2007 - 2008

DEIXIS

THE DOE CSGF ANNUAL

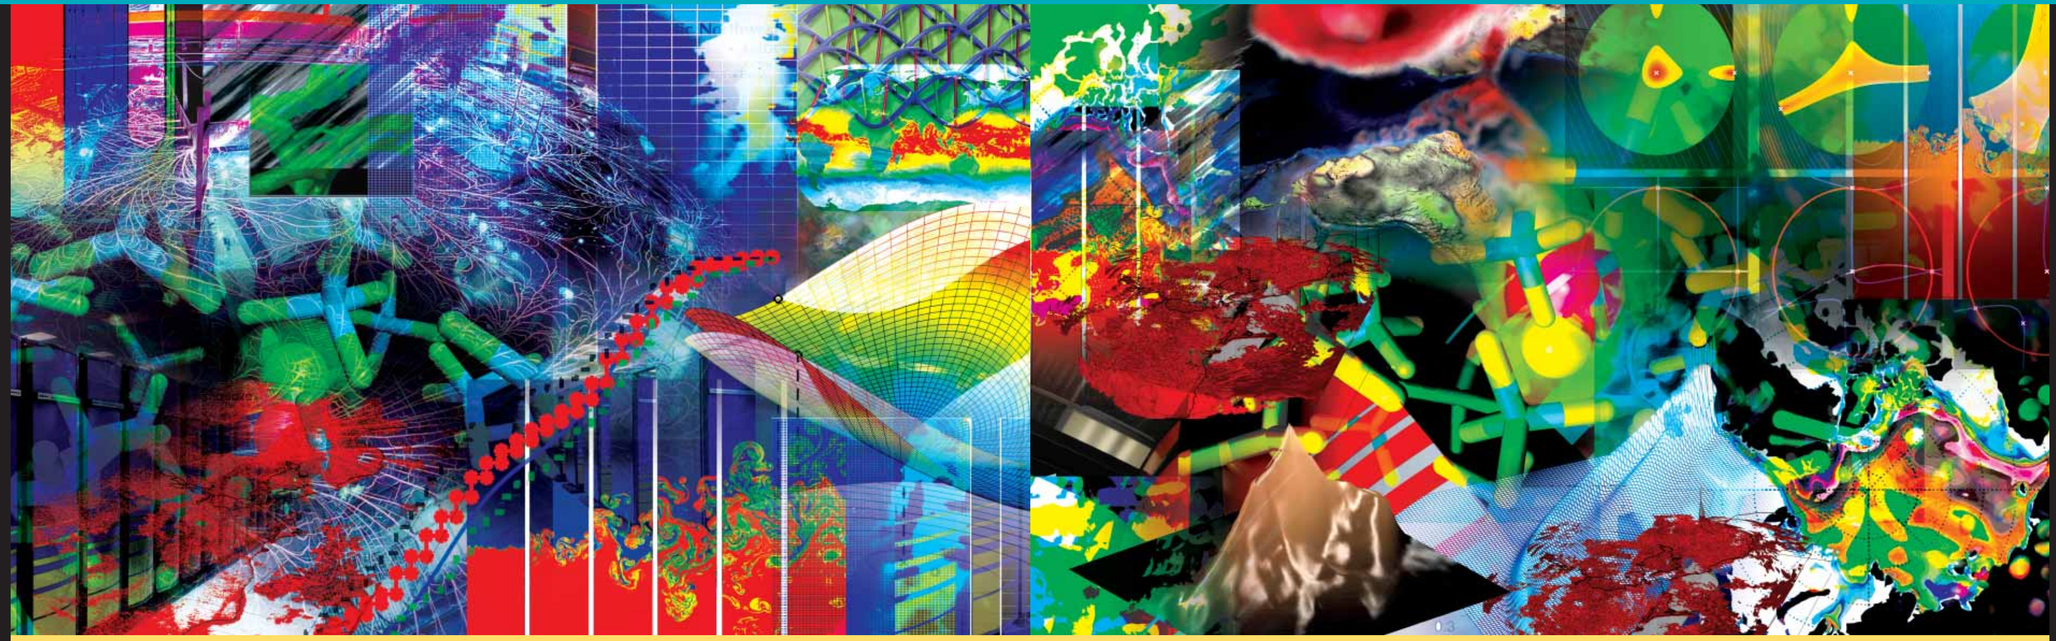

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

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

2007 - 2008

2007 - 2008

DEIXIS – THE DOE CSGF ANNUAL

DEIXIS – THE DOE CSGF ANNUAL

THE DOE CSGF ANNUAL

DEIXIS

TABLE OF CONTENTS

Department of Energy Computational Science Graduate Fellowship


PAGE 30


PAGE 22

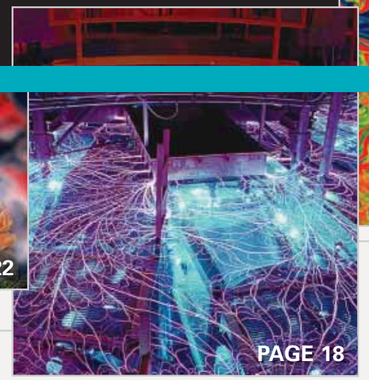
PAGE 18

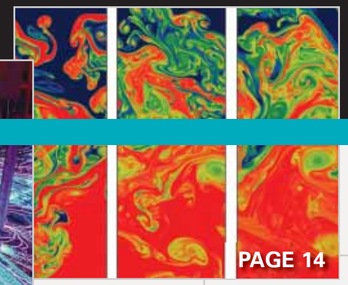
PAGE 14


PAGE 26

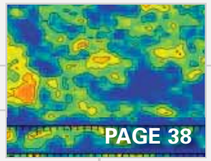
PAGE 38


PAGE 8

2007-2008

Editor
Shelly Olsan

Copy Editor
Ron Winther

Design
julsdesign, inc.

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

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

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

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

DOE CSGF

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

DEIXIS *is an annual publication of the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) program.* DEIXIS *illustrates work done at 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.*

This issue of DEIXIS *highlights research supported by the Department of Energy, Office of Science, Office of Advanced Scientific Computing and Research. This research represents an important part of the overall efforts by the Department of Energy labs that use computing to advance scientific discovery and engineering practice.*

The Summer Job That Stretches

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: Jeff Hammond, Stefan Wild, Ethan Coon and Jimena Davis.

BESIDES STUDIES AND RESEARCH at their home universities, each DOE CSGF participant must spend at least three months — usually in the summer — working with scientists at one of the Department of Energy's national laboratories. The practicum experience exposes the fellows to research distinctly different from their doctoral projects, stretching them with new challenges.

Some students return to their routines with new perspectives that improve their work. Others may take their doctoral studies in entirely new directions. And many decide they want to spend their careers at the national laboratories. Regardless, all of them are changed by their summer experience.

The Dilemma of Scale

ETHAN COON

Columbia University | Los Alamos National Laboratory | Story by Jacob Berkowitz

As every rock climber knows, survival is in the details. You can't think of the 3,000 foot-tall rock face while you're clinging to it. The only thing that matters is the half-inch-wide granite lip onto which your feet are wedged.

It's a perspective born of years of rock climbing that Ethan Coon brought to his summer practicum at Los Alamos National Laboratory in 2006. The DOE CSGF fellow took a page from his climbers' manual to make more physically accurate computer simulations of Earth processes, putting hand-grip level details into the big-picture view. And it turns out that the attention to detail that keeps climbers safe can help us better predict the consequences of earthquakes and squeeze more precious petroleum out of aging reservoirs.

As a student in applied mathematics at Columbia University, Coon works with those at the forefront of merging computational modeling and earth sciences. That includes his thesis adviser, Marc Spiegelman, an associate professor jointly appointed to Columbia's renowned Lamont-Doherty Earth Observatory and the Department of Applied Physics and Applied Mathematics.

In his doctoral research, Coon is creating a new mathematical model of the two-dimensional geometry of rock faults in order to improve computer models of earthquakes.

"Many standard physics-based computational models, including geological ones, are inherently bad at capturing small-scale effects," Coon says. "Like a postcard view, they can only show the broad strokes, while features smaller than these strokes are blurred out or completely missed."

It's not from lack of trying. Computational scientists continually work at putting the fine-scale into the big picture to create more detailed models. But they face a perennial hurdle: the increased time and cost it takes to run these more complex simulations.

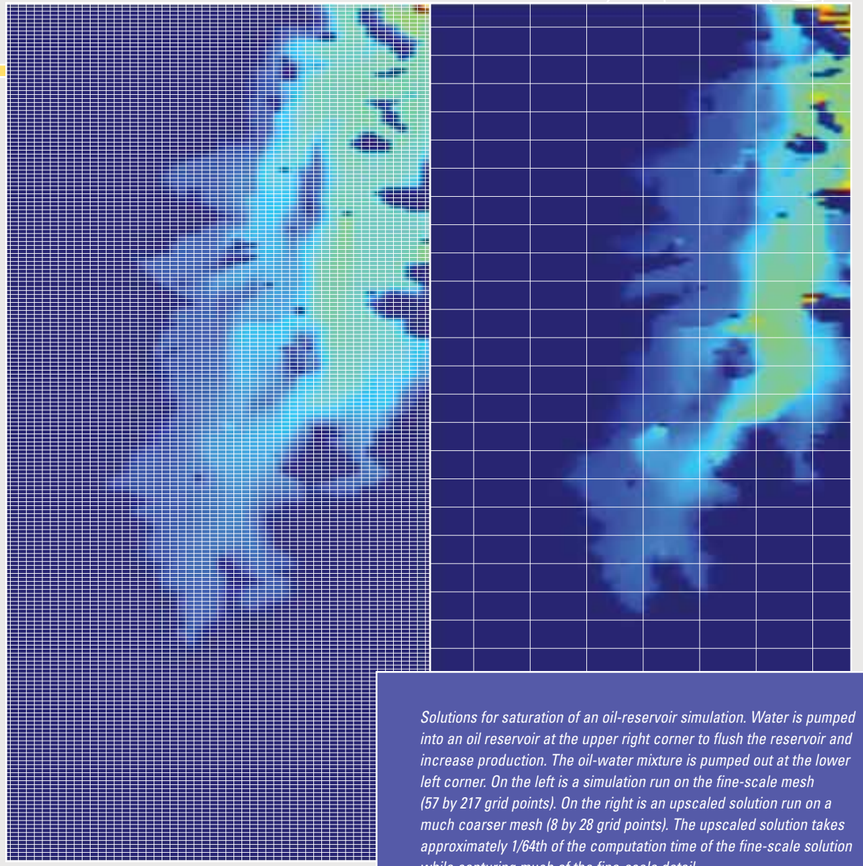

Solutions for saturation of an oil-reservoir simulation. Water is pumped into an oil reservoir at the upper right corner to flush the reservoir and increase production. The oil-water mixture is pumped out at the lower left corner. On the left is a simulation run on the fine-scale mesh (57 by 217 grid points). On the right is an upscaled solution run on a much coarser mesh (8 by 28 grid points). The upscaled solution takes approximately 1/64th of the computation time of the fine-scale solution while capturing much of the fine-scale detail.

Coon created an algorithm…to simulate the flow of two materials over time — an ability that's critical to simulating the interaction of oil and water in a reservoir.

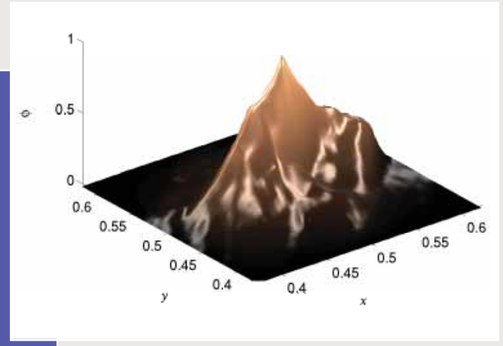

Our multilevel upscaling algorithm constructs a self-consistent hierarchy of coarse-scale models for single-phase saturated flow, as well as the corresponding multiscale basis functions, without solving any local or global fine-scale problems. The multiscale basis function for the center of the domain, shown in the figure, was generated using this algorithm. The fine-scale structure is clearly visible in the surface, which accurately represents the influence of this structure on the flow.

To tackle this challenge, computational scientists turn to approximations — they average out fine-scale details over the big-picture view. In modeling geological formations, this could mean averaging out tiny layers of sedimentary rock that compose a miles-long and miles-deep oil reservoir. Put into a climber's perspective, it's like averaging the number and locations of rocky hand-holds, without knowing exactly where each one is — a situation that can render the model useless as a real-life guide.

"This is one of the big challenges in trying to model real geological systems," says David Moulton, an applied mathematician with the Mathematical Modeling and Analysis Group (or T-7) at Los Alamos, and Coon's summer practicum adviser. "We obviously can't resolve everything, but we need to ensure that the approximations we introduce are useful. The key is to create approximations that capture the influence of unresolved features well enough to advance our understanding of the system and inform policy decisions."

For Moulton and Coon this means how to best extract an increasingly rare and valuable resource — oil. In a new well, the oil is under pressure from the overlying rock and spurts out of the ground like a geyser. However, there's still substantial oil left when the geyser stops. How to get it out? Oil companies often pump in water or carbon dioxide to flush out the remaining oil.

"Being able to increase the efficiency of these oil reservoirs has numerous advantages. It means more energy, it makes business sense, and it also minimizes the environmental impact of more drilling," Coon says.

To do this, it's critical to know how the oil and water flow in the rock, and that means understanding and modeling the fine-level details.

At Los Alamos, Moulton leads development of a state-of-the-art computational tool for improved modeling of flows through porous rock structures.

"Our Multilevel Upscaling (MLUPS) approach to modeling flow through heterogeneous media doesn't make the jump to a coarse-scale model directly from the fine-scale one, as do most existing approaches," Moulton says. "Instead it uses a multi-resolution approach, building a hierarchy of models, each one a baby-step up from the previous one. And in this way it is more effective."

When Coon arrived at Los Alamos, MLUPS was developed to the point it could model a single phase, or material, in a static state 15 times faster than the main competing technique. Coon created an algorithm enabling MLUPS to simulate the flow of two materials over time — an ability that's critical to simulating the interaction of oil and water in a reservoir.

In collaboration with Scott MacLachlan, a former Los Alamos summer student and (as of January 2008) an assistant professor at Tufts University, Coon also upgraded MLUPS so it incorporates mass conservation of the fluids, ensuring that the amount of material is the same at the start and end of a model run. While not always critical in initial model development, mass conservation is essential to making technology that can be practically applied.

Coon programmed the additional abilities into MLUPS in Python, a powerful computer language, and then added the components to the existing code, which was developed in the legacy language Fortran. Coon was introduced to Python through a collaboration with staff at Argonne National Laboratory as part of his doctoral research.

The work used real-world data from the Society of Petroleum Engineers' *Tenth Comparative Solution Project* — a baseline data set that exists as a testbed for developing and testing computational models. Moulton says MLUPS is now at the proof-of-concept, desktop stage for two-phase modeling. He's confident it will demonstrate significant advantages over competing models.

Coon presented the results of his summer research as an invited speaker at the March 2007 meeting of the Society for Industrial and Applied Mathematics (SIAM).

"It was a great experience. The other seven speakers in my session were all tenure-track professors, and I received great feedback," he says. Coon made contacts with researchers at the University of Texas at Austin and the University of Bath in England while there.

The DOE CSGF practicum also was a chance for Coon to see the fine scale in Los Alamos' scientific community.

"I really liked the ability to get out and talk with the engineers and scientists at the lab — people outside my immediate field — and to work on a larger, multidisciplinary project," says Coon, who notes that the T-7 group is noted for its highly collaborative approach.

While in New Mexico, Coon took time to get what geologists call "ground truth" with rocks. He scaled Sandia Mountain outside Albuquerque and also hiked to the summit of Mount Wheeler, the state's highest peak, near Taos. Both provided stunning vistas of the surrounding deserts and forests.

Similarly, from his scientific perch at Los Alamos, Coon was able to see the grander-scale picture that ties together his work using computational models to understand how the Earth — and things in it — moves.

"Talking with David Moulton about my thesis work really helped because he had a very different perspective on it," Coon says. "In the end I saw that both problems come down to a question of how you incorporate levels of detail that you can't completely resolve in your model."

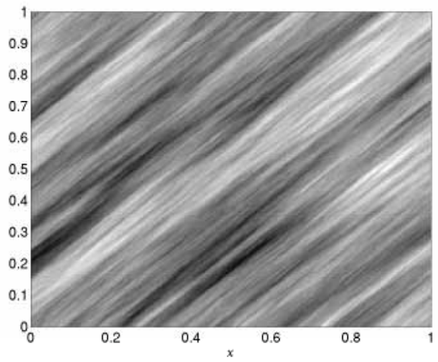

A geostatistical realization of a strongly heterogeneous permeability field with variation (from light to dark) of 6 orders of magnitude.

PROGRAM REQUIREMENTS

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

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

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

Prior to the third year of the fellowship, fellows must complete a practicum assignment at a Department of Energy laboratory. Currently, over 20% of program alumni work or have worked at a Department of Energy laboratory.

DISCIPLINES PURSUED

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

Applying Science

JEFF HAMMOND

University of Chicago | Pacific Northwest National Laboratory | Story by Jacob Berkowitz

Every day, while making the 15-minute drive to his summer practicum at the DOE's Pacific Northwest National Laboratory (PNNL) in southeast Washington State, Jeff Hammond passed the Hanford Site. Formerly used to produce plutonium for America's first nuclear weapons, the DOE's 586-square-mile reservation now is the nation's most contaminated nuclear waste site and home to the world's largest environmental cleanup.

The Hanford Site threw a new light on Hammond's summer internship. The University of Chicago doctoral student and DOE Computational Science Graduate Fellow had wanted to be a scientist since he was a kid growing up in Seattle. He was entranced by the pure intellectual intrigue of math, chemistry and theoretical physics. Later he found a love for big — really big — computers, which he could use to do chemistry without breaking lab glassware. But Hanford was the other face of science — where equations and algorithms meet the pavement.

As fate would have it, during his time at PNNL Hammond boosted the capabilities of DOE's premier computational chemistry tool, NWChem, so that chemists, biologists, nuclear scientists and even astrophysicists can more accurately, safely and reliably model not just nuclear wastes, but new medicines and even molecules light-years away.

"In academia, you don't have to think about the big picture all of the time — it's there when you write the grant but day-to-day it doesn't really matter," Hammond says. "But the DOE has a much more applications-driven mission. At PNNL I was inundated by the importance of applications, specifically in biology and nuclear sciences. When you're driving by tanks of nuclear waste each day that's quite the motivation to ensure that those tanks are still solid, decades from now, when your kids drive past."

Hammond, however, almost didn't make it to what he describes as "the best professional experience of my life." The job was to work on NWChem, but there was a problem: The 26-year-old Hammond wasn't familiar with the specific physics theory required for the project. And then there was the fact that he had almost no experience with computer coding for parallel supercomputers.

"No problem," said his adviser, Bert de Jong, a staff scientist in PNNL's Environmental Molecular Sciences Lab. "You can learn."

It was a heady introduction to the world of on-the-job training — one that Hammond relished.

"It was such an amazing experience to be in a work environment where, as opposed to academia, they actually have people whose job it is to help others do stuff they don't know how to do," Hammond says. "The computer people there, along with Bert and (PNNL staff scientist) Karol Kowalski, taught me basically everything I know about compilers, parallel computing, debugging and Fortran."

What he lacked in particulars, Hammond made up for in drive and broad quantum chemistry knowledge. After sadly leaving his new wife behind in Chicago, he settled into 90-hour work weeks fueled by lots of Mountain Dew.

His challenge was significant: Give NWChem the power to do highly accurate spectroscopy — to model the light emitted and absorbed by a molecule.

"Experimentalists study molecules spectroscopically," Hammond says. "But most computer codes can't calculate spectra directly. This is a big problem and makes the models less useful as predictive tools. What I did was give NWChem the ability to calculate the spectra of molecules based on their quantum chemical characteristics."

Hammond spent his first month at PNNL on a steep learning curve — reading scientific papers, deriving equations and talking with Kowalski to develop a solid working understanding of coupled-cluster theory. That's a well-established theoretical physics tool for approximately solving the Schrödinger equation; this enables scientists to predict a molecule's reactivity with other molecules or its response to a jolt from a laser beam.

"The math that we did can model the interaction of molecules with lasers," Hammond says. "But with a little more work it can also model the molecule's interaction with magnetic fields. This means modeling nuclear magnetic resonance, one of the most important tools biologists use to study the structure of large, complex proteins."

With Kowalski, Hammond developed the specific code to implement the coupled-cluster theory in NWChem and subsequently spent a month debugging it so it could run effectively on supercomputers, including the one located at DOE's National User Facility, the Environmental Molecular Sciences Laboratory at PNNL.

"The chemistry part of *computational chemistry* really does come last," Hammond says. "Half the time you're thinking about theoretical physics, another 45 percent of the time is spent getting the code to work correctly, and only then do you spend about five percent of the time actually thinking about chemicals."

When Hammond and others on the NWChem team did get to the chemistry, the results were impressive. They demonstrated the new code's ability to simulate the molecular properties of oxygen dichloride and carbon tetrachloride in solution. This is a crucial capability, because most pollutants and almost all biological molecules exist in solution, usually water.

"At the end of the day we're getting closer to giving biologists and nuclear scientists the ability to model things that are dangerous to study firsthand, or simply impossible to do so — for example, the interaction of uranium and other molecules in groundwater," Hammond says.

The researchers also modeled the spectra of a group of complex organic molecules that includes polyaromatic hydrocarbons, or PAHs. They ran the largest calculations of this kind ever performed and were able to determine the spectroscopic properties of these molecules at a level of accuracy not previously possible.

"The massively parallel codes developed by Jeff enable the full integration of the coupled cluster codes with NWChem's existing molecular dynamics module, and this is truly cutting-edge," PNNL's DeJong says. "For the first time ever, this merger will enable us to model the properties of the molecule with the full inclusion of surrounding environment. We believe that this will have a profound effect on the interplay between theory and experiment."

Hammond's contributions to NWChem will be available to the computational chemistry community in the next release of the software, expected in late 2007. The research resulted in at least four papers, with more to come as an ongoing collaboration expands in scope.

Back at the University of Chicago to continue his thesis work, Hammond says his time at PNNL and driving past Hanford "profoundly changed the way I'm going to do things in the future."

"By using existing methods and software, rather than starting from scratch, I can spend less time deriving equations and debugging code, and more time studying important chemical problems," he says.

In addition to his mostly–theoretical thesis research, Hammond also wants to talk to a University of Chicago spectroscopist about PAHs in space. These organic molecules produced by dying stars are plentiful in interstellar space, and are thought to be a key source of organic material for the origins of life. Hammond's work with NWChem could help astrobiologists identify PAHs in interstellar space — one more application of his science that could help all of us see the really big picture.

"The DOE has a much more applications-driven mission… When you're driving by tanks of nuclear waste each day that's quite the motivation to ensure that those tanks are still solid, decades from now, when your kids drive past."

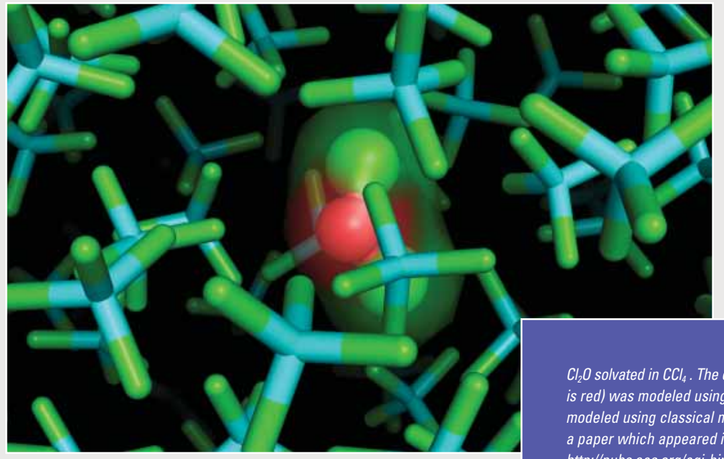

Cl_2O solvated in CCl_4. The oxygen dichloride molecule (in the center of the page, oxygen is red) was modeled using quantum mechanics, while the solvent molecules were modeled using classical mechanics. This picture represents the system studied in a paper which appeared in the Journal of Physical Chemistry A *on June 5th, 2007, http://pubs.acs.org/cgi-bin/abstract.cgi/jpcafh/2007/111/i25/abs/jp070553x.html.*

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.

Math Tools for Life

JIMENA DAVIS

North Carolina State University | Sandia National Laboratories – New Mexico | Story by Jacob Berkowitz

Seated on her apartment patio in Albuquerque, New Mexico, Jimena Davis looked west across the cactus-covered foothills and watched the setting sun splash pink light across the peaks of the Sandia Mountains. The sweet fragrance of desert sage wafted through the air. Compared to her usual surroundings, she might as well have been on a different planet, but she felt right at home.

It was the first time this Department of Energy Computational Science Graduate Fellow had visited or lived in part of the United States outside the East Coast. The summer practicum at Sandia National Laboratories that brought her to Albuquerque was like the mountains — they pulled her out of herself and made her think creatively, both personally and professionally. As a result, she's using math to contribute to computational biology tools that are enhancing our ability to mimic cellular processes, and thus better understand what makes life tick — including us.

Fished In

As a high school student in Mullins, South Carolina — population 5,000, with an economy rooted in centuries-old tobacco farms — it was clear Davis had a passion and a penchant for math. Her guidance counselors recommended a career in engineering and she agreed. But at Clemson University her freshman applied math professor pulled her aside one day after class and asked: *Have you considered majoring in math?*

"I was surprised. I didn't know I could major in mathematical science," Davis recalls. "The high school guidance counselors had never mentioned it."

Within weeks, Davis was a math major. During her last semester at Clemson her math path was further paved when, at a workshop sponsored by SAMSI, the Statistical and Applied Mathematical Sciences Institute, she met North Carolina State applied mathematician H.T. Banks. He spoke passionately about the use of applied mathematics on a variety of problems in science and engineering. It was exactly what she was looking for — the ability to make a difference with a subject she loves.

At NC State, Davis was quickly hooked on a long-standing computational biology model that Banks, now her advisor, had been angling for: understanding the population dynamics of mosquitofish.

"I'd never heard of a mosquitofish," laughs Davis, who's now contributed as much to understanding their population biology as almost anyone on Earth.

In rice fields in the United States and other major rice-growing countries, including India, mosquitofish are used as a natural mosquito control tool, reducing the need for pesticides. The challenge for biologists is to understand the mechanisms behind the growth of mosquitofish populations as a way to maximize the fishes' impact on mosquitoes. For example, when seeding a field with fish, what's the optimum fish-size distribution to use?

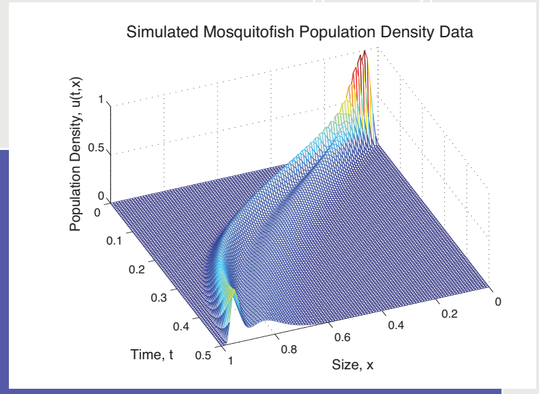

This figure shows the simulated mosquitofish population density data generated with a Bi-Gaussian growth rate.

"I'd never heard of a mosquitofish," laughs Davis, who's now contributed as much to understanding their population biology as almost anyone on Earth.

"What we know is that there isn't a single growth rate to describe the entire mosquitofish population but, rather, a distribution of growth rates. Individuals, and individuals at different sizes, grow at different rates, just as with humans," Davis says. "Estimating the distribution of the growth rates is impossible without computationally efficient approximation methods."

This problem, she notes, is an example of her specialty: inverse mathematical questions. These are ones in which the final solution is known and the challenge is to determine the parameters, or factors, that led to it — in this case the distribution of growth rates among the fish.

With the mosquitofish population model, she's not only improved the approximation methods for the growth rate distribution but also tested their reliability.

"We've been able to go a step further and compute confidence bands for these probability distributions," she says. "This is exciting because these techniques aren't just applicable to this problem, but to a whole range of probability estimation problems."

Bacterium Biology

At Sandia, Davis teamed with her practicum advisor, Elebeoba May, to tackle a more complex inverse biological computation problem — one that took her from fish populations to bacterium biochemistry.

May is leading an ambitious project to develop a large-scale systems biology simulation platform that can computationally model whole-cell, multi-cellular and host-pathogen systems at the molecular level. Called the *BioXyce Project*, it's based on modeling the flow of proteins and genetic molecules as if they were current flowing through an electrical circuit. This large-scale, parallel computing biocircuit simulation is one of the holy grails of computational biology — to mathematically model the hundreds of interacting genetic and protein pathways that constitute an organism's metabolism.

Davis' task was to improve a component of the system dedicated to computationally modeling the *E. coli* bacterium's central metabolic system. Once again, as with the mosquitofish work, Davis was thrown into deep waters.

"I had to hit the biochemistry textbooks to get a handle on this project. But I wanted to step outside the box of what I'd been doing at NC State and do something totally different," she says.

While the specifics were completely new, Davis saw that the essence of the problem was her specialty: an inverse problem — the mathematical modeling of biological systems in which there's a significant amount of data uncertainty.

As a first contribution, Davis reduced the level of unknown factors in the model by using a reformulation technique developed at the Lab to input known chemical reaction rates, derived from an on-line database.

Davis then coupled the reformulated model with DAKOTA, a computational optimization toolkit also developed at Sandia.

Initial desktop tests of the revised model show it's effective in parameter estimation using simulated data.

"The *E. coli* work was completely new to her, but in three short months she was able to significantly contribute to our work," May says. "Using DAKOTA and empirical data, she examined various computational approaches for the estimation of multiple-rate parameters for the *E. coli* central metabolic system. The results demonstrate the feasibility of this approach in determining reliable model parameters for systems biology applications."

Now completing her thesis at NC State, Davis hopes to move from fish and bacteria to humans.

"I want to work on problems such as HIV-AIDS or cancer research that will really be beneficial to people — that will make a difference in people's lives," she says.

And while she's helping shape the big picture in computational biology, she also wants to go back to high schools to talk about how math can directly shape lives.

"I want to mentor young female students," Davis says. "To tell them '*Yes, you can major in math. Yes, you can be a mathematical scientist.*' "

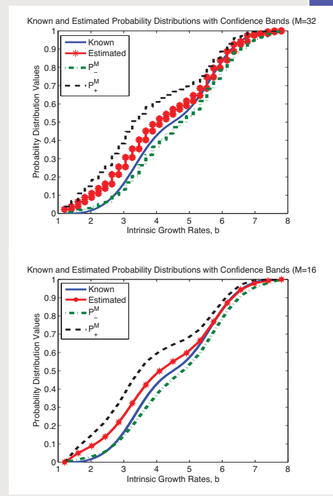

Top: Here is an example of the results obtained from the inverse problem for the estimation of the growth rate distribution using the delta function approximation method with 32 delta functions. The known distribution is Bi-Gaussian and is shown in the solid blue line. The estimated probability distribution is shown in red. Also shown in this figure are the confidence bands (lower in green and upper in black) quantifying the uncertainty associated with this estimated distribution.

Bottom: Here is an example of the results obtained from the inverse problem for the estimation of the growth rate distribution using the spline based approximation method with 16 piecewise linear spline elements. The known distribution is Bi-Gaussian and is shown in the solid blue line. The estimated probability distribution is shown in red, while the upper and lower confidence bands are shown in black and green, respectively.

Green Silicon

STEFAN WILD

Cornell University | Argonne National Laboratory | Story by Jacob Berkowitz

Along with being home to Cornell University, Ithaca, New York is famous for the Moosewood Restaurant, America's best-known vegetarian eatery.

The Moosewood's fame is largely based on its spectacular and creative cookbooks. The cookbooks are renowned for complex recipes that both taste and look great. The problem is that for busy cooks with 20 minutes to whip up something for a 5:30 family dinner, the recipes are impractical — fun to read, but often impossible to cook.

It's a conundrum Stefan Wild understands well — in a computational science sense — and he's rolling up his sleeves in the computational kitchen to help solve the problem. During a summer practicum at Argonne National Laboratory and in his doctoral research at Cornell, Wild, a Department of Energy Computational Science Graduate Fellow (DOE CSGF), is working to streamline cordon-bleu-complex computer models and make them stove-top fast. The results will have green impacts of another kind: environmental engineers use the computer models he's improving to find the best ways to stop the spread of pollution and to clean polluted groundwater.

Big, Not Fast

Even with advances in computer hardware and modeling software, computational scientists often face a central hurdle: when the model bites off more than computers can easily chew. Many models are so detailed, contain so much information, and are so complex that they become technically unwieldy. The mathematical recipes, or algorithms, that make up the model are beautiful, but take too long to run to get timely results within budget.

"We call these computationally expensive models," Wild says. "They often give researchers headaches because the model evaluation expense makes many of their favorite analysis tools ineffective in practice."

Enter the world of model optimization. Model optimizers are applied mathematicians, like Wild, who are passionate about getting the most from cumbersome computational models.

"It's not about making the models less detailed," Wild says. "The original modeler is thinking of replicating some physical phenomenon, not the larger goal of using the model to improve some system. If they were to make the model faster, they'd try and make it simpler. What we're doing is looking for mathematical ways to keep the detail, but boost the speed." For Wild, optimization is about bringing math to the people — to turn an unusable model into a dinner-table standard.

It's why he sought out his doctoral advisor, Cornell engineering professor Christine Shoemaker. She's a world leader in the application of sophisticated computations to solve environmental problems and, like Wild, takes a "math with a mission" approach.

It's also why Wild's DOE CSGF practicum stint at Argonne National Laboratory, outside of Chicago, was so valuable.

The Laboratory for Advanced Numerical Simulations in Argonne's Mathematics and Computer Science Division is a world leader in optimization technologies. The lab's Toolkit for Advanced Optimization (TAO) is a collection of high-quality, high-performance codes, primarily for distributed computing applications, used by hundreds of researchers within DOE, industry and academia in the United States and beyond.

Wild spent the summer of 2006 at Argonne, working with practicum advisor Jorge Moré to create a new code for the optimizer's toolkit: an algorithm specifically designed for engineers with computationally expensive models.

"The optimization technique that Stefan is developing is different from most," says Moré, an Argonne staff scientist who develops algorithms and software for large-scale optimization problems, such as modeling nuclear energy production.

"Most of the optimization techniques we're currently using require additional information from the user," Moré says. "The technique that Stefan's developing doesn't require additional inputs, and this makes it much simpler and more user-friendly."

Computationally expensive models involve as many as dozens of parameters, or variables. For example, in assessing the best way to clean contaminated groundwater in on-site remediation using wells, pumps and purifying techniques, the parameters include the rate of pumping, the changing levels of contamination, and groundwater movement in response to the pumping.

Some computationally expensive models can involve a hundred or more parameters. The more parameters, the more expensive and complex the model becomes, since each parameter requires its own algorithm, or part of the simulation code that must interact with all the other parts. This boosts the time it takes to run the model on a computer, often making it impractical to use.

"What we've done," Wild says, "is to build a faster mathematical surrogate that can replace key bottlenecks in expensive models and make them computationally inexpensive to evaluate."

Wild and Moré's optimization algorithm is performing well in a Matlab implementation, and the two are continuing to finesse the code to solve implementation details and make it user-ready for addition to TAO.

"It's already beating the competition," Moré says.

For Wild, it's an important step toward bringing applied math to the environmental engineering community. He says one of the most rewarding aspects of the practicum at Argonne was realizing the enormous impact of optimization codes.

"Argonne's Toolkit for Advanced Optimization is a facilitator for important science around the world, and it felt great contributing to this," he says.

Nowhere is this truer than in the realm of environmental engineering. There are tens of thousands of contaminated groundwater sites in the United States, from EPA-designated Superfund sites to smaller ones, involving pollutants from radioactive wastes to pesticides and petroleum residues. Remediating these sites often costs tens of millions of dollars and involves decades of work. Optimizing the cleanup approach before starting can save years of work and millions of dollars.

When Wild presented his preliminary results at the Society for Industrial and Applied Mathematics (SIAM) meeting on Computer Science and Engineering in September 2006, he was approached by scientists from the United States Geological Survey interested in applying the optimization techniques to their models.

It was a heartening response — one that told Wild he had the right recipe for long-term success.

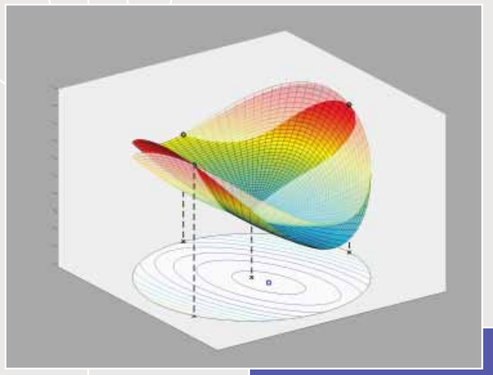

Shown is a quadratic approximation (transparent) of a function (opaque) using only 4 function values. The quadratic "surrogate" is easy to evaluate and hence it can be efficiently used to suggest a new point (the square) at which the expensive function should be evaluated.

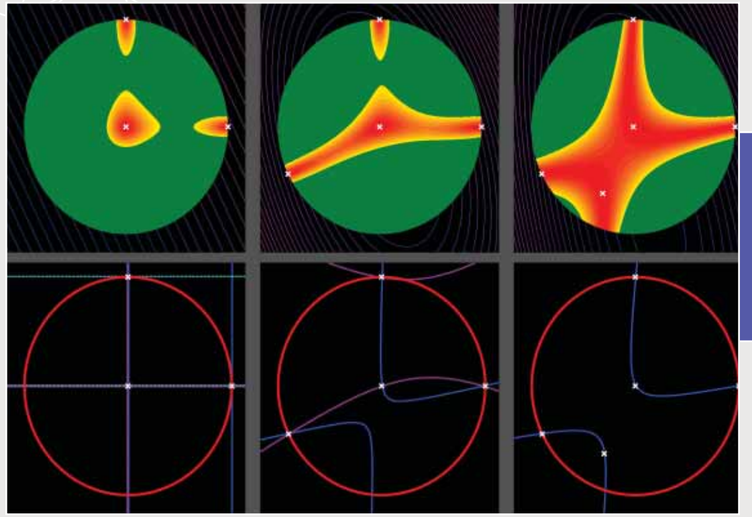

When constructing interpolation surrogates, one often encounters a battle between the optimization (points suggested by the surrogate) and the numerical conditioning of the resulting interpolation set. The optimizer's goal is to make good progress while respecting numerical stability concerns.

For Wild, optimization is about bringing math to the people — to turn an unusable model into a dinner-table standard.

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 $2,500 for a computer workstation purchase
- Opportunity to complete a practicum working with scientists and researchers at a DOE Laboratory
- Yearly fellows' conference with opportunities to meet other fellows and academic and government professionals
- Renewable up to four years

For more information: www.krellinst.org/csgf

Power Play

NERSC Makes High-Performance Computing Accessible

By Thomas R. O'Donnell

THE PERSONAL COMPUTERS in Lawrence Pratt's laboratory weren't cutting it. His research on the structures and interactions of lithium compounds was hindered because the Pentium 4-type machines he uses at Fisk University in Tennessee couldn't keep up with the demands of modern computational chemistry.

His work finally kicked into high gear with a grant of 150,000 high-performance computer processor hours — and the help of the National Energy Research Scientific Computing (NERSC) Center, based at the Department of Energy's Lawrence Berkeley National Laboratory (LBNL) in California.

"I use every last minute" of computer time, says Pratt, who got onto the powerful NERSC machines through the DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. He's since qualified for another grant of 150,000 processor hours and "I'm burning through it like crazy, but I'm also publishing a lot of papers — five so far this year."

As INCITE projects go, Pratt's is small. Most INCITE projects, which hold potential for major scientific breakthroughs, are awarded millions of processor hours — but they consume up to 20 percent of the hours available on NERSC's massively parallel computers, says center director Horst Simon. The remaining 80 percent is divided among 300 or so projects, each using tens of thousands to millions of processor hours per year. A small amount of time is set aside for "startups," or researchers who are still preparing their software for massively parallel processing.

Simon would be happy to have more such projects use NERSC computers as a gateway to high-performance computing. In particular, he's working to provide more students — especially DOE Computational Science Graduate Fellowship (DOE CSGF) participants — more opportunities to try their programs. "We want to get students interested in using NERSC or any of the other DOE computational resources, so they have a positive experience and become used to integrating scientific computing into their work," especially after graduation, says Simon, who is leaving his post soon to concentrate on other roles at the Berkeley lab.

The center is a perfect place for a first taste of science on massively parallel computers. For more than 30 years, it's been DOE's main production center for scientific computing, and it hosts some of the department's largest, fastest systems for unclassified research. More than 2,500 users from dozens of universities, private research institutions and DOE laboratories work on around 300 projects each year. Yet, users rarely visit the NERSC facility. Most connect to and use center computers via ESnet, DOE's high-speed network, and the Internet. Their work produces mountains of results — for 2006, researchers published more than 1,400 papers related to calculations on NERSC computers.

The NERSC Center's role as a DOE service facility means the projects running on its computers cover virtually every strategic theme pursued by DOE and its Office of Science — "Everything from astrophysics down to nanoscience," Simon says. Fusion energy, materials science, chemistry, climate, genomics, computational biology, applied math and computer science are just some of the disciplines with research on NERSC Center machines. "Our mission is really to be the high-end production resource for the Office of Science, so general purpose and diverse applications have been part of our mission," Simon adds.

In-house expertise is part of what makes the NERSC Center popular with researchers. The staff works with computational scientists to tune applications for the best performance, visualize their results and make their research more effective. There's a strong culture focused on assisting users, and the staff is experienced and stable, Simon says. The center stages regular training sessions, encourages user communities to exchange information, and hosts databases of user questions.

"If there's a difficult project that NERSC staff can work with the scientific users on, it often turns into a scientific collaboration," with the staff member listed as a joint author on research, Simon adds.

Pratt, who scaled his mathematical models up from the single processor in his desktop computers to eight processors, says NERSC experts were helpful. Without access to high-performance computing, "There are projects I would not have been able to complete," he says. "We were able to find out a lot about the chemistry through computational mechanisms that would not have been easily obtained by experiment."

Large allocations of NERSC computer time generally are awarded competitively through annual requests for proposals, with DOE program managers making the decisions. Computing experts at NERSC, Argonne National Laboratory, Oak Ridge National Laboratory, and Pacific Northwest National Laboratory judge whether the codes and algorithms described in the proposals are ready to run on massively parallel machines. Peer review panels scrutinize the proposals for their impact on science. "This is sort of self-selecting, because the people who apply know what the reviewers are looking for," Simon says.

NERSC is making it even easier for DOE CSGF fellows to get on its machines. It's allocating 40,000 to 50,000 startup hours to between 50 and 70 projects fellows put forward.

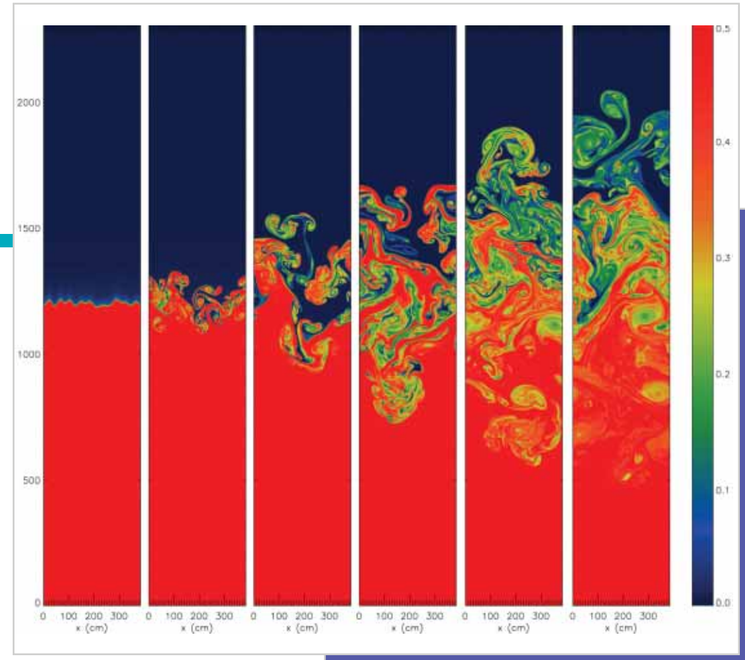

Accelerating a thermonuclear flame to a large fraction of the speed of sound (possibly supersonic) is one of the main difficulties in modeling Type Ia supernova explosions, which likely begin as a nuclear runaway near the center of a carbon-oxygen white dwarf. The outward propagating flame is unstable, which accelerates it toward the speed of sound. Bell et al. investigated the unstable flame at the transition from the flamelet regime to the distributed-burning regime through detailed, fully resolved simulations. At the low end of the density range, the instability dominated the burning.

Image illustrates a carbon mass fraction for a 384 cm wide, 6.67 x 106 g cm–3 C/O flame shown every 1.6 x 10–3 s until 8.12 x 10–2 s. The fuel appears red (carbon mass fraction = 0.5), and gravity points toward increasing y. At this low density, the instability dominates over the burning, and a large mixed region develops.

Image and caption courtesy of www.nersc.gov

COLLABORATORS

Horst Simon was named Associate Laboratory Director for Computing Sciences at Berkeley Lab in 2004. He represents the interests of the lab's scientific computing divisions — the NERSC Center and Computational Research — in the formulation of laboratory policy, and leads the overall direction of the two divisions. He also coordinates constructive interactions within the computing sciences divisions to seek coupling with other scientific programs. Simon joined LBNL in early 1996 as director of the newly formed NERSC Division, and was one of the key architects in establishing NERSC at its new location in Berkeley. Simon also is the founding director of Berkeley Lab's Computational Research Division, which conducts applied research and development in computer science, computational science, and applied mathematics. His research interests are in the development of sparse matrix algorithms, algorithms for large-scale eigenvalue problems, and domain decomposition algorithms for unstructured domains for parallel processing.

Simon's recursive spectral bisection algorithm is regarded as a breakthrough in parallel algorithms for unstructured computations, and his algorithm research efforts were honored with the 1988 Gordon Bell Prize for parallel processing research. He also is one of four editors of the twice-yearly "TOP500" list of the world's most powerful computing systems.

David Skinner was the lead high-performance computing (HPC) consultant for the Department of Energy's first six Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects before starting the SciDAC Outreach Center. The INCITE program has since blossomed into a multi-lab allocation process for large-scale computing. In that work and other projects Skinner often focuses on making scientific applications run fast and scale well. The core areas of Skinner's present HPC research include improving application performance, characterizing scientific workloads, and analysis of emerging architectures.

Skinner heads the NERSC Center's Open Software and Programming Group at NERSC, which is active in making software deliver for HPC centers and users and in promoting software development practices that enhance reliability and performance in the overall HPC process. His group works on a variety of software related to parallel computing applications themselves as well as HPC center infrastructure software for system monitoring, allocation banking, and Web services. Skinner also publishes scientific research in the areas of molecular dynamics, chemical quantum dynamics and kinetics.

Further Reading:
NERSC Center Web page:
http://www.nersc.gov/

NERSC Center strategic plan:
http://www.nersc.gov/news/reports/LBNL-57582.pdf

NERSC Center newsletter archive:
http://www.nersc.gov/news/nerscnews/

Contact:
Horst Simon HDSimon@lbl.gov
David Skinner deskinner@lbl.gov

PRACTICUM COORDINATOR
Daniel Martin . . . . . DFMartin@lbl.gov

Named "Franklin" after Benjamin Franklin, America's first scientist, the Cray XT4 will consist of more than 19,000 processor cores when fully installed. It will deliver sustained performance of at least 16 trillion calculations per second, with a theoretical peak speed of more than 100 teraflop/s. Franklin the computer was powered up for the first time on January 17, 2007 thus celebrating its birthday with Benjamin Franklin, who was born January 17, 1706.

Image and caption courtesy of www.nersc.gov

In the past, requests often totaled more than 10 times the amount of available time. Now the deployment of ever-more-powerful computers allows NERSC to meet about half of the requests. While NERSC seeks to accommodate projects and investigators with little or no background in parallel processing or high-end computing, those are becoming rare as parallel machines become ubiquitous.

The process and facilities may seem daunting, but it's not hard to get a foot in the door. New users who want to try their codes or develop new ones on NERSC machines can apply online for a startup allocation. Startups must meet the Office of Science mission and require high-performance computing.

NERSC is making it even easier for DOE CSGF fellows. It's allocating 40,000 to 50,000 startup hours to between 50 and 70 projects fellows put forward. The allocations will let students see how well their codes scale in parallel, and enable them to work with NERSC consultants to improve their projects and code performance.

The idea is to get fellows to look beyond the computational resources they have at hand through their major professor or department, says David Skinner, leader of NERSC's Open Software and Programming Group and coordinator for the SciDAC Outreach Center (see sidebar). Some students may have discovered NERSC Center computer resources if their advisor has used them, but the goal is to attract even those who have not worked with the center before.

Fellows can go to a Web page and fill out a survey about their computing needs. About half of respondents through summer 2007 "were people who said 'This is where I'm generally headed, but I'm not there yet,'" Skinner says, but other fellows said they can't get their research done fast enough with their present computer resources.

"A couple…said 'My workstation is too slow for the work,'" Skinner says. "For those, my response was to get them onto some high-performance computing facilities," either at NERSC or another national laboratory, such as Oak Ridge or Argonne. The startup allocations should let fellows try out their codes for up to 18 months.

Students, Simon says, often are shy about asking for help getting their codes to run on NERSC Center computers. Many are accustomed to solving their own computer problems and working with university computing centers that often were staffed by their fellow students, Simon adds. They're "quite surprised when they come to NERSC, because we are a full-service organization."

"We don't discriminate against students," Simon says; they're often the people doing the nitty-gritty coding for their major professors' research. DOE CSGF fellows who get access to NERSC can work with some powerful computers:

- Seaborg, an IBM RS/6000 SP with 6,656 processors (6,080 of which are available to run scientific computing applications). The system has a peak performance of 10 teraflops, or 10 trillion operations per second.
- An IBM p575 POWER5 system with 976 processors named Bassi. Bassi's peak performance is 7.4 teraflops and it has 100 terabytes of disk storage.
- A 712-processor Opteron Linux cluster named Jacquard, with a peak speed of 3.1 teraflops.
- PDSF, a 700-processor Linux cluster dedicated to high-energy physics research.
- DaVinci, a 32-processor SGI Altix shared memory cluster devoted to visualization, data analysis, and long-running interactive work.

NERSC also has two data storage systems:

- The NERSC Global File System (NGF), with about 150 terabytes (trillion bytes) of user-accessible storage. NGF allows users to create and access a single file from any of the lab's high-performance computing systems.
- The High-Performance Storage System (HPSS), with a theoretical capacity of 22 petabytes (quadrillion bytes) for long-term archival data storage.

In the fall of 2007, NERSC also will bring its latest computer system on line. Dubbed Franklin, the Cray XT4 will have 19,344 compute CPUs, at least two gigabytes of memory per CPU, and a sustained performance of 16 teraflops, as opposed to a theoretical peak performance of 100 teraflops. With future upgrades, Franklin could have a theoretical peak of 1 petaflops — one quadrillion calculations per second.

Franklin will increase the number of NERSC computer cycles available for research by a factor of 16, Simon says — and every one of them is needed. "Researchers often ask for many more processor hours than we can actually accommodate. In the last couple of years we had requests that were more than six or seven times what we had available," he adds. Most researchers got only a fraction of the processor hours they wanted. With Franklin, "We expect we will have, for once, enough cycles to keep everybody happy" — but not for long. Demand for computing time is constantly growing.

That's why, as soon Franklin is stabilized and running, NERSC will begin preparing for the next system, called NERSC-6 for now. NERSC-6 is likely to start life as a 1-petaflops-capable machine. It's also likely to have even more processing cores on a single chip, a change that poses challenges for the future NERSC director.

"There's always enough work to do," Simon says. "There's exciting stuff to do as long as computers grow and become more powerful. We never stand still."

GETTING UNSTUCK
SCIDAC OUTREACH CENTER BRINGS SOLUTIONS TO SCIENTISTS

>> ***Even though it too has a toll-free number (866-470-5547) and e-mail address (help@outreach.SciDAC.gov) for obtaining assistance, the center that David Skinner oversees isn't quite like the help desks that computer makers run for perplexed users. For one thing, callers to the SciDAC Outreach Center get better service.***

The center is part of the second round of the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC-2) program, which applies computational science to projects in a range of disciplines. The center is a clearinghouse to disseminate computational tools and techniques SciDAC researchers develop throughout the program's entire community.

"SciDAC has tried to bring software solutions to scientists and make them easier to use," says Skinner, the outreach center coordinator and leader of the Open Software and Programming Group at the National Energy Research Scientific Computing (NERSC) Center. But "Some of the technologies are quite complicated, and having someone who can assist in deploying a high-performance computing software library or a different methodology" is the outreach center's role.

Skinner and NERSC staff member Andrew Uselton are the only people permanently assigned to the center, but they have access to other NERSC application experts. More often, the center is an intermediary to connect researchers to each other.

Skinner says the most common thing he hears from researchers is "I'm stuck."

"That's the most motivating factor to get people to seek out help, is that they're bottlenecked or can't get their research done," Skinner says. Researchers contact the center seeking information on new software or technology they might use. On the other side, researchers also contact the center to offer software solutions.

"We definitely are the matchmaker between producers and requesters of new technology," Skinner says, but he and Uselton don't just sit back and wait for calls. They're spreading the word, both in person and electronically, about available technologies. They attend meetings to learn what scientists need and to help them learn new software. For example, about 100 researchers attended a series of software tutorials the center staged in conjunction with the annual SciDAC conference in Boston in June 2007.

Electronically, the center provides tools for collaboration, including software testing and evaluation, software repositories, and even FAQs. For example, the center is developing a GForge website for software development activities, including news, downloads, bug reporting, feature tracking and other tools.

The outreach center builds on NERSC's strong track record, Skinner says. "We have a lot of experience in bringing people up to speed on parallel computing resources," he adds, but "The SciDAC Outreach Center has a broader scope than just NERSC." It will connect researchers with any laboratory or computing center that might have the answers or services they need.

By Thomas R. O'Donnell

From Flames to Fusion

Sandia Researchers' Methods Prove Their Versatility

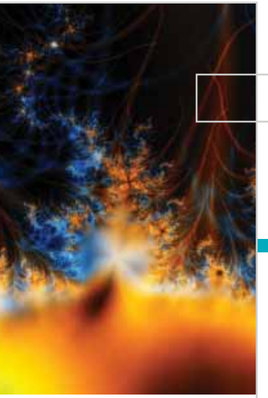

>> NERSC

WHETHER IT'S THE SUN, THE STARS, or the long-sought viable fusion reactor, they all involve plasmas, and researchers at Sandia and Los Alamos national laboratories want to help understand what happens in them.

They're developing computer algorithms that simulate how ions and electrons move and react in plasmas under the influence of magnetic fields. If they're successful, it could help physicists better understand the processes fueling the sun and stars, and how to harness that power for energy.

The project builds on nearly 20 years of Sandia's field-leading research in hardware, computer science, applied math and numerical algorithms to devise efficient methods for massively parallel computing. The results include Aztec, one of the first large-scale, parallel, iterative solver libraries; and Chaco, an early tool to balance the computational workload between parallel processors. A host of complex applications have used those and other codes devised by John Shadid, Ray Tuminaro, Scott Hutchinson, Bruce Hendrickson, Rob Leland and other Sandia researchers.

But Shadid had his own application in mind. He's worked with fellow Sandians Roger Pawlowski, Andy Salinger, Karen Devine, Gary Hennigan and Paul Lin to simulate the intricate physics of coupled fluid flow and complex chemistry. These transport/reaction (or chemically reacting flow) systems have a multitude of applications, including simulating combustion, cardiac cell activity, and even the spread of a biological material released in a busy airport.

In one case, Sandia researchers worked with Ford Motor Co. researcher Kevin Elwood to simulate a hydrogen solid oxide fuel cell. The researchers' algorithms and the MPSalsa simulation software helped model flow and reactions in the fuel cell to transform hydrogen and oxygen into water and electricity. The researchers also have collaborated with Dow Chemical Co. to optimize a new design for a chemical reactor to efficiently convert ethane gas to liquid ethylene, an important feedstock for plastics.

Stability Control

Such transport/reaction processes couple complex physical processes: Fluid flow, energy transfer, radiation effects, chemical species transport and chemical reactions — both in fluids and on surfaces — and more. They also occur across an expansive range of space and time scales.

"This myriad of coupled phenomena produce a highly nonlinear, multiple timescale behavior," Shadid says. It's a huge challenge to find stable, accurate, and efficient numerical methods that simulate these processes and run efficiently on ever-larger computers.

The usual algorithmic approaches have used semi-implicit, operator-splitting, or explicit time-stepping. In explicit time-stepping, the value of each quantity — density, momentum, energy, species concentration, and others — at the new time is calculated based on the old values. That approach produces a simpler solution method, but it can be unstable, meaning it's more likely to produce nonsense answers. "You have to go at the smallest time-scale of the individual physics to achieve stability, even if it doesn't provide extra accuracy in the simulation," Shadid says.

"This myriad of coupled phenomena produce a highly nonlinear, multiple timescale behavior," Shadid says. It's a huge challenge to find stable, accurate, and efficient numerical methods that simulate these processes and run efficiently on ever-larger computers.

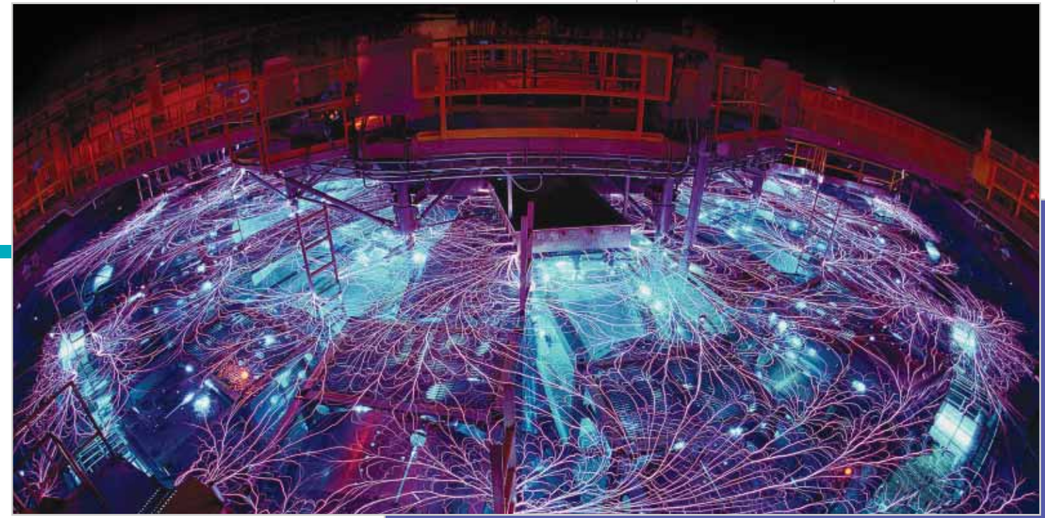

Time-exposure photograph of Sandia's Z machine firing. Z is the world's most efficient (15%) and powerful laboratory x-ray source, producing x-ray powers in excess of 200 trillion watts. (For more information: http://www.sandia.gov/pulsedpower/).

Photo provided by Randy Montoya

Shadid and his fellow researchers approach the problem from the other end of the spectrum: using fully implicit methods, which represent all the physics consistently at each new time level. "The advantage is you can take time steps that are associated with the physics you're interested in," rather than stability, Shadid says. The desired accuracy dictates the time steps, and fewer time steps can be calculated without losing accuracy. The disadvantage: When the governing partial differential equations (PDEs) are discretized — transformed into algebraic equation systems digital computers can solve — implicit time integration produces a system of strongly coupled, nonlinear equations with tens of millions or hundreds of millions of equations and just as many unknowns. These equations must be solved simultaneously on thousands of computer processors.

These huge algebraic systems must be attacked intelligently so parallel-processing computers can solve them efficiently. The methods Shadid and his fellow researchers have developed generally do that in two steps: applying a nonlinear solver; and applying iterative solvers with multigrid preconditioners.

Getting a Grid

In general, discretizing PDEs involves distributing points in a grid or mesh throughout the space being simulated, such as a chemical reaction chamber. The more points, the better the resolution and accuracy of the simulation — and the more demanding it is for computers to solve.

Multigrid methods use a hierarchy of grids of varying resolutions. Because they use information from a sequence of grids rather than a single grid, these methods scale optimally — the number of iterations is constant and the work grows in direct proportion to the problem size. Collaborators Ray

COLLABORATORS

John N. Shadid is a distinguished member of the technical staff at Sandia National Laboratories in the Computational Science R&D Group. He received bachelor's and master's degrees in mechanical engineering and a master's degree in mathematics from the University of Wisconsin-Milwaukee and his doctoral degree in mechanical engineering from the University of Minnesota. After graduation he joined Sandia National Laboratories as a senior member of the technical staff in the Parallel Computational Sciences Department and was named a distinguished member in 1999. At Sandia he has been lead principal investigator and co-PI on a number of large-scale computational science projects, including research and development of a parallel implicit transport/reaction simulation code, MPSalsa, and a parallel preconditioned Krylov solver library, Aztec. The Aztec library received an R&D 100 award in 1997. The MPSalsa simulation code has been honored twice as a Gordon Bell Prize finalist. His current research interests include high-performance computing; parallel algorithm development; numerical solution methods for multiple-time-scale nonlinear coupled PDEs; and the simulation of a wide range of complex transport/reaction systems that includes, most recently, magnetohydrodynamic (MHD) systems.

Luis Chacón is a member of the technical staff at Los Alamos National Laboratory. He obtained a master's degree in industrial engineering from the Polytechnic University of Madrid in 1994, and master's and doctoral degrees in nuclear engineering from the University of Illinois in 1998 and 2000, respectively. In 2000, he accepted a Director's Postdoctoral Fellowship appointment in T-division at Los Alamos National Laboratory to research the application of Newton-Krylov methods to resistive MHD. During this appointment, he has made important contributions in the context of implicit, nonlinear algorithms for two-dimensional resistive and Hall incompressible MHD.

Roger Pawlowski is a senior member of the technical staff at Sandia National Laboratories. He obtained his doctoral degree in chemical engineering from the State University of New York-Buffalo in 2000. He has contributed to a variety of projects at Sandia, including the development of solvers for large-scale circuit networks, coupled circuit/device solvers, catalytic oxidation reactor design, MEMS reactor design, multi-phase aerosol modeling, combustion, and fundamental studies of stagnation flows. His current interests focus on developing robust finite element discretization techniques for MHD and reacting flow physics and algorithm development for nonlinear systems, bifurcation analysis, and multi-physics coupling.

Further reading:

J. N. Shadid, A. G. Salinger, R. P. Pawlowski, P. T. Lin, G. L. Hennigan, R. S. Tuminaro, and R. B. Lehoucq. *Large-scale Stabilized FE Computational Analysis of Nonlinear Steady State Transport/Reaction Systems,* CMAME. 195, 1846-1871 (2006).

D. A. Knoll, V. A. Mousseau, L. Chacón, J. Reisner. *Jacobian-free Newton-Krylov methods for the accurate time integration of stiff wave systems,* Journal of Scientific Computing. 25, no. 1, 213-230 (2005).

Contact:

John Shadid jnshadi@sandia.gov
Luis Chacón chacon@lanl.gov
Roger Pawlowski rppawlo@sandia.gov

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

PRACTICUM COORDINATOR

Heath Hanshaw . . . . . hlhansh@sandia.gov

Tuminaro and Mazio Sala (from Swiss research institution ETH Zurich) have fully algebraic methods that accomplish this.

The preconditioning step is critical to solving these large systems, Shadid adds. On test problems with 10 million to 250 million unknowns and running on up to 8,000-plus parallel processors, the Sandia multilevel preconditioning techniques proved to be 10 to 100 times faster than one-level methods. The iterations necessary to reach a solution also stayed constant, even as the problem size grew by more than four orders of magnitude. That means the algorithms should scale well on even bigger computers.

The Sandia researchers' nonlinear solvers are robust — unlikely to blow up and produce nonsense — and allow them to crack elaborate problems. This includes bifurcation studies, which efficiently identify changes in stability — points where a change in model parameters leads to multiple steady states or oscillatory solutions. For example, the researchers' codes have been used to analyze chemical vapor deposition, which places films of semiconductor materials on wafers for microelectronics. Industry wants to scale up the process to get more chips out of each wafer, but reactor instabilities produce inconsistent results. "Some days you get a good uniform layer. Some days you get a really bad, nonuniform layer," Pawlowski says. Small changes in the system — whether a valve sticks during start-up, whether a temperature setting is slightly out of adjustment, or even whether someone bumped the reactor during the run — can lead to a fundamental change in system performance.

The group's code can determine the point at which parameter variations can cause destabilizing effects that destroy uniform, steady-state operation. That information can help improve designs for reactors and other devices. The code calculates the bifurcation point directly, compared to programs that must be run numerous times with different operating parameters. "We can set one run, let it go, and it will directly compute the parameter values of the bifurcation point," Pawlowski says. "You don't have to adjust the parameters by hand and try to hopefully figure out where it changes stability." It's "a powerful tool to design stable operation" of an experiment or process.

Similarly, the group's unique algorithms can help scientists and engineers choose optimal designs. "I could sit at a computer all day and set different parameters and run (a program) out, change the parameters and run it out again and again" to get an optimal answer, Pawlowski says. That may work if only a few parameters are changing, but the problem is tougher when many parameters are involved.

"You don't want to simulate for every possible combination of parameters. You want methods that will take you directly to the optimal solution," Pawlowski says. The Sandia codes do that quickly.

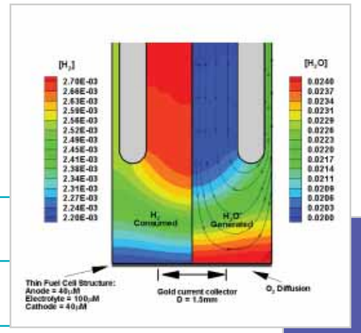

A contour plot of hydrogen and water concentrations in a 2D-cylindrical solid oxide fuel cell (SOFC). In the MPSalsa transport / reaction simulation hydrogen fuel reacts with oxygen ions at the anode interface (lower surface) to produce water and electrons for electrical power.

Courtesy of J.N. Shadid of Sandia National Labs and K. Ellwood of Ford Research.

Plasma Application

The group's nonlinear solvers and preconditioner algorithms have been disseminated to the computational science community as part of the Trilinos solver framework, a parallel computing software package assembled at Sandia. Now Shadid and his fellow researchers are extending their methods to plasmas — the stuff of stars and fusion energy. The work could apply to astrophysics and plasma processing of advanced semiconductor and microelectromechanical systems (MEMS). It also could help simulate plasma behavior in devices like ITER, the international fusion reactor project to be built in southern France; the National Ignition Facility (NIF) at DOE's Lawrence Livermore National Laboratory; or Sandia's own Z machine (see sidebar).

Sandia researchers Pavel Bochev and Jeff Banks work on the project with Shadid, Pawlowski and Tuminaro. They're collaborating with Luis Chacón of Los Alamos National Laboratory, Dana Knoll of Idaho National Laboratory, and DOE Computational Science Graduate Fellow John Evans of the University of Texas. The project marries methods the Sandia group devised with ones Chacón and Knoll have developed for magnetohydrodynamics (MHD) and extended magnetohydrodynamics (XMHD).

MHD studies how electrically conducting fluids move and are affected by magnetic fields. The physics has similarities with those governing chemical transport and reaction — but with some twists. "You add in electric and magnetic fields with Maxwell's equations," Shadid says. "It makes the solution processes significantly more difficult."

XMHD is trickier still. Standard MHD simulations track just one type of fluid behavior, and couple ions and electrons so they behave as a single fluid. XMHD brings in more complex electron dynamics, Chacón says. "Standard MHD is an approximation of what reality does, and the approximation involves several assumptions," he adds. XMHD makes fewer assumptions, but "The price you pay is that it becomes a much more complicated system" to solve.

Like the Sandia group, Chacón, Knoll and their fellow researchers have focused on multigrid preconditioning methods with applications to MHD simulations. In essence, their physics-based preconditioners determine which coupled physics equations govern the simulation's time scale.

"We know that the coupling produces the problems in time scales and so we address those," Chacón adds. "That's where the physics-based (approach) comes from. You need to have the insights of what couplings are producing the time scales" and address those directly. The preconditioning technique breaks the complex equation systems into smaller subsystems for multigrid solutions.

The Sandia group has used a similar physics-based preconditioner technique and applied it to low Mach-number fluid flows. In effect, they decouple the equations into subsystems that can be more easily solved by multigrid methods. "We don't decouple the actual iterative solver," Shadid says. "We decouple some of the physics in the preconditioners, which is just an approximation that gives us very fast, very good approximate solution."

The method Chacón's group uses also doesn't assume any particular grid or mesh structure to discretize the governing equations. The Sandia researchers use unstructured data grids, allowing them to discretize problems with complex geometries. The Los Alamos approach "readily generalizes to unstructured meshes," Chacón adds. "In principle there is a good marriage between the two. There is a direct translation to unstructured meshes."

Shadid agrees: "It's really an excellent collaboration."

FEELING Z PINCH

>> ***There's a homegrown brand of fusion science Sandia and Los Alamos researchers want to help model with their algorithms for extended magnetohydrodynamics.***

Sandia National Laboratory's New Mexico location houses the Z machine, the world's largest controlled X-ray generator. It's designed to test materials under extreme conditions of radiation and pressure as part of the U.S. nuclear weapons stockpile stewardship program, but it's also being studied as a route to creating the clean energy of nuclear fusion.

Through a system of huge capacitors and wires, and with 20 million amperes of electrical power, the Z machine can generate pulses of 200 trillion watts of power for short periods. The pulse vaporizes an array of metal wires around a housing the size of a spool of thread, turning them into plasma. The powerful magnetic field generated by the current compresses the plasma to about the thickness of a pencil lead. Under compression, the moving ions and electrons suddenly stagnate, releasing energy as X-rays and reaching temperatures of billions of degrees.

The current passing through the wires travels vertically, or along the Z axis. Since magnetic fields "pinch" the plasma, the process has been called Z-pinch confinement.

Sandia scientists have used the Z machine to fuse tiny amounts of deuterium, producing thermonuclear neutrons — a step toward creating a self-sustaining fusion reaction.

Sandia researcher John Shadid says some of his group's algorithms may be useful in helping model hydromagnetic Rayleigh-Taylor instabilities in the plasma implosion. Such instabilities are an important limiting mechanism for the amount of radiation energy produced by a Z-pinch and are therefore critical to understand and control. The algorithms the groups are developing may find their way into production simulation codes like Sandia's ALEGRA, which is used in modeling the Z machine.

"There's a large program already at Sandia in terms of modeling this kind of physics, but there are a number of open issues still, and one of them is handling these complex, interacting, multiple time scales" — a focus for his group's research, Shadid says. Luis Chacón, Shadid's collaborator at Los Alamos National Laboratory, says the powerful machine is ripe for simulations using the group's approach. It's been hypothesized that electron physics plays a role in the Z machine's operation "and that brings all the tools we're talking about in that regard." Computer models of the Z machine use operator-split time-stepping techniques. These methods can sometimes introduce instabilities and lead to significant error accumulation. Chacón and Shadid believe more effective implicit coupling could cut errors and increase efficiency, provided the faster time scales electron physics brings to the problem can be managed.

Measuring Up

On the Petascale

By Karyn Hede

"ANYONE CAN BUILD A FAST CPU. The trick is to build a fast system." — *Seymour Cray*

The words of Seymour Cray, founder of supercomputer maker Cray Inc., will soon be tested. The company is to deliver a system to Oak Ridge National Laboratory (ORNL).

Petascale computing has the potential to eclipse today's fastest computer by a factor of three. But how fast it really runs depends more on the lab than on the manufacturer.

As the Department of Energy's National Leadership Computing Facility (NLCF), Oak Ridge has been asked to take a dramatic step toward a new capability for complex, computationally intense science.

"Leadership Class Computing is a combination of world-leading computing power and the policies that allow it to be used to solve very large problems," says Doug Kothe, director of science at the lab's National Center for Computational Sciences (NCCS). Instead of hundreds or thousands of simulations running on the system at any one time, there may be only a handful. The idea is to push the limit on what simulation can tell us about about the natural world.

Jaguar, ORNL's current Cray XT4 system, comprises 124 cabinets containing more than 11,700 dual-core processors. The system has achieved 101.7 teraflops (trillion floating point operations per second) on the Linpack benchmark — more than 85 percent of its 119-teraflops theoretical peak speed. It's scheduled to receive new quad-core processors in late 2007, bringing its top speed to greater than 250 teraflops.

The new Cray "Baker" system will make that look like small potatoes. When it arrives in early 2009, it is expected to achieve petascale speed — more than a quadrillion operations per second, 10 times Jaguar's current Linpack rating.

But fast computers alone can't guarantee science breakthroughs. It's just as important to allocate large blocks of time on them to attack the most difficult problems. DOE does that through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, which provides large computer time grants to just a handful of peer-reviwed projects. The 2007 program made 28 allocations, but only a handful will be selected to have the first crack at Baker.

"It's an interesting and exciting new model, where you are trying to go deep into the science, and that can be done if you only have a few projects," Kothe says.

A third ingredient, besides fast computers and large allocations, also is necessary if these projects — some of the most complex calculations ever attempted — are to succeed. The NCCS works with the science teams through user assistance and scientific computing groups. A 5-million processor-hour INCITE allocation is equivalent to a multi-million dollar grant in terms of the support they receive, Kothe explains.

"To do breakthrough science at a user facility, we can't just be a 'cycle shop' and say 'OK, here's your account, good luck,'" Kothe says. "We intimately involve ourselves with each team."

"To do breakthrough science at a user facility, we can't just be a 'cycle shop' and say 'OK, here's your account, good luck,'" Kothe says. "We intimately involve ourselves with each team. We partner with them in ways that can help them by giving them our best and brightest people who are most aligned with their interests and needs."

That includes former DOE Computational Science Graduate Fellowship scholar Richard Mills. Mills worked in user support at Oak Ridge after completing his practicum with Peter Lichtner at Los Alamos National Laboratory. That experience led to a full-fledged collaboration with Los Alamos, and now Mills is a fully supported member of ORNL's computational earth sciences group.

"Richard spun up very quickly through our system," Kothe says. "It's a good example of what's possible for computer scientists who join our group."

Bridging Worlds

Climate modeling is one area expected to benefit from the great leap forward petascale computing offers.

There is a concerted effort to add the full carbon cycle to climate models, Kothe explains. Previous research has suggested ecosystems respond to increasing atmospheric carbon dioxide by taking up more carbon, but paradoxically also respond to increasing temperature by releasing carbon. Scientists would like to compare these responses in a full climate simulation, but it has been too expensive in terms of computer time and cost to add a fully functional carbon cycle to the models. Instead, models specified the amount of carbon dioxide in the atmosphere as a fixed input.

The advent of petascale computing opens up the possibility of simulating a full carbon cycle — including human-generated carbon — as well as simulating how the land and ocean will respond.

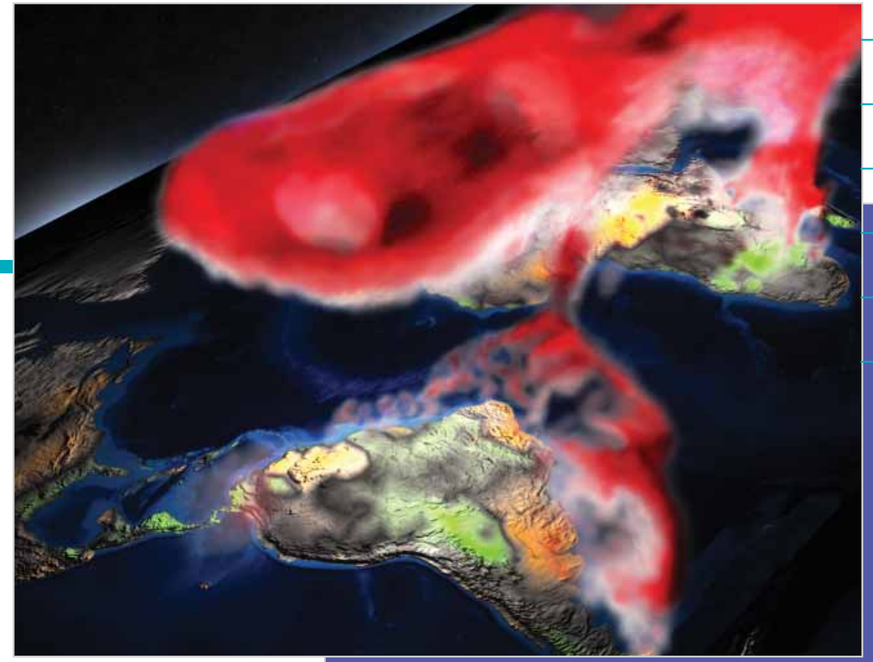

This image shows a snapshot of the simulated time evolution of atmospheric carbon dioxide (the red plumes) concentration originating from the land surface at the beginning of the industrial carbon cycle (around 1900). This CO_2 is a product of the net ecosystem exchange, the CO_2 flux due to respiration of vegetation and soil microbes (green areas on land) minus that taken up for ecosystem production (orange areas on land). The underlying simulation is one of a number of runs performed for Phase 1 of the Coupled Climate/Carbon Cycle Model Intercomparison Project on the Leadership Computing Facility.

Courtesy: Dr. Warren Washington, Principal Investigator.

COLLABORATORS

Doug Kothe is Director of Science for the National Center for Computational Sciences at Oak Ridge National Laboratory (ORNL). He is responsible for guiding the multidisciplinary research teams using the Center's leadership computing systems. Dr. Kothe has more than 20 years of experience in computational science research. His research interests and expertise have centered on developing physical models and numerical algorithms for simulating physical processes in the presence of incompressible and compressible fluid flow. Before joining NCCS, he was deputy program director for Theoretical and Computational Programs in the Advanced Simulation and Computing (ASC) Program at Los Alamos National Laboratory. He joined the technical staff at LANL in 1988 as a member of the Fluid Dynamics Group. Dr. Kothe received his bachelor's degree in chemical engineering from the University of Missouri – Columbia and his master's and doctoral degrees in nuclear engineering from Purdue University. He is author of more than 60 refereed publications and has written more than a half-million lines of source code.

Robert J. Harrison holds a joint appointment with ORNL and the chemistry department of the University of Tennessee, Knoxville. At Oak Ridge, he is leader of the Computational Chemical Sciences Group in the Computer Science and Mathematics Division. He has more than 75 publications in peer-reviewed journals in the areas of theoretical and computational chemistry and high-performance computing. He earned a bachelor's degree in natural science from Cambridge University, England, in 1981, and continued on there to earn a doctoral degree in organic and theoretical chemistry in 1984. He worked as a postdoctoral research fellow at the Quantum Theory Project, University of Florida, and the Daresbury Laboratory, England, before joining the staff of the theoretical chemistry group at Argonne National Laboratory in 1988. In 1992, he moved to the Environmental Molecular Sciences Laboratory of Pacific Northwest National Laboratory, conducting research in theoretical chemistry and leading the development of NWChem, a computational chemistry code for massively parallel computers. He started the joint faculty appointment with UT/ORNL in August 2002. In addition to his research into efficient and accurate calculations on large systems, he has pursued applications in molecular electronics and chemistry at the nanoscale. In 1999, the NWChem team received an R&D Magazine R&D100 award, and Dr. Harrison received the IEEE Computer Society Sydney Fernbach award in 2002.

Contact:
Doug Kothe kothe@ornl.gov
Robert Harrison harrisonrj@ornl.gov

Trey White trey@ornl.gov
Richard Mills rmills@ornl.gov

PRACTICUM COORDINATOR

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

The massive project, called "A scalable and extensible Earth system model for climate change science," includes input from several national laboratories, academic programs and DOE's Atmospheric Science Program and Terrestrial Carbon Program, among others. Its goal is nothing less than transforming an existing global climate model, the Community Climate System Model (CCSM), to create one that fully simulates coupling between the physical, chemical, and biogeochemical processes in the climate system.

John Drake heads the project's group at ORNL's Computer Science and Mathematics Division, and Peter Gent of the National Center for Atmospheric Research is chairman of the group's scientific steering committee. A group at Pacific Northwest National Laboratory contributes the component that calculates the chemistry of aerosols. Lawrence Livermore National Laboratory provides the chemistry component. Groups at Argonne National Laboratory work on coupling all of the components. The ORNL Leadership Computing team is contributing personnel from NCCS's Scientific Computing Group, which works directly with users, their software, and data, to get all the components ready for the jump to petascale computing.

"The people in our Scientific Computing Group have a lot of fingertip knowledge, meaning their hands are on the keyboard a lot in terms of writing code, and so they know the latest practices that can help these teams evolve their codes and algorithms to be more productive and more long-lived," Kothe says.

A background in physics and experience in high-performance computing allows Trey White, a research computer scientist at NCCS, to contribute to several projects. Presently he works on the Parallel Ocean Program (POP), a component developed at Los Alamos National Laboratory for the CCSM. He assists Pat Worley, a computer scientist in the climate group, to streamline POP and improve its scalability for petascale computing.

Worley, who has a long history of POP performance analysis, found a sluggish solver was one thing limiting its scalability. The code must find the sum of a value across all the processors, a step called an "all-reduce," which significantly slows the parallel operation. White is looking for a way to reduce the need for all-reduce queries by improving what's called the preconditioner. Meanwhile, Worley is making the entire operation faster by arranging the processes more efficiently.

"The Scientific Computing Group is really synergistic, since we are often members of the development teams for the applications and members of the operation team for the center," White says. "We act as advocates and proxies within the center for each project, and we bring detailed knowledge of the center computers, infrastructure, and plans back to the development teams."

White says POP could be among a handful of pioneering applications used to track performance when the Baker system first comes on line.

"The Scientific Computing Group exists because of the understanding that we are treading new ground; this isn't just using computers as is typically done, but really pushing the envelope on the scale of parallel computing," White says.

The Cray system is a "hybrid" XT3/XT4 (68 XT4 cabinets hooked to 56 XT3 cabinets, for a total of 124 cabinets).

Making the Transition to Petascale

"The NCCS is co-located at ORNL with many of DOE's premier science teams who use the computer systems," Kothe says. Several of these teams are already busy developing codes specifically for petascale machines. Two such ORNL projects received INCITE awards in 2007.

Robert Harrison, group leader of the Computational Chemical Sciences Group at Oak Ridge, is working with NCCS in an attempt to stave off program obsolescence by working ahead of the game. He's developing a quantum chemistry program that incorporates petascale capability into the initial design (see sidebar).

"Many codes have been caught sort of flat-footed in the race to get to the petascale because they have to make a transition from running on a few hundred or thousand processors to running on hundreds of thousands or millions of processors," Harrison says. "Typically what happens with older codes is that programmers take an existing program, parallelize it, and get stuck there, limiting their ability to migrate to even larger machines." Likewise, Thomas Schulthess of ORNL's Computer Science and Mathematics Division heads a research team that is developing a program called DCA++, a new-generation program that will predict behavior of materials at the nanoscale. The project's goal is ultimately to develop new materials designed on the nanoscale.

When these and other researchers are ready to run their simulations, members of the Scientific Computing Group will be there to serve as liaisons by scheduling time on the machine and smoothing the rough edges of the programs.

"Setting up queuing systems to be effective is a black art, particularly when you have really large jobs," White says. "Fitting everything in the available space/time holes can be as much a social science as technical. When a project has a big deadline coming up, we try to make sure they've got the resources they need."

Kothe adds: "We've mastered the art of user support, while pushing the envelope of the science that can be done on these machines. That's what makes us unique."

IT'S MADNESS

>> ***"Software gets slower faster than hardware gets faster."***
— Nicklaus Wirth, 1995

Robert Harrison has a beef with Wirth's Law. He's not saying software evolution hasn't kept up with the hardware in some cases, just that it doesn't have to be a self-fulfilling prophecy. Harrison, group leader of the Computational Chemical Sciences Group at ORNL, is working ahead of the game. He is part of a team developing MADNESS (Multi-resolution ADaptive Numerical Evaluation for Scientific Simulation), a next-generation software code intended from the outset to run on petascale computers still under construction.

The idea behind MADNESS is to harness the full capacity of petascale machines to solve problems that require complex quantum mechanical calculations. It is designed to be particularly useful for solving multi-scale problems in not just two or three dimensions, but up to 9 dimensions — calculations that would be unthinkable without massively parallel systems and a new application of quantum wave methodology.

"Our goal at the outset was to be very open-ended," he says. Therefore, MADNESS can run on existing systems, but should not be limited by ever-expanding hardware capacity.

One of the first applications to use MADNESS will be a simulation of the chemistry of heavy elements. The goal, explains Harrison, is to understand how to separate and reprocess spent nuclear fuel, thereby reducing waste and increasing its useful life. Exploring new fuel separations plants is a strategic goal at DOE, but doing the thousands of required experiments on spent fuel would be a practical impossibility. Realistic simulations of the behavior of uranium, cesium and other radioactive isotopes in a separation scenario could reduce experimental work and increase understanding of the fundamental behavior of this class of elements, says Harrison.

"People say nanoscale is different," says Harrison. "Well, petascale is different. It allows you to harness the power of complexity. We are just at the beginning to learning what it can do."

By Victor D. Chase

Modeling an Earth-Shaking Event

>> NERSC
>> Computational Science
>> Using LCC
>> Applied Math
>> Climate Modeling
>> Algorithms
>> Activities for LCC

"EARTHQUAKES DON'T KILL PEOPLE," seismologist Arthur Rodgers says. "Buildings do." "There are casualties in earthquakes because buildings collapse, freeway sections collapse, and bridges go out," says Rodgers, a member of an earthquake modeling team at the Department of Energy's Lawrence Livermore National Laboratory (LLNL). In essence, "We are vulnerable to earthquake damage because we choose to build and live near places where earthquakes occur."

To help understand and prevent such devastation, a team of computational scientists, applied mathematicians and seismologists at LLNL has created an earthquake simulation model as part of a larger Serpentine Wave Propagation (SWP) project. The project, headed by applied mathematician Anders Petersson, looks at the propagation of waves in nature. Whether they're seismic, electromagnetic, or sound waves, they're all governed by essentially the same mathematical equations.

The team spent several years developing the advanced mathematics and algorithms necessary to run computer models of wave propagation. For its first practical application, the team used the software to model what happened during the most famous and damaging earthquake in U.S. history.

It began at 5:12 a.m. on April 18, 1906. San Francisco Bay area residents were awakened when the San Andreas fault, a 296-mile fissure beneath the Pacific Ocean a few miles along the California coast and off shore, slipped. The displacement of a few meters along that fault line, where the Pacific and North American tectonic plates meet, was enough to set off one of the most monumental quakes in recorded history. After a 20-second foreshock, the full power of the quake was felt for about one minute. It would have measured 7.9 on the Richter scale — if the scale had existed then.

The shock was felt from Coos Bay, Oregon, to Los Angeles, and as far east as central Nevada. In all, the area of devastation was about 400 miles long, and 30 miles on either side of the fault zone. The quake and the resulting four-day San Francisco fire killed about 3,000 people, left 225,000 homeless, and destroyed about 28,000 buildings.

Poorly Understood

Earthquakes were poorly understood and little studied before the San Francisco disaster. The 1906 event put an end to that and marked the beginning of the science of seismology in the U.S. and gave rise to a more quantitative approach, applying physics and mathematics to the problem. Shortly after the quake, damage throughout the region was studied and quantified on the Modified Mercalli Intensity Scale. Unlike the Richter scale, the Mercalli scale does not require instrumentation. It rates a witness's impressions and physical damage to structures. Scientists can use Mercalli scale information to backtrack and determine the kind of ground velocities corresponding to the reported destruction.

The report has proven invaluable to later earthquake investigators, including the SWP team, which developed its computer simulation quake to mark the quake's 100th anniversary. Conducted under the leadership of the U.S. Geological Survey (USGS), the research also involved scientists from Stanford University, the University of California at Berkeley, and URS Corp., a worldwide engineering firm. DOE's Office of Advanced Scientific Computing Research supported the SWP group's modeling work.

Each participating group created its own model of the San Francisco earthquake using its own methods. The results were quite consistent, proving the exercise's value, Rodgers says.

Each participating group created its own model of the quake using its own methods. The results were quite consistent, proving the exercise's value, Rodgers says.

The exercise took almost two years to complete because of the complexity involved in creating a computer simulation of an earthquake. The rarity of major earthquakes in the 7.0 to 7.9 magnitude range, means there is less empirical data about them.

The findings of the centennial study were presented at the 2006 meeting of the Seismological Society of America, which was held in San Francisco to commemorate the100th anniversary of the famous quake — and of the society's founding.

To develop their simulations, each of the participating groups began with a USGS-created geological model of the Greater San Francisco Bay area. The model characterized rock and soil properties and was developed from years of study of seismic data, drilling, and tomography up to a depth of 50 kilometers. The data was crucial to creating a computerized picture of the earthquake in progress because different types of earth have dramatically different responses to the spreading shockwaves.

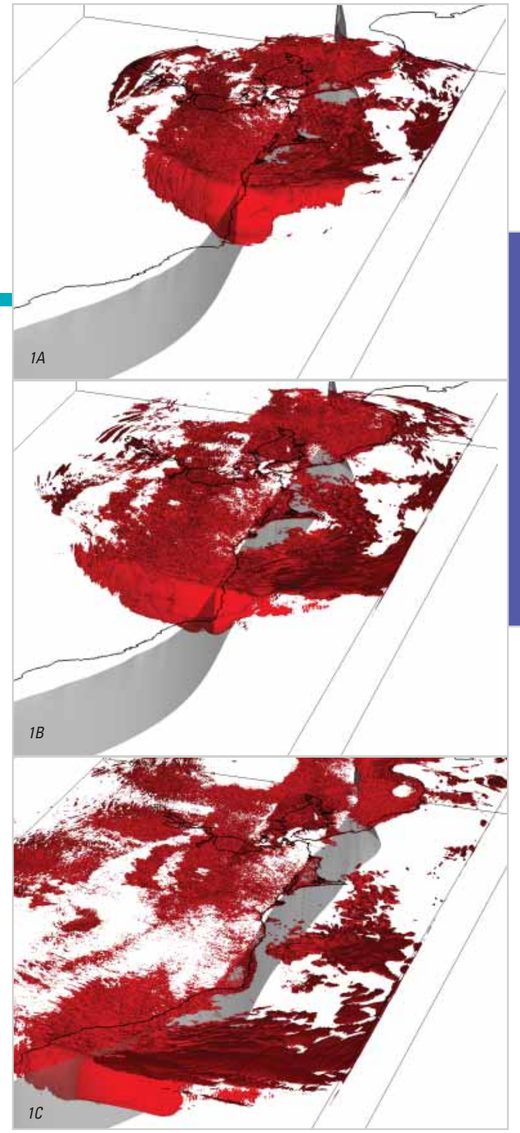

The propagation of shear waves (in red) through the three-dimensional model of the greater San Francisco bay area, looking south-east towards the bay area from the Pacific ocean. The San Andreas fault surface is shown in gray and the coast line of northern California is shown in black.

FIGURE 1A: 22.5 seconds after the start of the earthquake

FIGURE 1B: 30 seconds

FIGURE 1C: 60 seconds

COLLABORATORS

Anders Petersson is an applied mathematician in the Applied Mathematics group in the Center for Applied Scientific Computing (CASC). His research interests lie in the areas of grid generation and numerical solution of partial differential equations. Dr. Petersson earned his doctoral degree in Numerical Analysis from the Royal Institute of Technology in 1991. He joined the Lawrence Livermore National Laboratory in 1999.

Arthur J. Rodgers is a physicist and Group Leader of the Seismology Group in the Atmospheric, Earth and Energy Department at Lawrence Livermore National Laboratory. He received his Ph.D. and M.S. in Physics from the University of Colorado. Dr. Rodgers' research interests include computational seismology, earthquake ground motion simulation and nuclear explosion monitoring. He joined Lawrence Livermore National Laboratory in 1997.

Contact:
Anders Petersson
petersson1@llnl.gov

Arthur Rodgers
rodgers7@llnl.gov

PRACTICUM COORDINATOR

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

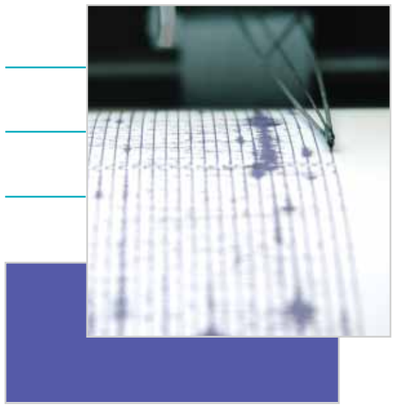

To create its model, the SWP team used two LLNL-based supercomputers:

- MCR incorporates 2,300 processors and has a peak performance of about seven teraflops (trillion floating point operations per second).
- Thunder has 4,096 processors and a peak performance of about 21 teraflops.

Doing the Math

One reason creating an earthquake model is difficult is that "a continuous elastic body, such as the earth, has an infinite number of degrees of freedom corresponding to motion at each point in space," Petersson, the SWP team leader, says. "So before the motion of a system with an infinite number of degrees of freedom can be calculated in a computer, we need to reduce the number of degrees of freedom to a finite, large number through a process called discretization."

The group used a finite difference discretization method. It broke the area being modeled, known as the computational domain, into a series of equally spaced grid points.

"In our calculations we used a grid spacing of 125 meters, which corresponds to something like 2.3 billion grid points," Petersson says. "At each of these grid points you have three degrees of freedom, so you get about 6.8 billion degrees of freedom. You reduce your infinite number of the degrees of freedom to 6.8 billion," which "is still fairly large."

The earthquake's motion also is integrated in time, which also must be discretized into a manageable number. The group took equal steps in time to simulate the first 300 seconds of earthquake motion. "Each of these steps is 0.01 seconds, so you take about 30,000 time steps in the calculation," Petersson says.

In discretization, "All the derivatives in the partial differential equations are replaced by what are called divided differences, and that converts the original mathematical equation to a set of algebraic equations, and those can be solved in the computer," Petersson says.

Without discretization, a computer cannot solve the original equations because they are too complicated — even for a supercomputer. "The computer can only deal with simple operations like addition, subtraction, multiplication and division. The partial differential equation involves very complicated relations between how the solution varies in space and time," says Petersson.

The Freedom is in the Details

But there is another factor to complicate the discretization process: stability. "If the method is unstable then perturbations due to round-off errors in the computer will accumulate during simulations, and they can make the computed result completely useless," Petersson says.

That's where the mathematician's expertise comes in. The key is to develop a mathematical theory that guarantees the stability of the process "so you know before you start your calculation that it is not going to go unstable, or 'blow up,' as we also say," Petersson adds. To do so, "Mathematically, you study how perturbations propagate through such a calculation without actually computing the solution. You can then estimate how large these perturbations can become in a calculation."

This is accomplished the old-fashioned way. "With pen and paper you can analyze the properties of your numerical method. And there are various ways you can modify your finite difference method so there's not just one prescription, there's a lot of freedom

in how you do all the details. So you've got to figure out the way to deal with all the details such that you can guarantee that it is stable. And that," Petersson says, "is the challenge."

Earthquake in a Box

Rodgers describes the SWP modeling effort from a seismologist's perspective. "Think of it as though we built a box that represents the Bay area in three dimensions," he says, "and we put an earthquake in the box, and the earthquake sets the box in motion. By doing so we were able to put a simulated seismic station that measures the ground motions as a function of time at any place, which allowed us to compute the ground motion anywhere."

Modeling the San Francisco earthquake has significance beyond commemorating the event. The ability to create a computer simulation of such a complicated occurrence enables scientists to model other earthquakes before they happen. Architects and civil engineers can use data gathered from those models to design structures that withstand tremors.

What DOE's supercomputers and scientists can't do is predict exactly where or when earthquakes will strike.

"We know that earthquakes are going to happen," Rodgers says. "The problem is that we have only been looking at earthquakes in detail for about 100 years. The return times of large earthquakes are hundreds, if not thousands of years, so we haven't got a statistical sample to allow us to do meaningful statistics."

Nonetheless, "In the Bay area the most likely next earthquake will be along the Hayward fault," he says. This supposition is based on geological studies indicating there have been 11 earthquakes along the fault at intervals averaging 140 years. The last such quake occurred in 1868, making 2008 the 140th year.

Using modeling, the researchers "can put in a hypothetical Hayward fault magnitude 7 earthquake, and see what happens," Rodgers says. Although models cannot tell experts precisely where along the fault the earthquake will start or in what direction it will run, "We can do lots of simulations to look at how the ground motion might vary depending on those types of factors."

Such scenarios are valuable because "Our data set of actual large earthquake shaking is limited," Rodgers say. "So this modeling effort is very important because it allows us to, in the safety of our computer, compute the shaking that would occur if an earthquake were to happen on a specific fault of a certain size within a certain geology."

From the Soil Up

Such projections are especially needed to avoid accidents at the many nuclear power plants being considered to meet increasing worldwide energy demands. Seismic safety of nuclear power plants is guided by observed, as well as computed, ground motions. The same computer modeling can also be used to simulate potential damage should an earthquake impact nuclear storage facilities, such as the controversial Yucca Mountain site. "Then you would know how to design containers to withstand the possible motions of the Earth," says Petersson.

Now that the SWP team has created a program that models how earthquake waves propagate from the source through rocks and soil to the foundations of buildings, the next logical step will be to follow those waves up from the soil through complicated structures, such as nuclear power plants, airports, and bridges, to learn how they will respond to the shaking of an earthquake.

"The foundations of buildings are embedded within soil so they need to be modeled together," said Rodgers. "We would like to be able to model what's called the soil/structure interactions."

Meanwhile, the SWP team recently made some waves of its own by receiving an internal award from the DOE's Energy and Environment Directorate for its work on the 1906 earthquake.

Model Mixes Ice, Heat, Water and Salt to View the Ocean's Future

Looking to the Sea for Clues to the Climate

By Victor D. Chase

>> NERSC
>> Computational Science
>> Using LCC
>> Applied Math
>> Climate Modeling
>> Algorithms
>> Activities for LCC

WITHIN THE NEXT SEVERAL DECADES, ice over the Arctic will completely disappear during the summer. That's just one of the clear and dramatic predictions to come from models developed by the Climate Ocean and Sea Ice Modeling (COSIM) program at the Department of Energy's Los Alamos National Laboratory.

Additionally, "The models now predict globally an average surface temperature rise of 2 to 4 degrees Celsius over the next 100 years," says the program's project manager, Philip Jones. And the oceans account for more than 70 percent of the Earth's surface.

Jones and his group of 12 researchers use the power of supercomputer simulations to study how Earth's oceans are affected by and influence global warming. Their specialty is modeling ocean circulation and how it affects both heat transport and sea ice up to 100 years into the future.

Results from his team and other research lead Jones to conclude that some aspects of climate change could take place more rapidly than previously anticipated. "There are amazing changes going on right now," he says.

The loss of summer Arctic sea ice provides a case in point. Satellite observations since 1978 show that the average annual sea ice shrinkage has been some 2.7 percent per decade with summertime average losses of 7.4 percent. At some point in the not-too-distant future, the ice will become so thin that one warm summer — or as Jones puts it, "one pulse of warm water" — will be enough "to cause a rapid transition to an ice-free state." This will happen, he anticipates, within 30 to 40 years.

The loss of Arctic sea ice, dramatic as it may be, will not directly raise the sea level, since the ice already displaces water as it floats on the surface. The disappearance of sea ice can, however, indirectly affect sea level since the water temperature will rise, causing thermal expansion. This is significant because the oceans' huge mass makes them act like a heat sink, absorbing 80 percent of the climate's temperature increase.

The oceans also are slow to respond to a rise in global temperature, especially at depths below the top few hundred meters.

"What that means is even if we stop putting greenhouse gases into the atmosphere today, we are committed to another half-degree or so of warming just because the oceans haven't caught up," Jones says.

Thermal expansion is expected to increase ocean levels by some 30 centimeters over the next 100 years. That's not a large increase, but it will affect some areas and will have secondary effects on things like storm surges. A much more dramatic rise in sea level would occur if huge land-based glaciers on Greenland and Antarctica melt — a phenomenon that's already being observed.

Results from his team and other research lead Jones to conclude that some aspects of climate change could take place more rapidly than previously anticipated. "There are amazing changes going on right now," he says.

Glacial Decline

A recent report notes that the decline of glaciers and mountain ice caps in both hemispheres has already led to sea-level increases — and that doesn't include melting from Greenland and the Antarctic. That conclusion is based on the increased flow speed of ice draining from the interior of the ice sheets, with a corresponding increase in ice loss.

Because "Current climate models don't represent ice-sheet changes very well, we are in the process of putting together some new ice-sheet models to look at the sea-level rise issue," Jones says. The goal is "to figure out how rapidly that ice is going to melt, because if Greenland melts you get about six meters of sea level rise, and that is pretty significant."

Though not a cause for complacency, the time frame for such occurrences is on a geologic scale. Because Greenland's ice sheet, for example, is several kilometers thick, "The question is whether a complete melting will take a thousand years or a few hundred years," Jones says.

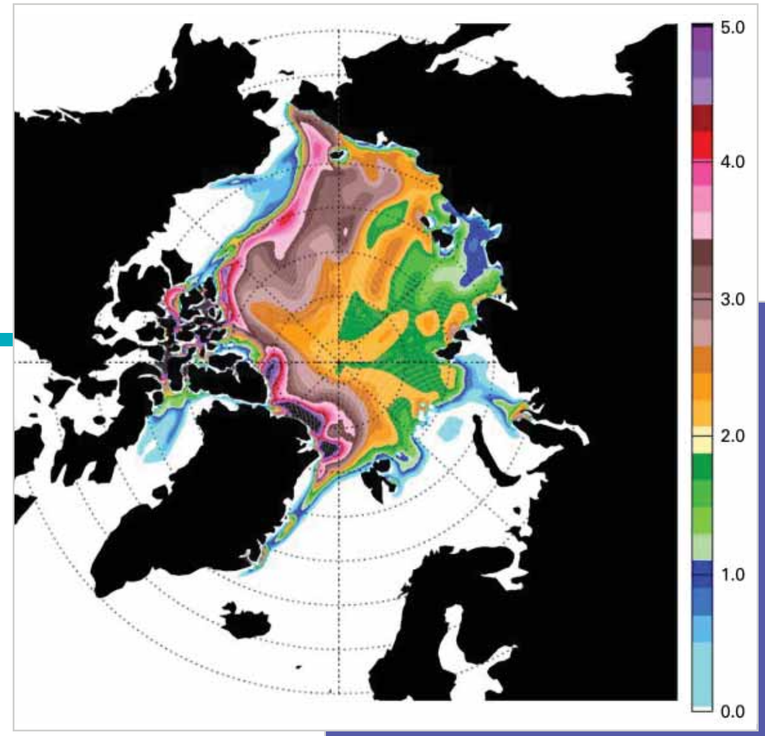

Sea ice thickness (m) in September 1998 from an ocean-ice simulation.

Graphics provided by Elizabeth Hunke.

The Cause Is Us

Analyses of ice core samples dating back eons, combined with modern-day computer simulations, have firmly established the cause of global warming — and it is us.

The power of the Cray high-performance computers the COSIM researchers use allows them to turn on or off various factors that may or may not contribute to global warming. That's let them confirm that greenhouse gases generated by burning fossil fuels cause climate changes. "We can turn off the industrial, human-caused carbon dioxide emissions and artificially pretend that humans didn't exist, and see what the natural forcing is," Jones says.

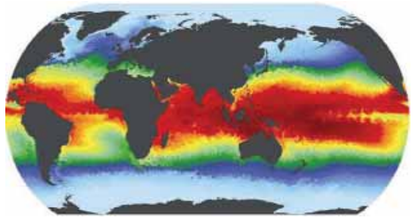

Sea surface temperature at 15m depth from an eddy-resolving ocean simulation.

Graphics provided by Mathew Maltrud.

PHILIP JONES

Philip Jones is leader of the Climate, Ocean and Sea Ice Modeling Project within the Theoretical Fluid Dynamics Group at Los Alamos National Laboratory. He also is the lead software developer of the POP ocean model. His research interests include coupled climate modeling, ocean modeling, remapping and interpolation and computational performance of climate models. He holds a doctoral degree in astrophysical, planetary and atmospheric sciences from the University of Colorado and a bachelor's degree in physics and math from Iowa State University.

Further reading:

IPCC, 2007: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. B. Averyt, M. Tignor and H.L. Miller (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 996 pp. Also available at: http://ipcc-wg1.ucar.edu/wg1/wg1-report.html.

W. D. Collins, C. M. Bitz, M. L. Blackmon, G. B. Bonan, C. S. Bretherton, J. A. Carton, P. Chang, S. C. Doney, J. J. Hack, T. B. Henderson, J. T. Kiehl, W. G. Large, D. S. McKenna, B. D. Santer, and R. D. Smith. *The Community Climate System Model: CCSM3,* J. Climate, 19, 2122-2143 (2006).

J. B. Drake, P. W. Jones, and G. Carr Jr. *Overview of the Software Design of the Community Climate System Model.* Int. Journ. High Perf. Comput. Application, 19, 187-201 (2005).

Contact
Philip Jones
pwjones@lanl.gov

PRACTICUM COORDINATOR

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

"When you do that, in the early part of the century you can explain some of what we see via natural warming, but anything from about 1975 onward can't be explained by any kind of natural forcing — either solar variability or any other natural process. Unless you include the human contribution to greenhouse gases you can't match the observations," he says. Based on that data, current climate change scenarios show a 1 to 2 percent increase in carbon dioxide per year, or a doubling over the next 100 years.

No Ice Age

Improved ocean and ice models have let climate scientists put to rest at least one of the more ominous global warming scenarios, in which changes in the Atlantic Ocean's currents have a dramatic cooling effect on Northern Europe. Such cooling brought on a cataclysmic ice age in the popular 2004 movie, *The Day After Tomorrow*, which was replete with scientists rushing to model climate changes on supercomputers.

Real-life scientist Jones, however, says this scenario will not take place, at least not due to the driving force, known as thermohaline circulation, depicted in the movie.

This circulation pattern acts like a heat pump, transporting heat from the equator to the North Atlantic. It is primed by water cooling as it moves northward. As it does, water evaporates and ice forms. Each of those actions rejects salt, leaving a higher salt content in the ocean as it flows northeast. As it approaches Northern Europe the cold, heavier, salt-laden water sinks, completing the circulation and drawing more warm water northward, warming the nearby land mass.

The concern has been that as the ocean surface warms and ice melts rather than solidifying, the added fresh water at the surface, having less salt content, would not sink. This could slow the ocean circulation and ironically have an ominous chilling (rather than warming) effect on Europe.

Ironically, climate modeling shows that though melting sea ice does slow thermohaline circulation, the cooling effect is largely counterbalanced by global warming. As Jones puts it, "The fact that the northern latitudes are warming faster than the rest of the globe tends to overwhelm the fact that you are missing a little bit of ocean circulation. So we are less worried about that scenario than we used to be. I guess that's good news in a sense, although I would not want people to get complacent."

Impact Scenarios

The COSIM group receives support from DOE's Office of Biological and Environmental Research and from DOE's Scientific Discovery through Advanced Computing (SciDAC) program, which matches mathematics and computational science with physical scientists in various areas. Climate modeling has been going on at Los Alamos for some 16 years, and Jones has been a part of it almost from the beginning.

Its work is part of the larger Community Climate System Model (CCSM), headquartered at the National Center for Atmospheric Research (NCAR) which is based in Boulder, Colorado, and supported by the National Science Foundation. The two groups jointly develop the model and use it to project climate change scenarios.

As part of the CCSM collaboration, the Los Alamos group also feeds data to the Intergovernmental Panel on Climate Change (IPCC), which uses them in its influential climate change reports (see sidebar).

Jones says that since his group's inception, models have gotten good enough "that we feel much more confident in believing in what we are predicting." That's partly due to a collaboration with computational scientists at DOE's Oak Ridge and Lawrence Livermore national laboratories, who help ensure the COSIM team gets optimum performance from the Oak Ridge-based Cray computers they use for modeling.

The DOE is interested in climate change because of its impact on the nation's energy portfolio. "The DOE is tasked with determining what are 'safe levels of carbon dioxide' in the atmosphere," says Jones. This information helps determine what will be used to generate power in the future — "Whether it's still going to be coal and oil or whether we need to put more emphasis on other forms of energy," he adds.

From Stars To Earth

Jones is trained as an astrophysicist, and comes to modeling global warming's effects on oceans and sea ice from studying how heat is transported from the inside of a star to the surface. Though the temperatures are quite different, the fluid dynamics are essentially the same.

When he joined Los Alamos as a post-doctoral fellow, Jones applied his knowledge of fluid dynamics to study fluid motions inside the Earth. "The similarity among these is that they all require high-end computation and the same fluid equations," he explains.

When he isn't modeling the oceans and sea ice or spending time with his family, Jones entertains the public with his trombone, which he has played since fourth grade. He has an eclectic repertoire, and performs with the Los Alamos Symphony, the Los Alamos Big Band (which plays '40s and '50s dance tunes), and a brass quartet that plays everything from baroque music to Dixieland jazz.

Clearly, Jones is a man for all seasons and temperature climes.

DOE CONTRIBUTES TO CLIMATE'S BIG PICTURE

>> ***To put the current global climate into "paleoclimatic perspective," it has been about 1,300 years since it was as warm as during the last half century; and one needs to look back some 125,000 years to find a time when the polar ice regions were significantly warmer than they are now.***

These estimations are based on studies of indicators such as tree ring width, ancient ice samples, and computer simulations. They're reported in "Climate Change 2007," issued by the Intergovernmental Panel on Climate Change (IPCC). The World Meteorological Organization and the United Nations Environment Programme created the panel in 1988.

The IPCC combines the computational findings of the joint National Science Foundation/Department of Energy Community Climate System Model (CCSM) program with data from 17 other groups to create its influential periodic report. The 2007 version is the panel's fourth, and is divided into three sections. CCSM's contribution is included in the "Physical Science Basis" section, prepared by Working Group I. The other two sections are "Impacts, Adaptation and Vulnerability" and "Mitigation of Climate Change."

In its "Summary for Policymakers," Working Group I writes that carbon dioxide is the most important greenhouse gas created by human activity: "Global atmospheric concentrations of carbon dioxide, methane and nitrous oxide have increased markedly as a result of human activities since 1750 and now far exceed pre-industrial values determined from ice cores spanning many thousands of years. The global increases in carbon dioxide concentration are due primarily to fossil fuel use and land use change, while those of methane and nitrous oxide are primarily due to agriculture."

The summary adds: "Warming of the climate system is unequivocal, as is now evident from observation in global average and ocean temperatures, widespread melting of snow and ice, and rising global average sea level."

Even if greenhouse gases held constant at 2000 levels, the report says, global warming would continue for the next two decades due primarily to a lag in the response of the oceans. If greenhouse gas emissions continue at or above current levels, continued warming would "induce many changes in the global climate system during the 21st century that would very likely be larger than those observed during the 20th century." These changes would likely include more frequent heat waves, some heavy precipitation, and more intense tropical cyclones.

It's certain the effects of greenhouse gas emissions will be felt far into the future, even if humans immediately stopped adding significant amounts to the atmosphere, but that need not spell doom. The warning signs are more than abundant, yet some scientists feel there is still time to act. How much has yet to be modeled.

By Alan S. Brown

Harnessing Hundreds of Thousands of Processors

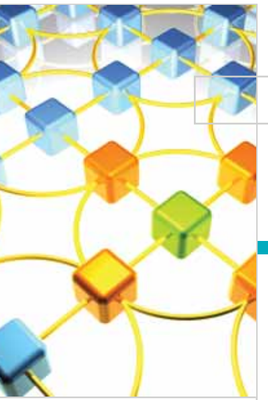

>> NERSC

WHAT DO ANCIENT CHARIOTS and modern high-performance supercomputers have in common? Both started with a single engine pulling the load. With chariots, that motor was a horse. To go faster, drivers bred larger, stronger horses.

With supercomputers, it was a high-performance processor. To boost performance, engineers created bigger, faster, and more expensive processors.

The big breakthrough in chariots came when drivers harnessed teams of horses. By distributing the load among two, four, or even six horses, chariots could go faster and farther than any one horse could pull them. For this to work, however, drivers had to train their horses to start, stop, and turn together. Otherwise, each animal would bolt in a different direction during the tumult of battle.

Supercomputer architects made a similar breakthrough. They learned to break big calculations into smaller parts that multiple processors could solve simultaneously. Instead of a single processor, computers were composed of many networked processors, pushing calculation speeds into the stratosphere.

Parallel computers also slashed costs, since they replaced expensive custom processors with the same off-the-shelf chips used to run servers and workstations. The world's fastest computer, IBM's BlueGene/L, uses more than 130,000 IBM Power PC processors. A new model unveiled in 2007, the BlueGene/P, will use almost 900,000 processors.

This is where the similarity between chariots and high-performance supercomputers breaks down. The mind easily wraps around the task of training a handful of willful horses. It's simply a matter of training one horse to lead and the others to follow. But how do you yoke together hundreds of thousands of computer processors in a single harness? How do they share data, store information, or recover when one node fails?

"Our goal is to find a practical solution for those problems," says Jarek Nieplocha of the Department of Energy's Pacific Northwest National Laboratory (PNNL) in Richland, Washington.

"In academia, most people do research, write a paper, and then forget about it and move onto another problem," he explains. "In DOE national laboratories we're developing real applications here, and our goal is to create software that has an impact on science. We want to create something that benefits everyone."

Managing Memory

Memory is a key issue in supercomputers. In conventional PCs, each processor has its own local memory. In a parallel supercomputer, however, every processor has its own local memory, but also shares memory with the hundreds or thousands of other processors in the system. In large simulations, where the results of one set of calculations drive another and another and another, processors are constantly reading, writing, and exchanging data in memory.

Most supercomputers share information using an approach called Message Passing Interface (MPI). It's a complicated technique that Nieplocha compares to postal delivery.

"Let's say I want to send data," he begins. "So I put it in an envelope and drop it in the mailbox. The post office comes, picks up my letter, and delivers it to a post office in another city, which puts it in the right mailbox. Then someone goes to the mailbox and picks up the letter."

In computers, "This two-sided process requires handshaking: both the sender and the receiver need to know what's being moved and agree on when to do it. It is laborious, expensive, and difficult to program this type of communication."

Nieplocha's alternative is an approach called Global Arrays. "You don't have to send data because it is visible to every processor without handshaking," he says.

This works because users first define where all their data will be within a matrix, or array. The array itself is a logical object — that is, software sees it as a single entity even though portions of the array (and the data it holds) may reside on hundreds or even thousands of different memory chips. As long as users know which part of the array holds their data, they can access it without complex handshaking and tracking routines.

Shared memory, however, comes at a price. "One-sided sharing is like accessing somebody else's mailbox without involving the post office, so everybody has to know where everything is," Nieplocha says. "That means Global Arrays must maintain indexes to track the physical location of the data, and employ communications techniques that optimize how the data flows between processors."

The payoff, however, is huge. Instead of spending time describing handshaking routines for thousands of memory locations, programmers can access shared memory the same way they would on an individual PC. They don't even need to know the underlying mechanics of memory manipulation.

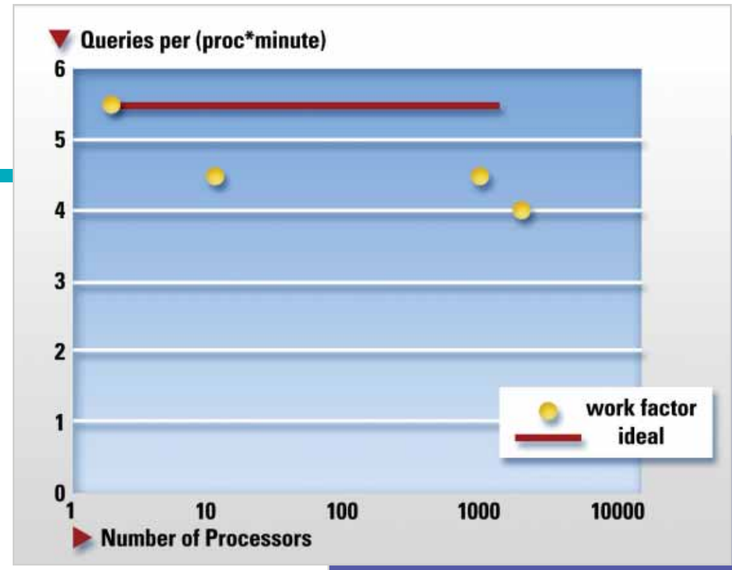

On a recent benchmark, we demonstrated that the per-processor throughput for analyzing biological sequence data using ScalaBLAST is essentially constant whether one uses 8 or 1,800 processors. This scaling characteristic is due, in part, to efficient memory management enabled by Global Array features for hiding latency and replacing repeated global file I/O with memory operations. This same approach has been used to perform trillions of pairwise homology calculations in a single day. Figure from C. Oehmen, J. Nieplocha "ScalaBLAST: A scalable implementation of BLAST for high-performance data-intensive bioinformatics analysis." IEEE Transactions on parallel and distributed systems, 2006:17(8), 740-749.

The mind easily wraps around the task of training a handful of willful horses. It's simply a matter of training one horse to lead and the others to follow. But how do you yoke together hundreds of thousands of processors?

JAREK NIEPLOCHA

Jarek Nieplocha is a Laboratory Fellow and Deputy Director of Computational Sciences and Mathematics Division of the Fundamental and Computational Science Directorate at Pacific Northwest National Laboratory (PNNL). He is also the Chief Scientist for High Performance Computing in Computational Sciences and Mathematics Division. He leads Advanced Computing Technology Laboratory at PNNL.

His area of research has been in collective and one-sided communication on modern networks, runtime systems, parallel I/O, scalable programming models, multithreading architectures and, and component technology for scientific computing. He has led development of Global Arrays, a portable shared memory programming toolkit widely used in scalable computational chemistry applications and other areas, and the ARMCI runtime system for global address space programming models.

Further Reading:

S. Benson, M. Krishnan, L. C. McInnes, J. Nieplocha, J. Sarich. *Using the GA and TAO toolkits for solving large-scale optimization problems on parallel computers.* ACM Trans. Math. Softw. 33(2): 11 2007.

C. Oehmen and J. Nieplocha. *ScalaBLAST: A scalable implementation of BLAST for high performance data-intensive bioinformatics analysis.* IEEE Trans. Parallel Distributed Systems. Special issue on high-performance computational biology, Vol. 17, No. 8, 2006.

M. Krishnan, S. J. Bohn, W. E. Cowley, V. L. Crow, and J. Nieplocha. *Scalable Visual Analytics of Massive Textual Datasets*, Proc. IEEE International Parallel and Distributed Processing Symposium, 2007.

J. Nieplocha, B. Palmer, V. Tipparaju, M. Krishnan, H. Trease, and E. Apra. *Advances, Applications and Performance of the Global Arrays Shared Memory Programming Toolkit.* International Journal of High Performance Computing and Applications, 20(2), 2006.

D. P. Scarpazza, P. Mullaney, O. Villa, F. Petrini, V. Tipparaju, D. Brown, and J. Nieplocha. *Transparent System Level Migration of PGAs Applications using Xen on Infiniband.* Proceedings of the IEEE International Cluster Conference 2007.

Contact:
Jarek Nieplocha
Jarek.nieplocha@pnl.gov

PRACTICUM COORDINATOR

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

Nieplocha and colleague Chris Oehmen recently put Global Arrays to work in BLAST, a bioinformatics program. BLAST lets researchers match snippets of newly-sequenced DNA or protein structures with known genetic information. In August 2006, Nieplocha and Oehmen published a paper showing that their new global array-based ScalaBLAST software is the most scalable and highest-performing parallel implementation of the BLAST algorithm.

Smarter Memory

With so many processors yoked together in modern supercomputers, the researcher explains, the speed of an individual processor hardly matters. Instead, the most common impediment to high-performance computing is the speed at which the network shuttles data between memory, processors, and hard disk storage.

By doing away with time-consuming handshaking, Global Arrays let data flow smoothly along the network to tens of thousands of thirsty processors. But what if supercomputers could process some types of data without even taking them out of storage?

Storage on computers means hard drives, which act like file cabinets to stockpile information. When computers need the data, they go to the cabinet, open a file, read it, and do something with the data. Then they write the results onto another file in the cabinet.

But what if the filing cabinet were smarter? What if you could write measurements in feet and inches, but pull out the same data in meters and centimeters when you needed metric units?

Better yet, suppose you wanted to study rainfall at certain elevations. Ordinarily, the supercomputer would have to send all the rainfall data to a processor to sort out the information you want. But suppose your hard drive did the sorting for you? This would reduce the flow of unwanted data over the network and also cut the time processors waste on doing menial calculations.

Visionaries first suggested the concept, called active or intelligent disks, 20 years ago. They proposed using small processors to run simple calculations on stored data. Hard drive manufacturers rebelled at the idea. They did not want anyone tinkering with drive electronics, since it might cause unexpected failures.

In 2005, the PNNL team found a way to do those calculations using another part of the computer that's usually better left alone: the kernel. The kernel contains the core functions of the computer operating system. Like brain surgery, tinkering with the kernel is both delicate and costly. "In the long term, such modifications are too complex to be practical," Nieplocha says.

One year later, however, the group developed a way to achieve the same results with software that intercepts requests for data as they move through the network.

"Now we don't have to modify the operating system or any system software," Nieplocha explains. "And we can also do some types of processing, such as floating point computation, that you cannot do in the kernel." The group is testing the approach on bioinformatics software. "We still have a few issues, but there are no major research challenges," he concludes.

Soft Landings

In addition to harnessing the power of healthy supercomputers, Nieplocha tries to prepare for the unavoidable failures. While the individual components of any supercomputer — processors, circuit boards, hard drives, power supplies — are highly reliable, the sheer numbers of components used in high-end systems make failure a certainty.

Supercomputer failures carry a high price in time and money. A single simulation may take months of preparation. Researchers may wait weeks for computer time. And some big simulations take days or even weeks to run.

"Most applications running on supercomputers are not fault-tolerant," Nieplocha says. "If a hard drive or memory chip fails, users typically lose their data and have to restart from scratch."

One solution is to save the work periodically. PC word processors and spreadsheets do this automatically by making mirror images of documents. Saving a simulation that sprawls across thousands of different nodes and processors is far more complex. "You have to think about what to save and where and when to save it. The applications are very complex, going through different types of stages, and some points simply cannot be saved," Nieplocha explains.

To resolve these issues, Nieplocha's team uses virtualization, an approach that traces back to mainframe computer days. It involves slipping a layer of software, called a hypervisor, between the hardware and the operating system that controls it. Instead of delivering instructions to the computer, the operating system talks to the hypervisor, which whispers those instructions to the hardware.

This sounds inefficient, but recent advances have made hypervisors much more economical and, for the first time ever, worthwhile, Nieplocha says. "Today, using a modern hypervisor, we will periodically stop each node and map the application state and operating system memory. Our software writes the image to a hard drive, then moves on to the next node. If a node fails, we will sense the problem, retrieve the saved image, and mount it on a healthy node so that the application can continue from the last checkpoint," Nieplocha says.

"People have worked on this problem for some time, but the results were specific to individual computers or basic academic research," he adds. "People describe how they made it work, write a paper and move on to other challenges. We're creating a more practical solution that can manage multiple nodes automatically and reconfigure the system if there is a failure."

Nieplocha's team understands that virtualization degrades computer speed, though they believe their technique makes the slowdown negligible. Given the high cost of failure, they believe most users would prefer a minor delay to a failure that leaves them without critical results for months at a time.

Fault tolerance was never much of a problem in a chariot, but then, chariots were relatively simple machines. Modern high-performance supercomputers are anything but simple. With their cluster-based architecture, they promise an unprecedented combination of flexibility and power. One day they may help us produce pollution-free fusion energy or develop new genetically based medications.

The goal of Jarek Nieplocha and his team is to make that power and flexibility more easily accessible to those who need it. It's just a matter of harnessing all those processors so they pull in the same direction.

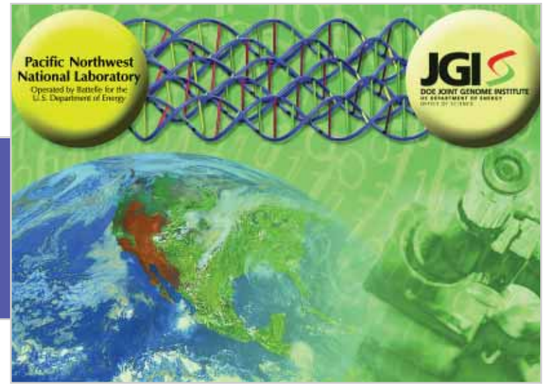

ScalaBLAST gives a sizeable performance boost over BLAST, a conventional sequence analysis tool.

http://picturethis.pnl.gov/PictureT.nsf/webpages/Search

STANDARDS MAKE ADVANCED COMPUTING FASTER — AND EASIER

>> ***Nearly all PC software conforms to standards. This is why software components written using a wide variety of languages — from C++ and Java to Visual Basic and Perl — all run and exchange data with one another on the same PC.***

"PC components work with each other all the time," notes Jarek Nieplocha of Pacific Northwest National Laboratory (PNNL). "If you open a Word document with an embedded graph, the underlying component technology understands that the graph is an Excel object. When you click on the graph, it opens Excel. These programs know about each other and work together."

In recent years, the Common Component Architecture (CCA) community has been adopting this technology for scientific supercomputing.

"There are supercomputer chemistry applications that need to perform some sort of optimization, such as a minimum energy function," Nieplocha explains. "These optimization routines are readily available in science libraries, but in the past programmers had to struggle with data structure and code to adapt them to their chemistry program.

"Today, we're seeing common interfaces that let us glue those two programs together so they work like a single application. If someone develops a better routine, they could swap it for the old one in a completely transparent way and it would work," Nieplocha says.

The PNNL team is working with ways to extend Common Component Architecture to specialized hardware accelerators, such as field programmable gate arrays (FPGAs) and graphics processors. While not as flexible as conventional processors, these specialized chips excel at accelerating repetitive calculations.

"We want to hide the complexity of moving data onto the accelerator, so programmers have a simple way to take advantage of its speed," Nieplocha says.

By Karyn Hede

The New Paradigm of Petascale Computing

WHAT WOULD HAPPEN if a magnitude 8.0 temblor wrenched the earth beneath a nuclear power plant? Can nano-scale technology be used to build a new generation laser?

From planning for the worst case to designing for the best case, there are some questions that we would like to answer without having to find out the hard way. No one wants to experience a nuclear reactor failure firsthand, or find out that the expensive nanoscale materials they designed won't work at the macroscale. Scientists would like to be able to use computer simulations to predict behaviors of large systems. But simulating a devastating earthquake or a nanostructure assembly has been hampered by the sheer complexity of the problems.

These kinds of multi-scale simulations require tremendous computational power, says Rick Stevens, associate laboratory director for Computing and Life Sciences at the Department of Energy's Argonne National Laboratory. Stevens leads Argonne's advanced computing initiative, targeting the development of petaflop computing systems capable of handling the kinds of computations multi-scale problems require.

The Department of Energy has tapped Argonne to become its second Leadership Computing Facility by early 2008, when the world's first IBM BlueGene/P class system comes on-line. The system, capable of petaflop performance, will be available for open science and engineering for applications granted computing time through DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program.

Stevens says the Leadership Computing program represents more than another leap in computational power. It is, quite simply, a new paradigm for doing science.

"This is different from the way supercomputing centers traditionally are run, where you might have dozens or more users on the system simultaneously," he says. "With the Leadership Computing Systems, there might be only one or two groups using the system at any one time. But what they are doing is very hard. They might be running across the whole machine, which is always interesting, because of the sensitivity to machine 'hiccups,' and these big runs generate terabytes of data so real-time data management is essential. It requires a full-time team working around the clock to manage it all. Then at the end, another team helps the users analyze their data."

The support team at Argonne will form long-term partnerships with the scientific teams as projects go to petascale on BlueGene/P and researchers begin to analyze results. But, equally important, teams of Argonne researchers will work to ensure that when the scientific applications are ready, the accompanying software will be up to the task.

"There's a multistep process to get codes ready for the BlueGene/P," says Stevens. "If they are already running on the BlueGene/L then they are in pretty good shape." He says there are currently about 100 applications that run on BlueGene/L, ranging from astrophysics to chemistry to economics. But some application scientists will need to modify their software to get ready for the BlueGene/P.

"The structure of leadership computing means that you can have close collaborations between teams over time," Stevens says. "That's really the vision and the promise of this approach."

BlueGene/P Short List

Stevens has a short list of applications that are likely to be the first to access the BlueGene/P system. One such application is the earthquake simulations led by Arthur Rodgers, Anders Petersson and David McCallen of Lawrence Livermore National Laboratory (LLNL).

As the U.S. searches for new sources of energy, there is renewed interest in building new nuclear power plants. But safety concerns about earthquake hazards, particularly in the western U.S., require detailed studies of potential nuclear plant sites.

"Ideally, one would like to be able to model how an earthquake fault might rupture, how those seismic waves may propagate through the earth, and how they may arrive at your site," says McCallen, LLNL division leader in nonproliferation, homeland and international security. Historically, such predictions have been made based on past earthquakes and then extrapolated to predict future earthquake damage. But these methods have limitations. "You can make better predictions about the ground movements at a given site from physics and what we call 'first principles.'" The problem is that combining data from subsurface geology and earthquake fault models has been prohibitive because there simply has not been enough computational power available to do the simulations.

RICK STEVENS

Rick Stevens is Argonne's Associate Laboratory Director for Computing and Life Sciences.

The Computing and Life Sciences directorate is made up of the Biosciences Division, the Leadership Computing Facility, and the Mathematics and Computer Science Division, along with the Computation Institute and the Institute for Genomics and Systems Biology.

Stevens has been at Argonne since 1982, and has served as director of the Mathematics and Computer Science Division and also as Acting Associate Laboratory Director for Physical, Biological and Computing Sciences. He is currently leader of Argonne's Petascale Computing Initiative, Professor of Computer Science and Senior Fellow of the Computation Institute at the University of Chicago, and Professor at the University's Physical Sciences Collegiate Division. From 2000-2004, Stevens served as Director of the National Science Foundation's TeraGrid Project and from 1997-2001 as Chief Architect for the National Computational Science Alliance.

Stevens is interested in the development of innovative tools and techniques that enable computational scientists to solve important large-scale problems effectively on advanced scientific computers. Specifically, his research focuses on three principal areas: advanced collaboration and visualization environments, high-performance computer architectures (including grids) and computational problems in the life sciences, most recently the computational problems arising in systems biology. In addition to his research work, Stevens teaches courses on computer architecture, collaboration technology, virtual reality, parallel computing and computational science.

Contact:
Rick Stevens
stevens@anl.gov

PRACTICUM COORDINATOR

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

"This is an application that demands as much computer crunching and computer power as you can muster up," he says. "It's a heavily three-dimensional computationally intensive problem. Even today, our biggest computers will be taxed by this application. That's a good example of why there is a motivation to use the emerging leadership computing facilities at Argonne."

With the additional computing power, McCallen would like to construct fault models of an area of 50 to 100 kilometers across by 80 kilometers deep and then combine that data with subsurface geologic models to simulate what the ground motion might be at a given site for different earthquake scenarios.

To do that, the group uses Wave Propagation Program (WPP) developed through laboratory-directed research at Lawrence Livermore, a program that is expected to scale up to petascale relatively easily, since it has already proven it runs efficiently on a massively parallel computing architecture.

An Argonne group, led by Bob Hill, is working on next-generation nuclear power plant design, so there is a natural link between the groups, says McCallen.

Similarly, Argonne is teaming with François Gygi at the University of California at Davis and his research team to examine the fundamental physical nature of nanomaterials that gives them unique properties.

One of the biggest challenges in nanoscience, according to Gygi, is learning how to control nanostructure assemblies. The scientists want to understand and control the interacting of nanoscale building blocks used to design new materials. The computational challenge, then, is to predict interface properties at the nanoscale.

Gygi, in collaboration with Giulia Galli (also at UC Davis) wants to use BlueGene/P to study silicon dots embedded in silicon nitride and carbon nanotubes in solution. The investigation will focus on whether it is possible to build a silicon laser out of dots assembled on a matrix. The application runs on Qbox, a program built on an Message Passing Interface (MPI) framework and calculating first principles molecular dynamics. It is one of a handful of programs that have already demonstrated the ability to take full advantage of Leadership Computing systems.

In 2007, Qbox set the world record for floating-point performance by achieving a sustained performance of 207 teraflops on the LLNL BlueGene/L 65,536 node system. That corresponds to 56.5 percent of the theoretical peak using all 128,000 CPUs. Gygi's team simulated the electronic structure of molybdenum, a high-Z or heavy metal, that is important to the National Nuclear Security Administration's (NNSA) Stockpile Stewardship Program. In this case, the study was the first step toward simulating the effects of aging on nuclear materials.

Future Programming

Capitalizing on their strengths in scientific computing, teams of Argonne researchers are already working on next-generation programming. Current linear solvers, a key component of scientific software, are reaching their limit in scalability, says Stevens. As part of DOE's SciDAC-funded project "Towards Optimal Petascale Simulations (TOPS)." Argonne researchers, led by Lois McInnes, are working to relieve the linear solver bottleneck. Similarly, through SciDAC's Center for Technology for Advanced Scientific Component Software, Argonne scientists are working to expand the software toolbox for scientific discovery to include programs that act as discrete components with unique capabilities, but that also work together seamlessly. The idea is to build a component ecosystem, a community of interacting units that work together and provide feedback to each other in a self-regulating way.

Stevens says that most of these projects are in their beginning phases and there are many places where a DOE CSGF fellow could plug into the organization.

"It's an exciting time to be starting a career in scientific computing," Stevens says. "Depending on the career interests of the person, there are opportunities to become a co-developer, a discipline scientist, a numerical analyst, a visualization expert. It's all open right now. The structure of leadership computing means that you can have close collaborations between teams over time. That's really the vision and the promise of this approach."

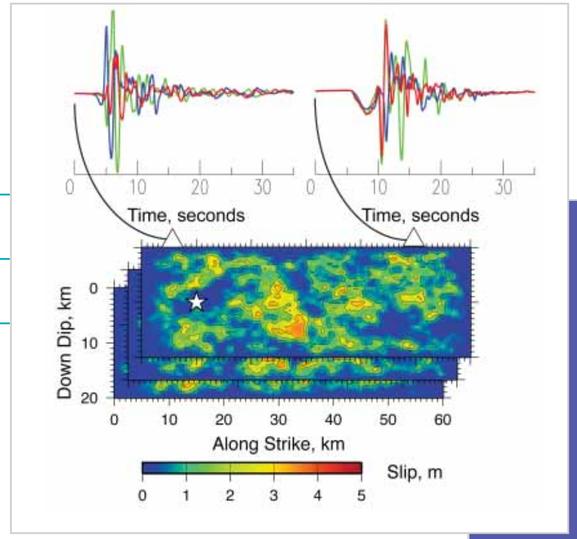

A great challenge to modeling large damaging earthquakes is the fact that it is difficult to know where slip will occur along the fault before the event happens. The figure shows ground motions computed with WPP code along two points (indicated by triangles) above a hypothetical magnitude 7.0 earthquake for three different earthquake rupture models (star indicates the hypocenter, fixed for all three models). The ground motions (colored for different rupture models) show variability due to variations in the slip along the fault. The LCF will allow us to run these simulations at higher resolution, resulting in more realistic higher frequency results and to sample the range of possible rupture models.

Image provided by Arthur Rodgers, Lawrence Livermore National Laboratory.

CHECKING OUT CODE AT ARGONNE'S NUMERICAL LIBRARIES

One of the thorniest issues in petascale computing is how to ensure that the software applications running on the machines measure up. Code developers typically don't write new code from scratch when a new machine arrives on the scene, says Rick Stevens, associate laboratory director for Computing and Life Sciences at Argonne National Laboratory.

"We think of these codes as having a family tree," he says.

"We've been working on the software problem for parallel computing for over 20 years," says Stevens. "We've created scalable numerical libraries that solve a lot of the problems scientists run into in getting these codes to work in a parallel computing environment."

What's developed over the years is a suite of dependable open-source software that is available for scientists to plug into their problem without having to "reinvent the wheel." And because the code developers are in many cases still available for consultation, Argonne's numerical libraries repository has become one of the go-to places when a homegrown solver isn't working.

Former DOE CSGF fellow Allison Baker, computational scientist at Lawrence Livermore National Laboratory, spent a summer at Argonne in 1999 working with an Argonne-developed library, PETSc, a suite of data structures and routines for the scalable, parallel solution of scientific applications modeled by partial differential equations. Her summer learning experience set the stage for a portion of Baker's thesis work at the University of Colorado, where she developed a new linear solver algorithm that reduced data movement through machine memory to gain efficiency.

"My experience with PETSc taught me a lot about the importance of thinking carefully about implementation details when developing software and about writing good code in general," she says. "It also made me realize the importance of having quality software tools available to the scientific community so that researchers can spend more time on their specific area of expertise and less time on code development."

Mary Biddy

Mary Biddy's assignment sounds like a task from MISSION: IMPOSSIBLE: understand the inner workings of a "black box" — a device that takes a known input, does something mysterious to it, and yields a different output.

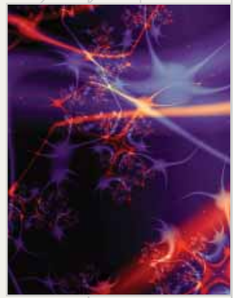

The "black box" is third-party software with proprietary modeling algorithms that are hidden. Biddy feeds in data and scans the results.

"If I'm trying to run a simulation and the data is not converging, then I look at the variables and try to figure out why," she says. "But there is always a part of the code that I can't see."

A year ago, Biddy received her doctoral degree in chemical engineering from the University of Wisconsin after researching vegetable-based lubricating oils. Today, she is a senior research engineer in the Aromatics and Acetyls Division of BP plc, one of the world's largest oil companies.

The black boxes are part of mammoth models used to optimize the operation and profitability of nearly 20 aromatic chemical plants worldwide.

It's a complex problem. Each plant operates somewhat different equipment. The cost of utilities fluctuates daily. So does the composition of chemical raw materials, which vary with the type of oil the refinery processes that day. Sometimes equipment breaks down, so plants need to find alternate routes to make the same products.

Most dauntingly, BP's optimization model works in real time, adjusting to each new fluctuation in customer orders, prices, chemical inputs, and available equipment.

The model's backbone was supplied by Aspen Technologies, which specializes in oil and chemical simulations. It also includes BP-developed models. Biddy has full access to BP's models, but Aspen's models are a trade secret.

When confronted with simulation problems, Biddy falls back on skills she learned in graduate school. "You can use logic to figure out why you're not getting the right answer," she says, "but you have to do it methodically and try to understand what's important for each component of the model."

Tickling the secrets out of massive models is a far cry from Biddy's original plans. Although she studied chemical engineering at Texas A&M University in Kingsville, she planned to become an environmental lawyer.

Then she took an internship at Johnson Polymers, a subsidiary of SC Johnson & Son Inc. "My mentor thought I was naturally curious and would enjoy the research side of engineering," Biddy recalls.

In graduate school, she pursued her environmental interests by looking at ways to make lubricating oils from renewable plants. "I found several ways to improve their physical properties, but the modifications reduced the oils' biodegradability," she says.

Biddy used molecular modeling to predict how modifications changed the physical properties of biolubricants — an approach she might not have taken were it not for her DOE CSGF.

"The fellowship required that I take math and computer science courses, such as a class on algorithm development, that would not have been part of the normal curriculum," she explains. "It gave me a more rounded education."

Biddy eventually chose to join BP because of its strong commitment to green technology and the challenging work it offered — even if her mission meant finding the possibilities in a black box.

"You can use logic to figure out why you're not getting the right answer, but you have to do it methodically."

Eric Held

It is easy to imagine Eric Held as a farmer. "I like BEING ALONE with my thoughts," says the Utah State University physicist — one of the few people on campus who does not own a cell phone.

Held often uses the silence to contemplate fusion, which promises to produce clean, emission-free energy by binding hydrogen nuclei. Held's simulations of magnetically confined plasmas may help scientists better understand this complex process.

This is a long way from Held's original plans. His grandfather and uncles were farmers, and his father was director of the South Dakota Farm Bureau. He wanted to farm after high school graduation, and didn't even apply to college until a few weeks before the fall semester. "My parents leaned on me to do it," he recalls.

At South Dakota State University Held discovered the joy of unraveling physics and math problems. "I thought it would be cool to spend all day thinking about a problem," he says. Such musings led him to the DOE CSGF while pursuing a doctoral degree in plasma physics at the University of Wisconsin.

Held honed his computer modeling skills and developed a new perspective on his work during his DOE CSGF practicum at Oak Ridge National Laboratory. "My advisor, Jean-Noel Leboeuf, was a bright, humble guy who thought carefully about what he said and did. I learned how to carry myself as a physicist by watching him," Held says.

The DOE CSGF stipend gave Held the freedom to consider what problem he would choose for his doctoral thesis. He decided to model how heat escapes from magnetically confined plasma as it races around a donut-shaped torus, and he's wrestled with the problem ever since.

"The torus is like a heating duct bent back on itself, and the magnetic field acts like insulation to prevent the heat from escaping through the walls," Held explains. If too much heat escapes the magnetic field, the plasma loses energy and cannot sustain fusion.

Unfortunately, heat travels up to 10 billion times faster along the torus than it does when it escapes. Like a gear that has a ratio of 10 billion to one, it takes 10 billion simulation cycles of the "duct" gear to turn the "escaped heat" gear forward one tick. That extreme difference in scale makes the system very hard to model.

Held uses hybrid models to simplify the calculations. He starts with a relatively simple fluid model to describe the plasma's density, flow, and temperature. He then adds elements of more complex kinetic models, which describe how individual particles interact with one another and with the electromagnetic fields around them.

"Imparting kinetic physics to fluid equations captures more of the real physics and enables us to do really big simulations," Held says.

The approach has already achieved at least one payoff. "Heat normally flows from hot to cold, the way a radiator heats up a room," Held explains. "But under some conditions in a magnetic field, heat will actually flow to even hotter areas."

Will those results help create a clean fusion reactor? Perhaps, Held says. Regardless, they promise better insights into a process that has baffled researchers for two generations.

It's nothing like farming, but it's just as rewarding.

"I thought it would be cool to spend all day thinking about a problem."

Ahmed Ismail

What do STICKY PROTEINS have to do with safely storing nuclear waste? On the surface, Ahmed Ismail says, they involve considerably different scales of time and size.

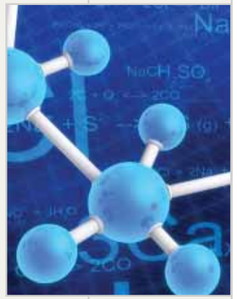

A nuclear waste repository "covers tens of square kilometers and its time scale is 10,000 years," he notes. "Protein interactions take place in nanometers and nanoseconds."

What ties them together for Ismail, who works for Sandia National Laboratories, are the complex models of molecular-scale interactions that he uses to predict behavior in both regimes.

While a doctoral student and Department of Energy Computational Science Graduate Fellow at the Massachusetts Institute of Technology, Ismail worked on ways to simplify simulations of molecular interactions so they eat less supercomputing time. "This is a widely researched problem, but the applications have not kept up with the theory," he says.

> **"Our group is trying to make sure that centralized storage is not a threat to the environment or to other people. I'm comfortable with that."**

After his MIT thesis defense — for which the lifelong Red Sox fan prepared while sneaking peeks at Boston's first World Series win since trading Babe Ruth to the Yankees — Ismail moved to Sandia as a post-doctoral researcher.

At Sandia, Ismail used models to explain why proteins sometimes stick to the non-stick polyethylene oxide (PEO) coatings used to prevent fouling in medical devices. "Proteins like to bond to regular surfaces," Ismail explains. "Coat with too little PEO and proteins bond to the exposed metal below. Use too much PEO and its molecules will crowd together to form a regular 'surface' to which proteins can stick.

"Between those two extremes, water slips between the PEO molecules, creating an irregular surface that proteins ignore. I think of this as the 'Goldilocks effect,' because the coating has to be just right to work," Ismail concludes.

After finishing his post-doc, Ismail joined Sandia's Carlsbad Programs Group, which supports the world's only operating nuclear waste repository. DOE's Waste Isolation Pilot Plant (WIPP) was built to safely hold clothing, test tubes, and other radiation-exposed items for 10,000 years.

WIPP stores sealed waste containers in rooms dug into an underground salt deposit. Over time, pressure from thousands of feet of rock above will buckle the rooms and isolate the containers. The salt rock resists chemical infiltration.

In addition to continuing his PEO work, Ismail tries to imagine how worst-case scenarios might affect the waste. "In 10,000 years, people might forget WIPP is there and drill down to a reservoir of briny water several thousand feet below the repository," he says. "We want to know what would happen if the brine infiltrated the repository."

Using models that describe fluid mechanics, geochemistry, and nuclear physics, Ismail studies how radionuclides might react with rock and brine. "If the salt does its job, there's likely to be very little release to the environment," he says.

Yet Ismail continues upgrading Carlsbad's chemical equilibrium model to improve its accuracy, as well as studying sites for the safe storage of more highly radioactive wastes.

Ismail never imagined using his ability to model molecular-scale interactions to understand radioactive containment, but he's pleased about the added application.

"We're solving a problem that already exists because we're storing waste all over the place now," he says. "Our group is trying to make sure that centralized storage is not a threat to the environment or to other people. I'm comfortable with that."

Dr. Kristen Grauman
Dr. Jaydeep Bardhan

Howes Scholars

THE FREDERICK A. HOWES SCHOLAR

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

Dr. Kristen Grauman of the University of Texas at Austin and Dr. Jaydeep Bardhan of Argonne National Laboratory have been named the 2007 Frederick A. Howes Scholars in Computational Science.

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

One of Howes' responsibilities was to oversee the Department of Energy's Computational Science Graduate Fellowship (DOE CSGF) program. He was extremely committed to this program that supports graduate students in computational science. This program is unique, as it requires candidates to take courses in mathematics, computer science and an applications discipline, such as physics or engineering. The DOE CSGF program currently supports over 60 graduate students and is administered by the Krell Institute.

To honor his memory and his dedication to the Department of Energy's Computational Science Graduate Fellowship program, one DOE CSGF fellow will be chosen each calendar year as a Howes Scholar. But because there were so many outstanding nominees for the award this year, two winners were selected. This award will provide the recipients with a substantial cash award, a Tiffany paperweight, and the distinction of being named a Howes Scholar.

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

Dr. Grauman was a CSGF fellow from 2001-2005. She received her Ph.D. in Computer Science in 2006 from MIT's Computer Science and Artificial Intelligence Laboratory. She is currently an Assistant Professor in the Department of Computer Sciences at the University of Texas at Austin.

A DOE CSGF fellow from 2002-2006, Dr. Bardhan received his Ph.D. in Electrical Engineering and Computer Science, also from MIT, in 2006. He is currently a Wilkinson Fellow at Argonne National Laboratory.

Both award recipients were on hand at the DOE CSGF annual fellows' conference where they presented their research and received their awards. Daniel Hitchcock from the U.S. Department of Energy's Office of Science presented the awards and Fred Howes' widow and son, Mary Hall and Michael Howes, were in attendance at the presentation.

For more information on this program, contact the Krell Institute at (515) 956-3696 or email Rachel Huisman at Rachel@krellinst.org.

Dr. Grauman presents at the annual Department of Energy Computational Science Graduate Fellows conference.

Daniel Hitchcock, U.S. Department of Energy presents Jaydeep Bardhan his 2007 Howes Scholar award.

Fred Howes' son and widow, Michael Howes and Mary Hall, (back row) pose with 2007 Howes Scholars Jaydeep Bardhan (left) and Kristen Grauman (right) and Daniel Hitchcock from the U.S. Department of Energy, Office of Science (center).

Alumni Directory

A
B
C

A

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

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

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

Kristopher Andersen
University of California – Davis
Physics
Fellowship Years: 2001-2005
Current Status: 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: Student, Texas A&M University

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

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

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

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

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

William Barry
Carnegie Mellon University
Structural & Computational Engineering
Fellowship Years: 1994-1998

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

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

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

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

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

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

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

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

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

C

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

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

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

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

Kevin Chu
Massachusetts Institute of Technology
Applied Mathematics
Fellowship Years: 2002-2005
Current Status: Vitamin D, Inc.

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

Joshua Coe
University of Illinois
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: Staff, Trinity College, Dublin Ireland

William Conley
Purdue University
Mechanical Engineering
Fellowship Years: 2003-2007
Current Status: Student, Purdue University

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

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

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

Robert Cruise
Indiana University
Physics
Fellowship Years: 1997-2001

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

Joseph Czyzyk
Northwestern University
Industrial Engineering
Fellowship Years: 1991-1994
Current Status: Staff, Central Michigan University Research Corporation

D

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

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

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

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

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

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

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

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

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

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

E

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

Thomas Epperly
University of Wisconsin
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

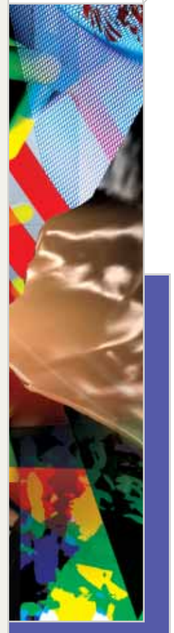

A
B
C
D
E

F

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

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

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

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

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

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

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

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

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

G

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

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

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

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

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

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

Matthew Giamporcaro
Boston University
Cognitive and Neural Systems
Fellowship Years: 1998-2000
Current Status: 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: 1996-1999
Current Status: Faculty, North Carolina Agricultural & Technical State University

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

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

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

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

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

H

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

Glenn Hammond
University of Illinois
Environmental Engineering & Science
Fellowship Years: 1999-2003
Current Status: Pacific Northwest National Lab

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

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

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

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

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

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

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

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

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

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

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

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

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

I

Eugene Ingerman
University of California – Berkeley
Applied Mathematics
Fellowship Years: 1997-2001
Current Status: General Electric

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

J

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

K

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

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

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

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

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

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: Staff, Virginia Commonwealth University

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

L

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

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

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: e-SciTek, LLC

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 Employee

M

William Marganski
Boston University
Biomedical Engineering
Fellowship Years: 1998-2002
Current Status: Ikonisys, Inc.

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

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

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

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

Matthew McNenly
University of Michigan
Aerospace Engineering
Fellowship Years: 2001-2005
Current Status: 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 and Mary
Computer Science
Fellowship Years: 2001-2004
Current Status: Oak Ridge National Laboratory

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

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

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

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

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

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

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

N

Heather Netzloff
Iowa State University
Physical Chemistry
Fellowship Years: 2000-2004
Current Status: 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

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

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

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

Joyce Noah-Vanhoucke
Stanford University
Theoretical Chemistry
Fellowship Years: 2001-2003
Current Status: Staff, University of California – Berkeley

Catherine Norman
Northwestern University
Applied Mathematics
Fellowship Years: 2000-2004
Current Status: Center for Naval Analysis

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

O

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

P

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

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

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

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, University of California – Los Angeles

Christina Payne
Vanderbilt University
Chemical Engineering
Fellowship Years: 2003-2007
Current Status: Student, Vanderbilt University

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

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

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

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

Q

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

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 and Planetary Sciences
Fellowship Years: 2003-2006
Current Status: Arete Associates

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

Clifton Richardson
Cornell University
Physics
Fellowship Years: 1991-1995

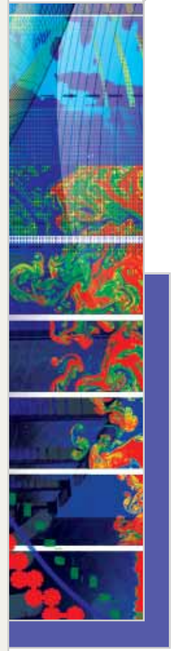

Christopher Rinderspacher
University of Georgia
Chemistry
Fellowship Years: 2001-2005
Current Status: Staff, Duke University

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
Biology
Fellowship Years: 1996-1997
Current Status: ActiveSight

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

S

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

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

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

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

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

Jason Sese
Stanford University
Computational Materials Science
Fellowship Years: 2003-2005
Current Status: U.S. Patent and Trademark Office

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

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

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

Steven Smith
North Carolina State University
Chemical Engineering
Fellowship Years: 1992-1994
Current Status: Invista

Eric Sorin
Stanford University
Chemical Physics
Fellowship Years: 2002-2004
Current Status: Faculty, California State University – Long Beach

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

Samuel Stechmann
New York University
Applied Mathematics
Fellowship Years: 2003-2007
Current Status: Student, 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: Scalable Display Technologies

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

T

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

Brian Taylor
University of Illinois
Engineering Mechanics
Fellowship Years: 2003-2007
Current Status: Student, University of Illinois

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
Mechanical Engineering
Fellowship Years: 1997-2000
Current Status: Staff, Pennsylvania State University Applied Research Laboratory

U

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

V

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

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

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

W

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

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

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

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

Collin Wick
University of Minnesota
Computational Chemistry
Fellowship Years: 2000-2003
Current Status: Faculty, Louisiana Tech University

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

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

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

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

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

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

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: Student, University of California – Berkeley

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

Z

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

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

Erik Allen

Massachusetts Institute of Technology
Chemical Engineering

Advisor:
Greg Rutledge
Practicum:
Sandia National Laboratories – New Mexico
Contact:
ecallen@mit.edu
Research Synopsis:
In the field of molecular simulation, many of the most interesting problems are currently computationally intractable. Despite significant advances in hardware and algorithmic efficiency, the combined requirements of a large number of atoms (10,000-1,000,000 or more) and a small integration time step (typically femtoseconds) relegates many of these molecular simulation problems to the category of future work. Coarse graining attempts to address many of the current weaknesses of molecular simulation by seeking to reliably replicate the behavior of an underlying atomistic simulation through a reduced representation.

The goal of our present research is to study the fundamental questions of coarse-grained experimental design by comparing multiple design options within a common particle system. The particle system we have selected is surfactants in water. Surfactant molecules are medium length chain molecules that, in sufficient concentration, will spontaneously assemble into spherical aggregates called micelles. Surfactant solutions are an interesting system in which to study coarse-graining techniques for a number of reasons: properties of interest span a large range of length scales, they are highly sensitive to the interaction model used, and significant experimental data are available as a basis for comparison.

Michael Bybee

University of Illinois
Chemical Engineering

Advisor:
Jonathan Higdon
Practicum:
Lawrence Livermore National Laboratory
Contact:
bybee@uiuc.edu
Research Synopsis:
Suspensions of colloidal particles (1nm to 1μm) are of great importance in a variety of industrial applications including paints, coatings, foods, drugs, and cosmetics. Additional applications arise in the fabrication and development of novel electronic materials, microscale biosensors, and nanostructured materials. These suspensions exhibit a wide variety of phase transitions and rheological behavior that can be far more complex than for normal fluids. The ability to understand, manipulate, and predict the behavior of these systems is important for designing and optimizing novel materials and manufacturing processes.

My research is in the large scale simulation of colloidal suspensions. More specifically, I am interested in the simulation of phenomena that arise from interparticle forces in the presence of hydrodynamic and Brownian forces. These phenomena include liquid-liquid phase separation, gelation, and crystallization. These phase behaviors in turn have a profound effect on the rheological behavior and material properties of the suspension.

Jimena Davis

North Carolina State University
Applied Mathematics

Advisor:
H.T. Banks
Practicum:
Sandia National Laboratories – New Mexico
Contact:
jldavis9@unity.ncsu.edu
Research Synopsis:
There are two types of measure-dependent problems: individual dynamics and aggregate dynamics. We are currently working on estimating growth rate distributions in size-structured mosquitofish populations, an aggregate dynamics problem. In order to achieve the properties of dispersion and bifurcation exhibited by the mosquitofish data, we use the Sinko-Streifer model as modified in the Growth Rate Distribution (GRD) model of Banks-Botsford-Kappel-Wang. Given data, we would like to estimate the growth rate distributions by considering a least-squares formulation for the inverse problem. We wish to minimize the cost functional over an appropriate collection of probability measures; however, finding an optimal probability distribution is not simple since we are in an infinite-dimensional setting. Also of importance in our work is the ability to be able to talk about the convergence and the continuous dependence of the estimated growth rate distributions on the data, but to do so we will need to talk about the distance between two probability measures.

Jeffrey Drocco

Princeton University
Biophysics & Computation

Advisor:
David Tank
Practicum:
Los Alamos National Laboratory
Contact:
jdrocco@princeton.edu
Research Synopsis:
Studies of the dynamics of mesoscopic systems have benefited greatly from advances in computational methodology. Also however, recent advances in many-body dynamics have made it possible for condensed matter physics to return the favor, in the form of potential new computer architectures.

While at present the size of microelectronics continues to decrease, allowing corresponding improvements in computing power, their size will soon reach a fundamental lower limit determined by quantum mechanics. One alternative to the standard field-effect transistor model is a cellular automata model. This uses the positions of particles in an array of cells to store state information. If cells are lined up next to each other, interparticle repulsion will correlate the positions of the particles and allow information to be transmitted through the array. Any viable computer architecture also requires a method of performing logic operations on the state information at a constant clock speed. This has recently been shown to be possible, numerically, when vortex cellular automata are used in conjunction with an AC current ratchet (Hastings et al., Phys. Rev. Lett., 6/20/2003). Moreover, a complete set of logic gates can be built with different arrangements of the pinning array.

The estimated maximum clock speed with this model is 800 MHz using conventional etching techniques, so further theoretical advances will be necessary to make the method practical for supercomputing applications. I plan to use my experience with computational studies of vortex dynamics to work on this problem. In this case, numerical methodology is useful to verify the feasibility of proposed solutions on a many-body scale without adding the unknown factors which accompany experiment.

Jasmine Foo

Brown University
Computational Fluid Dynamics

Advisor:
George Karniadakis
Practicum:
Lawrence Berkeley National Laboratory
Contact:
jfoo@dam.brown.edu
Research Synopsis:
My primary area of focus is the numerical simulation of fluid-structure interactions. To do this, I use a spectral/hp element Navier-Stokes solver coupled with a structural solver on serial and parallel systems. I am working on the development of numerical methods using polynomial chaos for modeling uncertainty in fluids. Currently, I am also studying the motion of a free rigid cylinder in a flow, subject to vortex-induced vibrations (VIV). I am investigating the effect of Reynolds number, mass ratio and damping parameters on the cylinder amplitude and mode of vortex formation. I am implementing a dual-level parallelism in this VIV code to increase the speed of computation. Other current and future projects include applications in biomimetics, such as flow around bat wings during flight and the energy-harvesting eel.

Bonnie Kirkpatrick

University of California – Berkeley
Computer Science

Advisor:
Richard Karp
Practicum:
Lawrence Livermore National Laboratory
Contact:
bbkirk@eecs.berkeley.edu
Research Synopsis:
My primary interest in Computational Biology is algorithms for human genetics, including disease association and haplotype phase resolution. My peripheral interests include regulatory interaction networks, phylogenetics, and biomolecular folding kinetics.

David Potere
Princeton University
Demography/Remote Sensing

Advisor:
Burt Singer
Practicum:
Oak Ridge National Laboratory
Contact:
dpotere@princeton.edu
Research Synopsis:
Space-born imagery when fused with existing census data offers a means of reducing the uncertainty in the spatial distribution of global population — a particularly urgent problem within the developing world. The DOE's LandScan project based at ORNL is one example of such an effort.

My aim is to use computational geography to improve such global demographic mapping efforts. There are correlations between population distribution and NASA's immense optical satellite remote sensing archive. The NASA archive is global in geographic extent, fine resolution, covers many spectral bands, has revisit times as frequent as every three days, and stretches back the early 1980's. Uncovering relationships between this imagery and global population will require a combination of machine learning, image processing, statistics, and data visualization.

Amber Sallerson
University of North Carolina – Chapel Hill
Applied Mathematics

Advisor:
Cass Miller
Practicum:
Lawrence Berkeley National Laboratory
Contact:
asalle1@email.unc.edu
Research Synopsis:
The problems that I am concerned with are of high complexity, involving multiple phases and species, multiple space and time scales, and which are inherently heterogeneous and three-dimensional in nature. The advancement of fundamental models and simulators for such systems requires high-performance computation.

Michael Veilleux
Cornell University
Computational Fracture Mechanics

Advisor:
Anthony Ingraffea
Practicum:
Sandia National Laboratories – New Mexico
Contact:
mgv5@cornell.edu
Research Synopsis:
My interests are in optimizing the accuracy of structural integrity prognosis for American Military and NASA aerospace vehicles. I want to advance the stochastic capabilities of computational fracture analysis programs for the purposes of developing more accurate, multi-scale structural damage state models of air vehicles. Given stochastic material geometries and properties, finite element analysis programs use extensive computation to produce multiple length-scale fatigue crack growth models in structures. Advanced technological sensors are currently being developed that will give local and global strain outputs from structural components on an aerospace vehicle during or after usage. By further enhancing the capabilities of finite element fracture analysis programs currently being used at Cornell, I will translate the sensor data into continually updated, multi-scale damage state models of the structural components. If component damage states can be accurately modeled throughout the existence of structures, then the safety, life expectancy, and performance of aerospace vehicles can be greatly enhanced.

Allan Wollaber
University of Michigan
Nuclear Engineering, Fission Concentration

Advisor:
Edward Larsen
Practicum:
Los Alamos National Laboratory
Contact:
wollaber@umich.edu
Research Synopsis:
My field of interest concerns the interaction of radiation with matter, or radiation transport. It is a multifaceted field, encompassing nuclear reactors, radiation cancer therapy, and industrial applications such as material thickness monitoring or oil well logging. Because working directly with radioactive materials can be hazardous or prohibitively expensive, scientists and engineers rely heavily on computational models of radiation transport. My research topic concerns a new technique by which we simulate radiation transport on computers.

Etay Ziv
Columbia University
Computational Biology

Advisor:
Chris Wiggins
Practicum:
Lawrence Berkeley National Laboratory
Contact:
ez87@columbia.edu
Research Synopsis:
Biological networks, such as cellular and genetic networks, are of particular interest because of their potential to elucidate design principles that nature has employed to perform computations, while maintaining robustness to noise and adaptability. Our lab studies biological networks using two approaches: statistical graph analyses and stochastic dynamic simulations. In both cases we aim to use computational technologies to discover underlying design principles.

John ZuHone
University of Chicago
Astrophysics

Advisor:
Donald Lamb
Practicum:
Oak Ridge National Laboratory
Contact:
zuhone@uchicago.edu
Research Synopsis:
Astrophysicists have known for many years that a non-luminous form of matter accounts for most of the mass in the universe. This mass resides in haloes in which the visible structures of the universe (galaxies, groups and clusters of galaxies, etc.) reside. Recent simulations of the formation of cosmological structure have suggested that the density profiles of these haloes are universal and are largely independent of the parameters of the halo, such as its mass. It is thought that structure in the universe has formed by the merger of smaller structures to form larger objects, but this does not by itself account for the universal density profile. I am performing a series of N-body simulations using the FLASH astrophysics code which set up controlled mergers of dark matter haloes. In doing so, I hope to isolate particular effects or circumstances which may contribute to the particular form of density profile that cosmological simulations have observed. In addition, the simulations may shed some light on the exact shape of the universal profile, as different simulations have given similar but not equivalent results.

Joshua Adelman
University of California – Berkeley
Biophysics
Advisor:
George Oster
Practicum:
Lawrence Livermore
National Laboratory
Contact:
jadelman@ocf.berkeley.edu

Zlatan Aksamija
University of Illinois
Electrical Engineering
Advisor:
Umberto Ravaioli
Practicum:
Los Alamos
National Laboratory
Contact:
aksamija@uiuc.edu

Jordan Atlas
Cornell University
Chemical Engineering
Advisor:
Michael Shuler
Practicum:
Los Alamos
National Laboratory
Contact:
jca33@cornell.edu

Christopher Carey
University of Wisconsin
Plasma Physics
Advisor:
Carl Sovinec
Practicum:
Lawrence Berkeley
National Laboratory
Contact:
cscarey@wisc.edu

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

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

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

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

Miler Lee
University of Pennsylvania
Genomics & Computational Biology
Advisor:
Junhyong Kim
Practicum:
Sandia National Laboratories –
New Mexico
Contact:
miler@mail.med.upenn.edu

Jeremy Lewi
Georgia Institute of Technology
Neuroengineering
Advisor:
Robert Butera
Practicum:
Lawrence Livermore
National Laboratory
Contact:
jlewi@gatech.edu

David Markowitz
Princeton University
Computational Neurobiology
Advisor:
David Tank
Practicum:
Los Alamos
National Laboratory
Contact:
dave@princeton.edu

Peter Norgaard
Princeton University
Computational Plasma Dynamics
Advisor:
Clarence Rowley
Practicum:
Lawrence Berkeley
National Laboratory
Contact:
norgaard@princeton.edu

Natalie Ostroff
University of California – San Diego
Bioengineering
Advisor:
Jeff Hasty
Practicum:
Oak Ridge National Laboratory
Contact:
nostroff@ucsd.edu

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

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

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
Contact:
abhatt@mit.edu

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

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

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

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

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

Ashlee Ford
University of Illinois
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
Contact:
bl76@cornell.edu

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

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

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

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

Ariella Sasson
Rutgers University
Computational Biology
Advisor:
Wilma Olson
Practicum:
Sandia National Laboratories –
New Mexico
Contact:
ariella@eden.rutgers.edu

Michael Sekora
Princeton University
Continuum Mechanics, PDE
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:
Salvatore Torquato
Contact:
bsonday@princeton.edu

FIRST YEAR FELLOWS

Matthew Adams
University of Washington
Computational Electromagnetics
Advisor:
Vikram Jandhyala
Contact:
adams.matthewr@gmail.com

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

Hal Finkel
Yale University
Physics
Advisor:
Richard Easther
Contact:
hal.finkel@yale.edu

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

Steven Hamilton
Emory University
Computational Mathematics
Advisor:
Michele Benzi
Contact:
sphamil@emory.edu

Joshua Hykes
North Carolina State University
Nuclear Engineering
Advisor:
Dmitriy Anistratov
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
Contact:
jmartin@ices.utexas.edu

Geoffrey Oxberry
Massachusetts Institute of Technology
Chemical Kinetics/Transport Phenomena
Advisor:
William Green
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
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
Cosmology
Advisor:
Paul Ricker
Contact:
psutter2@uiuc.edu

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

John Ziegler
California Institute of Technology
Aeronautics
Advisor:
Dale Pullin
Contact:
jackalak@caltech.edu