students and several faculty members. In the past three years, I have worked on implementing a high-order discretization, a multigrid solver, and parallelization, among other minor topics. My current work involves developing an output-based grid adaptation method for high-order finite element methods in two dimensions. Grid adaptation refers to changing the computational mesh, usually in order to improve the resolution by refining the mesh in a certain area, or to decrease computational cost by coarsening the mesh in an area. In particular, output-based grid adaptation refers to adapting the grid for improved prediction of a given engineering output, such as the lift or drag on a wing, for example. Such output-based error estimation is possible through the use of adjoint methods. A complete adaptation method then involves three basic ingredients: error estimation, mesh optimization, and a means of meshing the domain (i.e. a "mesher"). The goal of my research is to make all three of these ingredients automatic and robust. The motivation for an automatic adaptation method is to remove user involvement, the cost of which is very high in industry-level problems, from the adaptive solution process. Such a method is, of course, desired to be robust; that is, the failure rate of the adaptation should be kept to a minimum. An automatic and robust adaptation method is expected to make a significant impact in the area of aerospace design.

Tod Pascal
California Institute of Technology
Physicial Chemistry

Advisor:
William Goddard
Practicum:
Sandia National Laboratories – New Mexico
Contact:
tpascal@wag.caltech.edu
Research Synopsis:
My research involves the simulation of various materials as possible components for nano-machines. Using quantum mechanics to develop parameters for force fields, then using these force fields for molecular dynamics simulations, we examine the chemical and mechanical properties of these components, ranging from modified DNA molecules to nanotubes and buckyballs. Due to physical limitations, experimental techniques have difficulties characterizing the chemistry of these materials. Thus, in many instances, a computer simulation is the only tool capable of adequately describing these structures. Modeling and simulation, therefore, play a major critical role in nanotechnology, because of the difficulty in controlling nanostructures and in measuring reliable values for the properties of nanosystems. Additionally, simulations that lead to prediction of structural properties of nanostructures prior to experiments are of immense value since the experimental techniques for creating these nanostructures are usually time consuming.

Christina Payne
Vanderbilt University
Chemical Engineering

Advisor:
Peter Cummings
Practicum:
Sandia National Laboratories – New Mexico
Contact:
Christina.payne@vanderbilt.edu
Research Synopsis:
The focus my research is the simulation of a novel nanotechnology concept for rapid DNA sequencing. Through simulation, I examine the effect of electric fields, solvents, length and sequence of the DNA segment, and channel width on the transport properties of DNA through a nanoscale channel. These simulations will be performed using classical molecular dynamics techniques and high performance computing to obtain very detailed "data" that serves to provide insight into the behavior of the system. Additionally, ab initio methods will be used to develop forcefields (i.e., models for the ways atoms interact with each other both within the same molecule and in different molecules). Force fields are needed as input to the classical molecular dynamics simulations. The unusual combinations of material proposed for this screening device (e.g., DNA, ions, water, platinum, gold, etc.) mean that some of the force fields (DNA-platinum and DNA-gold) are either unknown or of questionable accuracy. This information would be used as an aid in experimental development of the nanodevice.

Mark Rudner

Massachusetts Institute of Technology
Physics

Advisor:
Leonid Levitov
Practicum:
Brookhaven National Laboratory
Contact:
rudner@mit.edu
Research Synopsis:
In my graduate research, I aim to combine ideas from quantum information science and condensed matter theory to investigate the problem of implementing quantum information processing in real physical systems. Currently, I am working on two projects along these lines.

1) Superconducting Persistent Current Qubit: Experimentalists at MIT and MIT's Lincoln Labs have created qubits based on persistent currents in superconducting rings. Loosely speaking, one can identify the clockwise and counterclockwise currents as the 0 and 1 states of the qubit, respectively. These researchers have demonstrated an unprecedented level of control over these devices, and have shown that they could maintain coherence even under strong driving from an external RF source. Their initial results in the high frequency, high power regime were described remarkably well by straightforward perturbation theory. They are now moving into a different regime where this approach will no longer be valid. Our aim is to understand the physics behind the behavior of the qubit in this new regime.

2) Quantum Dots: Quantum Dots are nanometer-scale semiconductor structures that can be easily fabricated and controlled, and tend to exhibit long coherence times. This makes them excellent candidates for future applications in quantum computation. Our task now is to better understand the properties of such devices and to figure out how they can best be engineered to allow reliable, coherent quantum manipulations. I am currently working on developing a model to understand recent experimental results pertaining to the spin-blockade effect and nuclear spin dynamics in two-electron double-dot systems.

Samuel Stechmann

New York University
Applied Mathematics

Advisor:
Andrew Majda
Practicum:
Los Alamos National Laboratory
Contact:
stechman@cims.nyu.edu
Research Synopsis:
When and where will it rain? To answer this question, countries all over the world have developed computer models to predict the weather. These computer models try to account for the rich variety of physics in the atmosphere to make their weather predictions better. My thesis work enters at the scientific level. We try to understand atmospheric physics better, so that weather prediction models can accurately represent the real world, so that we can better predict when and where it will rain.

For our models, and also for operational weather prediction models, computational science plays a central role. Observing the atmosphere is difficult and often expensive, and performing experiments is often impossible. Computer simulations can offer data with better resolution in both space and time, and they can offer data that is difficult to obtain with observations. To make sure the simulation data is believable, the models need to properly represent the real world. This is where my work fits in. We work with atmospheric models and numerical methods to make better simulations of the atmosphere, all from an applied mathematics perspective.

Brian Taylor

University of Illinois
Engineering Mechanics

Advisor:
Scott Stewart
Practicum:
Lawrence Livermore National Laboratory
Contact:
bdtaylo1@uiuc.edu
Research Synopsis:
My research is in the simulation of shock waves and detonations. Detonations are very high-speed (~3km/s) explosions initiated in materials that store vast amounts of chemical energy. Detonations are unique in that they are sustained by a thin chemical reaction zone immediately behind the shock interface.

To date, nearly all research involving detonation physics has been directed at large scale systems, with characteristic lengths on the order of a meter. My research is part of an effort to develop micro-detonic systems, using layers of explosives on the order of 100 microns thick. Detailed, accurate simulations are essential to this work, as they will provide the only time-resolved insight into the detonation process — the time scales are simply too brief to resolve by physical experiment.

William Triffo

Rice University
Bioengineering

Advisor:
Robert Raphael
Practicum:
Lawrence Berkeley National Laboratory
Contact:
triffo@rice.edu
Research Synopsis:
My general research interest involves the application of mathematical and engineering techniques to the analysis of physiologic processes. From a computational perspective, my goal is to develop quantitative models describing the behavior of neural and sensory systems and their associated pathology.

My current focus is the auditory system. In brief, sound waves traveling through the air give rise to vibrations of the eardrum, which are transmitted to the cochlea (the fluid-filled "inner ear") via the middle ear bones. The cochlear fluid oscillations then induce vibrations in the basilar membrane (BM) and Organ of Corti (OC). A cellular component of the OC, the inner hair cell (IHC), transduces these vibrations into neural signals which are then relayed up to the brainstem and cortex for further processing. The vibrations of the BM are spectrally tuned according to longitudinal location in the cochlear duct, and exhibit a transfer function left unexplained by passive mechanical models.

Michael Wolf

University of Illinois
Scientific Computing/Computer Science

Advisor:
Michael Heath
Practicum:
Lawrence Berkeley National Laboratory & Sandia National Laboratories – New Mexico
Contact:
mmwolf@uiuc.edu
Research Synopsis:
Much of my research has been in the field of computational electromagnetics, focusing on developing and improving high performance computing electromagnetic simulations. My previous work at the Stanford Linear Accelerator Center (SLAC) sparked an interest in the use of high-performance electromagnetic software to model and design next-generation particle accelerator structures. I developed a three-dimensional, time-domain electromagnetic field solver, Tau3P, used to model accelerator structures. I have spent a great deal of time improving the parallel performance of Tau3P. I used combinatorial algorithms to greatly improve the parallel mesh partitioning of Tau3P, greatly reducing the parallel communication. I also researched and determined a more optimal parallel matrix/vector multiplication algorithm which was the most significant computational step in Tau3P.

More recently, I have also been looking at a couple of different scientific computing problems. I have been collaborating with several people at Sandia National Laboratories, modeling an epidemiology problem. In particular, we are working on identification of the time and location of several small 'micro-releases' of pathogen based on the identification of infected individuals at medical institutions. I have been working on improving the efficiency of the forward disease propagation implementation. I am also working on developing parallel scalability analysis for several different parallel algorithms.

Brandon Wood

Massachusetts Institute of Technology
Computational Materials Science

Advisor:
Nicola Marzari
Practicum:
Lawrence Berkeley National Laboratory
Contact:
brandonw@mit.edu
Research Synopsis:
My work to date has focused on understanding superionic behavior in two materials: the archetypal fast-ion conductor alpha-AgI, and $CsHSO_4$, a promising solid-acid fuel cell electrolyte candidate. I have also recently begun using some more advanced dynamics techniques to study defect mobility in $NaAlH_4$, which is a complex metal hydride with potential application for solid-state hydrogen storage. The chemical characteristics of this material make it surprisingly similar to the superionics I have studied.

The basic understanding these simulations can provide is useful for exploring the structural, chemical, and thermodynamic factors that contribute to the stability and peculiarly high ionic conductivity of the superionic phase in known materials. It is my hope that such a fundamental understanding would prove to be of great value in designing and adapting new materials for promising energy technologies.

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
Practicum:
Sandia National Laboratories – New Mexico
Contact:
jldavis9@unity.ncsu.edu

Jeffrey Drocco
Princeton University
Biophysics & Computation
Advisor:
David Tank
Practicum:
Los Alamos National Laboratory
Contact:
jdrocco@princeton.edu

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

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
Demography/Remote Sensing
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
Practicum:
Lawrence Berkeley National Laboratory
Contact:
ez87@columbia.edu

John ZuHone
University of Chicago
Astrophysics
Advisor:
Donald Lamb
Practicum:
Oak Ridge National Laboratory
Contact:
zuhone@uchicago.edu

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

Zlatan Aksamija
University of Illinois
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
Practicum:
Lawrence Berkeley National Laboratory
Contact:
cscarey@wisc.edu

Ethan Coon
Columbia University
Applied Mathematics
Advisor:
Marc Spiegelman
Practicum:
Los Alamos National Laboratory
Contact:
etc2103@columbia.edu

Jeff Hammond
University of Chicago
Theoretical Chemistry
Advisor:
David Mazziotti
Practicum:
Pacific Northwest National Laboratory
Contact:
jhammond@uchicago.edu

Asegun Henry
Massachusetts Institute of Technology
Mechanical Engineering
Advisor:
Gang Chen
Practicum:
Sandia National Laboratories – New Mexico
Contact:
ase@mit.edu

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

Miler Lee
University of Pennsylvania
Genomics & 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
Practicum:
Los Alamos National Laboratory
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
Practicum:
Argonne National Laboratory
Contact:
stefan@orie.cornell.edu

FIRST YEAR FELLOWS

Mark Berrill
Colorado State University
Electrical & Computer Engineering
Advisor:
Jorge Rocca
Contact:
berrill@engr.colostate.edu

Arnab Bhattacharyya
Massachusetts Institute of Technology
Computer Science
Advisor:
Madhu Sudan
Contact:
abhatt@mit.edu

Jenelle Bray
California Institute of Technology
Computational Biophysical Chemistry
Advisor:
William Goddard III
Contact:
jenelle@caltech.edu

Julianne Chung
Emory University
Computational Mathematics
Advisor:
James Nagy
Contact:
jmchung@emory.edu

Tal Danino
University of California – San Diego
Bioengineering
Advisor:
Jeff Hasty
Contact:
tdanino@ucsd.edu

Jack Deslippe
University of California – Berkeley
Physics
Advisor:
Steven Louie
Contact:
jdeslip@berkeley.edu

John Evans
University of Texas
Computational & Applied Mathematics
Advisor:
Thomas Hughes
Contact:
JohnAEvans@mail.utexas.edu

Ashlee Ford
University of Illinois
Chemical Engineering
Advisor:
Richard Braatz
Contact:
anford2@uiuc.edu

Kristi Harris
University of Maryland – Baltimore County
Physics
Advisor:
Philip Rous
Contact:
kristi.harris@umbc.edu

David Ketcheson
University of Washington
Applied Mathematics
Advisor:
Randall LeVeque
Contact:
ketch@amath.washington.edu

Brian Levine
Cornell University
Transportation Systems
Advisor:
Linda Nozick
Contact:
bl76@cornell.edu

Oaz Nir
Massachusetts Institute of Technology
Computational Biology
Advisor:
Bonnie Berger
Contact:
oaz@mit.edu

Carolyn Phillips
University of Michigan
Applied Physics
Advisor:
Pam Derry
Contact:
phillicl@umich.edu

Alejandro Rodriguez
Massachusetts Institute of Technology
Condensed Matter Theory
Advisor:
Steven G. Johnson
Contact:
alexrod7@mit.edu

David Rogers
University of Cincinnati
Physical Chemistry
Advisor:
Thomas Beck
Contact:
rogersdd@email.uc.edu

Ariella Sasson
Rutgers University
Quantitative Biology
Advisor:
Wilma Olson
Contact:
ariella@eden.rutgers.edu

Michael Sekora
Princeton University
Continuum Mechanics, PDE
Advisor:
James Stone
Contact:
sekora@math.princeton.edu

Benjamin Smith
Harvard University
Experimental High Energy Physics
Advisor:
Masahiro Morii
Contact:
bcsmith@fas.harvard.edu

Benjamin Sonday
Princeton University
Plasma Physics
Advisor:
Salvatore Torquato
Contact:
bsonday@princeton.edu


Los Alamos