2008 - 2009

DEIXIS

THE DOE CSGF ANNUAL

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

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

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

DOE CSGF

2008 - 2009

2008 - 2009

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

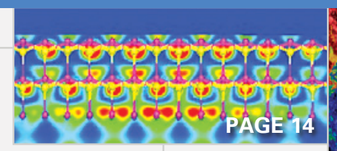
PAGE 14

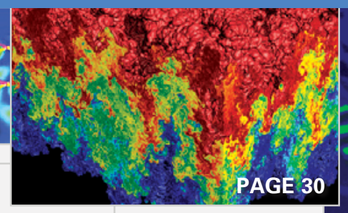
PAGE 30

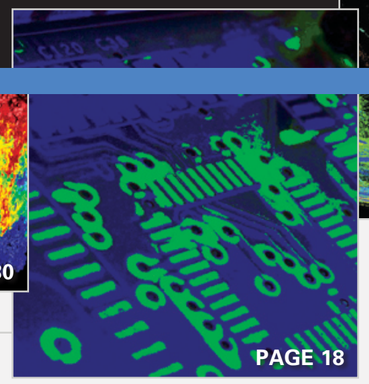
PAGE 18

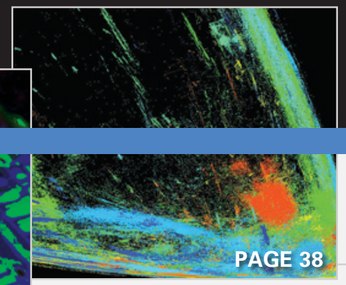
PAGE 38

PAGE 26

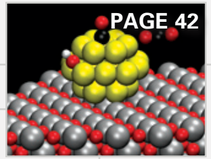
PAGE 42

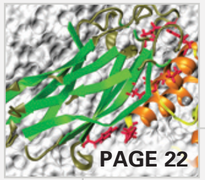
PAGE 22


PAGE 34

2008-2009

Editor
Shelly Olsan

Copy Editor
Ron Winther

Design
julsdesign, inc.

Publication Coordinator
Buffy Clatt

Contributors
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 some 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.*

R&R: Research and Revelation

THE DEPARTMENT OF ENERGY **Computational Science Graduate Fellowship (DOE CSGF) supports the nation's brightest science and engineering students, allowing them to concentrate on learning and research. The work of more than 200 DOE CSGF alumni has helped the United States remain competitive in a global economy.**

Left to right: Julianne Chung, David Ketcheson, Jordan Atlas and Ashlee Ford in Washington, DC at the annual DOE CSGF meeting.

SINCE 1991, the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) has trained the nation's brightest science and engineering students to apply computational science to complex problems. More than 200 DOE CSGF alumni now work in government, academia and industry, helping maintain American competitiveness.

Sometime during their graduate careers, participants in the DOE CSGF pack up for three months — usually in the summer — and head to a national laboratory in the DOE system. The fellows set aside their usual research and work with lab scientists on a new project. It's a chance for fellows to try something new, to experience lab life and to work with some of the world's most powerful computers.

As the experiences of these four fellows demonstrate, the practicum summer is no easy vacation filled with make-do work. It often leads to deep discovery — both professional and personal.

Practicum Proves Profitable for Image Reconstruction

JULIANNE CHUNG

Emory University | Lawrence Berkeley National Laboratory | Story by Thomas R. O'Donnell

Julianne Chung has an appreciation for versatility, starting with her academic interests. She showed an aptitude for mathematics while attending an all-girls high school in her hometown of Chattanooga, Tennessee, but Chung has a love for dancing that dates back even further. She followed both interests and graduated in 2004 from Atlanta's Emory University with a major in mathematics and a minor in dance and movement studies.

It's not just her interests that are varied, however. Chung also loves that math and dance are versatile pursuits. Dancers spend years learning techniques that influence their movements, says Chung, a third-year Department of Energy Computational Science Graduate Fellow. Then, "When you choreograph, you take those but you modify them and put in something new to make it your own."

It's the same when Chung develops computer algorithms for her doctoral research. "You have this structure of what all the previous researchers have done," she adds. "You take those different tools and things you've learned and try to put them all together in a way to get the best outcome."

Versatility also attracted Chung to applied mathematics and computational science. When she started college, "I thought there was no way on earth I was going to math graduate school."

Then Emory Mathematics and Computer Science Professor James Nagy suggested Chung do her honors thesis on image filtering algorithms. He showed her how the image reconstruction algorithms are applicable in everything from astronomy to microscopy. Chung followed Nagy's suggestion, then stayed on for graduate school at Emory so she could work with him.

Image reconstruction, especially from complex sources like medical PET and CAT scans, often requires the capability only high-performance computers can supply, Nagy says.

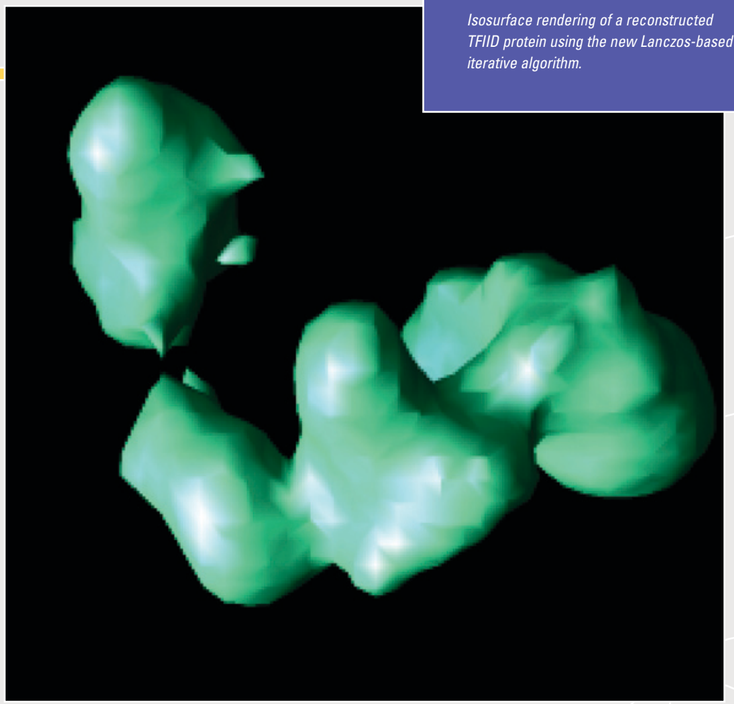

Isosurface rendering of a reconstructed TFIID protein using the new Lanczos-based iterative algorithm.

Each image can contain huge amounts of data and sometimes several must be compared and averaged to create clearer, more detailed pictures out of motion blur and background noise. "The mathematics behind that is quite complicated," Nagy adds. Nagy recognized that Chung's research had interesting parallels with a project involving Chao Yang, a staff scientist in Lawrence Berkeley National Laboratory's computational research division. Nagy introduced them at a 2006 Society of Industrial and Applied Mathematics (SIAM) imaging conference.

Versatility also attracted Chung to applied mathematics and computational science. When she started college, "I thought there was no way on earth I was going to math graduate school."

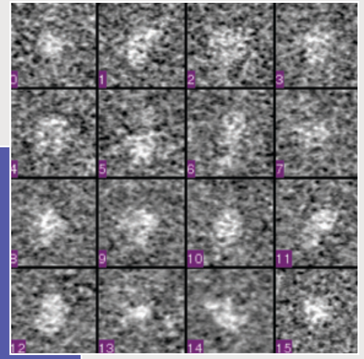

Sample projection images from real cryo-EM data. Usually researchers need tens of thousands of these images to reconstruct a good volume.

Yang and his group were working with biochemists Bob Glaeser of University of California – Berkeley, Pawel Penczek of the University of Texas – Houston Medical Center and others to develop algorithms to reconstruct images of complex particles like proteins and viruses. Biologists study these tiny materials because their structure affects their function, including how and whether they cause disease.

The images come from a process called single-particle cryo-electron microscopy. In cryo-EM, as it's called, a tiny purified sample of a biological substance is suspended in a water solution, then flash-frozen in liquid ethane kept near the temperature of liquid nitrogen. The sample is placed in the electron microscope, yielding thousands of two-dimensional images of single particles.

The outcome poses multiple problems. First, scientists must set the microscope's electron beam at a low level to avoid damaging the sample. As a result, "What you get … is very fuzzy and low-contrast. The signal-to-noise ratio is very low," Yang says.

Second, proteins can fold into many different conformations, making their structure look different from particle to particle.

The images also show molecules oriented in random directions, like jacks scattered across the floor. "You have to determine the orientation as well as the three-dimensional structure simultaneously," Yang says.

Lastly, thousands, if not millions, of images must be analyzed to find the particle's "signature" and decipher its structure from the many varied orientations. "This is a very computationally intensive job, mainly because the volume of data involved is huge," Yang says. Up to one million images must be analyzed to achieve atomic-level resolution, and each image can have as many as one million pixels of data.

That takes more than just powerful computers to solve in a reasonable time. Yang and his group work on algorithms that accelerate image reconstruction.

The technique they've developed first chooses a few good images, groups them according to their apparent orientations and averages them to mine the molecule's shape from amid the noise. The information is used to make preliminary three-dimensional "seed" reconstructions.

Each two-dimensional image then is compared with the seed reconstructions and grouped according to how much they differ from them. "Once you put each image into a different 'bucket,' each bucket is associated with a different orientation parameter," Yang says. The algorithm then merges images in each group.

A multi-reference reconstruction algorithm continually updates each three-dimensional seed structure based on two-dimensional images that correspond to its orientation. After the structures are updated, the algorithm recalculates how much each two-dimensional image differs from them and regroups the images. The process is repeated until the two-dimensional images no longer change groups or a maximum number of iterations is reached.

The group's algorithms significantly improved the speed and accuracy of cryo-EM image reconstruction. Chung's job, when she arrived at the Berkeley lab for her practicum in summer 2007, was to make it even better.

Building the seed structures for large proteins or viruses is the most computer-intensive step, Chung says. "This huge, 3-D volume has to reside on the memory of every single processor and the memory requirement is way too much," she adds. If each image is 1,000 pixels, a 3-D reconstruction is 1,000 cubed — one billion pixels.

The group's algorithm took advantage of parallel processing by parceling out the two-dimensional images over many processors. "There are so many images, it's natural to divide them into parts and assign each part to a processor," Yang adds. Imagine the processors spread out in a single column and the images divided among the processors.

Chung added another level of parallelization by organizing the processors into a two-dimensional grid with multiple columns. That allowed her to partition the three-dimensional volume among multiple processors along with the two-dimensional images, cutting the memory demand. "The whole volume didn't have to be on every processor," Chung adds.

She didn't stop there. Chung realized she could improve the cryo-EM implementation with research she did at Emory on regularization, which is designed to control errors in the image reconstruction algorithms.

Regularization is necessary because of the fundamental nature of image reconstruction, Chung says: "We solve the backward problem, or the inverse problem. Someone gives you this blurry image. You know a little bit about the image, and your goal is to work backward" to find the original picture.

Yang says with inverse problems "if you choose a very fast algorithm you may think it's converging rapidly" on a solution. In fact, "As the algorithm moves along, the noise tends to be amplified … if you don't do anything you just get garbage." To avoid that, the Berkeley group chose a slow algorithm and ran it for as many as 100 iterations.

Chung's solution was to combine the Lanczos bidiagonalization algorithm — a standard iterative method — with a filtered singular value decomposition (SVD) approach. In essence, Yang says, the Lanczos method projects the big problem into a subspace to make the problem smaller. The projected problem then is solved with appropriate regularization by the SVD approach, a technique that is highly effective for small problems.

Chung tested her implementation on Jacquard, the 712-processor cluster at Berkeley Lab's National Energy Research Scientific Computing Center (NERSC). She found that errors stabilize rather than "blow up" into garbage. Yang says the number of iterations needed to reconstruct a three-dimensional protein image dropped to just 30. Chung's improvements mean researchers can reconstruct high-resolution images of large structures like viruses, he says.

Yang has used Chung's code for other applications, including reconstructing X-ray tomography images of cell structures. He's also run it with an adenovirus data set on Franklin, NERSC's most powerful computer.

The code scales perfectly, Yang adds, so it takes full advantage of added computing power: "It's amazing to see it run on Franklin on 15,000 processors."

Chung presented the research at a SIAM conference in March 2008 and was expected to present it again later in the year. She and Yang are preparing a paper.

"Julianne made a tremendous contribution to our project," Yang says, not just improving algorithms but also implementing them. "That's usually rare to do in three months."

For her part, Chung loved her practicum. "To be able to … have a real-life application for something I was working on — that was a neat experience," she says.

Chung expects to graduate in spring 2009. After that, she hopes to obtain a postdoctoral fellowship and focus on research before returning to academia as a faculty member and investigator.

"She's an amazing student," Nagy says. "She has this great desire and curiosity. Sometimes it's unquenchable."

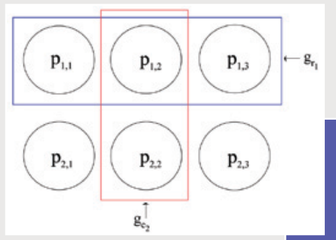

Processor layout for volume data distribution. Each processor is identified by its row and column group numbers.

PROGRAM REQUIREMENTS

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

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

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

Prior to the third year of the fellowship, fellows must complete a practicum assignment at a Department of Energy laboratory. Currently, approximately 25% 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 computational science in 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, materials science and engineering, and computational mechanics.

Cells and Uncertainty

JORDAN ATLAS

Cornell University | Los Alamos National Laboratory | Story by Alan S. Brown

Jordan Atlas wants to take a little of the doubt out of an uncertain world — the world of microscopic cells and the biology that drives them.

The Cornell University doctoral candidate and DOE Computational Science Graduate Fellow is researching ways to improve computer simulations of biological models. His main quest is to help build a minimal cell model — a simulation to find the fewest genes a cell needs to survive in a benign chemical environment. Identifying this core set promises powerful insights into the machinery of life — insights researchers could use to synthesize better medicines or to harness bacteria for industrial processes.

Atlas did similar research during his DOE CSGF practicum in summer 2007, as he worked with Los Alamos National Laboratory scientist Jim Faeder to refine cell-signaling models. Faeder focuses on how proteins in cells react with one another in order to respond to changes in the cells' environment.

Protein interactions are remarkably complex. Faeder likens each protein to a house with many doors. The right chemical key opens a door into a room with more doors. Behind each door are more rooms with more doors. Each door that opens changes the shape of the room, exposing some doors and hiding others.

Even the simplest cellular interactions involve dozens, or even hundreds, of proteins. With millions of possible reactions, complexity soon overwhelms conventional models that try to calculate every one of them. To simplify, Faeder looks only at the bonding sites — the doors — without considering the rest of the protein. He then uses supercomputers to develop rules that describe how and when these sites bond.

During his practicum, Atlas studied a conventional "heuristic" model — one that proceeds to a solution by trial and error or by following loosely defined rules — for the epidermal growth factor receptor system, which is involved in many cancers. He compared that with a rule-based model derived from the heuristic model. Faeder says the researchers believed the models would generate differences even though they portrayed the same process. Those differences could be resolved and understood through random variations of the model parameters as they were constrained by data like experimental results.

To see if the models behaved differently when constrained by information, Atlas fit them to both the data used to build the original model and to synthetic data. He used a new computational toolbox called SloppyCell to analyze the models.

It wasn't easy, Faeder wrote in an evaluation. "There turned out to be errors or ambiguities in both models that required detailed understanding of each model to find and correct," he wrote. Atlas excelled at finding and correcting the problems. He discovered and fixed several bugs in the published models and in the analysis program.

Atlas' results suggest both models were similarly constrained by data, but the rule-based model seemed to generate more error. "We just cracked the surface of that problem," Atlas says, but he started — and nearly finished — a second project to examine how the models are mathematically related. He's still working on it in his spare time.

The experience produced interesting and significant results, but just as importantly, Atlas saw how a major laboratory applies computational science on a big and important problem. He used many of the techniques he learned at Cornell, but in new and different ways.

"My philosophy is to try to take things from different instructors and fields. It's a little bit like running," he says, recalling an interest that obsessed him in high school. "You can run every day, but if you want to excel as an overall athlete, you have to sprint, jump and even throw. To do the best computer science, I have to learn new techniques. The best part of the practicum is learning something that I can apply somewhere else."

Just as he used his Cornell training at Los Alamos, Atlas has taken what he learned at the lab and is applying it to his doctoral research on the minimal cell model.

Sorting out which genes are the bare essentials for life isn't easy. All organisms carry useful but non-essential genes. The bacterium *Escherichia coli* (better known as *E. coli*), for example, has genes that let it infect other organisms. Without those genes, it couldn't thrive in the wild any more than a cheetah without legs. But are those genes essential for life? "Knock out those genes, provide a source of food, and E. coli will keep growing and dividing. It can survive without them," Atlas says.

"To do the best computer science, I have to learn new techniques. The best part of the practicum is learning something that I can apply somewhere else."

So how do we sort out the essential genes? Atlas likens the problem to figuring out what is essential for human life by examining what people pack in the luggage they carry onto a plane. "If you assume your checked-in bag will get lost, you will fill your carry-on with what you need to live," he explains.

"But what you need may differ from what I need. If you look into carry-ons, each one will have a toothbrush but not all of them will have a hair dryer. If you look at enough of them, though, you can see what people really need."

Atlas and other researchers take a similar approach when sifting through genes. "If two or more distantly related strains of bacteria have the same gene, then that must be something they need or nature would not have conserved it," Atlas says. On the other hand, he explains, different bacteria might use unlike genes to achieve the same results.

To simplify their task, researchers focus on one-celled organisms. Yet even a simple bacterium like *E. coli* has roughly 5,000 genes. Working with the smallest known free-living bacterium, *Mycoplasma genitalium*, makes the task more manageable. It has fewer than 600 genes. "Researchers who have done experimental studies on *Mycoplasma* think we can create a minimal cell with 200 to 300 genes or even less," Atlas says.

A few hundred genes may not sound like much, but each gene exists in a chemical soup that determines its behavior. Different chemicals in the environment "turn on" genes so they manufacture proteins and RNA, which regulate the cell as well as turn other genes on and off.

Modeling these interactions is a mountain of work and there are huge gaps in our knowledge. As a result, model builders sometimes make assumptions about reactions, just as skilled laborers eyeball the location of a door from a raw sketch.

Atlas is out to reduce those assumptions and make results less uncertain. He's developing novel methods to add new information on cell chemistry to the minimal cell model. This will make it easier to revise the model as researchers learn more about cellular chemistry. Atlas also uses high-performance computers to statistically determine how different estimates of reaction speed, frequency, or efficiency change model results.

The payoff is a model promising new insights into cell behavior. "It wouldn't replace experiments, but it could guide which experiments we chose to do," Atlas says. "We could make and test predictions to see if we truly understand how the cell works."

"If we understand the essential genes in a cell, we could identify targets for antibiotic drugs that could help cure bacterial infections," he adds. "Or we could speed up drug production by growing bacteria faster in a bioreactor."

Atlas has the perfect background for such work, thanks to his undergraduate degrees in chemical engineering and biochemistry from the University of Massachusetts in Amherst. He was drawn to the double major when he took classes with Susan Roberts, a recent doctoral graduate who was trying to increase the availability of a rare anti-cancer drug by growing it from cells. He joined Roberts' emerging lab in his sophomore year. "I was growing cells, doing experiments and running equipment," he says.

Although he spent five years at the university, Atlas took more than 20 credits in each of his semesters. He found it difficult to pick a graduate program because he found so many of them interesting. Eventually, he chose to work with Cornell's Michael Shuler on minimal cell models.

While the models he creates may be uncertain, Atlas' future is anything but.

Kholodenko, et al. (1999) JBC

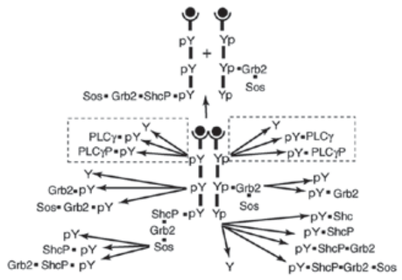

18 species, 37 reactions

Kholodenko, et al. (1999) JBC

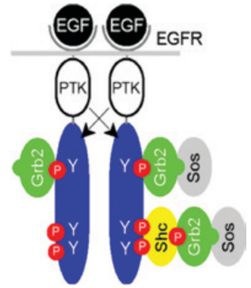

356 species, 3749 reactions

Even very simple models of protein interactions require tracking many different chemical species and reactions. Researchers use rule-based models to make it possible to model systems with millions of proteins. The models could help researchers develop drugs that target critical reactions in cancer and other disease cells.

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 over 60 students in 18 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.

A Matter of Scale

ASHLEE FORD

University of Illinois at Urbana-Champaign | Brookhaven National Laboratory | Story by Alan S. Brown

For Ashlee Ford, it's all about scale.

On one hand, Ford's doctoral research involves events that occur on vastly different scales of time. Her cutting-edge computational techniques show how interactions measured in picoseconds — one-millionth of one-millionth of a second — determine events weeks and even months later.

On the other hand, Ford has experienced similar differences of scale in her personal life. She grew up in Snyder, Oklahoma, a town of 1,500 where her graduating class of 60 was the high school's largest ever. It sits 100 miles southwest of Oklahoma City, surrounded by farmland in every direction.

Yet, in summer 2007 the chemical engineering doctoral student at the University of Illinois found herself in greater New York City. She was there to complete her DOE CSGF practicum at Brookhaven National Laboratory, located on Long Island.

Ford went there to work with James Davenport, whose research focuses on using Newton's laws of motion and quantum mechanics to explain how molecules bond with one another. "I was not familiar with it at all," Ford says.

Ford used the techniques to model how enzymes make sugars from cellulose, one of the substances that give plants their stiffness. The sugars can be used to make ethanol, making it possible to use woody materials like corn stalks and husks instead of grain for biofuel production.

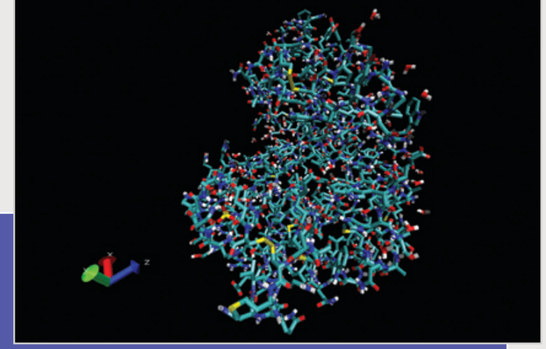

This natural enzyme breaks down cellulose plant fibers into sugars that can be converted into ethanol for fuel. Using cellulose, which is found in plants with little commercial value like switchgrass and in the stalks and stems of corn, would let us make fuel from renewable resources without cutting into food supplies. Models help illuminate the mechanisms involved in these reactions, and suggest how researchers can modify the enzyme and reaction conditions to make the process more efficient.

It's tricky, however. Cellulose doesn't yield its sugars readily, making the process energy-intensive and expensive. That's why numerous researchers are focusing on improving the enzymes' performance.

Sampling Time

At Brookhaven, Ford created molecular dynamics simulations of small glucose clusters representing cellulose. The models showed how the glucose interacted with molecules of benzene, which stood in for protein residues that play a role in enzyme-cellulose interaction. Other researchers have used molecular dynamics to study similar models, but for a span of only 150 picoseconds. Davenport and Ford wanted to increase the simulation's sampling time to see if the system's thermodynamic properties remained stable.

In the simulation, the researchers calculated the free energy potentials as a function of the separation between the glucose and benzene molecules. Free energy is a measure of the energy in a system that can be converted to do work. Davenport and Ford planned to compare plots for potential free energy over distance for two different simulation lengths to evaluate the effects of computational sampling time.

Ford contributed much to the project, including constructing and testing model structures, running the molecular dynamics simulations and analyzing the results. She even got to use New York Blue, Brookhaven's IBM Blue Gene/L supercomputer. "I was there for opening day, and it was really cool to be part of that," she says.

The Brookhaven researchers gave her plenty of time and latitude, Ford says. "They were great with encouragement and giving me individual attention if I needed help," she adds.

Davenport is just as enthusiastic about Ford's abilities. "I was impressed with her hard work and quick learning ability," he wrote in an evaluation. "I would be happy to have her back."

The practicum experience "helped make me more multidisciplinary, and that's the point," Ford says. "It also showed me that I can apply many of the things I'm working on in graduate school to any problem in engineering and modeling."

School

Ford discovered mathematical modeling as an undergraduate chemical engineering major at the University of Oklahoma. It interested her for the most practical of reasons: "When you model reactions, you're not wasting chemicals and you're not wasting time by running experiments over and over again," she says.

Her undergraduate research project involved fluid dynamics to model the flow of water over carbon nanotubes. "Water rolls off carbon nanotubes, and my professor had a grant to look at whether we could use them as a coating to make submarines quieter and faster," she says.

Ford thought she might work in industry after graduation, but her advisor warned that she might not return to school for an advanced degree. Although Ford applied to doctoral programs, she spent the summer after her senior year working as a process engineer at an oil refinery. "I wanted an industry internship so I could convince myself that grad school was a good option," she adds.

Ford chose to study at the University of Illinois because of its reputation in fluid dynamics. Eventually she joined professor Richard Braatz's research group, which was using multiscale models to study the bioproduction of butanol, an environmentally safer ethanol substitute.

Multiscale models tie together events that occur on vastly different time and size scales. Braatz thought he could apply the technique to drug delivery systems that slowly release medicine into the body. Ford jumped on the challenge.

The benefits of slow-release medicines are obvious. Instead of worrying about which pills to take when or about taking too little or too much medicine, patients would receive an injection that releases precise amounts of medication for weeks or even months. It would make it easier for everyone to comply with doctors' orders and simplify treatment of people with memory disorders like Alzheimer's disease.

Formulating precise slow-release systems involves delicate chemistry. "The drugs are encapsulated in polymer microspheres that degrade in the body's fluids," Ford says. "Once I have a model that explains how drug release occurs, I can say, 'I want to release this amount of drug in a patient for four weeks, so what size microspheres do I need to encapsulate them?'"

Creating that model is an imposing task. The body starts breaking down the polymer spheres as soon as they're injected. These reactions occur at picosecond speeds, rapidly transforming solid spheres into a mesh of polymer strands. The larger drug molecules then work their way through the gaps in the dissolving sphere over weeks or months.

The model must be precise to be useful. That means Ford must find a way to bridge the gap between picosecond reactions and prolonged drug release — a major computational challenge. It would take years of supercomputer time to calculate how each picosecond-long reaction alters the flow of medicine into the body. Ford relies on computational shortcuts to simplify the problem, but each one builds errors into the model. Too many errors and the model will not be precise enough to predict medication levels.

It's not a problem Ford will solve in weeks, months, or even years. In fact, Braatz calls the development of accurate multiscale models one of the grand challenges of computational science.

Ford embraces such challenges and with the research she has chosen, she will have plenty of them to keep her busy.

The practicum experience "helped make me more multidisciplinary, and that's the point," Ford says.

DOE CSGF HIGHLIGHTS

- Payment of tuition and required fees
- Yearly stipend of $32,400
- A $1,000 yearly academic allowance
- Matching funds of up to $2,475 for a computer workstation purchase
- Opportunity to complete a practicum working with researchers at a DOE Laboratory
- Yearly fellows' conference with opportunities to meet other fellows and academic and government professionals
- Renewable up to four years

For more information: www.krellinst.org/csgf

Mouse Model Traps Wave Interaction

DAVID KETCHESON

University of Washington | Sandia National Laboratories – New Mexico | Story by Thomas R. O'Donnell

During internships at the Department of Energy's Sandia National Laboratory and a practicum at its New Mexico neighbor, Los Alamos National Laboratory, David Ketcheson has worked on projects ranging from spectroscopy to verification and analysis of physics codes.

In his latest stint at Sandia, however, Ketcheson worked on a Mickey Mouse problem.

The University of Washington doctoral candidate spent summer 2007 on his Department of Energy Computational Science Graduate Fellowship (DOE CSGF) practicum, developing an algorithm to model electromagnetic (EM) wave dispersion. The test simulation was a waveguide — a tube that guides EM energy — with a cross-section shaped like the Disney mascot's head.

It's not because Ketcheson's practicum supervisor, Sandia Truman post-doctoral fellow Hung "Jacques" Loui, loves "House of Mouse." Any shape would work, and the mouse's eyes and other features could be "made" of different dielectric materials to test how algorithms portray waves' interactions with them.

Loui uses computational tools to research "frequency selective surfaces": materials and shapes that reflect or absorb specific frequencies of electromagnetic waves, like radio. They're like your microwave oven's glass door, Loui says: "Light passes through and you can see it cooking, but it confines the microwaves so it doesn't destroy your eyes."

One such surface Loui studies is a metal plate with periodic perforations, like a cheese grater. In most previous modeling research, the holes, or apertures, were filled with homogeneous material. "What we're trying to do is fill the holes with a complex material" that is selectively permeable to radio waves, Loui says.

"We hope to control transmission and reflection (of waves) or to tailor it to what we want," he adds. For instance, frequency selective surfaces could enclose radar antennae to mitigate interference when several are in close proximity, like in a naval convoy.

To get at what materials and shapes would work best for different purposes, Loui uses computers to model the typical wave modes they create.

Ketcheson sought algorithms to model EM modes of a Mickey Mouse-shaped tube filled with various materials. "That's much more difficult because it's not hollow," Ketcheson says. "You have dielectric materials with different properties."

Ketcheson used a vector-based finite element method approach. It's one way to replace the partial differential equations (PDEs) that describe continuous processes like wave propagation with algebraic equations to approximate a solution — a process called discretization. The algebraic equations calculate values like pressure or radiation at data points distributed in a grid throughout the space or material being modeled.

Through much of his career Ketcheson has used a different discretization technique, called the finite volume method.

"I think what made it work really well between me and David was he was very interested in learning a new language," Loui says. "Not only did he learn a new language — finite element method — very quickly, he was able to implement a solver that was quite unique."

Ketcheson adds: "Part of what I did was just implementing algorithms that existed in the literature — finding the best known algorithms. What we ended up developing is more capable in different ways than everything that was available in the literature."

Using the finite element method meant the problem could be discretized on an unstructured grid, giving researchers the ability to calculate wave dispersion in different geometries.

Using the MATLAB computing language, Ketcheson implemented algorithms designed to model heterogeneous materials like those that made up the Mickey model's eyes, nose and mouth. Although it's designed for just two-dimensional applications, it could be coupled with other codes to model properties in a third direction.

The code also is arbitrarily high-order accurate, meaning users can choose the order of accuracy they want. For example, doubling the number of data points in a second-order accurate code generally results in an error one fourth as big. Doubling the data points with a fifth-order accurate algorithm generally results in error that's just one thirty-second as big. However, other factors often must be accounted for to accurately depict the physics involved, Loui says.

Nonetheless, "That's some of the innovation that was inherent" in Ketcheson's practicum, Loui adds. "With most commercial code you're stuck in second or third order."

The two researchers hope to finish a paper on the algorithms and Loui included some of the work in a Sandia colloquium presentation.

For Ketcheson, applying the finite element method to wave mode problems is something of a reverse

"What we ended up developing is more capable in different ways than everything that was available in the literature."

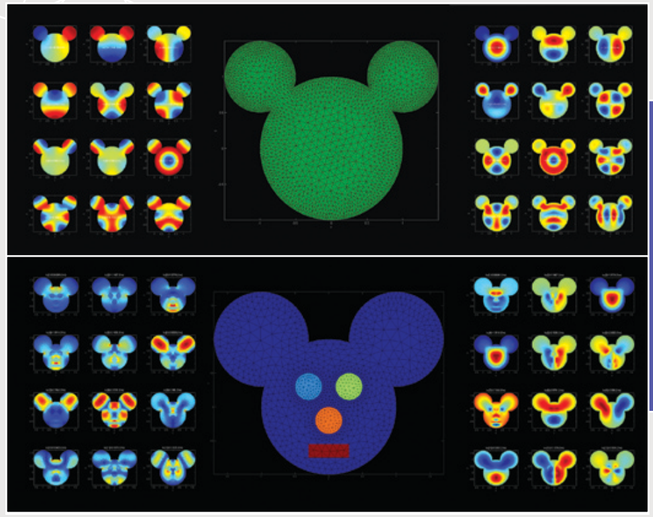

Top: This visualization shows a homogeneously filled waveguide of "arbitrary" cross-section and the data mesh used to simulate electromagnetic wave dispersion. The colors show normalized electric and magnetic field intensities of the various modes.

Bottom: In this simulation, the waveguide has been filled with inhomogeneous materials to calculate their effect on wave dispersion. The colors show normalized electric and magnetic field intensities of the various modes.

of an earlier Sandia experience. While earning his undergraduate degree from Brigham Young University, he helped develop and implement a finite volume fluid dynamics solver in ALEGRA, a finite element-based multiphysics code at Sandia.

The experience stoked Ketcheson's interest in computational science and prompted him to finish a second major in mathematics at BYU, in addition to astronomy and physics. He also learned about the finite volume research of Randall LeVeque at the University of Washington. LeVeque now is Ketcheson's academic advisor.

LeVeque created CLAWPACK (Conservation Law Package), a software package for solving wave propagation problems described by hyperbolic PDEs. Hyperbolic PDEs can depict a multitude of kinds of waves.

Waves are incredibly complex, Ketcheson says. They can show dispersion, with waves of different frequencies moving at different speeds. They also can be nonlinear, with high-amplitude waves moving faster than low-amplitude waves.

"The main difficulty in simulating the kind of waves I do — the nonlinear waves — is that the nonlinearity makes waves steepen" until they break like ocean surf, Ketcheson says. In other cases the steepening becomes a shock wave.

"To simulate these waves you need numerical methods capable of handling discontinuous functions — the points where their values suddenly jump to another value," Ketcheson says. Simple methods for solving equations depicting such waves create oscillations around the shock — "Little vibrations that can lead to all kinds of bad things," such as negative pressures in fluid dynamics problems and fictitious magnetic forces in MHD applications.

"Once that happens you can't have much confidence in the answer," Ketcheson says. "All the work I've done centers around capturing these discontinuities while avoiding the oscillations."

To cope with the tendency toward oscillations, Ketcheson developed WENOCLAW, a CLAWPACK extension named for the weighted essentially non-oscillatory (WENO) method for spatially discretizing hyperbolic PDEs with discontinuous solutions. Its main difference from CLAWPACK is that WENOCLAW, like the code Ketcheson developed with Loui, is arbitrarily high-order accurate.

"For certain kinds of problems you can benefit from high-order methods," LeVeque says. "Incorporating that whole class of methods into CLAWPACK allows people to try out different methods on their problems more easily."

Ketcheson is involved in so many projects LeVeque has difficulty tracking them. For instance, Ketcheson also studies strong stability-preserving methods (SSP), which are applied to ordinary differential equations after PDEs are discretized in space. His work on the methods as an undergraduate got the attention of Sigal Gottlieb of the University of Massachusetts, Dartmouth, and she brought Ketcheson east for a visit. Gottlieb and Ketcheson have since collaborated on papers and presentations with Chi-Wang Shu of Brown University and Colin Macdonald of Canada's Simon Fraser University. Ketcheson's independent SSP research earned him a 2008 Society of Industrial and Applied Mathematics award for one of the most outstanding student papers of the year.

Ketcheson "has two small kids, too," LeVeque marveled. "I don't know how he does it."

Scaling the Nanowire

By Karyn Hede

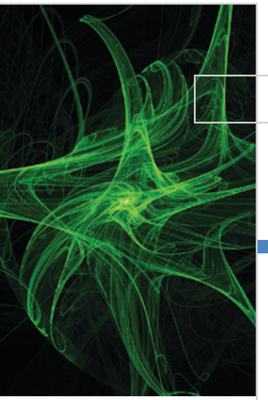

THE PREEMINENT PHYSICIST-FUTURIST Richard Feynman famously declared in a 1959 address to the American Physical Society that "there's plenty of room at the bottom." He then invited them to enter the strange new world of nanoscale materials, none of which had actually been invented, except in Feynman's fantastical imagination.

It took another generation of scientists before nanotechnology emerged, but Feynman's assertion still rings true. There's plenty of room at the nanoscale and scientists at Lawrence Berkeley National Laboratory (LBNL) in California are at the forefront in constructing new materials there.

Paul Alivisatos, director of LBNL's Materials Science Division, is a world leader in nanostructures and inventor of many technologies using quantum dots — special kinds of semiconductor nanocrystals. Quantum dots, which are one ten-millionth of an inch in diameter, fluoresce brightly, are exceedingly stable and don't interfere with biological processes because they are made of inert minerals. Alivisatos and his colleagues have constructed dozens of variations in which the fluorescent color changes with the dot's size. Today life-science researchers use quantum dots as markers, allowing them to visualize with extreme accuracy individual genes, proteins and other small molecules inside living cells and fulfilling a prediction Feynman made in his famous lecture.

LBNL physicist Lin-Wang Wang likes to say that some day we will view the 21st Century as the "nanostructure" age, much as we associate Neolithic humans with the Stone Age and their descendents with the Bronze and Iron ages.

Despite their obvious usefulness, however, the behavior of materials built from nanometer particles is still not completely understood or fully predictable. Part of the problem is that electron wavelengths also are on the nanometer scale. Electrons' quantum mechanical properties — the consequence of their wave-like behaviors — are changed by the sizes and geometries of the quantum dots, a subject that still generates heated discussion among physicists. These size and geometric changes allow electrons in semiconductor nanostructures to generate new energetic and optical properties.

Wang and his LBNL colleague Andrew Canning, a computational physicist who helped pioneer the application of parallel computing to material science, want to use computational methods to understand the emergent behaviors of novel materials, such as quantum dots, built from these exceedingly small blocks. "There are a lot of challenges and there are still many mysteries to be solved," Wang says. "For example, we still don't quite understand the dynamics of the electron inside a quantum dot or a quantum rod. There is a lot of surface area in a quantum structure, much more than the same material in bulk. So how the surface is coupled with the interior states and how this affects the nanostructure properties is not well understood."

The Rules are Different in Nanoscale

The research team is not starting from scratch, of course. There are established equations that predict the behavior of the electron wave function in these materials. The devil lies in the size of the problem.

"In terms of computation the nanostructure is challenging. For example, if you have a bulk material the crystal structure is a very small unit cell, just a few atoms, that repeats itself many, many times," Wang says. "So computationally, you can treat bulk structures by calculating one unit cell — you only deal with a few atoms. With only a few atoms, you can represent the whole, much larger structure of the material. However, for a quantum dot or a quantum wire you have to treat the whole system together. These systems usually contain a few thousand to tens of thousands of atoms, and that makes the computation challenging."

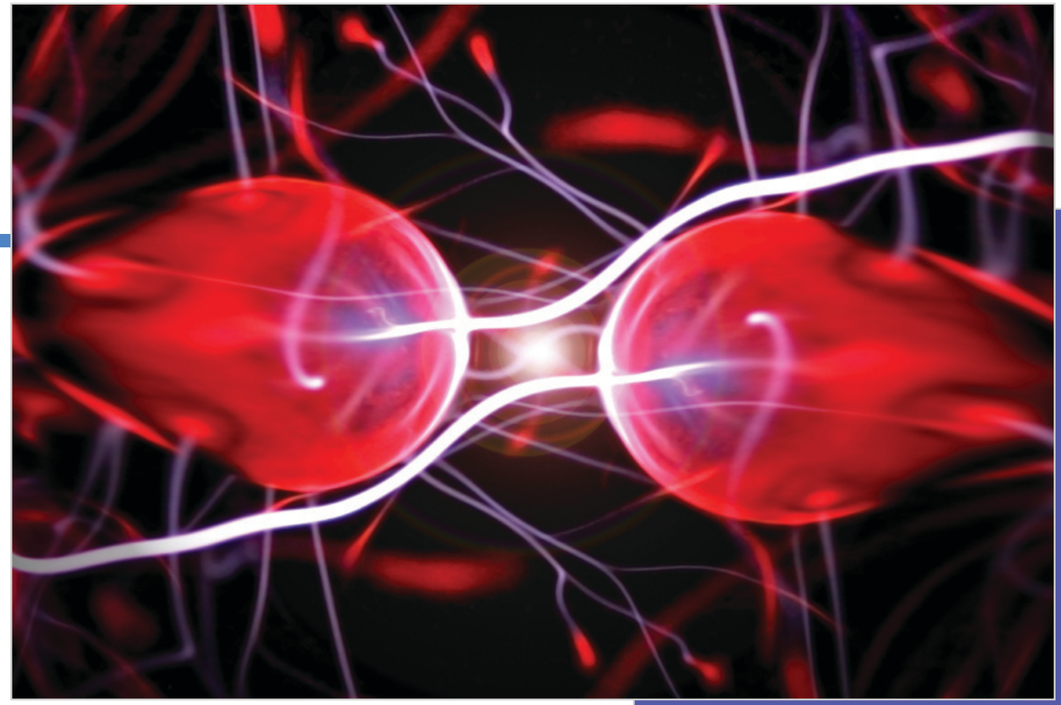

To solve a problem containing thousands of atoms requires new algorithms that handle the physics differently without compromising accuracy and parallel computing on a massive scale. That's where Canning's expertise came in.

"We know we need to solve the Schrödinger equation for these problems, but to do so fully is exceedingly computationally expensive," Canning says. "What we did was make advances to approximate, solve the problem, and still get the physics right."

Canning collaborated with Steven Louie's group at the University of California – Berkeley, to improve the Parallel Total Energy Code (Paratec), an ab initio, quantum-mechanical, total energy program. The program runs on Franklin, the Cray XT4 at LBNL's National Energy

LBNL physicist Lin-Wang Wang likes to say that some day we will view the 21st Century as the "nanostructure" age, much as we associate Neolithic humans with the Stone Age and their descendents with the Bronze and Iron ages.

COLLABORATORS

Andrew Canning is a staff scientist in the Computational Research Division at Lawrence Berkeley National Laboratory (LBNL). He was a pioneer in introducing parallel computing into the materials science community and developed many new algorithms for parallelizing codes, allowing new physics to be done on larger and more complex systems such as large nanostructures. Canning helped create many codes used by the materials science community, including PARATEC and ESCAN. His projects won the Gordon Bell prize in 1998 and received an honorable mention for the prize in 2001 and won the Smithsonian Computerworld Award in 2000. He has published more than 60 papers in physics and computational journals. Canning received his doctoral degree in computational physics from Edinburgh University in 1988 and joined LBNL in 1997 to work in the Scientific Computing group at the National Energy Research Scientific Computing Center.

Lin-Wang Wang has been a staff scientist at Lawrence Berkeley National Laboratory since 1999. His research interests include large-scale electronic structure calculations of nanosystems using supercomputers. He's devised or collaborated on many widely used methods and algorithms, including the kinetic energy functional, the folded spectrum method, linear combination of bulk bands method, the generalized moments method and many others. His simulations have pushed atomistic electronic structure nanosystem calculations from a few hundred to a million atoms. Wang has more than 140 publications and has been honored as an overseas outstanding young researcher by the National Natural Science Foundation of China and with the overseas scientist team of excellence award from the Chinese Academy of Science. He earned his doctoral and master's degrees in physics from Cornell University and his bachelors in physics from Shanghai Jiao Tong University. Wang is a fellow of the American Physical Society.

Further Reading:

A. Canning and D. Raczkowski. *Scaling first-principles plane-wave codes to thousands of processors.* Computer Physics Communications, Vol. 169, p. 449 (2005).

A. Canning, L. W. Wang, A. Williamson and A. Zunger. *Parallel Empirical Pseudopotential Electronic Structure Calculations for Million Atom Systems.* Journal of Computational Physics, 160, p. 29 (2000).

L.W. Wang, Z. Zhao and J. Meza. *A linear scaling three dimensional fragment method for large scale electronic structure calculations.* Phys. Rev. B 77, 165113 (2008).

Contact:

Andrew Canning
acanning@lbl.gov

Lin-Wang Wang
lwwang@lbl.gov

PRACTICUM COORDINATOR

Daniel Martin . . . . . DFMartin@lbl.gov

Research Scientific Computing Center (NERSC). The massively parallel system has 9,660 compute nodes, but is due to receive an upgrade, increasing its processing capability to a theoretical peak of about 360 teraflops.

"Paratec enables us to calculate thousand-atom nanosystems," Canning says. "The calculation is fast and scales to the cube of the system, rather than exponentially, as a true solution of the many-body Schrödinger equation."

Besides massive parallelization of the codes, the researchers also developed many new algorithms for nanostructure calculations. For example, Wang devised a linear scaling method, called the folded spectrum method, for use on large-scale electronic structure calculations. The conventional methods in Paratec must calculate thousands of electron wave functions, but the Escan code uses the folded spectrum method to calculate only a few states near the nanostructure energy band gap. That means the computation scales linearly to the size of the problem — a critical requirement for efficient nanoscience computation. Wang and Canning recently worked with Osni Marques at LBNL and Jack Dongarra's group at the University of Tennessee, Knoxville, to reinvestigate and significantly improve the Escan code by adding more advanced algorithms.

Wang and his colleagues also have recently invented a linear scaling three-dimensional fragment (LS3DF) method, which can be hundreds of times faster than a conventional method in calculating the total energy of a given nanostructure. The code has run at 107 teraflops on 137,072 processors of Intrepid, Argonne National Laboratory's IBM Blue Gene/P. The researchers have in essence designed a new algorithm to solve an existing physical problem with petascale computation. Wang says the LS3DF program is designed for materials science applications such as studying material defects, metal alloys and large organic molecules.

Within a nanostructure, the physicists are interested mainly in the location and energy level of electrons in the system because that determines the properties of a nanomaterial. For example, Wang says, electrons within a quantum rod or dot can occupy a series of quantum energy states or levels as they orbit the atomic nucleus and interact with each other. The color emitted by the material typically depends on these energy states.

Specifically, the scientists focus on two quantum energy levels: the highest occupied molecular orbital (HOMO) and the lowest unoccupied molecular orbital (LUMO), which the Escan code can calculate. The energy difference between these two levels determines the material's color.

The color also changes with the quantum dot's size, providing one way to engineer its properties. In principle, knowing the electronic properties of a given material lets the researchers predict how a new nanostructure will behave before actually spending the time and money to make it. It's a potentially less expensive way to experiment with new nanomaterials, Wang says.

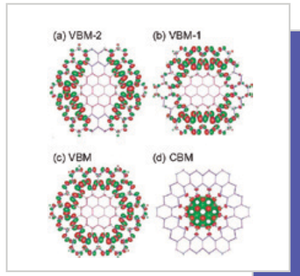

These figures show calculations of the lowest electron state (CBM or conduction band minimum) and highest valence band states (valence band maximum or VBM) for a ZnS/ZnO nanowire. The wavefunctions themselves are shown, with red representing positive and green representing negative.

Solar Cells on the Cheap

That kind of predictive power led to a discovery that may contribute to energy independence for the United States. In collaboration with scientist Yong Zhang, a senior scientist at the National Renewable Energy Laboratory in Golden, Colorado, Wang and his colleagues predicted a new kind of solar cell that could be manufactured inexpensively and from environmentally friendly materials.

The prediction is based on a new kind of nanostructure architecture that takes advantage of the unique electronic properties of materials constructed from nanometer-sized units. In order for a material to generate electricity efficiently from sunlight it must have an electron excitation potential around 2 electron volts (eV). Silicon is one such material, and most commercial solar cells use silicon that is cut into thin sheets from bulk material. These solar cells are expensive and it takes years to recoup the cost of manufacturing and installation.

Wang and Zhang collaborated with Joshua Schrier, Denis Demchenko and Paul Alivisatos to propose using zinc oxide (the white stuff often found in sunscreens) and zinc sulfide (an abundant, easily produced mineral) in a novel solar cell. The two materials by themselves could never generate solar energy efficiently. The key is how they are combined.

The researchers designed an architecture with a zinc sulfide nanostructure core surrounded by a thin shell of zinc oxide to form a nanowire. Using a new code developed by Wang and a team of LBNL computational scientists, the researchers simulated the electronic wave properties of the proposed solar cell.

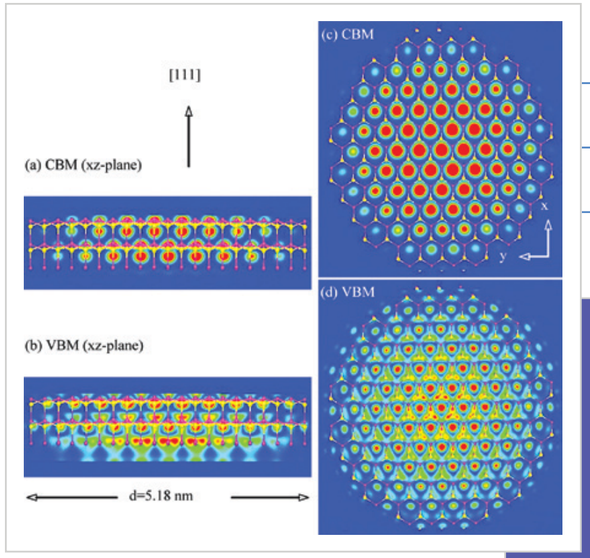

These figures show calculations for electron states for an InP quantum wire, including two different cross-section contour levels of the lowest electron state (CBM or conduction band minimum) and highest valence band state (valence band maximum or VBM). The squares of the wavefunctions are shown, with red representing a higher value while green represents a lower value.

"We wanted to reduce the band gap," Wang says. "By staggering the band energy alignment of the materials we calculated the overall band gap would be 2 eV, which produces a high efficiency limit of 23 percent."

The researchers published their finding in the journal *Nanoletters* and got immediate response. The publication was one of the top 10 most-viewed *Nanoletters* articles in 2007. Wang says several groups worldwide now are working on building versions of the nanowire and testing their efficiency. If they function as predicted, the group will have invented a safe, abundant, stable and environmentally benign solar cell that can be built immediately using existing manufacturing methods.

Looking forward, the scientists are using these newly developed computational methods to help guide other new nanostructure materials. Even now, when the experimental group creates a new nanostructure, the properties are not always easy to predict.

"When you do the experiment, you don't really know what result you are getting," Wang says. "How to explain experiments is sometimes challenging. That's the first task of simulation, to explain experimental results. Then the second task is to get some guidance for experimental design — to figure out what kind of experiment to do. That's theory guiding the experiment."

Both kinds of computation keep the group busy. Experimentalists at LBNL and collaborators at Washington University, NREL and worldwide are busy creating new nanomaterials: rod-shaped semiconductor nanocrystals that could be stacked to create tiny electronic devices; nanowires of various semiconductor materials; and quantum dots that can track tumors and help physicians diagnose and treat cancers more specifically.

There appear to be few limits to the number of nanostructures that can be created, proving once again that Feynman was right — there's lots of room in the nanosphere.

Invoking a Pharaoh's Name to

Model Nuclear Weapon Electronics

By Victor D. Chase

THE UNTHINKABLE HAS BECOME A REALITY. A nuclear strike has hit the U.S., and a retaliatory long-range ballistic missile loaded with a nuclear warhead streaks toward its target. As it begins reentry an enemy defensive missile explodes nearby creating extremely high temperatures and radiation aimed at disabling the incoming weapon.

Should this doomsday scenario ever occur, the U.S. must ensure that its nuclear weapons arsenal will be able to withstand such defensive missile blasts and successfully strike their targets.

But how to do so?

Until U.S. ratification of the Nuclear Non-Proliferation treaty in 1970 nuclear tests could be conducted to confirm the viability of nuclear weapons, but the treaty put a halt to such tests.

Prior to 2005, the job of creating a hostile neutron environment similar to what a nuclear weapon system might encounter was given to DOE's Sandia Pulsed Reactor (SPR), housed at Sandia National Laboratory, Albuquerque, New Mexico. When researchers wanted to determine if the electrical components of a nuclear weapon system could withstand being bombarded by a defensive missile explosion they placed the components in the SPR, hit them with short-pulsed neutrons and gamma rays and then examined their condition.

But then came 9/11 and everything changed.

There was concern that if the nuclear material used in the reactor fell into terrorist hands it could be used to produce "dirty bombs," or other kinds of nuclear weapons. To eliminate even the remotest possibility of that occurring, the pulse reactor was deactivated. That left those within DOE who are responsible for maintaining the viability of the nation's nuclear weapons stockpile without an important testing tool.

Refurbish and Requalify

Compounding the situation is a freeze on the building of new nuclear weapons that is also part of the Nuclear Non-Proliferation treaty. As a result, the U.S. must rely on its existing nuclear weapons stockpile, which is not immune to the effects of aging, and must, therefore, be periodically refurbished. And, "If you refurbish the old systems you have to requalify them," said Robert Hoekstra, who heads the 13-person Electrical and Microsystems Modeling group at Sandia that has taken on the task of ensuring that the new electrical components used in the refurbishing process can survive extreme environments.

Their goal is to use DOE's high-powered computing capabilities to model the effects of hostile environments on the electrical devices used in nuclear weapons as a substitute for testing in the SPR. In anticipation of the possible eventual need for such computational modeling, Hoekstra's department began creating the necessary computer software some two years prior to 9/11. As part of this effort, Hoekstra's team has been working in conjunction with other groups within Sandia that are providing data necessary to do the modeling, including codes representing material properties and radiation sources. Due to the complexity of the task, and the crucial need for proven accuracy, the development efforts are ongoing, with recent test results that have been promising.

Their goal is to use DOE's high-powered computing capabilities to model the effects of hostile environments on the electrical devices used in nuclear weapons as a substitute for testing in the SPR.

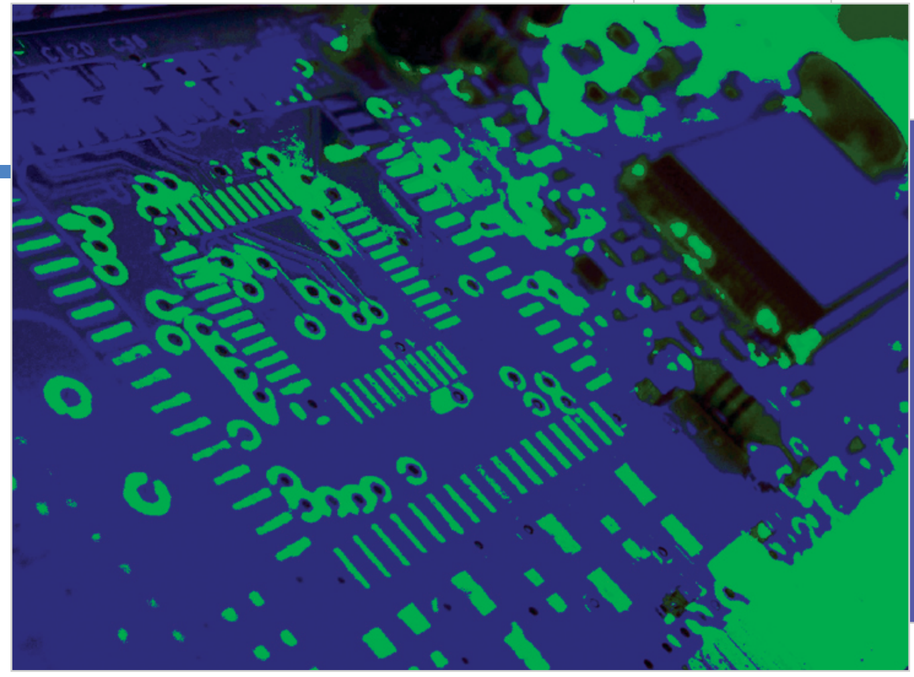

"We had the first prototype demonstration earlier this year and that was very successful," said Hoekstra. "It was a big milestone toward proving that we can use this modeling for new systems qualification." But the actual use of this software to qualify new nuclear weapon components is still several years away.

The project designed to accomplish this monumental task is called Qualification Alternative to the Sandia Pulsed Reactor — the acronym being QASPR (pronounced Casper). QASPR consists of several computer codes, with two of the primary codes being developed by Hoekstra's group. They are known as Xyce (pronounced Zeiss), and Charon, the more powerful of the two.

The River Styx

Because of its ability to solve a broad range of transport problems, including semiconductor physics, Hoekstra's team named the program Charon, after the Greek mythological boatman who transported the dead across the river Styx.

The name Xyce is a variant of SPICE (Simulation Program with Integrated Circuit Emphasis), a widely used circuit simulation program first developed at the University of California – Berkeley, around 1970. Today there are numerous

ROBERT HOEKSTRA

Robert Hoekstra joined Sandia National Laboratories in October 1998 after completing his doctoral work at the University of Illinois at Urbana-Champaign. His dissertation involved the simulation of low-temperature plasmas for semiconductor processing, and specifically the development and application of plasma surface interactions in ion enhanced etching tools. After coming to Sandia, Hoekstra initially was involved in Direct Simulation Monte Carlo (DSMC) techniques for modeling neutron generator behavior. Following this he became one of the lead developers of the Xyce parallel circuit simulator, for which development began in 1999. In 2005, Hoekstra became the project lead for the Charon project, and has continued in that role ever since with a focus on support for the QASPR program. Hoekstra has continued to contribute to the Xyce project as well the distributed linear algebra and graph algorithms in Trilinos. Hoekstra's research interests have included algorithmic issues of massively parallel circuit and device simulation with a focus on load balancing and linear solvers as well as semiconductor device physics. Earlier this year, Hoekstra was promoted to manage the Electrical and Microsystems Modeling department at Sandia.

Contact
Robert Hoekstra
rjhoeks@sandia.gov

PRACTICUM COORDINATORS

Heath Hanshaw (New Mexico) ..... hlhansh@sandia.gov
Alex Lindblad (California) ..... alindbl@sandia.gov

variants of SPICE that are used commercially in the electronics industry, but, says Hoekstra, none is as powerful as Sandia's Xyce.

The prototype of Xyce was written some nine years ago by Eric Keiter at Sandia shortly before Hoekstra joined the department. He and Scott Hutchison then worked with Keiter to create the first production version. The team is now working with Xyce version 4.1.

A high-fidelity code in its own right, Xyce can take a computerized snapshot of the electronic forest by modeling the logic and timing of multiple circuits, while Charon can bore down into the individual trees by examining the workings of a single device to the point that it can model the movement of electrons inside a semiconductor material.

"For a single device there may be as many as millions of finite elements to model, so to explore the physics of the materials, to learn how the device changes if it is hit with radiation or some other effect, requires extremely high-fidelity, which is what Charon provides," explains Hoekstra.

With it "we develop a detailed physics understanding of what the radiation effects are," he said. "We need to go to a lower-fidelity tool when we want to simulate the full electrical system, because we can't simulate every device in the electrical circuit at the Charon fidelity: it wouldn't fit on the largest computers. So we take what we learned with the higher-fidelity Charon and use that to inform the lower-fidelity Xyce, which simulates the full system."

Going Commercial

While addressing DOE's nuclear weapons verification needs, Hoekstra's group is also working with commercial vendors to make Xyce technology available to them. "It will give them the ability to do very large-scale simulation of circuits which is needed as integrated circuits get bigger and bigger," he says. The companies involved are major electrical design and automation tool companies that develop software and manufacturing processes for integrated circuits. They in turn supply semiconductor manufacturers.

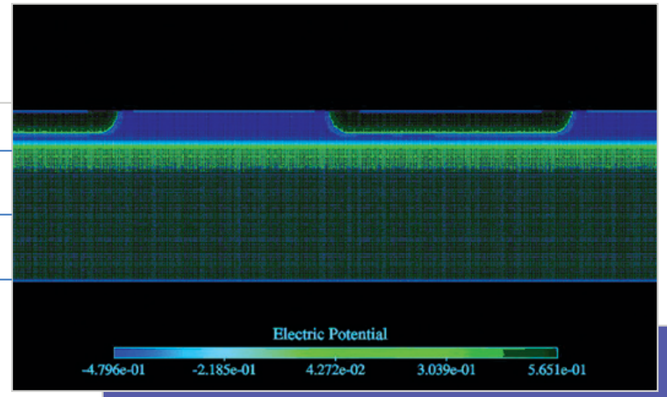

Charon model of bipolar junction transistor (~2 million elements in mesh).

"There are two features of our tools that differentiate them from the commercial tools," said Hoekstra. "One is that we can go to greater fidelity physics in the models. What goes in hand with that is our ability to scale up to the really large scale computing platforms of ASC (Advanced Simulation and Computing). Our code is the only version that is capable of massively parallel execution, which means we can run dramatically larger problems," he said.

What makes that possible is the supercomputing power made available through ASC, a DOE program created in 1995, aimed at developing supercomputer capability to simulate the performance, safety and reliability of nuclear weapons and to certify their functionality. It involves the collaboration of three DOE laboratories: Sandia, Los Alamos, and Lawrence Livermore, in conjunction with numerous university researchers.

For the higher-fidelity Charon code Hoekstra's team has run large-scale calculations using Livermore's Purple, an ASC supercomputer made up of 12,544 processors. Utilizing over 8,000 of these processors, it takes some two weeks for Charon to model 250 million variables. So far the group has used Purple for approximately 12 weeks during the past year. Because Xyce, powerful as it is, does not require such massive computing capacity, it is being run on a few hundred processors at Sandia on high-capacity computing Linux clusters, which themselves can consist of thousands of processors.

Until recently the suite of codes involved in the complex QASPR process were referred to individually by names such as NuGET, Cascade, and GRASP. Xyce and Charon are at the end of this chain of codes, and actually model the responses of electronic devices to hostile environments. To more tightly integrate the codes for effectiveness and efficiency they have now been coupled under a single umbrella term standing for Radiation Analysis Modeling and Simulation for Electrical Systems, or RAMSES.

Pushing the Limits

When QASPR will be ready for "real world" electronic component qualification is in large measure dictated by the needs of the Electrical and Microsystems Modeling group's customers — namely the military, who specify the requirements for refurbished nuclear weapons systems, and the DOE scientists and engineers who design the new components. Hoekstra expects that the software to model the devices' durability will be ready when the components themselves are ready for testing.

In the meantime, much needs to be done. "We have to improve our understanding of the physics and our ability to model that physics," said Hoekstra. For example, he noted, "The radiation effects are only partially understood, so we are trying to increase our knowledge of those effects and get that knowledge into the models. We are doing tightly coupled experiments and modeling to understand the physics."

Beyond that there is the issue of scale. "We are really pushing the limits of computing to model these very high-fidelity models. We need to improve our models and increase our computing horsepower," he said.

Another challenge is to convince those responsible for nuclear weapons qualification that they can rely on the QASPR methodology to accurately depict the response of electronic components to hostile environments. Doing so requires a two-step process known as verification and validation. Verification is more mathematical than empirical in that it verifies that the computer model gives the answer it is supposed to give in mathematical terms.

Validation answers the question, "Do the results represent reality?" To validate RAMSES results, physical tests are run on electronic components in radiation facilities other than the now defunct Sandia Pulsed Reactor. Though the conditions they create are not as close to an actual hostile environment as was achievable with the SPR, "We still have facilities that give us gamma and neutron irradiation. We put the devices and circuits into those test facilities and use the results to validate our models," said Hoekstra.

Two DOE Computational Science Graduate Fellows (CSGF) have been involved in the development of the suite of codes used by the QASPR group. Judith Hill, who did her DOE CSGF practicum at Sandia in 2000, worked on Charon when she returned to Sandia as a postdoctoral fellow. Similarly, David Ropp, who did his DOE CSGF practicum at Los Alamos in 1993, contributed to the ASC program while a postdoctoral fellow at Sandia from 2000 to 2004. During that period Ropp worked on Trilinos, a package of algorithms designed to run on the large-scale ASC computers such as Purple.

For its work, the Xyce team recently won a prestigious 2008 R&D 100 Award, given by R&D magazine. The magazines' Web site states that the "Award provides a mark of excellence known to industry, government and academia as proof that the product is one of the most innovative ideas of the year."

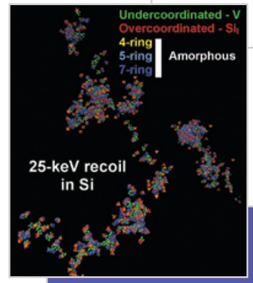

Molecular dynamics calculation of damage.

Helping Break

The Biomass Barrier

By Thomas R. O'Donnell

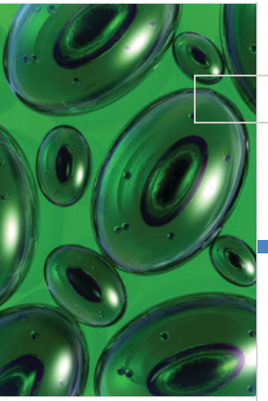

ANYONE WHO'S DEALT WITH AN UNCOOPERATIVE KID or combative employee is familiar with the concept of recalcitrance – stubbornness, disobedience and noncompliance.

It's easy to see why scientists would apply the same term to lignocellulose. The woody material found in plant cell walls makes trees tough and stems stiff, but also locks up sugars, making it difficult to ferment them into ethanol fuel.

"Plants don't want to be hydrolyzed and broken down into sugars. They develop these defense systems," says Jeremy Smith, director of the Center for Molecular Biophysics (CMB), an Oak Ridge National Laboratory (ORNL)-University of Tennessee joint project. In other words, the plant materials that make up biomass for energy are recalcitrant.

Smith, who also is the first University of Tennessee (UT)-ORNL Governor's Chair, and his fellow researchers are out to learn the atomic basis for this recalcitrance. With a 3.5 million processor-hour grant of computer time through DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, Smith is building computational models of lignocellulose and testing how enzymes interact with it and each other. They'll compare their models against neutron scattering experiments generated at ORNL's Spallation Neutron Source (SNS).

The Oak Ridge CMB participates in the Department of Energy's BioEnergy Science Center (BESC), an interdisciplinary coalition of experts from ORNL, UT and other universities, corporations and the National Renewable Energy Laboratory (NREL). The BESC mission is to achieve breakthroughs in biofuels from lignocellulosic biomass.

What they learn could help make cellulosic ethanol readily available and economically viable, helping the United States replace foreign oil with a renewable resource while curbing greenhouse gas emissions. As it stands now, making ethanol from biomass like plant stalks or wood chips is difficult and expensive, requiring substantial energy and chemical inputs. Cellulosic ethanol production could generate less carbon dioxide and interfere less with food supplies, but it's still easier and less expensive to make ethanol from corn and other crop-based feedstocks.

"We need to understand why (plant cell walls) are recalcitrant," Smith says. "To do that you need to understand the structures of plant cell walls and what's stopping them from being broken down by enzymes."

Cellulose, the source of fermentable sugars in biomass, is locked inside some pretty tough stuff, as researched by CMB graduate student Benjamin Lindner and postdoctoral research fellow Loukas Petridis. Cellulose is contained in compact, partially crystalline fibrils that block enzymes. Polysaccharides and lignin cover the fibrils, presenting another barrier that enzymes must overcome. The lignin also may inhibit enzymes by binding to their cellulose-binding components. In fact, removing lignin from biomass increases the cellulose-hydrolysis yield from 20 percent to 98 percent, the paper says.

What they learn could help make cellulosic ethanol readily available and economically viable, helping the United States replace foreign oil with a renewable resource while curbing greenhouse gas emissions.

Because of these barriers around lignocellulose, "You have to beat it up a bit before it will cooperate, and that's expensive," Smith adds. Computer simulation and neutron scattering experiments could provide clues on how to "kind of gently make (lignocellulose) cooperate in a way that's economical, so we don't have to heat it up to high temperatures or mix it with chemicals."

Smith and his fellow researchers are setting up molecular dynamics simulations that will help them understand the physics of the chemical reactions involved in enzymatic hydrolysis of lignocellulose. Petridis recently completed one of the first steps: a lignin "force field" model.

"You have to know what the forces on each atom of a system are" to make the lignin simulation accurate, Petridis says. "The force field is the parameters of the system that determine these forces." Using data from neutron scattering and other experiments, the researchers have focused on making the parameters closely represent lignin's properties.

"We have to make the model very specific to lignin," Petridis adds. "The forces acting on an atom of cellulose will be very different from those acting on the atoms of lignin and we need to understand the difference."

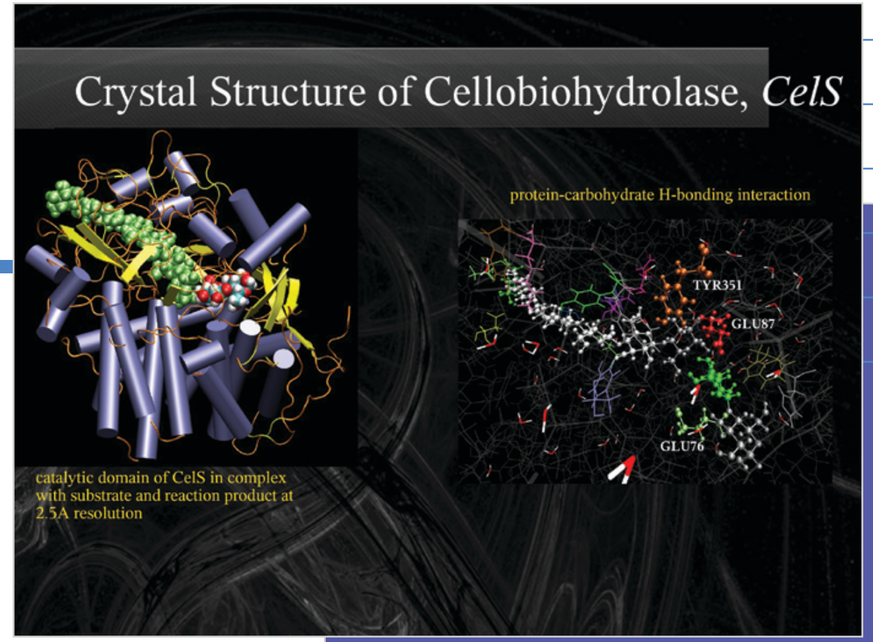

The model on the left is of the crystal structure of Cellobiohydrolase (CelS), the major enzymatic component of the cellulosome produced by the anaerobic, thermophilic bacterium Clostridium Thermocellum to degrade cellulose. It shows the CelS catalytic domain with substrate and reaction complex at 2.5 angstrom resolution. The model on the right is of the carbohydrate tunnel in CelS.

Once the lignocellulose model is complete, the research can turn to modeling how it interacts with enzymes that break lignocellulose down to sugars — and particularly how lignin keeps enzymes away. The enzymes of interest are cellulases, some of which are contained in spindly protein machines some bacteria and fungi produce to convert cellulose into sugars for energy. "If we understand how these enzymes catalyze the chemical reactions they do in breaking down cellulose, we will know how to improve them" and harness them to make ethanol, Smith says.

JEREMY C. SMITH

Jeremy C. Smith became the first Governor's Chair at the University of Tennessee and director of the UT/Oak Ridge National Laboratory Center for Molecular Biophysics in October 2006. Smith obtained his doctoral degree in Biophysics from the University of London and was a post-doctoral associate and lecturer at Harvard University. Before joining UT and ORNL Smith led research groups in biomolecular simulation at the Commissariat à l'Énergie Atomique (CEA) at Saclay, France (1989-1998) and as Chair of Computational Molecular Biophysics at the University of Heidelberg, Germany (1998-2006). Smith has performed and directed research in high-performance computer simulation of biological macromolecules, neutron scattering in biology, the physics of proteins, bioenergetics and the analysis of structural change in proteins. As of 2008 Smith had published more than 220 peer-reviewed scientific articles.

Further Reading:

L. Petridis and J. C. Smith, *A molecular mechanics force field for lignin*, Journal of Computational Chemistry, in press.

I. Daidone, M. Ulmschneider, A. Di Nola, A. Amadei and J. C. Smith, *Dehydration-Driven Solvent Exposure of Hydrophobic Surfaces as a Driving Force in Peptide Folding*, Proceedings of the National Academy of Sciences (U.S.A.), 104(39), 15230-15235 (2007).

L. Meinhold, F. Merzel and J. C. Smith, *Lattice dynamics of a protein crystal*, Physical Review Letters 99 (13):138101 (2007).

L. Meinhold, J. C. Smith, A. Kitao and A. H. Zewail. *Picosecond Fluctuating Protein Energy Landscape Mapped by Pressure-Temperature Molecular Dynamics Simulation*, Proceedings of the National Academy of Sciences (U.S.A.) 104(44) 17261-17265 (2007).

Contact

Jeremy Smith
smithjc@ornl.gov

PRACTICUM COORDINATOR

Dr. James J. Hack .....jhack@ornl.gov

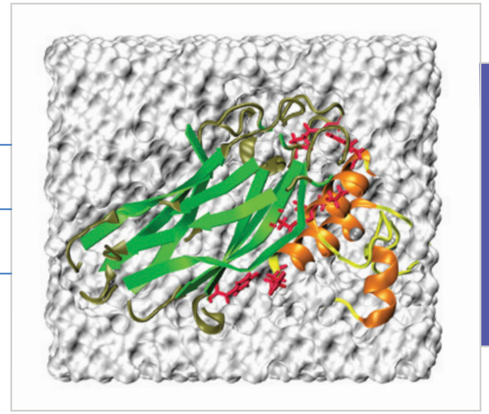

Researchers hope to overcome the recalcitrance of lignocellulose to enzymatic hydrolysis by understanding and modifying the properties of bacterial cellulosomes, large extracellular enzyme complexes that anaerobic bacteria produce to efficiently break down plant cell wall polysaccharides into sugars. The organization of the cellulosome is mediated by high-affinity protein-protein interactions between Type I cohesin domains within the scaffolding proteins and complementary Type I dockerin domains carried by cellulosomal enzymes. This image shows the cohesin-dockerin complex in a water box.

Cellulases actually are groups of proteins (one of Smith's specialties — *see sidebar on page 25*), each serving a different function but linked together. Researchers want to know how the proteins work together, why they are attached and how this makes them efficient. Another BESC member, Cornell University's John Brady, is refining a computer model of a cellulase molecule in action. He's collaborating with Mike Himmel of NREL and ORNL Genome Analysis and Systems Modeling Group researchers Ed Uberbacher, Phil LoCascio and Pavan K. Ghattyvenkatakrishna, who's also a University of Akron graduate student.

In Smith's lab and with NREL, postdoctoral research associate Jiancong Xu is working on a molecular dynamics cellulosome model. Meanwhile, Moumita Saharay, another postdoc in Smith's lab, is combining molecular dynamics and quantum mechanics to model how enzymes break and reform chemical bonds to break down cellulose. While the model will be more accurate, it also will demand more computing power, Petridis says.

The simulations explicitly recognize each atom in the molecules they portray, Smith says. Because the simulations can track a million or more atoms at once, "We need big supercomputers to perform these calculations," he adds.

At Oak Ridge, Smith and the rest of his group use Jaguar, a Cray XT4 with an aggregate system performance of 263 teraflops. Jaguar ranked as the world's fifth most powerful computer in the June 2008 TOP500 list. Quad-core processors have boosted Jaguar's power in recent years, but Smith and his fellow researchers want more.

"We are so hungry for computing power because we can never get enough. It's really what limits the usefulness of these simulations," he says. "A million atoms seems like a lot but it's only a tiny sliver of biomass. We'd also like to simulate it for milliseconds or a second. (Now) we can simulate it only a microsecond." Smith's group plans to integrate their simulations with results from another major facility at Oak Ridge: The Spallation Neutron Source, which generates the most intense pulsed neutron beams in the world. How neutrons scatter when they strike protein molecules provides clues about how atoms are arranged in those molecules.

"The simulations are a way of interpreting the neutron experiments, so the neutron experiments and simulations will be performed simultaneously," Smith says. Simulations will show how neutrons would scatter from a computer model of proteins. "If there's agreement of the scattering profile with the experimental measure then you can use the simulation to interpret the experiment in detail," he adds.

Having the two together will be key to unraveling biomass. "Lignocellulose is difficult to characterize experimentally," Smith says. "It's a heterogeneous mixture of different subunits with varying linkages and it's really a mess." Smith and his fellow researchers are only beginning to refine their computational models. "You can be sure that for the next year or two we'll get everything wrong," he joked, as models are tweaked and refined. "We'll do simulations that are not quite the real thing, but then eventually we'll get a model that seems to agree with everything," creating a powerful tool to help design new plants or enzymes that make lignocellulose more amenable to ethanol production.

Smith has little doubt they'll succeed. Although he doesn't claim to be an economist, he and other BESC researchers believe the recent turmoil in the energy markets brings further justification for the development of supply-stable, renewable and environmentally clean fuel sources such as cellulosic ethanol.

"Cellulosic ethanol may already be economical, but we want to make it even more economical," he adds. "I'm sure there will eventually be success in this, whether it's from us or other researchers."

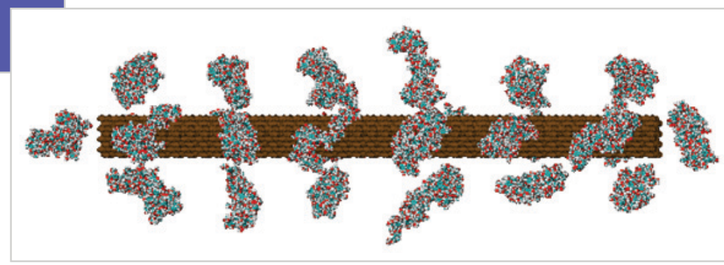

Atomic-detailed model of lignocellulose of softwoods. The model was built based on experimental data on the structure of cellulose (colored brown) and lignin (colored cyan and red).

RESEARCHER'S MODELS ARE HIGH IN PROTEIN

>> Jeremy Smith turned down offers of professorships in six countries to move his research from Germany's University of Heidelberg to Oak Ridge National Laboratory and the University of Tennessee in Fall 2006. The draw: A top-five supercomputer and the Spallation Neutron Source (SNS), generator of the world's most intense pulsed neutron beams.

Smith, a native Briton, uses high-performance computer simulations to interpret the results of neutron-scattering experiments. The experiments, meanwhile, provide data to refine and improve the computer models. Smith often focuses on proteins — with interesting results.

Running simulations on Jaguar, ORNL's Cray XT4 supercomputer, and on computers in Italy and Germany, Smith and a team of UT and Oak Ridge researchers found that the way parts of a protein interact with water can influence the way the whole protein folds into its final form.

"Protein folding is the key to how genes get translated into biological systems and functions," Smith says. How amino acid chains fold into a protein determines that protein's function, and understanding the process is vital to treating and preventing disease. Drugs often work by binding small molecules to proteins, and misfolded proteins have been linked to conditions like Alzheimer's and Parkinson's diseases.

The group's simulations looked at peptides — smaller chains of amino acids. They found that the way water wets hydrophobic areas of peptides determines their shape and behavior. Hydrophobic areas about the size of a water molecule have little effect on water around them. Larger hydrophobic areas of a peptide, in contrast, repel water and determine the peptide's structure.

Smith's group continues to work on the simulations and on interpreting experimental results, such as fluorescence tests to examine the kinetics of protein folding, including the time it takes for the ends of peptides to come together.

In another project, Smith led a team that performed the first atomic-detail simulation of how proteins vibrate in a crystal. Proteins carry out most of the cell's functions, so understanding how they interact is a key to deciphering basic biology.

The simulation depicts how phonons — units of vibrational energy — disperse through the repeating units of protein molecules locked in a crystal lattice. The simulations predict that the proteins "move relative to each other a bit like beads on a string," Smith says. "How they move, the frequency of these vibrations and the patterns of vibrations determine how the proteins interact and what the forces are" that affect them.

It will take neutron scattering experiments on the SNS to test the simulation's conclusions. Experimentalists must build the right instruments to track the vibrations and also grow appropriately sized crystals, "but the challenge is out there for them," Smith adds. "I'll be keeping an eye open to make sure they rise to it."

By Victor D. Chase

Going Big to Study Small

IT TAKES A BIG COMPUTER to model very small things. And, like its namesake state, New York Blue is big. Made up of 36,864 processors, the massively parallel IBM Blue Gene/L is housed at DOE's Brookhaven National Laboratory (BNL) on New York's Long Island, where, among other things, it's used to model quantum dots, or nanoparticles, just a few atoms in size.

Overseeing New York Blue's care and feeding is Brookhaven's Computational Science Center, headed by James W. Davenport. Davenport is spearheading an effort to model nano-sized slivers of metallic material, such as gold and palladium, on the supercomputer.

Modeling nano-sized particles is necessary because conventional diagnostic tools have not yet been able to accurately determine the makeup of such tiny bits, which differ from larger chunks of the same materials. When large pieces of gold and palladium are X-rayed their crystal structure becomes apparent. "The atoms are arranged in a certain periodic array that is well defined. But small particles are not the same. It is not known exactly what the atomic structures of most of these particles are," Davenport says.

Learning the physical properties of these nanocrystals is the first step in determining how they may be used as catalysts — materials which manipulate chemicals — and in turn create new applications in DOE-related fields such as advanced energy technology.

Hydrogen Storage Turns Golden

Gold, for example, is a noble inert metal that normally does not react with other chemicals. Yet when gold particles made up of fewer than about 1,000 atoms are placed in the presence of certain chemicals, catalytic reactions take place.

A particular catalytic interaction involving gold nanoparticles may prove useful for producing hydrogen, most of which is currently made from natural gas. Hydrogen separated from natural gas often also contains carbon monoxide, which poisons the fuel cells used to produce electricity. Gold nanocrystals may be used as a catalyst to scrub carbon monoxide from the hydrogen stream. This could greatly expand clean electricity production, since water is the only byproduct of fuel cell operation.

There also is evidence that nanoparticles might be a good hydrogen storage medium. Hydrogen storage currently requires considerable space, presenting a major obstacle to its use as a significant alternative energy source. Storing hydrogen in a large tank, for example, makes it less practical as an automotive fuel. Cutting the space requirement by getting hydrogen to adhere to tiny particles of metal would be a major step toward creation of a hydrogen-based economy.

All of this makes it important "to understand how the chemical bond of hydrogen to these nanoparticles differs from the chemical bond of hydrogen to a large chunk of metal," Davenport says.

Learning the physical properties of these nanocrystals is the first step in determining how they may be used as catalysts — materials which manipulate chemicals — and in turn create new applications in DOE-related fields such as advanced energy technology.

View of New York Blue along with a portion of the 560 terabyte disk storage facility.

To accomplish these goals first requires understanding how the properties of the very small differ from the big. And that is where the modeling on New York Blue comes in.

Searching For Low Energy

"We are studying the shapes of the different particles as a function of their size by calculating the energies of these different structures and seeing which one is the lowest," Davenport says.

The energy in this case is potential energy — the force drawing the atoms together — as opposed to kinetic energy, which involves actual movement of the atoms. "If I allow two atoms to come close together the bond energy increases and that's what makes chemical bonds — that's what holds our world together," Davenport says. "We are trying to calculate that chemical bond energy."

Researchers look for structures with the lowest energy because that's how nature tends to arrange things. To conduct this search, Davenport's team

JAMES W. DAVENPORT

James W. Davenport is a senior physicist and director of the Computational Science Center at Brookhaven National Laboratory. He's also interim technical director of the New York Center for Computational Sciences and an adjunct professor in the Stony Brook University Physics Department. Davenport's research interests include computational science, electronic structure and magnetism of nanoscale clusters, applied mathematics, massively parallel computing and protein molecular dynamics. Davenport is author or coauthor of more than 100 publications dealing with the electronic structure of atoms, molecules and solids. He invented LASTO, the linear augmented Slater-type orbital method. Davenport earned his doctoral degree in physics from the University of Pennsylvania in 1976, his master's degree in electrical engineering from Princeton University in 1968 and his bachelor's degree in electrical engineering from Brown University in 1967.

Further reading:

N. Kocharian, G. W. Fernando, T. Wang, K. Palandage, and J. W. Davenport, *Exact thermodynamics of pairing and charge-spin separation crossovers in small Hubbard nanoclusters*, Physics Letters A 364, 57 (2007).

J. C. Zheng, Y. Zhu, L. Wu and J. W. Davenport, *On the sensitivity of electron and X-ray scattering factors to valence charge distributions*, J. Appl. Cryst. 38, 648 (2005).

J. C. Zheng and J. W. Davenport, *Ferromagnetism and stability of half-metallic MnSb and MnBi in the strained zinc-blende structure: Predictions from full potential and pseudopotential calculations*, Phys. Rev. B 69, 144415 (2004).

Contact:

James W. Davenport
jdaven@bnl.gov

PRACTICUM COORDINATOR

James Davenport . . . . . jdaven@bnl.gov

computationally arranges atoms into different shapes and models their potential energy.

"If two atoms are far apart then the energy is two times the sum of the isolated atoms," Davenport explains. "If I allow those two atoms to come close together they will form a chemical bond. It is the chemical bond energy that we are really looking for."

As part of their search, the investigators use quantum mechanics to simulate the arrangement of atoms in different configurations and then attempt to determine which has the lowest energy.

"If I have three atoms I can put them into a triangle. If I have four I can put them into a tetrahedron, or if I have eight I can put them into a cube," Davenport says. "I can then vary the cube edge to see how the energy varies. If the cube edge is very small, usually the energy will be high, because these atoms are too close together and repel each other. If I take them far apart they also have a high energy because they like to bind together. So someplace in between there is an optimum and we search for those optima."

Up to now, the group has verified its results with cutting-edge optical experiments that are slowly becoming available. "There are molecular beam experiments in which these particles can sometimes be created and can be stuck on surfaces and examined with electron microscopes, but it is difficult," Davenport adds. The researchers also have looked at hydrogen as though it was stuck on a palladium nanostructure, but do not yet have experimental confirmation of their results. They're now working to scale up their models to examine the structures of particles consisting of hundreds to thousands of atoms.

The researchers chose to model metals first because of previously existing evidence that they would make good storage media and because of past experience at BNL, where other groups have modeled metals for other potential applications.

Davenport's team has used a number of DOE-developed computer codes in its work, including NWChem, which was developed at Pacific Northwest National Laboratory; GAMESS, a product of Ames Laboratory; and a home-grown code dubbed LASTO.

Scaling Up

Researchers at BNL's Computational Science Center have modeled nanoparticles for several years and succeeded in calculating the crystalline structures of small clusters of up to 150 atoms of gold and palladium. Their results agree with some empirical evidence.

New York Blue, an IBM Blue Gene/L supercomputer, occupies some 2,000 square feet at DOE's Brookhaven National Laboratory.

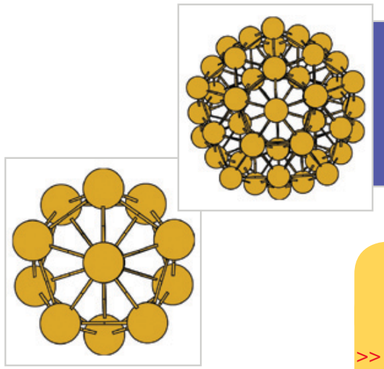

Atomic arrangement in icosahedral gold clusters with 13 and 55 atoms, view from the top.

Figure provided by M. McGuigan, BNL.

In addition to the nanoscience work, other researchers at BNL's Computational Science Center are working on magnetohydrodynamics, with a particular focus on ITER, a fusion research project undertaken by an international consortium including the United States. The ITER reactor, being built in France, will be used to create and study conditions required to produce electricity through nuclear fusion, rather than fission, the process behind present nuclear power plants. Nuclear fusion, the process that takes place in the Sun, holds the potential for solving the world's energy problems because it would be a virtually limitless energy source independent from finite fossil fuels.

Brookhaven researchers also are involved in modeling the intense collisions produced in Brookhaven's Relativistic Heavy Ion Collider (RHIC). Activated in 2000, the RHIC is used to study what the universe may have looked like in the first few moments after its creation. It does so by crashing beams of gold ions head-on in a subatomic collision. Davenport's department is helping increase the intensity of the beam in the collider by simulating the motion of particles in the accelerator.

Some of this computational science work is done on other computers at the Computational Science Center, including the 500-processor Galaxy, a conventional Beowulf cluster. But it is New York Blue that has become the shining star of a very small universe.

GETTING BLUE AT BROOKHAVEN

>> Brookhaven National Laboratory (BNL) and New York's Stony Brook University have formed a joint venture that revolves around New York Blue, an IBM Blue Gene/L supercomputer owned by the state and housed at the Department of Energy lab.

Known as the New York Center for Computational Sciences, or NYCCS — pronounced Knicks, as in the New York Knicks — the partnership was formed to foster the cooperative use of New York Blue (NYB) between DOE, Stony Brook and some 200 other users throughout the state, including universities and various private organizations. The machine became operational late in 2007.

Though Stony Brook — part of the State University of New York system — owns NYB, the supercomputer was placed at Brookhaven because of DOE's supercomputing expertise, because BNL had the 2,000 square feet required to house NYB, and because there is a strong connection between the lab and the university. "We have an especially strong interaction with the applied mathematics and statistics department," says James W. Davenport, head of BNL's Computational Science Center, which maintains and operates NYB.

New York's goal in obtaining the computer and sharing it with the DOE is "to upgrade the computational science capabilities of researchers within the state. It also wants to upgrade the high-end computing capabilities of industry. One of the interests is high-tech job growth," Davenport says.

"The state also recognizes that having high-end computing capabilities attracts people to universities, so there is an increase in federal grants, and it also builds a high-tech culture that allows companies to attract engineers and scientists even if they are not directly using it," Davenport says.

In return for overseeing NYB's care, Davenport's group of some 10 people also gets to run much of its modeling on the supercomputer, including simulating nanoparticles, molecular biology and fluid dynamics. Other BNL groups also use NYB to run high-energy physics programs (specifically as they relate to the lab's Relativistic Heavy Ion Collider) and climate modeling programs.

New York Blue also is used to train DOE scientists and other researchers from throughout New York state on massively parallel machines. "We anticipate that our users will then migrate to even larger DOE machines as they develop the skills and the codes that are suitable for such large machines," Davenport says.

The other supercomputers are used in a variety of DOE programs, including Advanced Simulation and Computing (ASC), which is designed to simulate the performance, safety and reliability of nuclear weapons; and the Leadership Computing Facilities at Argonne and Oak Ridge national laboratories and at the Lawrence Berkeley National Laboratory's National Energy Research Scientific Computing Center.

A Small Space

The 2,000 square feet NYB occupies is a relatively small amount of space for a 100-teraflops machine — a single teraflops being one trillion floating point operations per second. This is made possible by a low-power design IBM incorporated into the supercomputer, enabling the processors to run cooler than conventional ones. This also reduces electricity consumption and allows more processors to be packed into less space. With a relatively low clock frequency of 700 megahertz, the processors require some six times less electricity than conventional ones. The slower processors do compromise speed somewhat, but as Davenport explains, "It is nonlinear, so you get less performance per processor, but you can make up for it on volume."

The Blue Gene design is similar to what is known as a Beowulf cluster — a large number of off-the-shelf computers linked in parallel. Though the NYB concept is much the same, it's different from Beowulf computers because IBM designed it from scratch and it uses a special high-speed torus-interconnecting network. It runs a Linux operating system, standard for cluster computers, and uses MPI — message-passing interface, commonly used to enable communications among processors in a cluster.

By Alan S. Brown

Small Splats Help Models
Predict the Untestable

SOMETIMES, JOE SEFCIK'S JOB **comes down to measuring little splats. These are not just any little splats, but the patterns formed by blasting a small piece of metal with the world's most powerful (and expensive) laser.**

He compares them with predictions made by one of the world's most elaborate mathematical models. If the model cannot predict the pattern to within a few microns, it's probably not accurate enough to describe the short but extremely energetic life of a nuclear weapon.

Sefcik heads Lawrence Livermore National Laboratory's Verification and Validation (V&V) team, whose researchers determine the accuracy of models that simulate the detonation of nuclear weapons. In other words, their job comes down to answering one critical national security question: How do we know if we can trust our nuclear weapons models?

The question is more than theoretical. Sefcik's group works for the Department of Energy's National Nuclear Security Administration (NNSA), which is charged with ensuring the safety, security, and reliability of U.S. nuclear weapons. As long as NNSA can guarantee that the stockpile will perform as designed, it will remain the strongest possible deterrent against attack by a foreign power.

Certifying weapons, however, is an extraordinarily complex undertaking. Since the end of the Cold War, the United States has ended underground nuclear testing, stopped the production of new nuclear weapons and reduced the size of its nuclear arsenal by 80 percent.

Many of the remaining weapons are approaching the end of their lifespans. They may require repairs and new parts to remain functional or modifications to support changing national security missions.

Models

Since the tests have been discontinued, computer modeling is the chief means for the Department of Energy to predict the reliability of aging or modified weapons. Teams working for NNSA at Livermore and other national laboratories have created sophisticated mathematical models of nuclear weapons and run them on the world's largest supercomputers. By simulating nuclear weapons in action, researchers hope to judge how potential repairs and modifications will affect their performance.

It takes highly accurate models to answer these questions because thermonuclear weapons are unlike other weapons. Conventional bombs work by igniting explosive substances. Thermonuclear weapons are far more complex and must move through a series of phases prior to detonation. They start with exotic materials and complex mechanical systems that produce a controlled nuclear explosion. The weapon utilizes this energy to compress hydrogen isotopes until they produce a fusion explosion. Structures within the warhead have to deform in a controlled pattern for fusion to occur.

"The physical behavior — the temperatures and pressures and energy densities — are typical of what you would find in the center of sun," Sefcik says. "These are not regimes human beings have access to. And we have to simulate them in the model from the time you press the button through the final explosion."

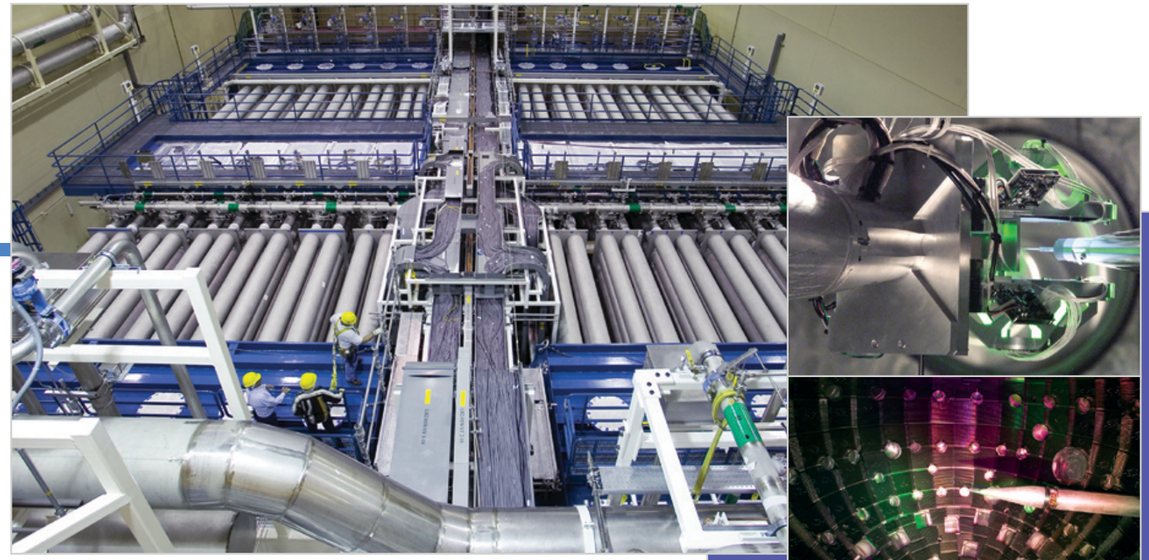

The National Ignition Facility uses a bank of 192 powerful lasers beams to recreate the temperature and pressure regimes found in stars and in thermonuclear explosions. Testing materials in this environment enables researchers to validate models by checking predictions against test results.

Graphics credit is given to Lawrence Livermore National Security, LLC, Lawrence Livermore National Laboratory, and the Department of Energy under whose auspices this work was performed.

It takes 8 million lines of code for all of Livermore's models to describe the process. The models do it from the ground up, starting with chemical reactions and the deformation of metal crystals through the trillionths-of-a-second interactions of subatomic particles during the nuclear explosion. Often, those different steps rely on different types of theories. One may use quantum mechanics, another electrodynamics and a third classical thermodynamics. Yet each of those steps produces data that drives the next step of the model.

The V&V team's job is to ensure the entire model fits together and makes sense. On one hand, they have to verify the code to make sure it contains no programming or math errors. On the other, they also must validate that the model uses the right theories to predict events. "In other words," Sefcik says, "our model has to give the right results for the right reasons."

Validation can be tricky. Often, more than one competing theory does a good job of explaining phenomena in a lab. There's also the question of how an explanation will fare under the extreme conditions of a nuclear blast. To determine which theory or combination of theories fits best, Sefcik's team needs to find a way to reproduce those conditions in a laboratory.

Splats

It turns out that the perfect laboratory is just down the road from Sefcik's office. It's Livermore's National Ignition Facility (NIF), the largest, most expensive, and most powerful laser system ever built. When completed sometime in 2009, its 192 individual laser beams will harness an incredible 500 terawatts of power — roughly 125 times the planet's total electrical output — for a billionth of a second to produce an intense pulse of light.

Certifying weapons, however, is an extraordinarily complex undertaking. Since the end of the Cold War, the United States has ended underground nuclear testing, stopped the production of new nuclear weapons and reduced the size of its nuclear arsenal by 80 percent.

JOSEPH A. SEFCIK

Joseph A. Sefcik is program leader for the Verification and Validation Program in the Advanced Simulation and Computing (ASC) Program at Lawrence Livermore National Laboratory (LLNL). He received his bachelor's degree in applied and engineering physics from Cornell University and his doctoral degree in nuclear engineering from the Massachusetts Institute of Technology. He joined the laboratory in 1981 to work on issues in nuclear nonproliferation, export control and strategic defense. With the moratorium on nuclear testing, he focused on applying laboratory-developed defense technologies to commercial industry. Dr. Sefcik assisted the Department of Energy in forming the Advanced Design and Production Technologies (ADAPT) program to bring modern fabrication and integration technologies into aging plants in the DOE complex. He then led the Nuclear Materials Technology Program at LLNL, where he was responsible for the research and development, facilities and security associated with special nuclear materials and tritium. Sefcik is the recipient of two Department of Energy Awards of Excellence and shares an R&D 100 Award for the development of femtosecond laser materials processing. He is the author of numerous publications on nuclear nonproliferation, nuclear design and directed-energy technology.

Contact:
sefcik1@llnl.gov

PRACTICUM COORDINATOR

Jim McGraw .....mcgraw1@llnl.gov

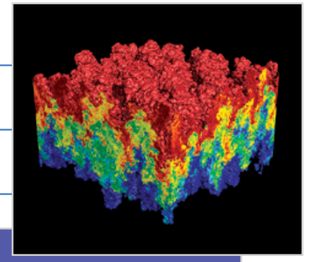

Multiphysics models, such as those used to predict weapon performance, can provide powerful insights into the physical world. Here, a simulation of the Rayleigh-Taylor instability at the interface of a lighter fluid pushing into a heavier fluid.

Graphic provided by National Nuclear Security Administration

The laser beams vaporize anything in their path. In fact, they add so much heat so fast, they cause anything they light up to explode and create an enormous shock wave. Researchers at NIF plan to use these shock waves to compress hydrogen isotopes in small canisters until they fuse and produce energy.

The shock waves NIF lasers generate are similar to those produced by the controlled blast of a thermonuclear weapon. In fact, even a small fraction of the facility's lasers provide enough energy to test how metals behave under the extreme temperature, pressure and energy density regimes of a thermonuclear explosion.

This enables Sefcik's team to get around a fundamental problem. To build models, researchers start with theories and then fine-tune the math so the output matches test data. Once a model is tuned, they use it to predict results. Ordinarily, the model's predictions are close; but how close is close enough?

Ordinarily, scientists would vary some parameters and run additional tests to see how well the model forecasts the results. "Unfortunately, the regimes we needed to test were only available from underground testing, which we discontinued in 1992," Sefcik says. "We could simulate the old tests and see if simulation gets the correct answer, but sometimes the answer is not accurate enough. Then we don't have enough certainty to say for sure whether the model is working properly or not."

When that happens, the first thing Sefcik's team does is trace the result backwards through the program. "Sometimes it's easy, and we find something anomalous and we can trace it back to a bug," he notes.

"Other times, we do not predict something to the level of accuracy needed to match the data," he continues. "So we step through the model to see where it began to diverge from the accuracy we need. Then we look at the physical regime where this occurred and design an experiment to see if we can test the code at that point."

The NIF splatter test was one such experiment. "We wanted to see how well we could predict hydrodynamic jets," Sefcik explains. "These are jets of hot plasma that occur when we focus a powerful laser on a spot of metal. The spot gets so intensely hot that you get bursts of metallic plasma jetting out the other side. Our models let us calculate the position of all those jets within tens of microns. What was critical was where the splats go. When we looked at the code's predictions, the results were very close."

Flops and Teamwork

It takes some of the world's most powerful supercomputers and a great deal of teamwork to produce such accurate predictions. The Department of Energy and NNSA have been building that infrastructure since the earliest days of the nuclear weapons program.

Today's highly precise models trace their genesis back to 1995, when the 10-year Accelerated Strategic Computing Initiative (ASCI) began to change the face of supercomputing at the national laboratories. The program's two-fold goal was to produce a 100 teraflops supercomputer and use it to create more realistic three-dimensional simulations.

"Flops" are a sophisticated measurement of such common math operations as multiplication and division. To appreciate how many calculations that is, imagine that each one of the Earth's 6.7 billion inhabitants did one math problem per minute without stopping or sleeping. It would take 10 days for all of us to complete the same number of calculations a 100-teraflops computer does in one second.

ASCI's powerful supercomputers *(see sidebar on this page)* brought revolutionary changes, Sefcik says. "It let us go from calculating behavior in two dimensions to calculating them in all three dimensions," he adds. "We could build models based on the fundamental principles of physics and ASC researchers have used them to discover new types of metal failure that were later confirmed in the laboratory. The simulations also run much faster and produce higher-fidelity results."

The Advanced Simulation and Computing (ASC) program, which began in 2005, is building on those successes. It ultimately will push supercomputer speeds well into the quadrillion operations per second (petaflops) range and make simulation software easier to design and use.

Besides hardware, people comprise an important part of the NNSA's infrastructure. Livermore is one of three major national laboratories, along with Los Alamos and Sandia, that model the nuclear stockpile. Sefcik estimates that 1,300 people are involved in the Stockpile Stewardship Program at Livermore alone.

His V&V team of 36 researchers is part of the laboratory's larger ASC program and consists of many different specialists who leverage the work of the larger Stockpile Stewardship Program. Some are quality assurance experts who comb the code for errors or bugs. Others focus on understanding which of the model's adjustable parameters will have the most influence on its output, so researchers can focus on reducing uncertainty in those assumptions. One group consists of computation experts who seek to improve model efficiency, while another tries to determine whether the errors inherent in any model, such as rounding off very small numbers, will produce very large errors in the results.

NNSA also has reached out to the academic community. After an intense competition, the Predictive Science Academic Alliance Program (PSAAP) chose five university partners: the California Institute of Technology, Purdue University, Stanford University, the University of Michigan and the University of Texas at Austin.

The five schools will receive $17 million each over the next five years to establish centers of excellence in predictive science — the use of verified and validated models to predict the behavior of complex systems for which routine experiments are not feasible. The universities will build unclassified predictive models for phenomena as diverse as material deformation caused by high-energy impacts and the flow of hypersonic shock waves.

"We've had a fair number of students from similar programs in the past come to work at the labs because they were fascinated by our research or by building and running some of the world's largest computers and models," Sefcik says.

And why wouldn't they? Those researchers will enter an exciting new world of discovery. Sefcik's team has only scratched the surface when it comes to validating models through high-energy testing at NIF and other laser and accelerator facilities around the nation.

For the first time, models are accurately describing phenomena that we cannot easily observe or test in a lab. Such models are already ensuring the safety and reliability of the nuclear stockpile. They also have led to discoveries about fracture mechanics in metals. One day, the techniques NNSA researchers pioneer may allow us to understand the behavior of other regions where we cannot venture, from the cores of stars to the heart of the Earth's interior.

POWERFUL HARDWARE

>> NNSA laboratories operate the world's most powerful supercomputers, including Los Alamos National Laboratory's Roadrunner, the first computer to break the 1 petaflop barrier. Built with commercially available parts, including souped-up versions of the same processors used in Sony PlayStation video systems, the $100 million computer hit 1,026 trillion floating point operations per second on a common benchmarking program.

Roadrunner, which began operation in 2008, is nearly twice as fast as the previous top-rated system, the 596-teraflops (trillion calculations per second) Blue Gene/L at Lawrence Livermore National Laboratory. Blue Gene/L itself is nearly three times faster than the next fastest supercomputer, comparisons compiled by the TOP500 organization show.

Livermore also operates the world's 33rd fastest supercomputer, ASC Purple. At 93 teraflops per second, Purple is not as fast as Roadrunner or BlueGene/L. Yet its large processor memory makes it ideal for more memory-intensive codes. Another NNSA national laboratory, Sandia, houses Red Storm, the world's 9th fastest computer

All four of these high-performance supercomputers have massively parallel architectures. This means that instead of doing all their calculations on a single powerful processor, they apportion the computing workload among thousands of individual processors. Many of the processors actually contain two or four processing cores that can operate independently of the others. Roadrunner, for example, runs calculations in 122,400 individual cores.

Massively parallel supercomputers operate much like the oarsmen aboard a Roman galley. Individually, none of the rowers could move the ship. By dividing the work, however, they were able to cruise and maneuver at high speeds.

NNSA's complex simulations typically contain many smaller models. Some describe atomic-level phenomena, while others simulate events that occur as larger devices and forces interact. Parallel supercomputers enable researchers to run many of these smaller models at the same time, with each part feeding information to the others as they solve a problem.

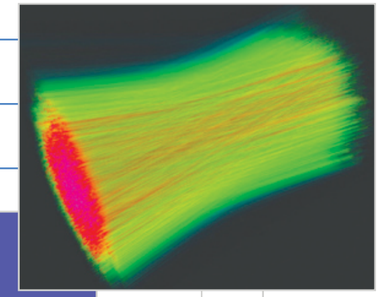

ASC Computers are being used to model complex phenomenon such as this simulation of laser filamentation produced by laser-plasma interaction at the National Ignition Facility.

Graphic provided by Lawrence Livermore National Laboratory

Finding the Right Answer

For the Right Reason

By Alan S. Brown

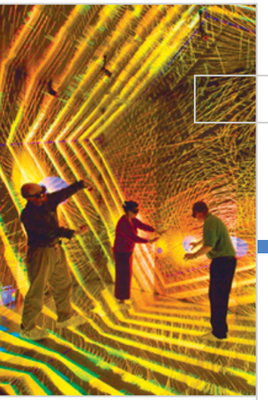

WHEN PEOPLE THINK ABOUT NUCLEAR WEAPONS, **most probably picture warheads sitting silently on missiles inside dimly-lit silos or submarine launch tubes, or perhaps within secure hangers near airbases. The bombs themselves appear remote, cold and unchanging.**

Nothing could be further from the truth. Nuclear weapons are changing all the time, says Scott Doebling of Los Alamos National Laboratory, who recently took leave from managing the lab's Verification and Validation (V&V) program to take an assignment with the National Nuclear Security Administration (NNSA).

"You're dealing with systems that have radioactive isotopes in them. As the isotopes decay over time, they emit radiation that changes the properties of the materials surrounding them in the warhead. If it is a polymer or a metal, for example, it might make it brittle and less strong. It could impact the performance of that material in some way."

Just how does prolonged exposure to nuclear radiation change a material's properties? How do those changes alter the way a weapon performs? Will an older bomb still perform as intended? Will it remain safe and intact if it falls off a truck, crashes in an aircraft or burns in a fire?

At the height of the Cold War, the United States had easy answers to those questions. First, it never let nuclear weapons age, replacing them every 10 to 15 years. Second, it tested weapons to ensure they performed as intended. Finally, America kept adding to its stockpile. Planners could count on sheer numbers alone to deter any potential attack.

That changed when the Cold War ended. In 1992, the United States halted underground nuclear testing, reduced the size of the existing nuclear stockpile and stopped making nuclear weapons.

Since the nation no longer built new warheads, it needed to extend the life span of the ones it had. It also needed to assess the performance of those aging and reconditioned weapons — and it had to do it without actually testing the weapons.

Today, the Department of Energy relies on computational simulations to predict weapons performance, employing several of the world's largest supercomputers and some of the largest and most sophisticated computer models ever created.

Yet the models raise questions of their own. They are large, complex combinations of many smaller models. Often, those models are based on different understandings of how physics works. They sometimes deal with temperatures and pressures where the physics are not well understood. Running such complex software also requires compromises that could change a model's ability to predict future warhead behavior.

The task of the Los Alamos V&V team is to quantify those uncertainties and improve our confidence in the simulations.

Today, the Department of Energy relies on computational simulations to predict weapons performance, employing several of the world's largest supercomputers and some of the largest and most sophisticated computer models ever created.

Inherent Problems

Computational simulations have always played a role in managing the uncertainty surrounding nuclear weapons, dating back to the Manhattan Project, when scientists did calculations by hand. Since then, the Atomic Energy Commission and now NNSA have supported breakthroughs in computing power and in our ability to model complex science *(see sidebar on page 37)*.

Thanks to those advances, computational simulation was ready to move from a supporting role into a starring role when researchers could no longer test weapons for safety and reliability. Today's models produce incredibly detailed results that appear to precisely simulate the behavior of weapons and materials as they age.

Yet do they? The uncertainties start with construction of the models themselves.

"Many of the mathematical models we use to understand physics are limited in what they explain," Doebling says. One theory, such as electrodynamics, may elucidate the behavior of a few atoms individually. Yet it cannot portray the transfer of heat or force through a block of material made of those atoms. That would take a second theory, such as thermodynamics or mechanics.

Physicists describe this as a problem of scale. A model that describes the behavior of atoms does not necessarily scale up to describe the behavior of heat moving through a block of those atoms. Moreover, it isn't always obvious how to connect the results of such different models to describe the behavior of a single material.

Sometimes, just picking the right model is hard. Thermonuclear weapons operate under extreme conditions that aren't easy to understand. As Doebling notes, weapons models must simulate a broad range of environments, from room temperatures and pressures to conditions in the center of the sun. A theory may work in one regime but prove less reliable in another.

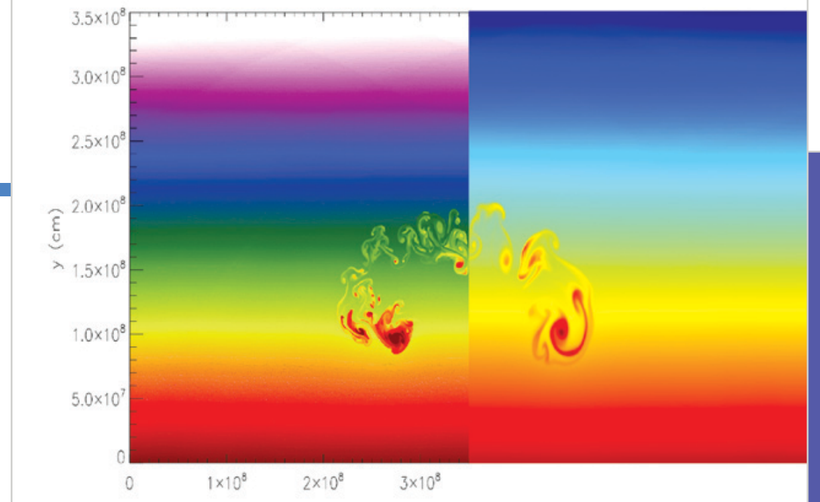

Sometimes comparisons of calculations from one physics code (left) with calculations from another physics code (right) can provide insight as to whether either code is solving the equations correctly.

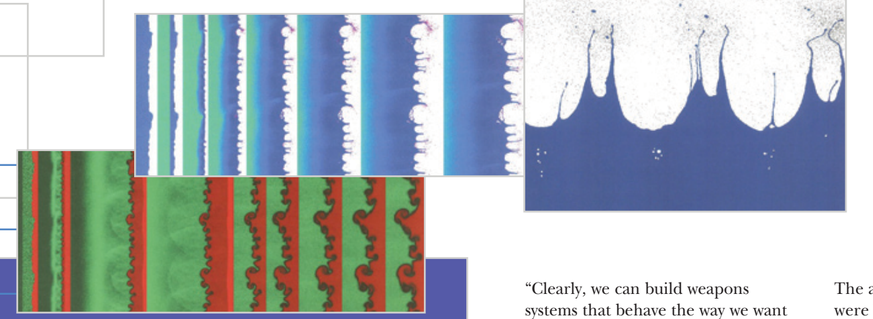

Simulating nuclear weapons requires a detailed understanding of material behavior. Here, a large shock will cause solid copper to eject material. The simulation, which ran for 88 hours on the world's second-fastest supercomputer, Blue Gene/L, described the behavior of 800 million atoms over 1 nanosecond.

Graphics provided by Los Alamos National Laboratory

SCOTT DOEBLING

Scott Doebling is a manager and researcher in the Applied Physics Division at Los Alamos National Laboratory. He has served as a program manager in the Advanced Simulation and Computing program since 2004 and currently is assigned to the National Nuclear Security Administration in Washington, D.C. His interests center on the assessment of errors and uncertainties in complex physics and engineering calculations, including the comparison of simulation results with experimental data. His previous research includes the study of shock propagation through mechanical interfaces, the detection of damage in civil and aerospace structures and assessment of uncertainty in vibration measurements. Doebling earned a doctoral degree in Aerospace Engineering Sciences from the University of Colorado in 1995, after earning bachelor's and master's degrees in Aero & Astro Engineering from Purdue University. He lives in Los Alamos, N.M. with his wife and three children.

Further Reading

S. W. Doebling, as member of the ASME Committee PTC 60, ASME Publication V&V 10: *Guide for Verification and Validation in Computational Solid Mechanics*, October 2006.

S. W. Doebling, F. M. Hemez, A. N. Robertson, R. D. Maupin, J. F. Schultze, A. L. Cundy, and J. E. Hylok, *Validation of the Transient Structural Response of a Threaded Assembly: Phase I*, Los Alamos National Laboratory report LA-14104-MS, May 2004.

S. W. Doebling, F. M. Hemez, and M. C. Anderson, *A Brief Tutorial Introduction to Model Verification & Validation for Structural Dynamics*, Proceedings of IMAC XXII: A Conference & Exposition on Structural Dynamics, 26-29 January 2004, Dearborn, Michigan, USA.

Contact:
doebling@lanl.gov

PRACTICUM COORDINATOR

Aric Hagberg ..... hagberg@lanl.gov

"Clearly, we can build weapons systems that behave the way we want even if we do not understand every detail that goes into them," Doebling says. "The question is, can the models we build give us a high degree of confidence that a system will perform as designed even though we are not sure of every little detail in the model? In other words, how good of a simulation is good enough?"

Knobs

There are several ways to get a handle on how aging and reconditioned warheads will perform. One is to learn more about how they age. NNSA does that by dismantling weapons and examining the pieces as part of its stockpile stewardship program.

Deciphering the physical evidence is harder than it sounds. Radiation may have turned a polymer a different color, but its other properties may remain unchanged. A metal may look new but behave differently.

Those changes are critical only if they affect the bomb's performance. To predict whether they will, model builders need to understand how those materials behave when exposed to the extreme heat, temperature and radiation of a nuclear explosion. In the past, there was no way to run those experiments without actually detonating a nuclear bomb.

Researchers finessed the problem by including what they call "ad–hoc parameters" in the mathematical models. These parameters accommodated both the unknowns of aging materials and unknowns in the models themselves. "We saw a problem with the accuracy of our mathematical models and built a factor into the equation to account for those discrepancies," Doebling adds.

The ad–hoc parameters, however, were based on data accumulated during underground tests that ended more than 15 years ago. Researchers adjusted their materials models until the model outputs matched the test outputs.

Historically, this approach has always been necessary when designing weapons. Engineers added a margin of error to their warhead plans to make up for ad–hoc parameters as well as any manufacturing variations.

"It's similar to what engineers do when they build a bridge," Doebling says. "They use extra concrete and steel in order to avoid unanticipated failures. We do the same with our designs because we don't really know exactly how those materials will behave as they move away from the design point."

Modelers call ad–hoc parameters "knobs." "From the weapon design perspective, they are necessary to get the job done in design engineering," Doebling says. "Now we're going back and trying to understand the physics of those knobs."

Verification

Verification and validation, Doebling says, is about ensuring the calculations yield the right answers for the right reasons, which means reducing and eventually eliminating dependence on knobs. "If we want our models to predict unknown behaviors, they cannot be right just because we tuned the model to our experimental data," he adds. "They have to be right because the physics embedded in the models represents how the system truly behaves."

The Los Alamos V&V team approaches that issue in two ways. First, it attempts to verify that its models solve equations properly. This takes both finely honed math and programming skills and judgment.

The team starts by looking for programming errors. Since complete models contain millions of lines of codes, it's a vast undertaking. The team also re-checks the math. "Many of the code's algorithms and equations were written before the guys checking them were born, so they have to make sure they really understand why the math was done the way it was," Doebling says.

Equally important, verification ensures that the model provides enough precision without sacrificing run speed. Part of the problem, Doebling says, is that computers cannot solve the continuous equations that describe the fundamental conservation laws of physics. Instead, they digitize those equations into smaller segments.

For example, a differential equation might generate a curve. A computer, however, will break that curve into minute segments and calculate the length and angle of each segment. The smaller the segments, the closer they resemble the original curve, but the more time they take to calculate. "We have to make sure the model's segments are small enough to capture the important details and large enough to produce results in our lifetime," Doebling says.

Validation

The Los Alamos team's second approach is to validate a model's underlying physics. In other words, it ensures the model reflects reality.

New testing facilities make this possible. NNSA now operates several sites that can produce conditions of extreme temperature, pressure and radiation. These include the newly renovated Sandia National Laboratories' Z accelerator, the University of Rochester's expanded OMEGA laser and Lawrence Livermore National Laboratory's new National Ignition Facility (NIF).

NIF is the most powerful of them all. Its lasers generate so much heat they explode small objects on contact, generating heat and shock waves that resemble those created by thermonuclear devices. For the first time, scientists actually can observe material behavior under extreme conditions without setting off a thermonuclear device.

The experiments are costly and time-consuming to set up and run. So instead of testing everything, the V&V team identifies areas where precision counts most, such as the changes in plutonium crystals prior to fission and the behavior of solid hydrogen isotopes before fusion.

"Experts typically have several hypotheses about how something works," Doebling says. "Using our simulations, we try to gain insights into which theory might be the most likely. Then we focus our experimental resources on testing that hypothesis."

Doebling pauses for a moment, then adds, "These are incredibly complex projects, but verification and validation has to be tightly linked to testing. Decisions about the safety, security and reliability of nuclear weapons have pretty high consequences. These weapons are operational military hardware. We have to know that they will survive a fire or a plane crash without accidentally detonating. For them to act as a deterrent, we have to know that after aging and modifications, they will still operate as designed."

The sophisticated and detailed models NNSA has developed accurately predict how materials behave under unknown conditions, yet they still contain gray areas. At Los Alamos, the verification and validation team is making sure those gray areas grow less uncertain all the time.

ADVANCED SIMULATION & COMPUTING

>> The National Nuclear Security Agency (NNSA) launched the Accelerated Strategic Computing Initiative — later renamed the Advanced Simulation and Computing (ASC) Program — in 1995. Its mission was to develop the infrastructure needed to predict nuclear weapon performance without detonating a bomb.

The program brings together three national laboratories — Los Alamos, Lawrence Livermore and Sandia — and their academic partners. Under the latest Predictive Science Academic Alliance Program, NNSA awarded five universities $17 million each over five years to develop unclassified aspects of predictive modeling.

In its first 10 years, the ASC program built four of the world's fastest supercomputers. These include Los Alamos' Roadrunner, the first computer to break the petaflops barrier of 1 quadrillion calculations per second; and Blue Gene/L, the world's second-fastest computer. Many other top supercomputers are based on ASC designs.

These supercomputers are called "massively parallel" because they break problems down into smaller segments that run simultaneously on tens of thousands or even hundreds of thousands of individual processors. Duplication of effort enables massively parallel supercomputers to run thousands of times faster than supercomputers using a single processor with a higher clock speed. ASC also supported new ways to program models to take advantage of massively parallel processing power.

During its second 10 years, ASC plans to push computer speeds well into the petaflops range. On the software side, it will back the development of more accurate codes based on a better understanding of the underlying science. The program's Verification and Validation teams will ensure those codes truly reflect that understanding and produce the right results for the right reasons.

The Los Alamos National Laboratory CAVE uses 33 projectors to create a 43-million-pixel virtual reality environment. Researchers can simulate a problem, then walk through its visualization to detect trends, correlations, anomalies and unusual events.

Graphic provided by Los Alamos National Laboratory

By Thomas R. O'Donnell

Research Faces Up to Unstructured Data

TYPE A PERSON'S NAME INTO THE POPULAR IMAGE SEARCH **engine**

Google and the results are likely to vary wildly. You may find pictures of the person you're seeking, but you're also likely to see completely irrelevant images just because their name appears on the same web page.

You might have better luck if your computer could analyze a picture of the person you want, then search through millions of other images — even hours of videotape — to find someone who looks identical or similar. Ideally, the computer could match the faces regardless of whether the subject is in bright or low light, is only partially facing the camera or is near or far.

That's exactly what two Pacific Northwest National Laboratory (PNNL) researchers have done. Their algorithms analyze millions of video frames, pluck out the faces and quantify them to create searchable databases for facial identification. It's also fast, with codes shown to run at nearly optimal efficiency on everything from laptops to supercomputers.

"We're measuring the information content of a face much like Google" analyzes written web material, says Harold Trease, a PNNL computational physicist. "What they do for text searching we're trying to do for video and image processing."

A program that picks faces out of streaming or recorded video and identifies them regardless of conditions could be useful in many areas, but for Trease and Rob Farber, a PNNL senior research scientist, it's just a test case.

"It doesn't have to be webcams," Farber adds. "This is 'a first toe in the water' work" to prove the concept on massive amounts of unstructured data and high-performance computers. The algorithms could be generalized to work with almost any set of digital images, including X-ray and infrared pictures, and to identify a variety of objects, including hidden roadside bombs, features captured in satellite photos and tumors.

In fact, Trease had a biomedical application in mind when he delved into image processing. He wanted to create a "virtual lung" simulation and planned to "stack" two-dimensional CAT scan and MRI images into a "data cube" from which he could extract lung geometry data.

"We wanted to make that (extraction) process as fast as possible, but … the biomedical people couldn't supply us enough data, so we started to look at other images that would give us more data to practice on," Trease says. The researchers chose video and settled on faces because they were good targets for extraction.

It wasn't easy. "For computers, it's very difficult to recognize a face, first off, and then pull that out of unstructured data," where the light levels, size and angles of the faces change constantly, Farber says. For instance, humans typically have few problems recognizing people regardless of whether they're close or somewhat distant, but computers aren't as adept. So facial recognition algorithms must have "scale invariance" — the ability to pick a face out of video regardless of its distance from the camera.

Likewise, a successful algorithm must have a degree of "rotation invariance" — the ability to distinguish faces that aren't facing the camera head-on. And it must have "translational invariance" — the ability to extract faces or other target objects in a video even if they're moving within the frame.

"We have to have all three degrees of invariance so we're able to operate correctly in an unstructured environment," Farber says.

Then there is the mountain of data to sift. "If you want to look at all of YouTube and identify the faces, you're talking about a huge number of videos, each of which is comprised of 24 to 30 frames a second" and runs for minutes or hours, Farber says. "That translates into an astronomically large number of frames. To process that data it's very important that the algorithm scale well."

To do the job, the researchers combined Trease's work on data and pattern extraction with Farber's work on machine learning and pattern recognition.

The first part of the algorithm, largely Trease's work, starts with a raw red-green-blue (RGB) format video frame and transforms it to concentrate on the qualities of hue, saturation and intensity. The intensity parameter is discarded, allowing the algorithm to work regardless of lighting conditions in the image.

Next the algorithm sifts out facial "blobs" based on hue. "It turns out that skin color occupies a very narrow band in the hue dimension" regardless of race or ethnicity, Trease says, making them relatively easy to distinguish.

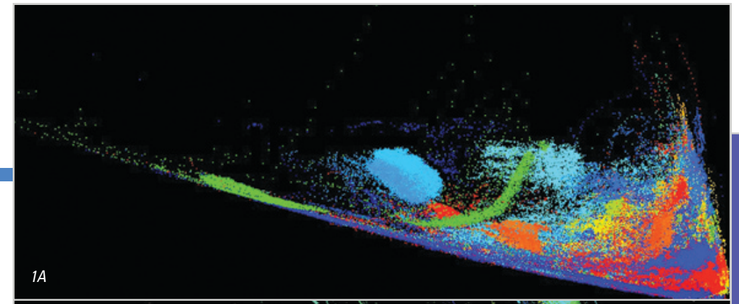

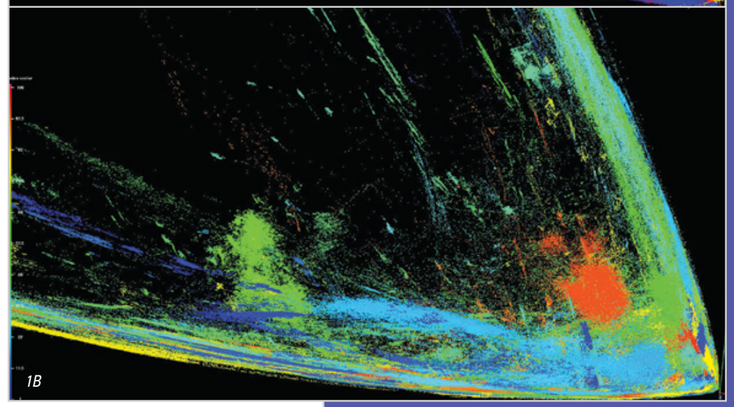

The algorithms identify patches of skin pixels, apply edge detection filters to separate overlapping faces, and compute the two-dimensional geometry of each blob based on the minimum and maximum number of pixels vertically and horizontally. The aspect ratio of the face inside this rectangle is computed and compared to the "golden ratio" artists use for faces. The algorithm computes the percentage of skin pixels that fills each rectangle.

FIGURE 1A: This image shows a larger view of 2.2 million frames of video data from the Supercomputing 2005 conference, plotted at their respective PCA coordinates. The color corresponds to the sequence number of the source video.

FIGURE 1B: This image shows a larger view of 22.6 million frames of YouTube data plotted at their respective PCA coordinates. The color represents the sequence number of the source video, ranging from blue (video #1) to red (video #513).

"We're measuring the information content of a face much like Google" analyzes written web material, says Harold Trease, a PNNL computational physicist. "What they do for text searching we're trying to do for video and image processing."

COLLABORATORS

Rob Farber is a senior scientist in the Environmental Molecular Sciences Laboratory at Pacific Northwest National Laboratory. He is recognized for his work in massively parallel computation, machine learning, complex dynamical systems and high-energy physics. Farber has more than 40 publications in books, peer-reviewed journals and other media. He has been a staff scientist in the theoretical division at Los Alamos National Laboratory (LANL), a member of the external faculty at the Santa Fe Institute, a staff member at the National Energy Research Scientific Computing Center (NERSC) and a successful company co-founder and consultant. Farber has been an invited speaker at both national and international conferences. He recently was credited with achieving peak performance on Ranger, the University of Texas supercomputer, in a feature article on the Texas Advanced Computing Center main page (http://www.tacc.utexas.edu), and was recognized in Science magazine for his role in laying the theoretical groundwork (along with Alan Lapedes) at LANL that enabled Derek Smith to produce insightful maps on evolution of the influenza virus (Science, 320 pp 310-311, 18 April 2008).

Harold Trease has more than 26 years of technical research experience in the design, implementation and application of high-throughput, high-performance computer software used for the characterization, classification, clustering and prediction of continuum and information physics. He currently leads the P3D Code Development Project at Pacific Northwest National Laboratory. P3D is a large-scale computational physics/information modeling, simulation and prediction framework with major software components that include digital image processing, computational geometry, multi-dimensional Voronoi/Delaunay mesh generation, time-dependent computational predictive solvers, and GUI-driven graphics/visualization.

Further Reading:

R. M. Farber, H. Trease, *Massively Parallel Near-Linear Scalability Algorithms with Application to Unstructured Video Analysis*, Proceedings of TeraGrid '08, June 2008, Las Vegas, Nevada, http://teragrid.org/events/teragrid08/Papers/papers/6.pdf.

R. M. Farber, A. S. Lapedes, R. Rico-Martinez and I. G. Kevrekidis, *Identification of Continuous-Time Dynamical Systems: Neural Network Based Algorithms and Parallel Implementation*, Proceedings of the 6th SIAM Conference on Parallel Processing for Scientific Computing, Norfolk, Virginia, March 1993, http://arxiv.org/PS_cache/comp-gas/pdf/9305/9305001v1.pdf

H. E. Trease, T. Carlson, R. Mooney, R. Farber, and L. L. Trease, *Unstructured Data Analysis of Streaming Video Using Parallel, High-Throughput Algorithms*, IASTED Signal and Image Processing, 2007:300-305, PNNL-SA-57622.

H. E. Trease, R. Farber, A. Wynne, L. Trease, *High-Performance Video Content Analysis Using Hybrid, Multi-Core Processors*, IASTED Signal and Image Processing, 2008. PNNL-SA-61871.

Contact:

Rob Farber
rob.farber@pnl.gov

Harold Trease
het@pnl.gov

PRACTICUM COORDINATOR

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

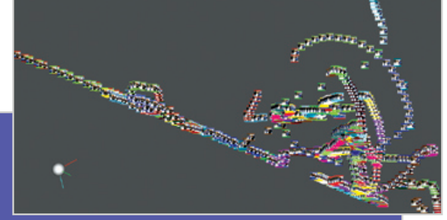

The three dominant principal components (derived from the PCA-based curve data) are used as the respective x-, y-, and z-coordinates to plot the frame in which each face appeared. Note that both trajectories (of the same face) and clustering are apparent.

Next, successive constraints based on Shannon entropy measures are applied to generate a 20-attribute signature characteristic of each face. The entropy measures calculate the images' information content. "We want to look at how much information each measure gives us, how many bits of information are contained between different pixels in the same region," Farber says. "We kind of slice and dice the picture in different ways, then we're able to combine all those slices together" into a signature.

It's impractical, however, to compare complex 20-dimensional signatures to identify faces or other objects extracted from a host of video frames. So the second part of the algorithm, largely based on Farber's research, reduces those 20 dimensions to just three.

"We crunch down to low dimensions and are able to do that with high accuracy — we don't introduce a lot of error between the different points," Farber says. That makes it possible to use similarity metrics that can distinguish between like faces. With high-dimensional data, every face can be considered similar in some aspect.

To do the job, the researchers start with principal component analysis (PCA), a tried and true technique for identifying the main directions and trends among a group of data points. "Principal components allows us to find the smallest set of linear combinations that lets us interpret the data," Farber says. "Using PCA, we can map out to a neural network type of architecture."

The process forces the 20-dimensional vector through a "bottleneck" of three linear "neurons" to find the three principal components that can accurately reconstruct the original signature. Data passes through the bottleneck to a set of output neurons that reconstruct the data. The process is repeated, reducing the error each time as the bottleneck neurons automatically learn the principal components.

Projecting from a high-dimensional signature down to three dimensions is "a wonderful situation to be in because we can actually plot those pictures in three dimensions," Farber says. With the three-dimensional coordinates, the program can create a trajectory tracking the subject's face as it moves in front of the camera. Faces can move out of the frame and back in, or from one camera to another, and the trajectory resumes, Farber says. It also could allow researchers to investigate social networks based on when and where faces appear together.

The algorithm could merely sample the frames in a video recording, rather than analyzing every one, "and do just fine," Trease says, "but as Google has found, the more data they get the better their searches are. It's true with us, too." The redundant information improves the program's search capability, letting it handle a variety of different situations, like a face viewed from an angle. "All of those different views let us develop a probability of detection which can range from 100 percent to whatever else," Trease adds. Researchers can apply uncertainty analysis, stating a threshold for the accuracy of matches they want to see. The program recognizes each three-dimensional signature under different conditions and forms a database, Farber says.

"That is an absolutely wonderful characteristic that allows us to do this facial recognition and answer that question of 'Have we seen this person's face before?'" he adds.

The approach has proven to be extraordinarily accurate and efficient. In one test using 2,000 pictures with known identities, the algorithm correctly identified all but two faces — although Trease notes that accuracy varies with image quality, resolution and other factors. And because Farber combined the neural network with a massively parallel mapping technique he pioneered in the 1980s, the program achieves high throughput with near-linear scaling — the amount of work the computer does rises in direct proportion to the number of processors employed. For instance, the researchers have run tests using a steadily increasing percentage of the 62,976 AMD Opteron Barcelona processor cores on Ranger, the Sun Constellation high-performance computer at the University of Texas. The data set included 2.2 million video frames captured at the Supercomputing 2005 conference and another 22.6 million captured from YouTube. To provide a fair measure of scaling behavior, the amount of data per core was kept constant as the number of processing cores increased.

With tweaking to optimize the code for Ranger and minimize communications between compute nodes, the algorithm has run at four flops per clock cycle per core — the theoretical peak performance for each core — on the part of the code dominated by floating point math operations. In late August 2008 the program ran at 363 trillion floating point operations per second (363 teraflops) using 60,000 of Ranger's cores. The researchers estimate the code will hit a blazing 380 teraflops when running on all of Ranger's cores.

Just because the algorithms run well on supercomputers doesn't mean they can't do as well on smaller machines, Farber notes. They also achieve near-linear scaling on inexpensive commodity hardware like NVIDIA graphics processing units (GPUs). Farber envisions low-power "smart sensors" that capture data, do initial data extraction and then compress the results and transmit them at low bandwidth for processing.

"For very low money we can get in the hundreds of gigaflops per GPU," Farber says. "We're able to do valuable work even on our laptops or workstations." The PNNL team's approach is readily generalized to a variety of computer architectures and also can handle "a huge number of different types of problems that may not be associated with vision recognition," Farber adds.

In fact, the potential applications are so wide that Farber seems a little overwhelmed.

"We have many different directions we're contemplating," he says. "It depends on what's going to be the best allocation of our available resources."

This montage shows the steps used to extract faces from video image data. Shown from the upper left to the upper right are the original frame, the RGB-to-HSI converted frame, the Sobel edge-detection filtered frame and an inverse color frame. The bottom row shows an RGB entropy frame, skin-colored pixel patches, the faces framed for easy visual identification and the final identified faces. These are added to the face database.

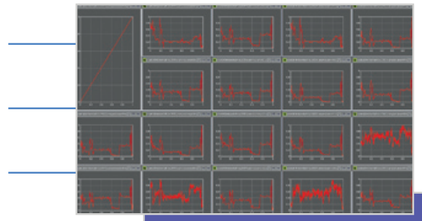

This figure illustrates 18 entropy vs. frame number plots (the first image in the upper left is not used) for a 2,107-frame segment of video from the Supercomputing 2005 conference.

PROGRAM MAY MEAN CUTTING THE TAGS

>> Image searches typically rely on tags – text humans have attached to the pictures to identify objects or people they depict. The algorithms PNNL scientists Rob Farber and Harold Trease have created could largely eliminate the need for tags because they recognize content automatically in massive amounts of data.

The application could make it as easy to index objects and people in hours of video as it is for search engines to find text on millions of World Wide Web pages.

So it's no wonder that Farber and Trease have been in touch with Google and with YouTube, the video-sharing web site Google acquired in 2006.

The PNNL codes could improve the sites' image search functions, Trease says, but the companies are too busy to consider the prospect at present. "They don't really get time to look inside the data," he adds. "It's all they can do to manage the text and the images," which are growing exponentially. "They're certainly interested, but it's going to take awhile."

In the meantime, Farber and Trease are refining their algorithms to better identify objects – more difficult, in some ways, than extracting and identifying faces.

"Faces have a built-in context. Objects don't," Trease says. To demonstrate the difference, he'll sometimes use the algorithms to search for round yellow objects in video archives. The results include everything from golf balls to the sun. Such objects need context to disambiguate them; faces are unambiguous by nature.

Trease also works on training the algorithms to recognize events in video by focusing on the specific signature each one creates. For instance, computer programs may be taught to recognize security risks like cars parking in restricted areas or making U-turns.

Computing Puts
Catalysts on Track

By Thomas R. O'Donnell

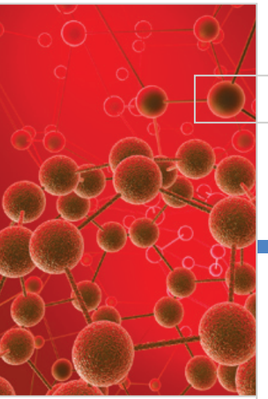

YOU COULD THINK OF **Larry Curtiss, Jeff Greeley and their Argonne National Laboratory colleagues as kind of a catalyst screening committee.**

Their computer models can quickly sift candidates for these important industrial and environmental materials, identifying the best ones for chemists to try out in time-consuming laboratory tests. They also can help scientists better understand results from such tests. The combination of computation and experiments is helping accelerate the development and improvement of catalysts.

Computational modeling of chemical processes like catalysis has grown in importance in the 33 years since Curtiss joined Argonne. "The capability has really tremendously increased" as the Department of Energy boosted high-performance computing at Argonne and other labs, says Curtiss, leader of the Theory and Modeling Group and an Argonne Distinguished Fellow. When that power and improved algorithms are combined, "It's unbelievable compared to 30 years ago … how much faster it can get done and the larger systems with more atoms (researchers can model), but also the accuracy with which we can calculate the properties of molecules and materials."

Now, "We can actually make predictions that can help to guide the experimentalists in making catalysts important to reducing the nation's energy dependency," Curtiss adds.

Catalysts increase the rate of a chemical reaction but are not consumed in the process. They're everywhere in industry, helping to efficiently produce a multitude of materials and chemicals. They're in your car's muffler, helping cut pollutants. They enhance reactions that generate power in a fuel cell, help convert grains, cellulose and animal and vegetable oils to biofuels and help clean polluted soil and water. Good catalysts cut the heat and energy needed to cause a reaction. That — and the fact catalysis can reduce the creation of wasteful byproducts in chemical reactions — also makes them easier on the environment than other chemical processes.

Scientists seek several properties in an improved catalyst, Curtiss says. It has to be selective, breaking certain chemical bonds but not others. If the catalyst isn't selective, it could generate unwanted byproducts that might have to be removed. A good catalyst also has to break those bonds in a way that uses less energy than the process scientists wants to replace. The new catalyst also must be stable, so it doesn't degrade, and be as inexpensive as possible.

The research falls under Argonne's Center for Nanoscale Materials almost by default, says Greeley, a materials scientist at the center. "In any industrial process, the catalysts are essentially on the order of nanometers," he adds. "The urgent question is how the size and shape of these particles affect their catalytic properties. In many cases you can obtain significantly improved properties by finding the right size or shape."

Chemists typically have used a time-consuming trial and error approach to finding good catalysts, Curtiss says. "Oftentimes, people don't really understand how they work," he adds. With the computational models he, Greeley and Argonne materials scientist Peter Zapol have devised, "We're trying to really understand at the molecular and atomic level what is happening for these catalysts. If you can understand what's happening, you can design new catalysts that are more efficient and more selective."

Computational modeling of chemical processes like catalysis has grown in importance in the 33 years since Curtiss joined Argonne. "The capability has really tremendously increased."

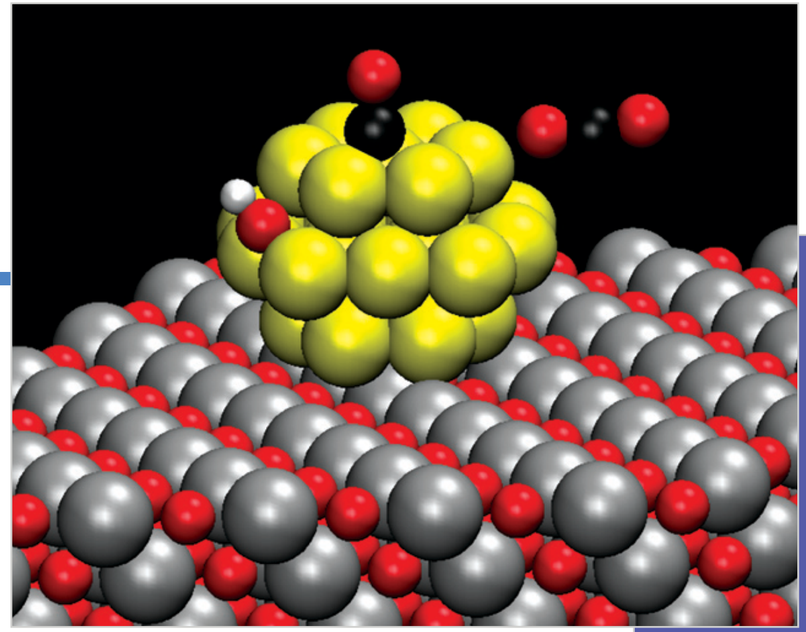

A supported catalytic nanoparticle of the kind the Argonne researchers model using first principles algorithms. Adding support material, like aluminum oxide, to the simulation greatly increases the computational cost.

The codes Greeley, Curtiss and Zapol develop and use are based on first principles: They calculate what's happening in a catalytic reaction based on fundamental physical properties, so they don't have to fit the model to experimental data and can make fewer assumptions or estimates about what's happening.

"That is an extremely powerful technique, in my opinion, but I admit I'm a little bit biased," Greeley jokes. With a first-principles approach, the computer model can calculate the energy needed for a catalytic reaction, known as the activation energy, and predict which catalysts lower that energy barrier the most and with the best selectivity. The first-principles methods can study all the elemental reactions between atoms "and figure out which steps are more energetically favorable," Greeley says. That's difficult to understand using purely experimental approaches.

"The first-principles strategy by its very nature is molecular-scale modeling," he adds. "It gives you lots of insight into what the molecules and atoms are doing on the surface of these catalysts."

First principles also helps understand the effect of reaction conditions, like temperature and pressure — another aspect of catalysis that is gaining more attention in research circles, Greeley says. When those environmental properties change, "That can sometimes cause a real change in the structure of the catalyst," he adds. First principles can help decipher how those changes affect the catalysis mechanism and the activation energy.

The level of detail first principles provides comes at a high cost in computational power, however, and it would be impossible without major computing access. Compute time "increases quite rapidly with the number of electrons that you're considering and the number of atoms you're considering in the calculations." Curtiss says. "As far as doing first principles, it limits the number of atoms that can be included" in the simulation.

COLLABORATORS

Larry A. Curtiss is a group leader and Argonne Distinguished Fellow in the Materials Science Division and Center for Nanoscale Materials at Argonne National Laboratory. He earned a bachelor's degree in chemistry from the University of Wisconsin-Madison in 1969 and a doctoral degree in theoretical chemistry from Carnegie Mellon University in 1973 and joined the lab in 1976. Curtiss' research focuses on computational chemistry, including development of new quantum chemical methods and applications to problems in materials science and chemistry. Specific areas of interest include nanocatalysis, nanocrystalline materials, computational thermochemistry, electron transfer processes and nanoporous materials. He has over 280 publications and is listed as a Highly Cited Researcher by the Institute for Scientific Information (ISI) in Chemistry for the period 1980-1999. His many awards and honors include the University of Chicago Distinguished Performance Award in 1995 and elevation to fellow of the American Association for the Advancement of Science in 1997. He has served on review panels, is on the Executive Committee of the Institute for Energy and Catalytic Processes of Northwestern University and is on the editorial board of the Materials Science and Engineering C.

Jeffrey Greeley did his doctoral research in the chemical and biological engineering department at the University of Wisconsin-Madison under Manos Mavrikakis. He was a postdoctoral researcher under Jens Nørskov at the Technical University of Denmark for two years before joining the Theory and Modeling Group at Argonne's Center for Nanoscale Materials as an assistant scientist. His current research focuses on the use of density functional theory techniques to understand the kinetics and thermodynamics of chemical reactions on nanostructured metal and oxide surfaces. His group's goal is to develop models of reactions relevant to heterogeneous catalysis and electrocatalysis. At the same time, the group will develop computational screening techniques that will permit the efficient design of new catalytic materials from first principles. Selected reactions of interest include the carbon monoxide electrooxidation reaction, the selective oxidative dehydrogenation of propane to propene and the electrooxidation of small organic molecules. A related interest is the development of new first principles-based techniques to model the dissolution and corrosion of metallic surfaces.

Further Reading:

D. Strmcnik, D. Tripkovic, D. van der Vliet, K. C. Chang, V. Komanicky, H. You, J. Greeley, V. Stamenkovic and N. M. Markovic, *Unique activity of platinum ad-islands in the CO electrooxidation reaction*, Journal of the American Chemical Society, 130, 15332 (2008).

P. Strasser, S. Koh and J. Greeley, *Voltammetric Surface Dealloying of Pt bimetallic nanoparticles: An experimental and DFT computational analysis*, Physical Chemistry Chemical Physics, 10, 3670 (2008).

S. P. Adiga, P. Zapol and L. A. Curtiss, *Structure and morphology of hydroxylated amorphous alumina surfaces*, Journal of Physical Chemistry C 111 7422 (2007).

P. C. Redfern, P. Zapol, M. Sternberg, S. P. Adiga, S. A. Zygmunt and L. A. Curtiss, *Quantum Chemical Study of Mechanisms for Oxidative Dehydrogenation of Propane on Vanadium Oxide*, Journal of Physical Chemistry, B 110 8363 (2006).

Contact:

Larry A. Curtiss
curtiss@anl.gov

Jeffrey Greeley
jgreeley@anl.gov

PRACTICUM COORDINATOR

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

For instance, the researchers may easily model a platinum catalyst composed of clusters of four or eight atoms. But if the simulation includes the support material, like aluminum oxide, on which the platinum is deposited, "the computational time really starts increasing," Curtiss says. Increasing the cluster to nanometer size or working with a two-element catalyst like vanadium oxide boosts the calculation time even further.

The researchers tap several computers to run their catalyst codes, including Carbon, a 1,000-plus-cores cluster at the Center for Nanoscale Materials capable of up to 10 teraflops (trillion floating point operations per second); Jazz, a 350-node cluster at Argonne; and computers at the National Energy Research Scientific Computing Center (NERSC) near San Francisco and the Pacific Northwest National Laboratory in Washington state.

First-principles codes are difficult to run efficiently on more than a few dozen processors, so the researchers have yet to make much use of the Argonne Leadership Computing Facility Blue Gene/P, one of the world's fastest for open computing. However, Greeley and colleagues at the Technical University of Denmark are working on a first-principles code that scales to thousands of processors *(see sidebar on page 45)*. "It looks promising," Curtiss says. "We hope in the near future to start using Blue Gene."

The researchers have plenty to do in the meantime. They're often sought out by scientists from Argonne's Catalysis Center, from other institutions or even from industry. The lab scientists sometimes have an experimental result they don't completely understand, Greeley says, and they want simulations to provide insights into the molecular reactions behind it. "That's something first-principles applications are very good at," he adds.

In one recent case, a group led by Argonne materials scientist Nenad Markovic was studying a fuel cell reaction that oxidizes carbon monoxide and converts it to carbon dioxide. They had found that tiny islands of five to 500 platinum atoms on the fuel cell anode were highly active in carrying out the reaction. With computational analysis, Greeley and his fellow researchers were able to show the islands stabilized a critical intermediate reaction. "We were able to explain the speedup of the carbon monoxide oxidation reaction on these important nanostructural features," Greeley says.

The "screening committee" approach comes into play when experimentalists have a list of proposed catalysts and want to find the best candidates for testing. "In some cases it could be hundreds," Curtiss says, with different combinations of metals and atomic configurations. Computer models offer a quick, relatively inexpensive way to narrow the list.

A subnanometer platinum cluster is shown on an oxide support. Such clusters may possess remarkable catalyic properties for catalytic reactions involving fuel cells and petrochemicals processing.

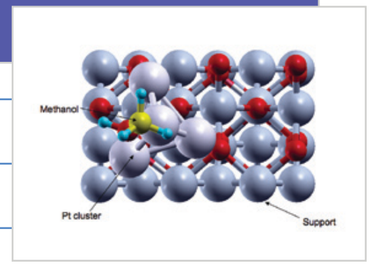

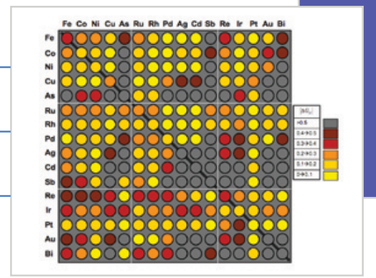

An example of how first-principles calculations are used to identify improved catalysts. Each circle represents an alloy and the color of the circle indicates how well the alloy is predicted to perform for the hydrogen evolution reaction, an important electrochemical reaction. The best catalysts are shaded yellow.

For instance, Greeley worked with Thomas Jaramillo of the Technical University of Denmark to study binary alloy catalysts for a reaction to produce hydrogen from protons and electrons. The computers were able to study the structures of around 750 alloys and predict which one could improve on the state of the art catalyst. In tests, Jaramillo's group found the predicted combination worked well, Greeley says.

More and more often, collaborations are flowing from theory to experiment, Greeley says. He and his fellow researchers are using their models to make predictions, then taking them to experimental scientists for testing.

Greeley frequently focuses on electrocatalysis, in which an electric potential like a voltage is applied to the catalytic process. "By changing the voltage you can … significantly affect the reaction," using it as a sort of "tuning knob" to control the rate and products produced, he adds. Electrocatalysis is especially important in fuel cell research.

The group also studies heterogeneous catalysis, including research on catalysts that would efficiently break down cellulose for ethanol production and would cut the energy needed to convert propylene into propylene oxide, a common ingredient in plastic packaging.

Curtiss says the group still is in the early stages of developing models with predictive power. To help the effort, the researchers also are compiling a database of simulation results.

"If we're looking for a new catalyst with certain properties, we would be able to access this database of calculations that's already been done to help find what that catalyst would be," Curtiss says. Additional calculations may be needed, but the database will provide insights to zero in on the best targets.

The Argonne researchers also are continuing to refine the algorithms behind their models. Combined with new, more powerful experimental tools like the bright x-ray beams at the lab's Advanced Photon Source, the simulations will provide new tools for improved catalysis.

"What's impressive is the tremendous progress that's been made in computational theory modeling, and how experimentalists are really looking for it to explain what's happening, as well as to guide them in designing new catalysts" and other materials, Curtiss adds.

A MATTER OF FIRST PRINCIPLES AND PARALLEL COMPUTATION

>> Jeff Greeley and Larry Curtiss have made strides in modeling catalysts using first-principles methods and teraflops-capable computers, but they have yet to fully deploy their biggest weapon: Argonne National Laboratory's Blue Gene/P.

Blue Gene may be rated one of the world's most powerful computers for open science (according to the TOP 500 benchmarking list), but the first-principles codes the researchers use to model catalysts generally don't run well on its tens of thousands of PowerPC processors.

"Typically, what people have done is to do first principles calculations on small systems. They might use 20 or 30 CPUs," materials scientist Jeff Greeley says. "What people have found is if you go beyond that size, performance declines rapidly." That limits the size of simulations the researchers can run to just a handful of atoms.

First principles calculations typically are done in reciprocal, or Fourier, space, Greeley says, but that approach often inhibits parallelizing the algorithms. To tackle the problem, he's collaborating with Jens Nørskov of the Technical University of Denmark to develop a "real space" algorithm that, in general, solves the differential equations associated with electronic structure calculations in a grid of points in space.

"Real space is very amenable to parallelization over lots and lots of processors," Greeley says. The approach requires developing new algorithms, but also sorting through libraries of computational methods already available. "It's not quite as sexy, in a scientific sense, to do those kinds of computational science optimizations," he adds. "Traditionally, electronic structure codes have been a little bit behind other codes" when it comes to parallelization.

In tests, the new first-principles codes have run well on more than 5,000 processors in some cases, Greeley says. "Our immediate goal is to actually do simulations of catalytic nanoparticles up to 3 nanometers in diameter — about 1,000 atoms or so," he adds. "That is beyond the limits of what anybody's been able to do so far, at least with first-principles calculations, but it's also exciting because that is a size range … where you see very interesting changes in the catalytic reactivity."

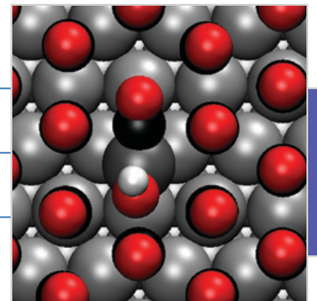

The oxidation of carbon monoxide (CO) is shown on a platinum adatom supported on a much larger platinum substrate. First-principles calculations have helped explain why these structures are much more active for CO oxidation than are normal platinum catalysts.

Ryan Elliott

The bizarre behavior of metals that can SHIFT SHAPE, only to metamorphose back to their original form, is by turns intriguing and, perhaps, slightly intimidating to the untrained eye. But understanding such alloys, often called "intelligent" or "shape memory" materials, poses a singular challenge to Ryan Elliott, a scientist who himself deftly shifts from discussions of materials science to physics to multi-scale simulation techniques.

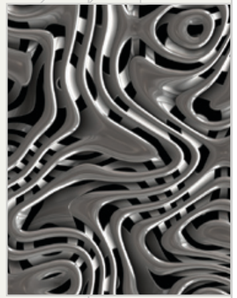

"As a graduate student at the University of Michigan I was struck by the properties of these exotic materials," says Elliott. "My advisors were interested in modeling these materials by simulating them from the atomistic level, and when I got the (DOE Computational Science Graduate) Fellowship, we were able to pursue that."

What started as a graduate research project became a full-fledged research program in 2005, when Elliott received a faculty appointment in the department of aerospace engineering and mechanics at the University of Minnesota.

"Fundamentally, my simulations try to describe the way that atoms exert forces on each other when they are brought together in a crystal," he says. "What I am predicting are the different crystal structures or phases that are possible, for a particular material, based on a model of how its atoms apply forces to each other."

Essentially, Elliott's computer models describe the physical properties that allow a metal — an alloy of titanium and nickel for example — to shift from one crystal shape to another at low temperatures and revert faithfully to the original shape when heated. His work is stretching the boundaries of elasticity theory to incorporate behaviors that current theory does not sufficiently support.

While the research is basic in the sense that Elliott predicts atomic-level forces, companies ranging from aerospace and automobile manufacturers to surgical and medical device makers are eager to develop applications for shape memory materials.

One of the most successful applications for shape memory materials has been in aerospace, where a "patch" was developed to repair leaks that appear in airplane hydraulic and fuel lines. When a line springs a leak, a cold shape-memory alloy sheath is installed over the line. As it warms to room temperature, the sheath contracts, forming a seal that efficiently repairs the leak.

"They work incredibly well," Elliott says. "To my knowledge, these patches have never failed."

Some shape memory alloys are biologically inert and have been used successfully for orthodontia, reducing the need for frequent adjustment of braces, and for self-expanding coronary artery stents — hollow mesh tubes that are used to keep arteries open after a balloon angioplasty procedure.

Despite their promise, the cost and time-consuming nature of devising reliable new materials is hampering development of shape memory alloys. Elliott's simulations are providing a theoretical underpinning to shape memory behavior and with it the potential to discover entirely new materials through computer simulation. The research is so promising that Elliott was awarded a five-year National Science Foundation (NSF) Faculty Development (CAREER) grant, the foundation's most prestigious early career award.

"Currently, materials scientists are good at predicting how these materials react when they make a small change," he says. "The goal of the computational work I am doing is to allow these scientists to make bigger changes and be able to accurately predict what will happen to the properties of the alloy."

Indeed, Elliott's work may one day reshape our understanding of these mysterious materials.

His work is stretching the boundaries of elasticity theory to incorporate behaviors that current theory does not sufficiently support.

Richard Mills

Richard Mills likes to SCALE MOUNTAINS. His web site, which shows a smiling Mills atop verdant Mount Taylor in the San Mateo Mountains of New Mexico, is a testament to his drive to conquer challenges, be they geologic or computational.

Now in his third year as a staff computer scientist at Oak Ridge National Laboratory (ORNL), Mills works at the leading edge of high performance computing, spending much of his time working with ORNL's Jaguar system.

Such computational power demands applications well suited to use it, and Mills contributes to several projects designed to make the most of Jaguar's capabilities. He is a core member of the multi-institutional PFLOTRAN team, which is developing code to simulate multi-scale, multi-phase, multi-component flow and reactive transport in porous media using machines ranging from laptops all the way to ultrascale computers.

In addition to work on PFLOTRAN, Mills occasionally contributes to development of the PETSc framework, and recently developed a universal interface that allows programmers to seamlessly transfer memory pointers between PETSc and Fortran 90.

"One of the first things I had to do to get PFLOTRAN to work on Jaguar was to develop a way to translate C pointers to memory allocated by PETSc into the Fortran 90 pointers that PFLOTRAN requires," says Mills. "It took me a few months, but working with the PETSc team at Argonne National Laboratory I came up with a universal PETSc/Fortran 90 interface that does not require you to know the internal representation that the Fortran 90 compiler uses for the pointer. It's not very glamorous, but it is an example of the kind of work that somebody has to do to make codes run on these low serial-number supercomputers."

Mills' code is now a part of Argonne's PETSc library and is available to anyone using the PETSc system.

In other work that involves scaling new heights, Mills is contributing to environmental data mining projects supported separately by the U.S. Forest Service and DOE's Atmospheric Radiation Measurement Program. He assists scientists who are trying to wade through terabytes of environmental data, works with climate scientists to help fine-tune climate models and works with ecologists on an early-warning system identifying forests that are under environmental stress from factors such as insect damage or climate change. Mills uses parallel data mining techniques to sift through information that would be impossible to interpret manually.

"With the Forest Service, we are trying to look at, in near real-time, satellite data to determine in an automated fashion places where forest ecosystem health might be threatened," Mills says. "It would be very helpful because the Forest Service can conduct only so many aerial surveys or send only so many personnel out in trucks to look for things such as parasite infestation. It's a difficult problem because of the sheer number of variables and the fact that we have to be able to differentiate between expected changes in ecosystem state – due to factors such as seasonal cycles – and unexpected changes that might indicate that the forest ecosystem is in trouble."

Mills uses similar techniques to compare climate model output and observations to determine which observed atmospheric states are never captured by a model, as well as to identify when a climate model is "hallucinating" — simulating climate states that are never observed in actual data. The goal of that project is to improve the Community Climate System Model used to predict future climate change and its impacts.

"There were a number of numerical steps my environmental scientist colleagues needed to do, such as principal components analysis, that had become a real bottleneck for them. I was able to develop parallel tools that enabled my colleagues to do in minutes what used to take hours or days using serial processing," he says.

Mills contributes to several projects designed to make the most of Jaguar's capabilities.

ALUMNI PROFILE

CATHERINE NORMAN

Catherine Norman

In graduate school at Northwestern University, Catherine Norman conducted COMPUTER SIMULATIONS to determine how gas bubbles behave in fluids moving under precisely defined conditions. After a few weeks on the job at the Center for Naval Analyses, a government-funded think tank in Alexandria, Virginia, her bubble had burst.

One of her first assignments was to analyze U.S. Navy ship movements and critique the efficiency of training exercises. Soon she was asked to devise a plan to reduce fuel consumption by optimizing travel between bases and refueling stations — all on deadline and working with inexact information.

"In graduate school I was doing my own little computational fluid model and I had complete control over absolutely every parameter that went into my model," she says. "I knew exactly what that data coming out meant, because I knew exactly what went into the model. This is very different. We are working with incomplete information and often times we don't understand where our information comes from. That's why we often go into the field, to understand how military operations work."

Before long, Norman was landing in Camp Victory, Baghdad, for a seven-month deployment to analyze Improvised Explosive Devices (IEDs) — roadside bombs — for the Joint IED Defeat Organization, a multidisciplinary group trying to stay one step ahead of the bomb-makers. Norman's job was to be a resource for military commanders, providing day-to-day updates based on information gathered from multiple sources.

"My job was to be available to answer questions about trends in tactics and materials so they could make decisions based on the best available information," she says. "It was amazing just being over there and listening to the people who were out in the field. If you just sit in an office looking at numbers all day, you really don't have any sense of perspective."

Norman's analyses are no theoretical simulations; they could literally be a matter of life and death. Heady stuff for a computational scientist, but Norman always had a bent toward foreign affairs. Her undergraduate major was international politics. She is fluent in German and Russian and has worked as a translator, transcribing reports from the German Ministry of Defense into English.

"I like working in an atmosphere of urgency, where you have time-sensitive analysis that can make a difference in the real world," she says. "My graduate training gave me the ability to face a large-scale problem and not be intimidated by it. And I am used to working in a collaborative group with people of differing backgrounds to solve complex problems."

Norman hasn't left her computational training behind. She has written a couple of scripts to solve linear optimization problems and put together a small Visual Basic script to parse data for one of her group projects.

"Some of the others in the group thought it was really impressive, but it was only a hundred lines of code," she says. "It wasn't like writing code to solve a partial differential equation."

"I like working in an atmosphere of urgency, where you have time-sensitive analysis that can make a difference in the real world."

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.

A researcher in drug design and analysis has been named the 2008 Frederick A. Howes Scholar in Computational Science.

Dr. Mala Radhakrishnan is an assistant professor of chemistry at Wellesley College, where she heads a research group that uses computational techniques to design and analyze drugs and other biologically relevant molecules. Radhakrishnan was a fellow from 2004 to 2007, when she graduated from the Massachusetts Institute of Technology with a doctoral degree in physical chemistry.

The annual award goes to at least one participant in the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) with an outstanding record of scholarship and character. They must have completed all the requirements for their degree under fellowship support or must have been supported by the program for the maximum number of allotted years.

The award honors Howes, who managed the Applied Mathematical Science Program in the U.S. Department of Energy from 1991 until his death in 1999. Howes was highly respected and admired for his energy, dedication and personal integrity.

One of Howes' duties was to oversee the DOE CSGF. He was extremely committed to this program, which supports graduate students in computational science. It takes a unique approach, requiring candidates to take courses in mathematics, computer science and an applications discipline, such as physics or engineering. The DOE CSGF currently supports more than 65 graduate students and is administered by the Krell Institute.

Radhakrishnan received her award, which includes a plaque and a substantial cash payment, at the 2008 DOE CSGF Conference in Washington, D.C. She also delivered a talk on her research, "The Many Roles of Computational Science in Drug Design and Analysis."

Lawrence Livermore National Laboratory's David Brown presents the Howes Scholar award to Dr. Radhakrishnan.

Dr. Radhakrishnan gives a presentation to the attendees of the DOE CSGF annual meeting.

Alumni Directory

A

Matthew Adams
University of Washington
Computational Electromagnetics
Fellowship Years: 2007-2008
Current Status: Alumnus

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

Erik Allen
Massachusetts Institute of Technology
Chemical Engineering
Fellowship Years: 2004-2008
Current Status: Svaya Nanotechnologies

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, Rutgers University

Kristopher Andersen
University of California – Davis
Physics
Fellowship Years: 2001-2005
Current Status: Faculty, Northern Arizona University

Matthew Anderson
University of Texas
Physics
Fellowship Years: 2000-2004
Current Status: Faculty, Brigham Young University

B

Teresa Bailey
Texas A&M University
Nuclear Engineering
Fellowship Years: 2002-2006
Current Status: Lawrence Livermore National Laboratory

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
Current Status: Faculty, Asian Institute of Technology

Paul Bauman
University of Texas
Computational & Applied Mathematics
Fellowship Years: 2003-2007
Current Status: Staff, University of Texas

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

Bonnie Carpenter Beyer
University of Illinois at Urbana-Champaign
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: Staff, New York University

Dean Brederson
University of Utah
Computer Science
Fellowship Year: 1996
Current Status, Staff, University of Utah

Paul Bunch
Purdue University
Chemical Engineering
Fellowship Years: 1994-1997
Current Status: Merck & Co. Inc.

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

Michael Bybee
University of Illinois at Urbana-Champaign
Chemical Engineering
Fellowship Years: 2004-2008
Current Status: Student, University of Illinois at Urbana-Champaign

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: AREVA

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: Vitamin D, Inc.

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

Joshua Coe
University of Illinois at Urbana-Champaign
Chemical Physics
Fellowship Years: 2001-2002
Current Status: Los Alamos National Laboratory

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: Faculty, University of Missouri – Columbia

William Conley
Purdue University
Mechanical Engineering
Fellowship Years: 2003-2008

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

Nathan Crane
University of Illinois at Urbana-Champaign
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

Aron Cummings
Arizona State University
Electrical Engineering
Fellowship Years: 2003-2007
Current Status: Student, Arizona State University

Joseph Czyzyk
Northwestern University
Industrial Engineering & Management
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: Los Alamos National Laboratory

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

Jimena Davis
North Carolina State University
Applied Mathematics
Fellowship Years: 2004-2008
Current Status: US Environmental Protection Agency

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

John Dolbow
Northwestern University
Theoretical & 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: Dataspora, Inc.

Jeffrey Drocco
Princeton University
Biophysics & Computation
Fellowship Years: 2004-2008
Current Status: Student, Princeton University

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

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

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

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

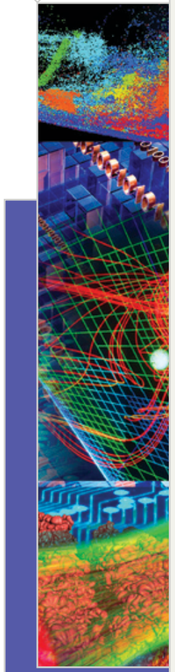

Thomas Epperly
University of Wisconsin
Chemical Engineering
Fellowship Years: 1991-1995
Current Status: Lawrence Livermore National Laboratory

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

F

Matthew Fago
California Institute of Technology
Aeronautical Engineering
Fellowship Years: 2000-2003
Current Status: Research Scientist, Colorado Springs

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

Matthew Farthing
University of North Carolina
Environmental Science & Engineering
Fellowship Years: 1997-2001
Current Status: USACE ERDC Coastal & Hydraulics Laboratory

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

Krzysztof Fidkowski
Massachusetts Institute of Technology
Computational Fluid Dynamics
Fellowship Years: 2003-2007
Current Status: Faculty, University of Michigan

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

Jasmine Foo
Brown University
Computational Fluid Dynamics
Fellowship Years: 2004-2008
Current Status: Sloan Kettering Institute

Gregory Ford
University of Illinois at Urbana-Champaign
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: Johns Hopkins Medical Institutions

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 & Neural Systems
Fellowship Years: 1998-2000
Current Status: GCI, Inc.

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

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

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: 1996-2000
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 at Urbana-Champaign
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 at Urbana-Champaign
Electrical Engineering
Fellowship Years: 1994-1997
Current Status: Intel

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

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

H

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

Glenn Hammond
University of Illinois at Urbana-Champaign
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: Oak Ridge National Laboratory

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 at Urbana-Champaign
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

Amber Sallerson Jackson
University of North Carolina – Chapel Hill
Applied Mathematics
Fellowship Years: 2004-2008
Current Status: Student, University of North Carolina – Chapel Hill

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: Staff, Carnegie Mellon University

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

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

Peter Kekenes-Huskey
California Institute of Technology
Computational Chemistry/Biology
Fellowship Years: 2004-2007
Current Status: Student, California Institute of Technology

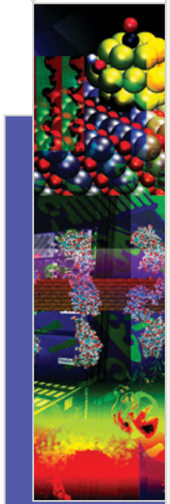

Jeremy Kepner
Princeton University
Computational Cosmology
Fellowship Years: 1993-1996
Current Status: Staff, Massachusetts Institute of Technology Lincoln Labs

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

Bonnie Kirkpatrick
University of California – Berkeley
Computer Science
Fellowship Years: 2004-2008
Current Status: Student, University of California – Berkeley

Justin Koo
University of Michigan
Aerospace Engineering
Fellowship Years: 2000-2004
Current Status: Air Force Research Laboratory

Michael Kowalok
University of Wisconsin
Medical Physics
Fellowship Years: 2000-2004
Current Status: Arizona Oncology Services

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: Boston Consulting Group

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: Krell Institute

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

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

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: Useful Bias, Inc.

Randall McDermott
University of Utah
Chemical Engineering
Fellowship Years: 2001-2005
Current Status: National Institute of Standards & Technology (NIST)

Matthew McGrath
University of Minnesota
Physical Chemistry
Fellowship Years: 2004-2007
Current Status: Peace Corps Volunteer, Cameroon

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

Matthew McNenly
University of Michigan
Aerospace Engineering
Fellowship Years: 2001-2005
Current Status: Lawrence Livermore National Laboratory

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

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

Julian Mintseris
Boston University
Bioinformatics
Fellowship Years: 2001-2005
Current Status: Staff, Harvard Medical School

Erik Monsen
Stanford University
Aerospace & Astronautical Engineering
Fellowship Years: 1991-1993
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: Ames Laboratory

Elijah Newren
University of Utah
Mathematics
Fellowship Years: 2001-2005
Current Status: Sandia National Laboratories – New Mexico

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

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-Vanhoucke
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: Staff, Princeton University

O

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

P

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

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

Ian Parrish
Princeton University
Computational Plasma Physics
Fellowship Years: 2004-2007
Current Status: Staff, Princeton University

Tod Pascal
California Institute of Technology
Physical Chemistry
Fellowship Years: 2003-2007
Current Status: Student, California Institute of Technology

Virginia Pasour
North Carolina State University
Biomathematics
Fellowship Years: 1998-1999
Current Status: Staff & Faculty, University of California – Los Angeles

Christina Payne
Vanderbilt University
Chemical Engineering
Fellowship Years: 2003-2007
Current Status: Staff, Washington Division of URS

Robert (Chris) Penland
Duke University
Biomedical Engineering
Fellowship Years: 1993-1997
Current Status: Biopharmaceutical & Medical Device Consultant

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

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

David Potere
Princeton University
Demography / Remote Sensing
Fellowship Years: 2004-2008
Current Status: The Boston Consulting Group

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

Q

Alejandro Quezada
University of California – Berkeley
Geophysics
Fellowship Year: 1997

Catherine Grasso Quist
Cornell University
Bioinformatics
Fellowship Years: 2000-2004
Current Status: Student, University of Michigan

R

Mala Radhakrishnan
Massachusetts Institute of Technology
Physical Chemistry
Fellowship Years: 2004-2007
Current Status: Faculty, Wellesley College

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

Nathan Rau
University of Illinois at Urbana-Champaign
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, Duke University

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, Stanford University

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

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

Mark Rudner
Massachusetts Institute of Technology
Physics
Fellowship Years: 2003-2007
Current Status: Staff, Massachusetts Institute of Technology

S

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

Samuel Schofield
University of Arizona
Applied Mathematics
Fellowship Years: 2001-2005
Current Status: 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
Computational Turbulence
Fellowship Years: 1997-2002

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

Jason Sese
Stanford University
Computational Materials Science
Fellowship Years: 2003-2005
Current Status: Environmental Consulting Company

Elsie Simpson Pierce
University of Illinois at Urbana-Champaign
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
Current Status: At Home

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: Faculty, California State University – Long Beach

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

Samuel Stechmann
New York University
Applied Mathematics
Fellowship Years: 2003-2007
Current Status: Staff, New York University

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 Corporation

Laura (Painton) Swiler
Carnegie Mellon University
Engineering & Public Policy
Fellowship Years: 1992-1994
Current Status: Sandia National Laboratories – New Mexico

T

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

Brian Taylor
University of Illinois at Urbana-Champaign
Engineering Mechanics
Fellowship Years: 2003-2007

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

William Triffo
Rice University
Bioengineering
Fellowship Years: 2003-2007
Current Status: Student, Rice University

Mario Trujillo
University of Illinois at Urbana-Champaign
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

Michael Veilleux
Cornell University
Computational Fracture Mechanics
Fellowship Years: 2004-2008
Current Status: Student, Cornell University

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 Atomic Power 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: Faculty, Louisiana Tech University

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

Stefan Wild
Cornell University
Operations Research
Fellowship Years: 2005-2008
Current Status: Staff, Argonne National Laboratory

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, Inc.

Allan Wollaber
University of Michigan
Nuclear Engineering, Fission Concentration
Fellowship Years: 2004-2008
Current Status: Argonne National Laboratory

Michael Wolf
University of Illinois at Urbana-Champaign
Scientific Computing/Computer Science
Fellowship Years: 2003-2007
Current Status: Student, University of Illinois at Urbana-Champaign

Matthew Wolinsky
Duke University
Earth Surface Dynamics
Fellowship Years: 2001-2005
Current Status: Staff, National Center for Earth-surface Dynamics

Brandon Wood
Massachusetts Institute of Technology
Computational Materials Science
Fellowship Years: 2003-2007
Current Status: Student, Jawaharlal Nehru Centre for Advanced Scientific Research

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

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

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

Etay Ziv
Columbia University
Computational Biology
Fellowship Years: 2004-2008
Current Status: Mount Sinai Hospital

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

John ZuHone
University of Chicago
Astrophysics
Fellowship Years: 2004-2008
Current Status: Student, University of Chicago

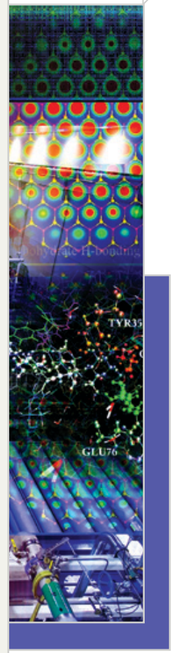

Joshua Adelman
University of California – Berkeley
Biophysics

Advisor:
George Oster
Practicum:
Lawrence Livermore National Laboratory
Contact:
jadelman@berkeley.edu
Research Synopsis:
During my graduate studies I intend to develop computational methods to elucidate the connection between structure and function in ring-shaped molecular motors, and in doing so provide a detailed understanding of the broader question of how mechanochemical coupling and cooperativity conspire to allow proteins to do work within the cell.

Zlatan Aksamija
University of Illinois at Urbana-Champaign
Biophysics Electrical Engineering

Advisor:
Umberto Ravaioli
Practicum:
Los Alamos National Laboratory
Contact:
aksamija@uiuc.edu
Research Synopsis:
My work focuses on examining the generation of heat in ultrascaled silicon MOSFET and FinFET devices using self-consistent Monte Carlo device simulation with full electron bandstructure from non-local pseudopotentials and a full phonon dispersion computed from the Adiabatic Bond Charge model. We have devised an efficient algorithm for the inclusion of full phonon dispersion in order to account for anisotropy and details of heat transport with great accuracy. We have also computed the density-of-states (DOS) and the lattice thermal energy numerically and used them to generate maps of local temperatures in a representative short-channel MOSFET device.

Jordan Atlas
Cornell University
Chemical Engineering

Advisor:
Michael Shuler
Practicum:
Los Alamos National Laboratory
Contact:
jca33@cornell.edu
Research Synopsis:
I am working on a dynamic modeling framework to integrate genomic detail and cellular physiology within functionally complete hybrid bacterial cell models. An initial step in this approach is the development of a whole-cell coarse-grained model which explicitly links DNA replication, metabolism, and cell geometry with the external environment. A hybrid model can then be constructed from chemically-detailed and genome-specific subsystems, called modules, inserted into the original coarse-grained model. We use the sensitivity analysis of the original coarse-grained model to identify which pseudo-molecular processes should be de-lumped into molecularly-detailed mathematical modules to implement a particular biological function.

Christopher Carey
University of Wisconsin
Plasma Physics

Advisor:
Carl Sovenic
Practicum:
Lawrence Berkeley National Laboratory
Contact:
cscarey@wisc.edu
Research Synopsis:
Many astrophysical bodies, such as supernovas and accretion disks, consist of magnetized ionized gas. Although scientists have been studying these objects for many years, due to a lack of direct access very few measurements of these bodies are available, and the mechanisms of their formation remain an unanswered question. I am interested in using computational models to relate direct measurements from laboratory plasmas to observational measurements of large scale astrophysical plasmas. Astrophysical plasmas exist at temperatures and densities that are starkly different from those of laboratory plasmas. Although they exist in a vastly different regime, laboratory plasmas do exhibit some of the same fundamental phenomena as laboratory plasmas. I will extrapolate experimental data to astrophysical applications using magnetohydrodynamical (MHD) models to understand how parameter scaling affects these observations.

Ethan Coon
Columbia University
Applied Mathematics

Advisor:
Marc Spiegelman
Practicum:
Los Alamos National Laboratory
Contact:
etc2103@columbia.edu
Research Synopsis:
Many interesting and important geophysical features at the Earth's surface are direct expressions of the dynamics of the solid earth below. In addition to the complex rheologies inherent to the material, current work suggests small-scale features such as localized faulting and melt transport in areas on plate boundaries can have important implications for heat and mass transfer. My research will look at computational and analytic techniques for modeling these multi-scale, multi-physics problems in Earth dynamics.

Jeff Hammond
University of Chicago
Theoretical & Computational Chemistry

Advisor:
Karl Freed
Practicum:
Pacific Northwest National Laboratory
Contact:
jhammond@uchicago.edu
Research Synopsis:
My current research focuses on molecular property calculations within the coupled-cluster approximation. Using the response formalism, which is just time-dependent perturbation theory, dynamic properties such as hyperpolarizabilities and Raman cross-sections can be calculated efficiently. By implementing these methods within the massively-parallel quantum chemistry program NWChem, molecules previously too large for study at this level of accuracy are now within reach.

Asegun Henry
Massachusetts Institute of Technology
Nanoscale Heat Transfer and Energy Conversion

Advisor:
Gang Chen
Practicum:
Sandia National Laboratories– New Mexico
Contact:
ase@mit.edu
Research Synopsis:
My research uses molecular dynamics (MD) to investigate phonon transport in nanostructures. MD simulations involve calculating the trajectories of particles within a system by numerically integrating classical equations of motion. Using the trajectories the properties of a system can be extracted. MD simulations can supply accurate descriptions of the heat transfer mechanisms within nanostructures because surface, interface and quantum effects can all be included, making it an excellent tool for analyzing nanoscale processes.

Kevin Kohlstedt
Northwestern University
Bio-Polymer/Soft Matter Computation

Advisor:
Monica Olvera de la Cruz
Practicum:
Lawrence Berkeley National Laboratory
Contact:
kohlstedt@northwestern.edu
Research Synopsis:
I hope to characterize and describe phenomena, in the context of self-assembled charged, bio-oligomers, that occur due to the competition between long (coloumbic) and short range (van der Walls) forces at a coarse grained level of description. My first project is to look at pattern formations on the surface of cylinders and describe new patterns with anisotropic Wigner cells.

Miler Lee
University of Pennsylvania
Genomics and Computational Biology

Advisor:
Junhyong Kim
Practicum:
Sandia National Laboratories – New Mexico
Contact:
miler@mail.med.upenn.edu
Research Synopsis:
I am working on developing general, low-bias RNA classification schemes based on features of the folded RNAs, in an attempt to define a structure space (and hence, function space) onto which to map the various classes of RNAs. I hope to apply this work to address several questions pertaining to noncoding RNAs using large-scale computational approaches. For example, what is the relationship among the different varieties of RNA-mediated silencing strategies? Or to what extent can a particular class of RNAs be derived from genomic sequences that are ostensibly non-RNA related? Careful attention to folded structure is sure to reveal various subtle aspects of RNA-mediated regulation.

Jeremy Lewi
Georgia Institute of Technology
Neuroengineering

Advisor:
Robert Butera
Practicum:
Lawrence Livermore National Laboratory
Contact:
jlewi@gatech.edu
Research Synopsis:
The primary application of my research is for adaptively optimizing neurophysiology experiments. The goal is to use active learning to significantly reduce the number of trials needed to characterize neural responses using parametric statistical models. The response functions of neurons are often estimated by measuring the neural responses to different stimuli and then applying techniques from system identification. The error of the estimated response function can potentially be minimized using fewer trials by choosing "optimal" stimuli. In this context, the optimal stimulus is the one which will provide the most information about the unknown response function. Information theoretic metrics allow the optimal stimulus to be defined rigorously.

David Markowitz
Princeton University
Computational Neurobiology

Advisor:
David Tank
Practicum:
Los Alamos National Laboratory
Contact:
dave@princeton.edu
Research Synopsis:
I am interested in the structure and function of microcircuits in neural systems. Currently, I am developing software to automate the 3D reconstruction of neurons from stacks of electron micrographs, with the long-term goal of mapping local network topology. My other research interests include dendritic information processing, neural coding and prediction.
Notable:
Lead author of an experimental paper accepted for publication in the June 17, 2008 issue of the prominent journal *Proceedings of the National Academy of Sciences (PNAS)*.

Peter Norgaard
Princeton University
Computational Plasma Dynamics

Advisor:
Clarence Rowley
Practicum:
Lawrence Berkeley National Laboratory
Contact:
norgaard@princeton.edu
Research Synopsis:
I am presently studying fast ion dynamics in the presence of large scale magnetohydrodynamic (MHD) instabilities. Experimental data indicates that the fast ion current redistributes during periods of high MHD mode activity. This influences the tokamak equilibrium profile and therefore affects the overall stability and confinement. Some of my other research interests include PDE numerical methods, model reduction via proper orthogonal decomposition /Galerkin projection, and applied dynamical systems theory for high dimensional systems.

Natalie Ostroff
University of California – San Diego
Systems Biodynamics and Computational Biology

Advisor:
Jeff Hasty
Practicum:
Oak Ridge National Laboratory
Contact:
nostroff@ucsd.edu
Research Synopsis:
The primary focus of my research in the field of Systems Biology is on the design, construction, and computational modeling of synthetic gene regulatory networks, with the ultimate goal of developing a quantitative and predictive model for cellular function. A gene regulatory network is the fundamental input-output device of a cell, responsible for determining which genes are expressed at a given time, how much product is made from each gene, and how the cell responds to external stimuli. A schematic of a gene regulatory network in many ways resembles a complex circuit diagram, and it is this analogy that drives the pursuit of a quantitative description of gene regulation. In the same way that an electrical engineer can describe the functionality of an electrical circuit with a set of equations, we are working towards a mathematical framework that will provide analogous information about a cell's functionality based on its gene network.

Christopher Schroeder
University of California – San Diego
High Energy Physics Theory

Advisor:
Julius Kuti
Practicum:
Brookhaven National Laboratory
Contact:
crs@physics.ucsd.edu
Research Synopsis:
My field of interest is particle physics, particularly Quantum Chromodynamics (QCD). QCD is the field theory of the Standard Model, which describes the interactions of the building blocks of the universe: electrons, neutrinos, quarks, and gluons. I intend to utilize lattice Quantum Chromodynamics (LQCD), the numerical study of QCD on a discrete spacetime lattice, to discover new laws of physics, both within the realm of the Standard Model and beyond.

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

Arnab Bhattacharyya
Massachusetts Institute of Technology
Computer Science
Advisor:
Madhu Sudan
Practicum:
Sandia National Laboratories – California
Contact:
abhatt@mit.edu

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

Julianne Chung
Emory University
Computational Mathematics
Advisor:
James Nagy
Practicum:
Lawrence Berkeley National Laboratory
Contact:
jmchung@emory.edu

Tal Danino
University of California – San Diego
Bioengineering
Advisor:
Jeff Hasty
Practicum:
Lawrence Berkeley National Laboratory
Contact:
tdanino@ucsd.edu

Jack Deslippe
University of California – Berkeley
Physics
Advisor:
Steven Louie
Practicum:
Oak Ridge National Laboratory
Contact:
jdeslip@berkeley.edu

John Evans
University of Texas
Computational & Applied Mathematics
Advisor:
Thomas Hughes
Practicum:
Sandia National Laboratories – New Mexico
Contact:
JohnAEvans@mail.utexas.edu

Ashlee Ford
University of Illinois at Urbana-Champaign
Chemical Engineering
Advisor:
Richard Braatz
Practicum:
Brookhaven National Laboratory
Contact:
anford2@uiuc.edu

Kristi Harris
University of Maryland – Baltimore County
Physics
Advisor:
Philip Rous
Practicum:
Sandia National Laboratories – New Mexico
Contact:
kristi.harris@umbc.edu

David Ketcheson
University of Washington
Applied Mathematics
Advisor:
Randall LeVeque
Practicum:
Sandia National Laboratories – New Mexico
Contact:
ketch@amath.washington.edu

Brian Levine
Cornell University
Transportation Systems
Advisor:
Linda Nozick
Practicum:
Sandia National Laboratories – New Mexico
Contact:
bl76@cornell.edu

Oaz Nir
Massachusetts Institute of Technology
Computational Biology
Advisor:
Bonnie Berger
Practicum:
Lawrence Berkeley National Laboratory
Contact:
oaz@mit.edu

Carolyn Phillips
University of Michigan
Applied Physics
Advisor:
Sharon Glotzer
Contact:
phillicl@umich.edu

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

David Rogers
University of Cincinnati
Physical Chemistry
Advisor:
Thomas Beck
Practicum:
Oak Ridge National Laboratory
Contact:
rogersdd@email.uc.edu

Ariella Sasson
Rutgers University
Computational Biology & Molecular Biophysics
Advisor:
Todd Michael
Practicum:
Sandia National Laboratories – New Mexico
Contact:
ariella@eden.rutgers.edu

Michael Sekora
Princeton University
Continuum Mechanics, PDE, Numerical Analysis
Advisor:
James Stone
Practicum:
Lawrence Berkeley National Laboratory
Contact:
sekora@math.princeton.edu

Benjamin Smith
Harvard University
Experimental High Energy Physics
Advisor:
Masahiro Morii
Practicum:
Lawrence Berkeley National Laboratory
Contact:
bcsmith@fas.harvard.edu

Benjamin Sonday
Princeton University
Applied & Computational Mathematics
Advisor:
Yannis Kevrekidis
Practicum:
Lawrence Berkeley National Laboratory
Contact:
bsonday@princeton.edu

Gregory Crosswhite
University of Washington
Physics
Advisor:
Dave Bacon
Contact:
gcross@u.washington.edu

Hal Finkel
Yale University
Physics
Advisor:
Richard Easther
Practicum:
Pacific Northwest National Laboratory
Contact:
hal.finkel@yale.edu

Robin Friedman
Massachusetts Institute of Technology
Computational & Systems Biology
Advisor:
Christopher Burge
Practicum:
Lawrence Berkeley National Laboratory
Contact:
robinf@mit.edu

Steven Hamilton
Emory University
Computational Mathematics
Advisor:
Michele Benzi
Practicum:
Los Alamos National Laboratory
Contact:
sphamil@emory.edu

Joshua Hykes
North Carolina State University
Nuclear Engineering
Advisor:
Yousry Azmy
Contact:
jmhykes@ncsu.edu

Milo Lin
California Institute of Technology
Physics
Advisor:
Ahmed Zewail
Contact:
miloiq@its.caltech.edu

Paul Loriaux
University of California – San Diego
Computational Biology
Advisor:
Alexander Hoffmann
Contact:
ploriaux@ucsd.edu

James Martin
University of Texas
Computational & Applied Mathematics
Advisor:
Omar Ghattas
Practicum:
Sandia National Laboratories – California
Contact:
duckofhealing@sbcglobal.net

Geoffrey Oxberry
Massachusetts Institute of Technology
Chemical Kinetics/Transport Phenomena
Advisor:
William Green
Practicum:
Sandia National Laboratories – New Mexico
Contact:
goxberry@mit.edu

Troy Perkins
University of California – Davis
Theoretical Ecology
Advisor:
Alan Hastings
Contact:
taperkins@ucdavis.edu

Matthew Reuter
Northwestern University
Theoretical Chemistry
Advisor:
Mark Ratner
Contact:
matthew-reuter@northwestern.edu

Sarah Richardson
Johns Hopkins University School of Medicine
Human Genetics & Molecular Biology
Advisor:
Joel Bader
Practicum:
Lawrence Berkeley National Laboratory
Contact:
notadoctor@jhmi.edu

Danilo Scepanovic
Harvard/Massachusetts Institute of Technology
Signal Processing/Modeling
Advisor:
Richard Cohen
Contact:
danilos@mit.edu

Paul Sutter
University of Illinois at Urbana-Champaign
Cosmology
Advisor:
Paul Ricker
Contact:
psutter2@uiuc.edu

Cameron Talischi
University of Illinois at Urbana-Champaign
Topology Optimization
Advisor:
Glaucio Paulino
Contact:
ktalisch@uiuc.edu

John Ziegler
California Institute of Technology
Aeronautics
Advisor:
Dale Pullin
Practicum:
Oak Ridge National Laboratory
Contact:
jackalak@caltech.edu

Carl Boettiger
University of California – Davis
Biology - Ecology & Evolution
Advisor:
Alan Hastings
Contact:
cboettig@gmail.com

Eric Chi
Rice University
Bioinformatics/Computational Biology
Advisor:
David Scott
Contact:
echi@rice.edu

Scott Clark
Cornell University
Applied Mathematics
Advisor:
Charles Van Loan
Contact:
sc932@cornell.edu

Curtis Hamman
Stanford University
Flow Physics and Computer Engineering
Advisor:
Parviz Moin
Contact:
cwhamman@stanford.edu

Ying Hu
Rice University
Biomedical Engineering
Advisor:
Rebekah Drezek
Contact:
hooying@gmail.com

Anubhav Jain
Massachusetts Institute of Technology
Materials Science & Engineering
Advisor:
Gerbrand Cedar
Contact:
anubhavj@mit.edu

Armen Kherlopian
Cornell University
Computational & Systems Neuroscience
Advisor:
Jonathan Victor
Contact:
ark2010@med.cornell.edu

Jeffrey Kilpatrick
Rice University
Computer Science
Advisor:
Luay Nakhleh
Contact:
Jeff.Kilpatrick@rice.edu

Kathleen King
Cornell University
Applied Operations Research
Advisor:
John Muckstadt
Contact:
kathleen.a.king@gmail.com

Eric Liu
Massachusetts Institute of Technology
Aerospace Engineering & Mathematics
Advisor:
David Darmofal
Contact:
ehliu@mit.edu

Brian Lockwood
University of Wyoming
Computational Fluid Dynamics
Advisor:
Dimitri Mavriplis
Contact:
blockwoo@uwyo.edu

Douglas Mason
Harvard University
Physics
Advisor:
Eric Heller
Contact:
douglasmason@gmail.com

Matthew Norman
North Carolina State University
Atmospheric Sciences
Advisor:
Frederick Semazzi
Contact:
mrnorman@ncsu.edu

Britton Olson
Stanford University
Computational Fluid Dynamics
Advisor:
Sanjiva Lele
Contact:
bolson@stanford.edu

Cyrus Omar
Carnegie Mellon University
Neural Computation
Advisor:
Brent Doiron
Contact:
cyrus@cmu.edu

Claire Ralph
Cornell University
Theoretical Chemistry
Advisor:
Garnet Chan
Contact:
ccr53@cornell.edu

Brenda Rubenstein
Columbia University
Theoretical Chemistry
Advisor:
David Reichman
Contact:
rubenstein.brenda@gmail.com

Anne Warlaumont
University of Memphis
Computational Developmental Psycholinguistics
Advisor:
David Kimbrough Oller
Contact:
anne.warlaumont@memphis.edu


DOE Staff