THE DOE CSGF ANNUAL

DEIXIS

2013

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

BURNING AMBITION

Leslie Dewan's New Nuclear Reactor Design Launches into a Promising Company

PAGE 5

SPECIAL FEATURE

LIGHT EXPERIMENTS, HEAVY DATA

At Light Sources and Other Facilities, Computation Sifts Mounds of Complex Results

PAGE 20

THE DOE CSGF ANNUAL

DEIXIS

DEIXIS, The DOE CSGF Annual is published by the Krell Institute. Krell administers the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) program for the DOE under contract DE-FG02-97ER25308.

For additional information about the DOE CSGF program, the Krell Institute or topics covered in this publication, please contact:

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

DEIXIS (ΔΕΙΞΙΣ — pronounced dāksis) transliterated from classical Greek into the Roman alphabet, 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 program that highlights the work of fellows and alumni.

The DOE CSGF is funded by the Office of Science and the National Nuclear Security Administration's Office of Defense Programs.

Editor
Shelly Olsan

Senior Science Writer
Thomas R. O'Donnell

Copy Editor
Ron Winther

Publication Coordinator
Lisa Frerichs

Design
julsdesign, inc.

TABLE OF CONTENTS

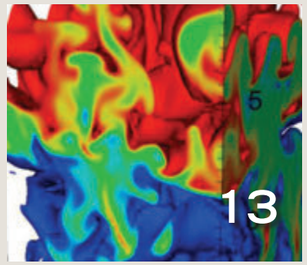

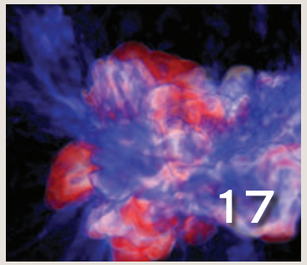

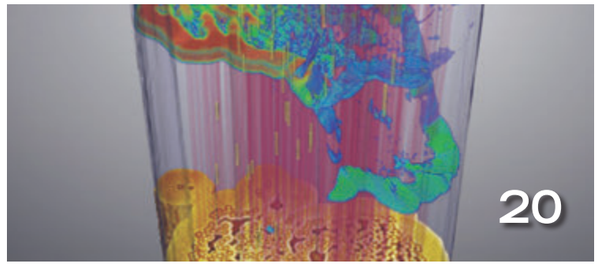

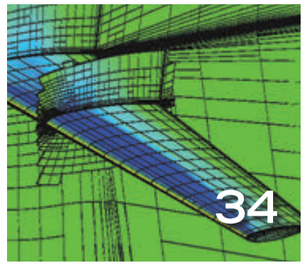

NO SMALL AMBITIONS

Students in the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) take on big challenges, in terms of difficulty or scale – or both.

Leslie Dewan, for example, promotes an approach to nuclear energy that concentrates on smaller reactors burning existing wastes more safely. Fellows Devin Matthews and Edgar Solomonik teamed up on a summer project to rewrite a key quantum chemistry algorithm and make the codes run faster on supercomputers.

Meanwhile, the small-scale mixing models Sanjeeb Bose helped improve during his Lawrence Livermore National Laboratory practicum could assist with experiments at the giant National Ignition Facility. And Samuel Skillman's Oak Ridge National Laboratory summer project addressed a programming problem for the heterogeneous architectures in today's largest computers.

These and other fellows can expect a future of big data sets, whether from simulations or experiments. This issue's special feature considers one example: results from advanced X-ray light sources.

Fellowship recipients graduate ready to address major challenges, as alumni profiled in this issue demonstrate. Aerospace engineering professor Krzysztof Fidkowski, programming language researcher Stephen Fink, computational scientist Timothy Germann and chemical and materials engineering professor Christina Payne are making significant contributions.

In these and other ways, the DOE CSGF advances technology and keeps America competitive. That's a big deal.

THE POWER OF THE PRACTICUM

SUMMER RESEARCH BROADENS FELLOWS' KNOWLEDGE

Front row: Edgar Solomonik
Back row (left to right): Samuel Skillman, Leslie Dewan, Sanjeeb Bose and Devin Matthews

SOME ADVICE for new Department of Energy Computational Science Graduate Fellowship recipients: Don't underestimate the power of the practicum.

For three months, usually in the summer, fellows leave their university labs and head to a Department of Energy national laboratory. They dive into subjects that are either new or tangential to their regular research. As the following profiles demonstrate, the results can be rewarding.

Fellow Leslie Dewan, for instance, worked on modeling fluids migrating underground from a leaking nuclear waste storage tank – a switch from her work commercializing a reactor design.

Devin Matthews and Edgar Solomonik stepped outside their comfort zones to collaborate during their practicums. Matthews got his first taste of supercomputer programming while Solomonik navigated advanced computational quantum chemistry.

On a break from his Stanford University research into swirling flows, Sanjeeb Bose worked on a different kind of turbulence. Meanwhile, fellow Samuel Skillman took on a challenging programming problem, setting aside his interest in galactic clusters.

Each fellow ended the summer with new skills and new perspectives – and a recharge of their creative batteries.

The Department of Energy Computational Science Graduate Fellowship supports the nation's brightest science and engineering students, allowing them to concentrate on learning and research. The work of more than 300 DOE CSGF alumni has helped the United States remain competitive in a global economy.

~~~~~

Credit: Savannah River Site

FOCUS ON NUCLEAR WASTE BRINGS FELLOW FAME

LESLIE DEWAN
Massachusetts Institute of Technology
Lawrence Berkeley National Laboratory

LESLIE DEWAN RARELY HESITATES to embark on ventures others might balk at. As a nuclear and mechanical engineering undergraduate at the Massachusetts Institute of Technology, she and some colleagues sailed the Charles River on a 3-meter-long balsa raft Dewan based on pre-Columbian Ecuadorian designs. She also rigged her dormitory room door to unlock when it sensed a radio frequency identification chip. Then she implanted the chip in her left hand – without anesthesia.

There have been other projects – a neutron interferometer, a tabletop cyclotron – leading to Dewan's biggest venture: a nuclear reactor design company. She and Mark Massie, a fellow MIT nuclear science and engineering doctoral student, will develop plants that consume nuclear waste, helping reduce storage needs and cut carbon emissions. The plan landed Dewan on a list of "30 under 30" leading innovators and entrepreneurs compiled by *Forbes* magazine.

Others might have doubts about entering a capital-intensive, highly regulated industry – especially after a 2011 earthquake and tsunami damaged Japanese reactors, releasing radiation and casting a pall on nuclear energy. But Transatomic Power's founders "felt there's so much potential in the nuclear industry and there are so many exciting technologies that can be invented," says Dewan, a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient who earned her doctoral degree this spring.

Transatomic's waste-annihilating molten salt reactor (WAMSR) will mix searing hot fluoride salts with radioactive waste from standard nuclear power plants. The waste will react, producing heat to keep the salts fluid and generate steam for turbines.

The design, conceived decades ago at Oak Ridge National Laboratory, requires no external coolant – unlike today's light-water reactors, which can overheat and melt down even after the reaction stops. The Japanese incident knocked out both the power and backup generators at three reactors. Without pumps to circulate cooling water, the fuel rods overheated and melted.

The WAMSR is designed to fail safely if power goes out. Electrically cooled fluoride-salt plugs are at the reactor's lowest points. As long as the power is on, the plugs are solid, even as nuclear reactions generate heat to keep the molten salts flowing. If power fails, the plugs would melt and the core would drain like a bathtub, the salts flowing into a secondary container shaped
~~~~~

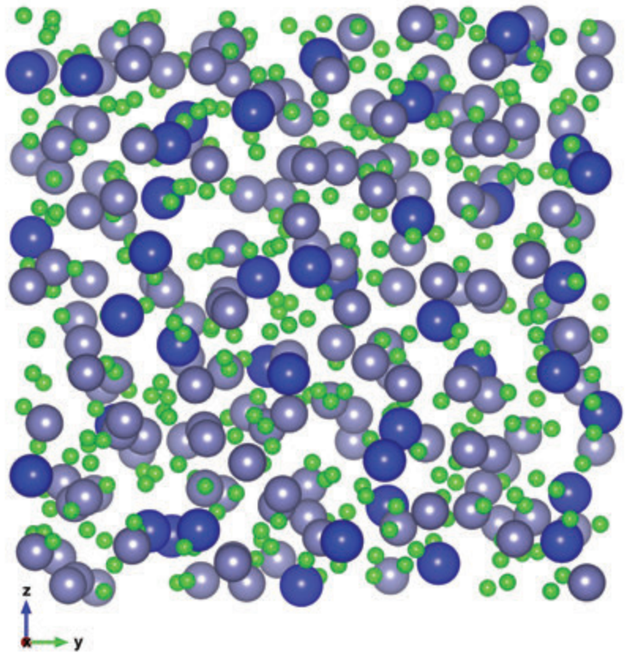

This visualization of a molecular dynamics computer simulation shows a cell of molten lithium thorium fluoride salt. Leslie Dewan's company is developing a nuclear reactor design that would rely on such molten salts to generate heat for power production.

to halt the heat-generating nuclear reactions. Like hot coffee poured into a saucer, the salts would naturally cool within hours.

WORKING WITH WASTE

Transatomic extends Dewan's doctoral research modeling how nuclear waste interacts with materials encasing it. In particular, she studied the products of vitrification, which dissolves waste such as thorium and uranium oxides in molten glass made of silicon, boron and other materials. The glasses solidify in canisters made of stainless steel and concrete.

About 10,000 tons of highly radioactive waste glass are stored around the world, says a paper Dewan, her MIT advisor, Linn Hobbs, and Jean-Marc Delaye of CEA, the French atomic energy agency, published last spring in the *Journal of Non-Crystalline Solids*. Because the materials are hazardous, it's difficult and expensive to test how they behave, but "with simulation, it's inexpensive and you can run it a thousand times," Dewan says. The goal: make nuclear waste storage media more stable.

In contrast, Dewan's summer 2011 practicum at Lawrence Berkeley National Laboratory (LBNL) focused on fluid nuclear waste and what happens if it escapes underground storage tanks. Working with John Bell, head of the lab's Center for Computational Sciences and Engineering (CCSE), Dewan helped improve a high-resolution algorithm for modeling radioisotopes as they move through and react with soil and rock.

DOE supervises underground storage at South Carolina's Savannah River Site and at the Hanford Site, a former weapons production facility in Washington State. Scientists are concerned about leaks of radioactive fission products, like cesium and strontium, but the tanks also hold actinide oxides, such as uranium oxide, that also are toxic.

Computers can simulate how this waste moves through and reacts with the soil. "It's an interesting problem because it involves a lot of different length scales," Dewan says, from the tens of meters groundwater crosses down to the millimeters over which chemical reactions occur. Detailed computations must capture sharp variations in chemical concentration and models must portray waste species in different phases, as each interacts differently with its surroundings.

Simulations generally represent materials or processes as a grid of data points. Computers calculate physical changes at each point and, taken together, they portray a complete picture, like pixels in a photograph. But such codes can demand considerable computer resources, especially when widely varying length scales are involved. CCSE's adaptive mesh refinement (AMR) approach conserves resources by concentrating computations in the most interesting areas, like the leading edge of a radioactive waste plume, and using fewer elsewhere.

Dewan tried to make the AMR subsurface flow code more accurate. "My first step was just poking at it and seeing where it broke," she says, by comparing results with data from research literature. With members of Bell's group, she worked to better incorporate chemical processes into each iteration.

ROBOTICS AND REACTORS

The practicum was Dewan's first taste of national lab life, but not her first non-academic engineering experience. After receiving her bachelor's degree, she spent a year at a Cambridge-based company, helping develop a module that would let a battlefield robot identify chemical, biological and nuclear weapons. "I couldn't imagine working anywhere else," Dewan says, but "if I stayed away for too long, it would just get harder and harder to come back to academia" for a doctorate.

And without returning to school, Dewan might never have started Transatomic Power. The company's chief executive officer is Russ Wilcox, the former head of E Ink, which developed electronic reader display technology. They've also lined up about a million dollars in financing – all before either Dewan or Massie finished their degrees.

At LBNL, Dewan tried a test problem: Modeling how far radioisotopes leaking from a Savannah River Site tank would spread via groundwater. She calculated how each species might flow through, become trapped in, or react with layers of bedrock, gravel, densely packed soil, and other material. Over the summer, she made the simulation more complex, incorporating multiple soil layers, various tank materials and different chemical reactions.

Dewan's tweaks helped the model more accurately reflect real-world data on how fission products move. She had less success modeling uranium and other actinides that often spread as colloids – undissolvable chunks that disperse through groundwater – making their chemistry more difficult to simulate.

In general, Dewan says, the radioactive species that travel farthest in groundwater often are the most important to contain. Fission products generally move intermediate distances. Soil absorbs dissolved uranium fairly quickly, but it travels farther as a colloid. "It depends a lot on the soil chemistry, which is part of what makes it such a tricky problem" and makes it important to tailor a simulation to the site, she adds. "It's hard to say that one particular species is going to be transported over a certain distance or will remain mobile for a certain amount of time."

Bell, who also heads LBNL's Mathematics and Computational Science Department, said Dewan helped improve the contaminant transport model and developed an interface to specify geochemistry in the

Dewan helped improve a high-resolution algorithm for modeling radioisotopes as they move through and react with soil and rock.

~~~~~

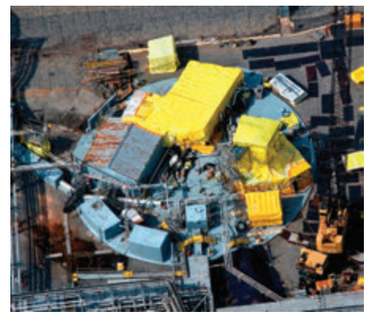

Credit: Savannah River Site

An aerial view of H Area Waste Tank 13 at the Savannah River Site in South Carolina. The tank was placed in service in 1956 and has a capacity of more than a million gallons. This and other older tanks at Savannah River must be emptied and closed by 2022.
~~~~~

simulation. "I was extremely impressed by Leslie's performance," Bell says. "She did a remarkable job given how different the project was" from her thesis research.

Hobbs, Dewan's doctoral advisor, says her talent for connecting apparently disparate subjects is what makes Dewan so valuable. "She moves between fields, because she can take and identify commonalities and common ideas that can be transferred." Her LBNL computational fluid dynamics experience, for example, could be useful for developing the WAMSR design.

MD AND MOLTEN SALTS

At MIT, Dewan used molecular dynamics (MD) algorithms, which calculate how atoms and molecules interact, to model molten fluoride salts carrying dissolved actinides and fission products, estimating the salts' viscosity, electrical and thermal conductivity, and other properties. Waste-laden molten salts are difficult to work with – blistering hot, corrosive and radioactive – so "in a lot of cases, this simulation work is filling the gaps in the experimental data set," she says.

Dewan also simulated how emissions from highly radioactive nuclear waste affect materials used to sequester them, a topic Hobbs has studied for decades. He wants to know why some crystalline materials lose their regular, periodic structure and become amorphous – turning to glass – under strong radiation. Scientists thought glasses containing radioactive waste should remain stable since they're already amorphous, but the materials' densities change as radioisotope atoms shed particles. "So the question is 'what's the nature of that change?'" Hobbs asks. "Can we predict swelling or shrinkage of these materials? Can we predict durability changes?"

Dewan combined MD with Monte Carlo mathematical methods that randomly choose distributions of model inputs. The last ingredient is connectivity topological analysis, which examines how atomic bonds break and reconnect as a material's structure changes under irradiation. The method mathematically examines how things are connected and quantifies amorphousness in a material – an inherently more difficult task than understanding crystalline structure, Dewan says. She developed an algorithm to count rings of atoms characteristic of amorphous substances. "That was something the fellowship helped me with significantly," she adds. In algorithm development courses, she learned methods that hadn't been used in materials science or nuclear engineering. "I was able to apply them to my ring-counting system and come up with something that was quite a bit faster" than other methods.

Left: A visualization of thorium tetrafluoride crystal structure from a simulation by Leslie Dewan. Her molecular dynamics-Monte Carlo models combine with connectivity topological analysis to consider how materials' structures change under irradiation.

Right: This visualization shows a collision cascade induced by an alpha recoil event in crystalline zircon, a material used to sequester radioactive wastes in a process called vitrification.

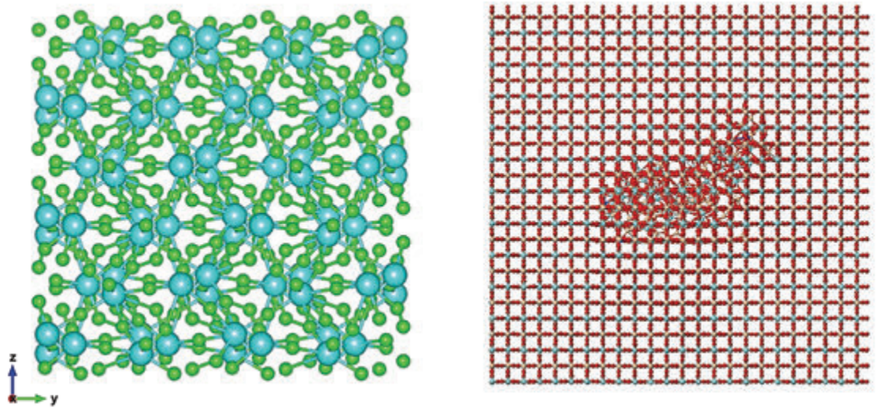

COUPLING SKILLS FOR CHEMISTRY QUEST

DEVIN MATTHEWS
University of Texas at Austin
Argonne National Laboratory

EDGAR SOLOMONIK
University of California, Berkeley
Argonne National Laboratory

"That this all came together – I did not actually think it was going to happen like this."

~~~~~

Credit: Argonne National Laboratory

AS DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP (DOE CSGF) recipients, Devin Matthews and Edgar Solomonik both attended the program's conferences, but only met when beginning practicums at Argonne National Laboratory (ANL), near Chicago, in summer 2011.

No surprise: With more than 70 fellows, getting acquainted takes time. More importantly, their fields differ significantly: Matthews studies theoretical and computational chemistry at the University of Texas at Austin; Solomonik's computer science research at the University of California, Berkeley, focuses on algorithms for science applications.

At ANL, the two convened on the common ground of high-performance computing (HPC), working with a third intern and DOE CSGF alumnus Jeff Hammond on a concept that could significantly influence computational quantum chemistry, computer science and applied mathematics. "It truly is a wonderful project," says Hammond, an assistant computational scientist at ANL's Leadership Computing Facility. "That this all came together – I did not actually think it was going to happen like this."

It's even more remarkable because Hammond assembled the project mostly as the fellows arrived. A short time later his wife gave birth and for two weeks he left them (and Martin Schatz, a Texas computer science graduate student) largely alone. "I came back and they had made a lot of progress," Hammond says. "I said, 'Oh, OK, this is really clicking. This is going to be fun.'"

It helped that Hammond carved the project into connecting parts that capitalized on each fellow's strength. Under advisor John Stanton, Matthews uses HPC to calculate the structures and interactions of atoms and molecules according to the strange physics that govern the tiniest scales. In the quantum realm, energy and matter interact as both particles and waves. Electrons and other particles absorb energy only in discrete amounts, or quanta, and one particle can influence another's behavior over a great distance.

Quantum calculations are important to understand details of interactions – especially the energy needed to form and break chemical bonds. They're vital to finding ways to efficiently convert plants into fuel, improve combustion and accomplish other nationally important goals. One of the most popular quantum chemistry codes for massively parallel computers is NWChem, developed under the auspices of DOE's Pacific Northwest National Laboratory. The package performs a variety of calculations, but the fellows focused on one: coupled cluster.

Coupled cluster methods describe the interactions of many-body systems, a notoriously difficult calculation because each body – electrons, in this case – influences the others in a complex dance. The calculations yield an approximate solution to the famed Schrödinger equation describing a system's quantum state.
~~~~~

GOING AFTER SCHRÖDINGER

The equation so fascinated Matthews in high school that he naively vowed to solve it, "which is what people have been trying to do for a hundred years. I basically got on Google and looked up how to do it," he says. "From what I could piece together I wrote a program to do the most basic thing. That was my start in theoretical chemistry."

Solomonik, meanwhile, took quantum physics while earning his computer science degree at the University of Illinois at Urbana-Champaign. It's a top computer science school, but also was close to home for Solomonik, whose parents brought him to Chicago from Russia at age 10. The ANL practicum was something of a homecoming from California, but Solomonik realized his physics education was insufficient to understand advanced computational quantum chemistry. Matthews talked him through it and Solomonik later supplemented that education with graduate chemistry and physics courses at Berkeley.

Coupled cluster starts with a single-particle wavefunction describing an electron's quantum mechanical orbital, Matthews says. "You put a certain number of electrons in each orbital, but that doesn't describe correlation, where the electrons interact with each other."

Correlation depends on excitations that push electrons into new orbitals. Coupled cluster methods account for excitations and products of excitations. Coupled cluster with singles and doubles (CCSD), an elementary form, calculates interactions of one or two excited electrons, producing a set of electron arrangements, each with a different weight. Adding triple excitations (CCSDT) generates a larger set of configurations, and quadruples (CCSDTQ) an even larger group. "It's expanding this set of things you're including in your wavefunction" but "it's also expanding the complexity with which you determine the weights."

High-level coupled cluster methods, Hammond says, are "central to this specific part of computational chemistry that cares about the highest possible accuracy with the most detailed physics."

Solomonik and Matthews, Hammond says, turned "an irregular problem into a very, very regular problem."

~~~~~

But those higher-level methods also demand great computer power and efficient algorithms.

Coupled cluster produces sets of nonlinear equations with multiple variables in the form of tensors, which describe relationships between data in multiple dimensions. (A matrix is a two-dimensional tensor.) In coupled cluster, tensor indices represent interchanges between electrons or orbitals. Hammond explains: "Each particle requires two dimensions, so two particles requires four dimensions, three particles requires 6-D, and that's how things get complicated fast." Cyclops (cyclic-operations) Tensor Framework, as the fellows call their method, supports up to 8-D tensors – CCSDTQ– "the holy grail of quantum chemistry for some people," Hammond adds.

CONTRACTION AND INTERACTION

Cyclops does the fundamental operation in quantum many-body calculations: tensor contraction, or computing the product of two tensors. It's complex math: summing the products of tensor components over one or more indices to reduce the answer's dimensions. Contraction represents an interaction between tensors, Matthews says, and is a more complicated version of matrix multiplication, a common operation.

"It gets interesting because tensors have symmetry: When certain indices are interchanged, the result (of multiplying) is the same or the same but with a minus sign." All indices may be symmetric – interchangeable with identical results – or antisymmetric – interchangeable but producing a negative – or they may be a combination. This is called index permutation symmetry.

When tensors get big it helps to only store distinct elements without duplicating symmetric indices, Matthews says. By preserving symmetry, the algorithm can reduce the amount of computation and data storage needed by a significant fraction that grows exponentially with tensor dimension. The summer students worked on exploiting index permutation symmetry to create a fast tensor contraction algorithm.

As Solomonik, Hammond and Solomonik's doctoral advisor, James Demmel, noted in a 2011 technical report, exploiting high-dimensional symmetry makes it hard to reduce contractions to matrix multiplication. The number of possible permutational symmetry types also grows exponentially with tensor dimension, complicating matters. "Traditional applications would try to, right away, take a tensor contraction and rewrite it as a matrix multiplication," Solomonik says, but that would surrender all symmetries.

Most contraction algorithms, like those in NWChem and ACES III, another popular quantum chemistry code, also assign tensors to processors sequentially, so each works with a neighboring piece of a tensor. The second processor starts work when the first finishes and the third processor starts when the second is done. "The problem with permutational symmetry is if you need a certain block … (it) may be in the wrong order for that symmetry," Matthews says. "You may have to transpose it, which means reaching across the network." The result: irregular, unbalanced work and time-wasting communication.

The key to countering these ill effects, Matthews says, is to avoid "unpacking" tensors. Packed tensor arrays consume less memory and are easier for processors to access. But "the algorithm has to know about the symmetry and has to do a lot of different permutations of the contraction to get the right answer."

Solomonik's solution: a cyclic layout that preserves the tensors' symmetry and decomposes them more efficiently and regularly. Instead of block scheduling, Cyclops assigns one tensor element to each processor in a specific order. If a tensor index is distributed among four processors, each "owns" every fourth element. "The local piece of the tensor each processor holds will end up having the same symmetry as the global tensor," Matthews says. "The parallel algorithm for moving the data doesn't have to know about the symmetry. It just has to know about the number of dimensions" in the tensor.

BETTER BALANCE

Cyclops inherently improves load balancing, so each processor does the same amount of work, Solomonik says. "Rather than processors asking for data … and waiting for it, they know exactly when data arrives. Everything proceeds synchronously."

Solomonik also extended his thesis research on communication-avoiding techniques for matrix computation into the realm of tensors. And the algorithm maps efficiently to each supercomputer's data-passing architecture, especially on IBM Blue Gene systems, such as Mira, the 10-petaflops (quadrillion scientific calculations per second) Blue Gene/Q at Argonne's Leadership Computing Facility. That's no accident: Solomonik worked with Blue Gene systems as an undergraduate, and Hammond does most of his research on them.

Solomonik and Matthews, Hammond says, turned "an irregular problem into a very, very regular problem" by splitting the parallel transpose permutation from matrix multiplication. "That's where this is first of a kind. Nobody, until we did this, I think, fully appreciated you could make this separation and get two very, very efficient, regular algorithms."

As Solomonik focused on making the code run well in parallel, Matthews mostly concentrated on the algorithm each processor uses for its tensor piece. It's "the same as if a single processor did the entire tensor," he adds. "I started with the dumbest possible implementation with the least capability and worked up to add more capabilities and make it more efficient." The hardest part was understanding cyclic computation – where data land and how they move. "Also, trying to think about tensors – you have to figure out what permutations are necessary and how you unpack things." He eventually realized a single piece of code could deal with tensor operations besides contraction.

In initial tests, Cyclops showed good scaling for matrix multiplication, hitting more than one petaflops on 16,384 nodes of Sequoia, a 20-petaflops Blue Gene/Q at Lawrence Livermore National Laboratory, where Solomonik did a second practicum in 2012. The test compared Cyclops with

These graphics compare load imbalances incurred or padding necessary for blocked, block-cyclic and cyclic distributions. Cyclops Tensor Framework employs cyclic distribution to preserve packed symmetric structure in subtensors, minimize padding and generating a regular decomposition.

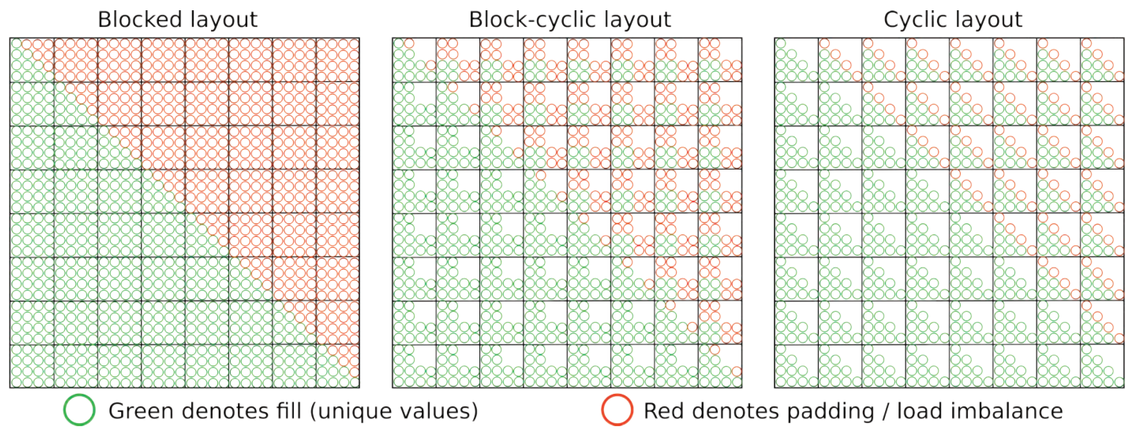
~~~~~

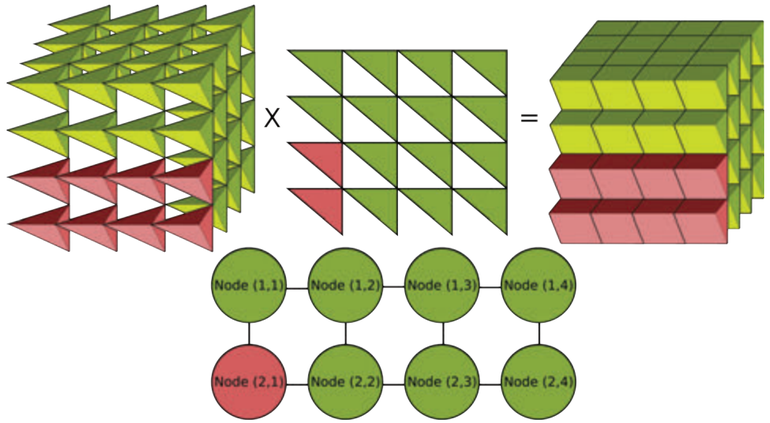

Cyclops Tensor Framework uses virtualization to decompose any given tensor contraction so the computational load is evenly balanced among processors. The virtual decomposition also is parameterized so it's effectively a multiple of the processor grid, insuring each processor owns the same number of sub-blocks. The scheme reduces the problem of mapping tensors with symmetry to mapping padded tensors with no symmetry. In this example, the three-dimensional virtualized mapping is decomposed among processors so each is contracting a matrix of symmetric tensors with a vector of symmetric tensors into a matrix of symmetric tensors. By the time the distributed contraction algorithm is executed, it need not be aware of the symmetry of the sub-tensors, but only of their size.

ScaLAPACK, a popular parallel linear algebra library, which ran the same problem at 100 teraflops.

Cyclops also compared favorably with NWChem on Hopper, a 1.28-petaflops Cray XE6 at Lawrence Berkeley National Laboratory. Although not tuned for Cray architecture, Cyclops computed a CCSD iteration twice as fast as NWChem and a CCSDT iteration three times as fast. On Mira, Cyclops efficiently scaled CCSD to 8,192 nodes, achieving 500 teraflops.

On Blue Gene, Solomonik says, "we'll be able to scale much further and do much bigger coupled cluster calculations than have been done before on any other architecture." The fellows, with Hammond and Demmel, describe their results in a paper in proceedings of the 27th IEEE International Parallel & Distributed Processing Symposium held this spring in Boston.

Hammond, an NWChem developer, says Cyclops is not designed to supplant it. The older code isn't just one computational procedure; it's a suite. Cyclops can plug in and work with NWChem's existing tensor contraction algorithm for added flexibility. Yet Hammond acknowledges NWChem performs poorly on CCSDTQ because it's tuned for lower-level calculations. "Once Cyclops gives the right answer for triples and quadruples, we will have something that we've really never had before." Then, Matthews adds, "we really get into doing some serious coupled cluster work and getting it to run very fast."

During the practicum, Matthews says, he learned several tricks and skills, like matrix multiplication algorithms and parallel programming techniques. The experience also introduced him to supercomputing, since he did most of his prior research on clusters using commodity processors.

The practicum, meanwhile, was one of the first times Solomonik has implemented a software library framework that is publicly available to scientists, rather than just studying algorithms' performance. Cyclops is the largest code he's worked on, partly because the team made it extensible to applications beyond coupled cluster. That makes coding "much harder, more painstaking, but in the end I find it more rewarding." Two years later, the project has virtually subsumed Solomonik's thesis research, while Matthews continues to refine the core algorithm.

With graduation approaching, Solomonik is leaning toward a university career, but also wants to expand his experience with postdoctoral fellowships. Matthews is unsure whether he'll go into academia or to a national laboratory. What's certain is Cyclops is part of their futures, Hammond says. "My son is two years old, and the fellows are still working together, so it was a lot more than a summer project."

IMPROVING INSTABILITY MODELS FOR FUSION'S FUTURE

SANJEEB BOSE
Stanford University
Lawrence Livermore National Laboratory

IT'S LITTLE WONDER Sanjeeb Bose's career combines engineering and computation. His father, Rathindra, is a cancer scientist and vice chancellor for research and technology transfer at the University of Houston. His mother, Anima, studies fuel cell technology. Both gave him gentle pushes toward technical fields, says Bose, a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient. They've always been interested in his work using computers to understand and predict turbulent fluid flow.

"They are primarily experimentalists by nature, so they're always very skeptical of anything that comes out of a computer," says Bose, a recent Stanford University doctoral graduate. "That skepticism is really nice to have."

It's also understandable: Because the underlying equations are impossible to solve precisely, Bose knows his computer models are subject to numerical errors. He's driven to reduce errors and isolate the remainder so engineers understand the uncertainty they introduce. His work contributes to research into uncertainty quantification (UQ): understanding error in models and putting a number on the accuracy of their outputs. It's also part of V&V: verification (whether equations are solved correctly and with what tradeoffs) and validation (whether the equations are the right ones for the simulation at hand).

Much of Bose's research focuses on large-eddy simulation (LES) of turbulent fluids. As the name implies, LES computes swirling flows like ocean currents. It's generally used to portray a single fluid or a well-mixed combination. But during his 2011 Lawrence Livermore National Laboratory (LLNL) practicum, Bose turned to modeling how different materials mix, particularly at the early stages, with implications for fusion, a potential source of virtually limitless, clean power.

Bose turned to modeling how different materials ... mix with implications for fusion.

~~~~~

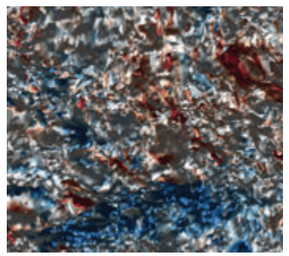

LLNL hosts the National Ignition Facility, a stadium-sized building where pulses from 192 laser beams converge in a gold container the size of a pencil eraser. The beams interact with the gold, generating X-rays that compress a BB-sized plastic shell containing a frozen mixture of hydrogen isotopes. If all goes well, the hydrogen nuclei merge in an inertial confinement fusion (ICF) thermonuclear reaction, generating more power than went in.

It's a big "if." Fusion depends on whether a spot in the fuel pellet gets hot enough, Bose says. "But you have several different materials that all exist at different densities inside this tiny capsule," including the plastic shell. Hydrodynamic instabilities form at the interface of the materials: Rayleigh-Taylor instability (RTI), in which a constant force like gravity pushes a dense fluid into a lighter one; and Richtmyer-Meshkov instability (RMI), in which a shock accelerates the interface. Instabilities can mix out the hot spot, limiting or stopping fusion.
~~~~~

"There isn't a prescription that says 'This is how you do it.'
It's a bit of an art."

~~~~~

BETWEEN THE EXTREMES

Scientists use computers to understand instabilities and help design experiments that minimize mixing. Precise RTI simulations already are available, Bose says, but they take too long to run, even on supercomputers, for practical use. Instead, scientists tap computationally inexpensive, but less accurate methods. Working under LLNL physicist Oleg Schilling on a project begun by a previous summer student, "the objective for my time was to see if you could live somewhere in between those two extremes," Bose says: a reasonably inexpensive but more accurate instability simulation.

The code, Schilling says, was two-dimensional and mainly modeled RTI and RMI by solving forms of the Navier-Stokes equations that describe the motions of fluids. Bose worked on extending it to three dimensions while limiting numerical dissipation and diffusion, which remove energy attributable to small-scale turbulence. Physics codes often use numerical dissipation as a surrogate for physical energy dissipation because omitting it can make algorithms crash and, on average, there's little energy at these small scales. Limiting numerical dissipation is vital, however, if the goal is high-fidelity simulations that resolve large-scale mixing and follow its evolution.

Schilling focuses on higher-order weighted essentially non-oscillatory (WENO) methods, which perform well when simulating shocked flows or when there are sharp differences between two species, such as a strong jump in the density gradient. But in smooth flows they are highly dissipative and lose accuracy, he says, so other methods must be used. Bose worked on creating a hybrid of WENO and central difference methods, which are better at simulating smooth flow.

Bose based hybridization on how much species were mixed in the simulation domain. "If they weren't mixed at all, you would clearly try to use WENO. If you only found one species or another in any part of the domain, you would try to use the central difference scheme. And then as you move from one end to the other, you would try to blend these schemes in some continuous manner." That's the hard part, Schilling says: "There isn't a prescription that says 'This is how you do it.' It's a bit of an art."

Bose's approach was somewhat successful but still vulnerable to crashes. In a test, he compared it with data from a Texas A&M University experiment in which a stream of air passes above a stream of helium. The researchers measured how much the two gases mixed as the dense air sank into the helium. The hybrid code inaccurately predicted the mixing rate early in the experiment, but improved in its late stages.

Bose added good capability to the code, Schilling says, and was a top-notch, quick programmer. During tests, "when things didn't appear right, he was quite creative in finding solutions or making improvements," and Bose accomplished more than expected during the practicum. Schilling set the code aside, but resumed work on it this year. Bose had hoped to help, but finishing his degree under Parviz Moin, head of Stanford's Predictive Science Academic Alliance Program (PSAAP) center, took precedence.

VERIFICATION, VALIDATION AND UNCERTAINTY

Stanford's center is one of five the DOE National Nuclear Security Administration (NNSA) sponsors. Each focuses on improving methods for V&V and UQ. NNSA also is interested in the applications each center uses to hone its V&V and UQ methods, such as Stanford's simulation of air-breathing planes capable of traveling at least five times the speed of sound – 3,400 miles per hour at an altitude of 60,000 feet. "We focus on multiphysics flow calculations in general, so it could have applications to jet engines, aircraft noise, you name it," Moin says, but the real goal is demonstrating the group's V&V and UQ methods.

Like most simulation techniques, LES uses data point grids or meshes to discretize, or divide, the physical domains it models. Computers calculate how properties like density and velocity change at each point. Drawn together, they portray a turbulent fluid the way pixels comprise a photograph. Unlike many techniques, LES adds filtering, removing from the solution a range of activity too small for the grid to capture. Instead, a mathematical model accounts for these fine effects, making the overall calculation less computationally demanding.

When LES doesn't agree with experiments or precise direct numerical simulations (DNS), Bose says, researchers want to attribute the error to models for small-scale turbulence or to numerical error. Codes most commonly are verified by running them with successively finer grids, each more closely spaced than the last. When differently spaced grids produce the same results, researchers know they've eliminated errors.

Most LES calculations, however, are sensitive to grid spacing: placing points closer together can cause results to differ from experimental or DNS data. The results converge only when grid resolution is so fine the turbulence model no longer matters. In essence, the calculation then is a DNS – something beyond the power of computers when simulating practical devices. That weakness makes it hard to say how much error is attributable to the subgrid-scale model.

Bose devised a way to solve the equations independent of mesh spacing, essentially eliminating errors due to numerical methods. "You could still refine a mesh and start to eliminate or reduce numerical errors, but the length scales you resolve don't change, because they're decoupled from the grid. You could

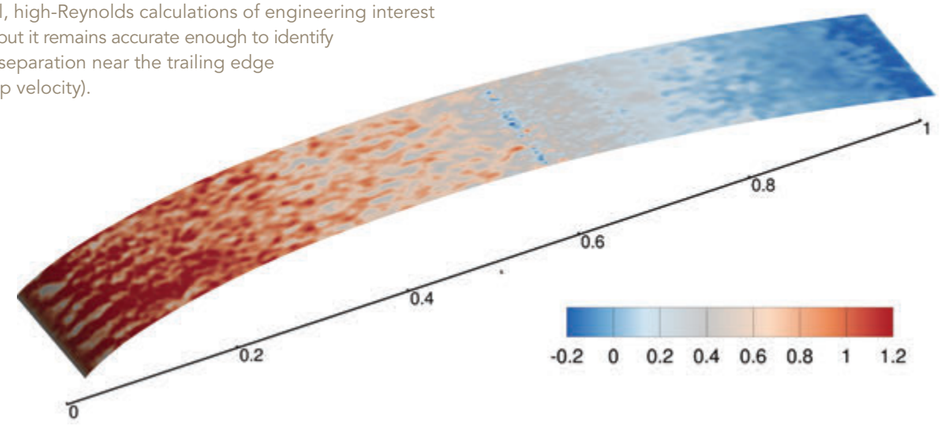

Instantaneous stream-wise slip velocity on the suction side of an airfoil at an angle of attack for a wall-modeled large eddy simulation (LES). Wall modeling in LES can make practical, high-Reynolds calculations of engineering interest computationally feasible, but it remains accurate enough to identify experimentally observed separation near the trailing edge (indicated by negative slip velocity).

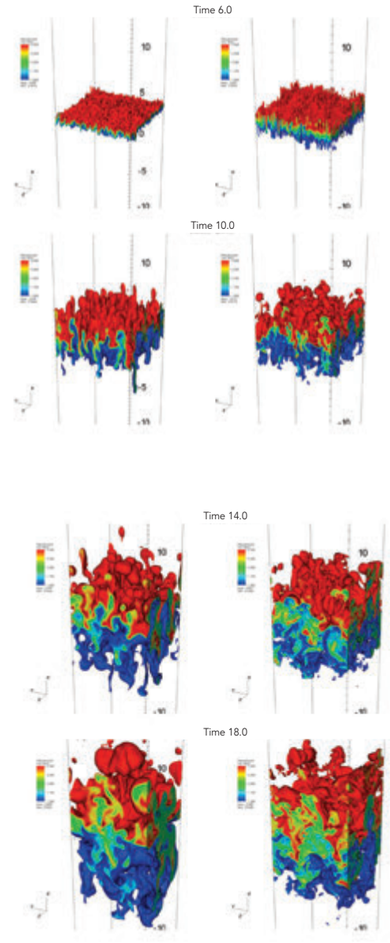

Credit: Kalen Braman, University of Texas at Austin

These visualizations show the time-evolution of density from direct numerical simulations of multimode Rayleigh-Taylor instability and mixing using the fifth-order WENO method at grid resolutions of 256 by 64 by 64 (left column) and 512 by 128 by 128 (right column).
~~~~~

actually get a converged solution that is not a DNS," Bose says. "If that converged solution fails to agree with either a highly resolved calculation or experiments, then you could say with some confidence that this is really due to the shortcomings of our subgrid-scale models."

OUT OF BOUNDARY

Knowing where the subgrid-scale model failed lets researchers refine the grid at just those spots, efficiently focusing computation to increase fidelity. Bose, Moin and Stanford Engineering Research Associate Frank Ham tested grid adaptation by modeling a Stanford flow diffuser experiment. Although diffusers are used in gas-turbine engines, this one was designed to create complex boundary-layer separation, in which fluid flow moves away from the wall, creating a bubble of slowly recirculating back flow. Diffusers increase pressure while slowing flow. Boundary layer separation stalls that, but predicting whether it will happen and where is a challenge, Bose says.

The researchers ran their grid-adapting LES of the diffuser on thousands of processors on LLNL's IBM Blue Gene/L and Hera, an Appro International cluster, and on Los Alamos National Laboratory's Mapache, a Silicon Graphics machine. Although their LES must span a wide range of scales, it agreed with pressure increase data from the experiment. The simulation also did well at predicting boundary separation location and extent.

The simulations are demanding because of their high Reynolds number, a measure of how a fluid's viscosity affects its ability to flow and generally taken as a measure of turbulence. The experiment and simulation reached a Reynolds number of 50,000 – high, but still less than needed to effectively model engines, cars and other real-world machines. Yet, "computers are getting fast enough and our abilities to model scales are slowly coming along," Bose says, so the solution of such problems may be within reach someday.

Bose can help drive that development, Moin says. "He is savvy in computer science – very, very savvy – he can program well, he knows his applied math very well and he has very good physical insight. These are the three components I look for in my students and he possesses all of them."

It's little wonder, then, that Cascade Technologies, a company Moin cofounded, hired Bose even before he graduated in December 2012. The firm, naturally, develops LES and computational fluid dynamics tools for industry. Bose plans to continue "interacting with what I hope will be increasingly exciting machine hardware. I fully intend to continue to study these problems."

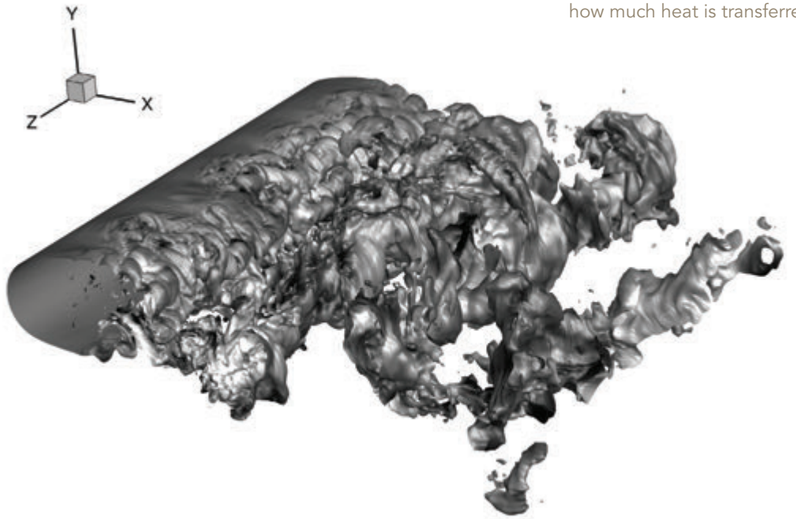

This shows an instantaneous isosurface of normalized temperature in the near wake of a heated cylinder in crossflow, used to visualize how much heat is transferred from the cylinder.

"I did want to take seriously the purpose of the practicum: to broaden your skill set beyond your own field."

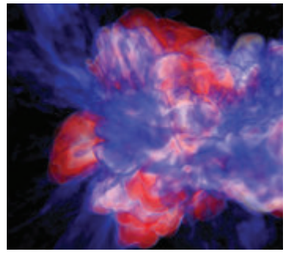

SMALL-TOWN GUY TACKLES BIG PROBLEMS

SAMUEL SKILLMAN
University of Colorado, Boulder
Oak Ridge National Laboratory

SAMUEL SKILLMAN IS FAMILIAR with the challenges of going big, whether in life, astrophysics or computer science.

Skillman, a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient, grew up in the northern California mountain hamlet of Sierra City (2010 population: 221). His high school, in nearby Downieville, had just 40 students. When Skillman expressed an interest in physics, a teacher directed him to custodian Mike Kelly, a science buff who loaned Skillman books like Brian Greene's "The Elegant Universe." Skillman says his talks with Kelly were crucial to deepening his interest in physics and cosmology.

With that background, scaling up to even a small institution like Harvey Mudd College near Los Angeles was a shock, Skillman says. He no longer was the smartest student, and after earning top high school grades, "dropping into Harvey Mudd, where I had to work my butt off," required adjustment.

Skillman went big for his doctoral research at the University of Colorado, Boulder, from which he recently graduated: galaxy clusters, the universe's most massive structures. Skillman's simulations predict what radio astronomers should find when seeking artifacts of shocks passing through clusters.

Skillman's Spring 2011 practicum at Oak Ridge National Laboratory (ORNL) addressed aspects of going big in computer science. He plunged into programming for supercomputers that incorporate two or more kinds of processors. In many of the world's top systems, chips like graphics processing units (GPUs) work alongside standard multicore processors. That includes ORNL's Titan, a Cray XK7 capable of more than 20 petaflops (quadrillion scientific calculations per second) and rated the country's fastest computer. It combines standard AMD multicore processors with NVIDIA Tesla GPUs.

These heterogeneous architectures will be key to reaching exaflops speeds – a million trillion scientific calculations per second, about a thousand times faster than today's best machines – without consuming as much electricity as a small city. "If scientists want to continue to run larger and more complicated and better simulations, they will need to learn how to take advantage of these complex systems," Skillman says. "That's what we'll all have to learn to deal with. Some of it might be painful."

STEPPING INTO ANOTHER FIELD

The chance to work with GPUs is partly why Skillman chose ORNL, even though the project had little to do with his area of expertise. "There were other practicums I could have done that were astrophysics related," but "I did want to take seriously the purpose of the practicum: to broaden your skill set beyond your own field."

Simulations on parallel processing computers typically represent materials or processes as a grid of data points, then use this framework to perform calculations. One technique, stencil computation, updates each point on a fixed pattern based on nearby elements'

values. Imagine a single dice: a stencil calculation updates the center based on exchanges with each of the six faces – a seven-point stencil. Stencils are key to a number of simulations, particularly molecular dynamics (MD), which can help scientists devise new plant-based fuels or improve combustion. Under ORNL Staff Scientist Wayne Joubert, "my piece of the puzzle was to write the fastest, most optimized code we could for doing this seven-point – and in the future, more points – stencil computation" on a GPU, Skillman says.

The main obstacle dealt with how GPUs work. They're best at taking small amounts of information, performing a huge number of floating point operations per second (flops) on them and holding the results until they're ready for output. Stencils require only a comparative handful of flops on each data point, so information must be moved from memory more often. This low ratio of flops to memory operations per second – flops per MOPS – makes stencils difficult to run at peak speeds. The chief brake is memory bandwidth – the quantity of information that can move to and from memory in a given time. "You have to optimize your memory movement a lot more than you have to optimize your floating point operations," Skillman says. "What you try to focus on is getting as close to the memory bandwidth limit as you can."

Another problem: Skillman didn't know CUDA, the standard GPU programming language, and it "really looked nasty."

HITTING TOP SPEED

Skillman's stencil code explicitly manages the memory hierarchy, doing most work at the level offering the fastest access: local memory or registers, where the computer performs calculations. It moves data in large chunks to limit communication with slow-access global memory, which the entire GPU processor shares. He tested the code on the NVIDIA Fermi C2050, which packs 14 multiprocessors with 448 computing cores on each chip to deliver 515 million flops. Skillman's code accelerated as the problem grew, as measured by grid size. At its peak, the algorithm ran at around 35 gigaflops (billion flops) and updated about 5 billion cells per second for a seven-point stencil.

That's "probably about as fast as you can get for that hardware," Joubert says, and shows that bandwidth is the main restraint. "We were very happy that no stone was left unturned in terms of maximizing performance."

Skillman also used the Python programming language to create a simple performance model that projects how the code would run under different conditions, including its requirements and the GPU's speed, memory and memory bandwidth. Its predictions nearly matched the code's performance on the Fermi C2050. The more advanced Fermi C2090, the model predicted, should boost the code's performance by about 15 percent – largely based on the GPU's larger memory bandwidth. A 125-point stencil, useful for some algorithms in large MD codes, should hit around 90 gigaflops.

Skillman lectured on his project at ORNL just before his practicum ended and displayed a poster at the lab's Fall Creek Falls Conference in September 2011 and at the 2011 CSGF Annual Conference.

Skillman and Joubert designed the stencil code to be incorporated into large simulations. But Skillman's biggest contribution may have been helping understand what's necessary to make stencils and other algorithms run efficiently on GPUs, both technically and in programmer time and skill. "With the lessons learned from that short and concentrated project, we can go back to these other projects and say … we have in-house expertise for how these kinds of computations work on GPUs," Joubert adds.

Skillman, meanwhile, notes that CUDA "ended up not being a horrible language to code in. It's actually pretty nice." He may even use it and GPUs for astrophysics research, including Enzo, the main program he uses.

SNIFFING OUT "RADIO RELICS"

With Enzo, Skillman and his thesis advisor, Jack Burns, probed galaxy cluster "radio relics": the radio-wave signatures of shocks generated when clusters merge or pull in gas from surrounding space, a process called accretion. The shock fronts accelerate electrons, which spiral around magnetic field lines, emitting radio waves that sensitive Earth-based instruments can detect.

Clusters contain gas clouds and hundreds of galaxies, all cemented by gravity. It takes powerful computers to simulate clusters' evolution because of the span of scales involved: Individual galaxies may be just 100,000 light years across, but the clusters are 10 million to 30 million light years wide.

"We were very happy that no stone was left unturned in terms of maximizing performance."

~~~~~

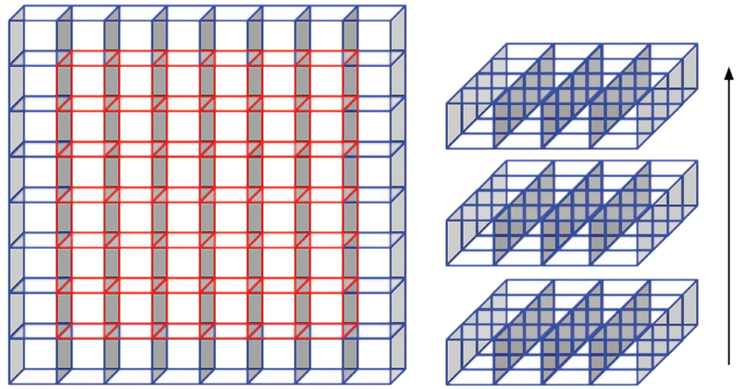

Far left: This schematic depicts stencil computation active zones in red, with a layer of ghost zones in blue. Each CUDA thread block is responsible for updating the redcells. Correctly setting up data access patterns to gather the needed values for the blue cells is a key optimization for the stencil computation.

Near left: Given a three-dimensional dataset, the algorithm Skillman and his colleagues developed decomposes first into two-dimension slabs. Since the center cell needs knowledge of the neighbors above and below (for a seven-point stencil), they found that keeping three layers in the shared memory space optimized data movement. These three active layers then move up through the domain, cycling out unneeded layers from the bottom at each step.

Enzo, originally developed by Columbia University's Greg Bryan and now by a team from across the country, is suited for the job, Skillman says. Its adaptive mesh refinement (AMR) approach performs detailed simulations of areas researchers are most interested in while leaving other parts less refined to save computer power. Like focusing a camera on part of a scene while leaving the rest fuzzy, AMR "captures the interesting physics that occurs on all these scales and couples them together."

Researchers combine galaxy cluster shock simulations with Earth-based observations to understand the physical processes that led to the radio emissions. Skillman's simulations predict what new radio telescopes might see, given the known physics. When observations don't agree with his simulations, it means something in the simulations' physics must be incorrect. "Now we have the task of going back to our models and trying to fix them so they better match" reality, he adds.

The computational tools Skillman added to Enzo are groundbreaking, Burns says, including an algorithm that's enhanced researchers' ability to identify shocks resulting from cluster mergers. Skillman later added models for diffusive shock acceleration, translating shock strength into an expected energy spectrum for radio emission-producing electrons.

A BRIGHTER SHOCK

For a 2011 *Astrophysical Journal* paper, Skillman, Burns and several colleagues simulated two random representative pieces of a universe like ours – "boxes" of space, one about 300 million light years and one about 930 million light years on each edge. Running on Ranger, a computational cluster at the University of Texas at Austin, the researchers evolved accretion and merger shocks over billions of years. Results indicate that interior shocks from mergers generate brighter radio emissions, even though accretion shocks are stronger. That's probably because interior shocks move through areas where matter is denser, yielding more accelerated electrons, the researchers say.

The paper also predicts how many radio relics newly sensitive observatories should see in portions of the sky out to a certain redshift – the distance from Earth, and thus the time into the past, as measured by the Doppler effect in light. There probably are many more than those already detected, Skillman says, even in previously surveyed clusters. "These galaxy clusters are very complex environments," with subclusters merging and swarming. "Those produce shocks and so far we've only seen the grandest, where there's a giant shock of two galaxy clusters coming together. There are a bunch of little guys and they're also producing radio emissions at a smaller scale."

Skillman's *pièce de résistance*, Burns says, was collaborating with other researchers to incorporate magnetic field effects using Enzo. "Now we can produce the most realistic radio maps," with the first results published earlier this year in *The Astrophysical Journal*. The radio source appearance and structure simulations "look scarily similar to the real radio observations."

Skillman is among the best graduate students and postdoctoral researchers Burns has supervised during 35 years in academia. "The last couple or three years, he's performed at the level of a postdoc," Burns adds. "It's wonderful and refreshing to have a student that operates at such a high level."

Now Skillman really is a postdoc, at the Stanford/Kavli Institute for Particle Astrophysics and Cosmology, where he plans to continue his radio relics simulations. New radio telescopes coming on line may provide details needed to find the "little guys" his models predict. No doubt even bigger problems and accomplishments await him.
~~~~~

Right: This view of Lawrence Berkeley National Laboratory's Advanced Light Source (ALS) looks head-on into the upper and lower rows of magnets in an undulator. The vertical distance between the magnets can be adjusted to determine the wavelength emitted. In this photo, a laser simulates the burst of light produced. *Credit: Lawrence Berkeley National Laboratory.*

SPECIAL FEATURE

BIG DATA HITS THE BEAMLINE

A Data Explosion is Driving a New Era of Computational Science at DOE Light Sources

By Jacob Berkowitz

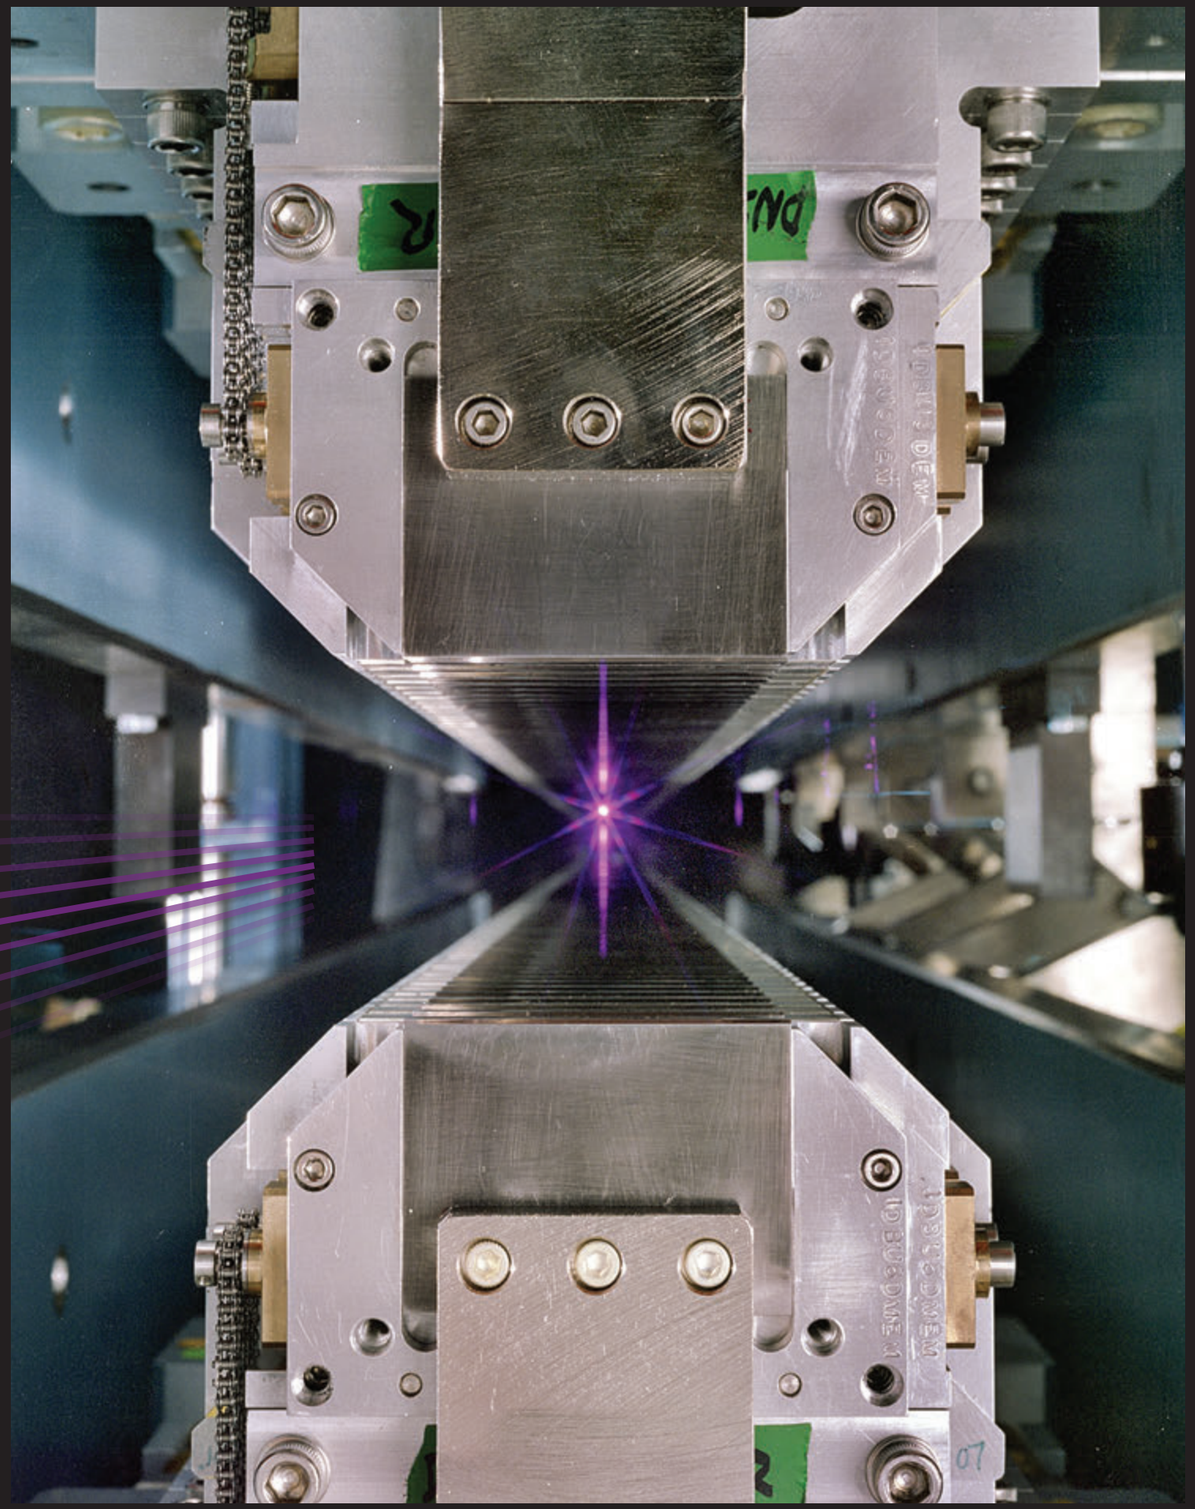

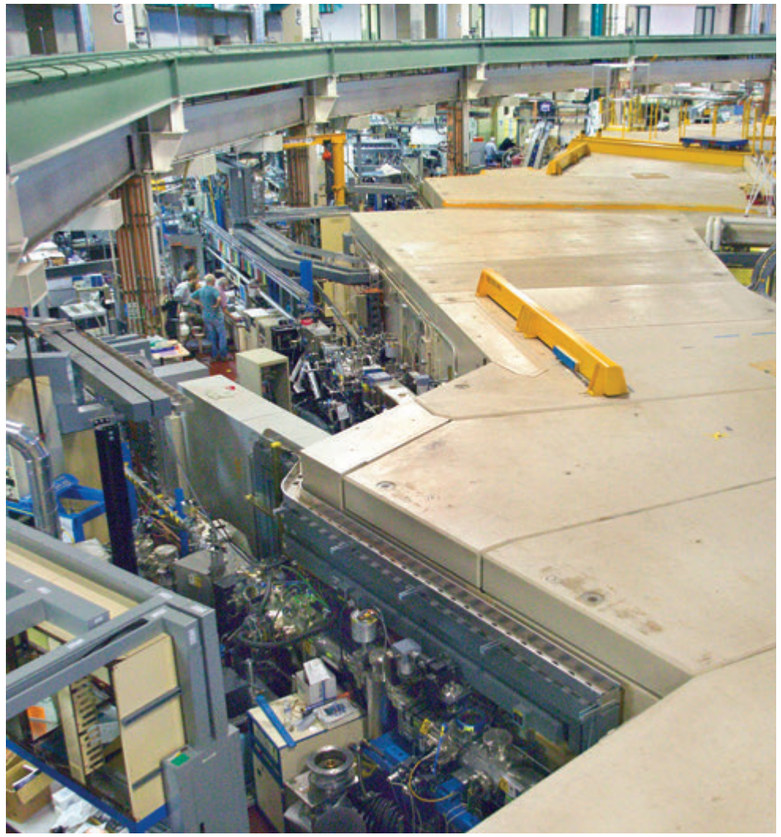

Right: An array of beamlines fans out from the storage ring (under concrete shielding) at Berkeley Lab's ALS. Each beamline can support a different experiment, each generating considerable amounts of data. *Credit: Lawrence Berkeley National Laboratory.*

When scientists from around the world visit Dula Parkinson's microtomography beamline at Lawrence Berkeley National Laboratory's Advanced Light Source, they all want the same thing: amazing, scientifically illuminating, micron-scale X-ray views of matter, whether a fiber-reinforced ceramic composite, an energy-rich shale, or a dinosaur bone fragment.

Unfortunately, many of them have left lately with something else: debilitating data overload.

"They're dying because of the amount of data I'm giving them," says Parkinson, a beamline scientist at the Advanced Light Source (ALS). "Often they can't even open up their whole data sets. They contact me and say 'Dula, do you have any idea what I can do? I haven't been able to look at my data yet because they crash my computer.'"

Data sets from light sources, which produce X-rays of varying intensity and wavelengths, aren't enormous by today's standards, but they're quickly getting bigger due to technology improvements. Meanwhile, other fields – astrophysics, genomics, nuclear science and more – are seeing even mightier explosions in information from observations, experiments and simulations.

If all this knowledge is to benefit the world, scientists must find the insights buried within. They must develop ways to mine mountains of elaborate information in minutes or hours, rather than days or weeks, with less-than-superhuman efforts.

To deal with the challenge, Department of Energy scientists – whether at light sources or particle colliders – are collaborating with computational scientists and mathematicians on data-handling and analysis tools.

"The data volumes are large, but I think even more importantly, the complexity of the data is increasing," says Chris Jacobsen, a 25-year veteran and associate division director at Argonne National Laboratory's Advanced Photon Source (APS). "It's getting less and less effective and efficient to manually examine a data set. We really need to make use of what mathematicians and computational scientists have been learning over the years in a way that we haven't in the past."

Big data management, analysis and simulation are driving users to new levels of high-performance computing at the ALS, the APS and three other DOE Office of Science light sources. Researchers are approaching their experiments from a data-intensive computational science perspective.

Those involved say there's enormous opportunity to stimulate

Above: A long hallway separates the Linac Coherent Light Source near hall from the far hall at the SLAC National Accelerator Laboratory. The X-ray laser beam travels through the pipe on the right, capturing images of atoms and molecules in motion. *Credit: SLAC National Accelerator Laboratory.*

Human Connections to Tackle Big Data

Even with leadership-class computers, the latest data-management software, and gigabyte-speed connections, light source and computational scientists say solving big-data issues comes down to a face-to-face formula: coffee and conversation.

"The key thing is for those of us who are working on a problem to be able to sit in a room together and discuss it, and hack out the base issues," says Craig Tull, leader of the Science Systems Software group at Lawrence Berkeley National Laboratory.

Tull heads the lab's involvement in SPOT Suite, an innovative collaboration capitalizing on the proximity to the National Energy Research Scientific Computing Center (NERSC). It brings together computational scientists and resources and provides high-performance computing capabilities for analyzing and simulating large datasets generated at the lab's Advanced Light Source (ALS).

At NERSC, development of SPOT Suite's Web portal is led by computational scientist Jack Deslippe, a 2010 DOE Computational Science Graduate Fellowship alumnus.

"The cool thing about working with a facility like the ALS is that there is an unlimited number of projects you might work on in any number of computational science fields," Deslippe says. He joined NERSC in the fall of 2011, fresh from earning his Ph.D. at the University of California, Berkeley.

This is music to the ears of ALS scientists like Dula Parkinson, whose beamline is the test case for SPOT Suite applications.

"As beamline scientists we haven't had the right language to even know how to ask the questions to get the computer people interested in our problems," says Parkinson, who supervised first-year DOE CSGF recipient Justin Lee for a summer 2013 practicum focusing on computational imaging. "It feels like we're finally overcoming that barrier more and more in this relationship."

Since SPOT Suite's Web portal went live in mid-April, conversations are fueling on-the-fly improvements.

"What we are finding is that it is useful for the users to give us real-time feedback on our portal," Deslippe says. "So the development has ended up being a back-and-forth process where research on our real-time analysis approach is coupled to real research problems at the beamline."

Left: This aerial view looks east down the SLAC National Accelerator Laboratory's 2-mile long linear accelerator toward Stanford University and the San Francisco Bay. *Credit: SLAC National Accelerator Laboratory.*

leapfrog advances in light-source science and to create innovative collaborations and a community of computational scientists specialized in such research. *(See sidebar: Human Connections to Tackle Big Data.)*

"Light sources are an archetype for the new data-intensive computational sciences," says Craig Tull, leader of Berkeley Lab's Science Systems Software group.

THE LIGHT SOURCE DATA DELUGE

If there's a symbol for the big data changes at the nation's light sources, it's the external computer drive. Even today, results from many beamline experiments can be loaded onto portable media like a thumb drive, just as it's been done for decades. Back at their home institutions, researchers can use workstations to process the information from X-rays that ricochet off a protein molecule, for example, and strike detectors.

Light sources have historically operated on this manual grab-and-go data management model, reflecting the nature of synchrotron experiments.

"Synchrotrons are big machines with traditionally bite-sized experiments," says Jacobsen, at Argonne's APS. As at all such facilities, the APS' stadium-sized, circular synchrotron accelerates electrons nearly to light speed, generating a cascade of bright X-ray photons. These photons are tuned and focused to feed 60 simultaneously operating beamlines with wavelengths ideal for resolving matter from the atomic to cellular level.

What's changed: On many DOE light source beamlines, a manageable bite of experimental data has ballooned into daily helpings of terabytes (TB – trillions of bytes). Four primary factors are driving this data volume spike, Parkinson says.

First, light source detectors are collecting images with unprecedented speed. At the APS, a new generation of detectors has turned what used to take 15 minutes of imaging into a 15-second job. Cameras already exist to capture even higher-resolution images in just milliseconds. That means at full use, APS could produce a staggering 100 TB of data a day – a rate comparable to that of the Large Hadron Collider, the giant European physics experiment.

Similarly, data output is doubling every year from the 40 beamlines at Berkeley's ALS. And Brookhaven National Laboratory's new National Synchrotron Light Source II (NSLS II, the sixth Office of Science light source) is expected to generate about 15,000 TB of data per year later this decade – not an enormous amount when compared to some other experiments, but a big jump for an X-ray source.

Second, light sources have gotten brighter at a rate even faster than Moore's Law of accelerating computer power. That means shorter exposure times and more data.

The shining case in point: the Linac Coherent Light Source (LCLS), the world's most powerful X-ray laser, which came on line at DOE's SLAC National Accelerator Laboratory in 2009. LCLS is the world's ultimate fast camera flash, illuminating samples for one-tenth of a trillionth of a second with X-rays a billion times brighter than previous sources.

The combination of faster detectors and greater light power is opening the door to time-resolved experiments, capturing the equivalent not of photos, but video – with the attendant exponential boost in data volumes. Today, detectors running at maximum output can generate a terabyte of data per hour, a study by Berkeley Lab's Peter Denes shows. By 2020, that could reach 1,000 terabytes – one petabyte – per hour.

Here's another example: Not long ago, it was common for a single tomographic X-ray scan, capturing hundreds or thousands of images as the sample rotates, to take an hour. Earlier in 2013, Parkinson says, a group at the ALS captured changes in a sample by collecting similarly sized scans once every three minutes for more than 24 hours, producing terabytes of data. "These are the users that bring the whole data system to its knees," he adds.

Finally, increasingly automated data management accelerates these factors. For example, Jacobsen is leading an APS pilot project to automate data transfer from the local beamline computer to networked machines, a process that has been done manually.

These big increases in information volume also are compounded, Jacobsen says, by the variety of materials studied and efforts to integrate data from experiments performed at different wavelengths.

With so much complex information to process, he adds, "what we really need is the intelligence about how you look at the data."

NEW ERA OF LIGHT SOURCE COMPUTATIONAL SCIENCE

With the results they produce rapidly growing in quantity, variety and complexity, light sources are laboratories for developing techniques to gather that intelligence. Users are turning to high-performance computing

hardware, software and, perhaps most importantly, computer scientists, launching a new era of data-intensive management, analysis and visualization at the facilities.

Two initiatives characterize the ways DOE light sources (and other DOE programs) are incorporating computational expertise and data-management capacity, whether on-site or by collaborating with a supercomputing facility – setting the groundwork for what many envision as a combination of both.

First, a Berkeley Lab project links beamline data in real time to some of the nation's most powerful open-science computers at the lab's National Energy Research Scientific Computing Center (NERSC), based in Oakland.

"We're trying to move the typical data-intensive beamline into a world where they can take advantage of leadership-class, high-performance computing abilities," says Tull, who leads the project. "In the final analysis what we're trying to do is drive a quantum leap in science productivity."

The software collaboration, dubbed SPOT Suite, unites computational scientists from NERSC and the lab's Computational Research Division with ALS beamline scientists and users. They plan to capitalize on NERSC's computational power to manage, analyze and visualize big data from ALS beamlines.

The initiative leverages DOE's advanced Energy Sciences Network (ESnet) to transfer data from ALS to NERSC at gigabytes per second.

SPOT Suite users access NERSC resources via a Web portal being built by Department of Energy Computational Science Graduate Fellowship (DOE CSGF) alumnus Jack Deslippe. A beta version went live in April.

With SPOT Suite, scientists using ALS' microtomography beamline for three-dimensional, time-resolved, micron-resolution images "can see their data being processed, analyzed and presented in visual form while they're at the beamline," Tull says. "This is something many of them have never seen before at any light source."

Second, in anticipation of the big output from the LCLS – adding a data dollop on top of that from the existing Stanford Synchrotron Radiation Lightsource – SLAC installed an on-site shared high-performance computing cluster. This new model of local computer support also will be implemented at Brookhaven's NSLS II.

"The data volumes and rates for LCLS are unprecedented for a light source," says Amber Boehnlein, SLAC's Head of Scientific Computing since 2011. She has a wealth of big data experience, including former responsibility for computing and application support for Fermi National Accelerator Laboratory experiments.

"We provide a robust computing and storage infrastructure common across all instruments at SLAC," Boehnlein adds, allowing experimenters to do a first-pass analysis on local SLAC resources, rather than waiting until they return to their home institutions.

One key to effective data management, Boehnlein adds, is using these resources for local data reduction and even to cut the information flow before it starts. SLAC computational scientists are collaborating with beamline users on algorithms for what Boehnlein calls "smart data reduction."

"People will start to develop much more sophisticated algorithms and they'll be using those during data collection to take better, higher quality data," she says.

Indeed, in many physics experiments, scientists set an initial filter that weeds out everything but the desired high-energy events, reducing the initial data collected by as much as 99 percent, says Arie Shoshani, who this year marks his 37th anniversary as head of Berkeley Lab's Scientific Data Management Group.

Shoshani provides the long-ball, bird's-eye perspective on DOE big data issues as director of the year-old Scalable Data Management, Analysis and Visualization (SDAV) Institute. SDAV's mission is to apply existing high-performance computing software to new domains experiencing data overload.

The need for large-scale simulations of things like nuclear fusion and climate has been the primary driver in supercomputer development, Shoshani says. Light sources – and some other DOE research facilities – need heavy iron for large-scale simulations less than they need the thinking that goes into high-performance computing, but that doesn't mean there aren't lessons to be learned.

"Whatever we know from doing large-scale simulations – from data management to visualization and indexing – can be useful now as we start dealing with large-scale experimental data," Shoshani says. Earlier this year he joined a discussion of hardware requirements for data processing at the 2013 Big Data and Extreme-Scale Computing meeting, jointly sponsored by DOE and the National Science Foundation.

For Parkinson, who runs the microtomography beamline, these high-performance computing tools change the experimental data he collects from a PC-crashing burden into a scientific bonanza.

"The beamline users are really excited about what's happening," he says. "Though it's early days, I think it's really going to accelerate what they can do and improve what they can do."

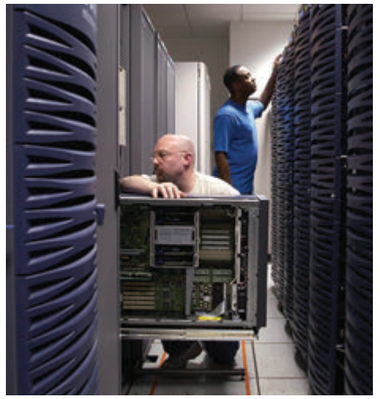

Left: Information technology workers check out systems at Argonne National Laboratory's Advanced Photon Source. Such systems are needed to handle a flood of data from the facility. *Credit: Argonne National Laboratory.*

Below: This three-dimensional rendering from computed microtomography data shows matrix cracks and individual fiber breaks in a ceramic matrix composite specimen tested at 1,750 C. Each of numerous ceramic samples is imaged with powerful X-ray scattering techniques over time to track crack propagation and sample damage, producing prodigious amounts of data. *Credit: Hrishi Bale and Rob Ritchie, Lawrence Berkeley National Laboratory.*

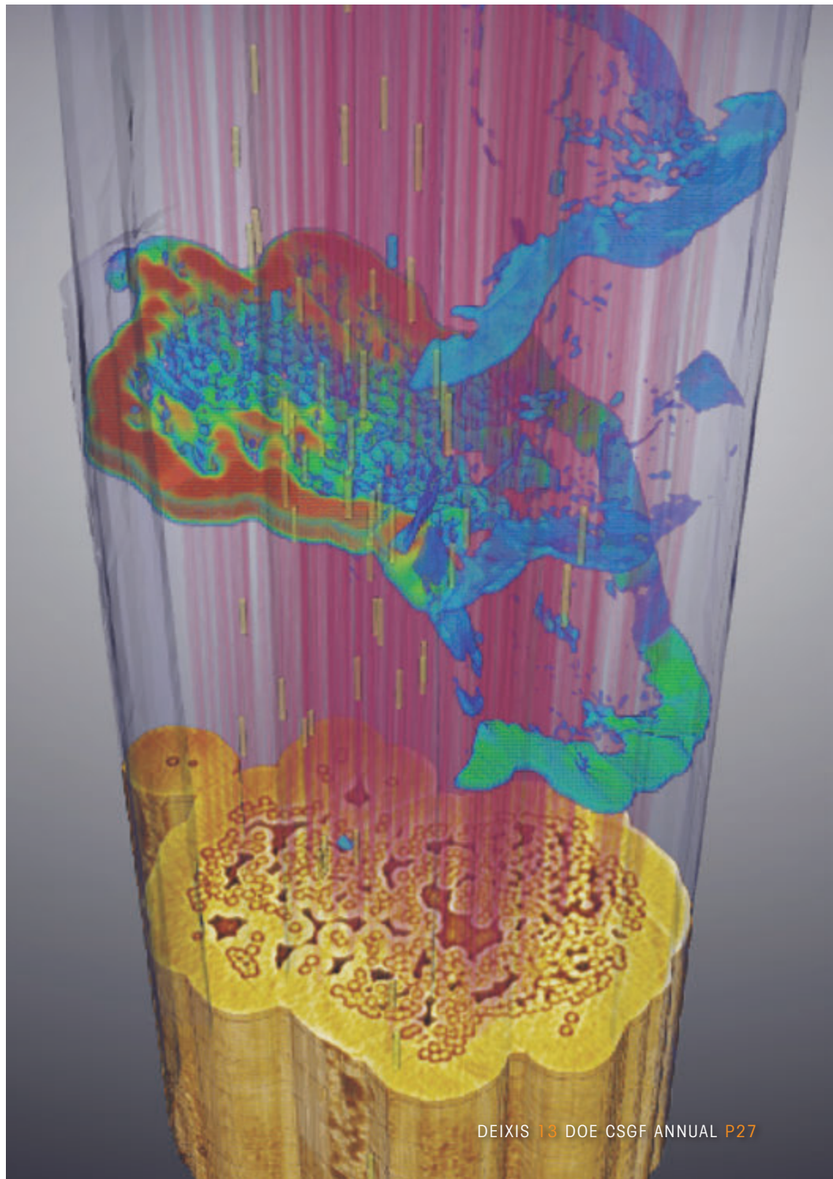

GRADUATES GO ON TO EXCEL

Going to Extremes in Materials and Disease

Timothy Germann
Los Alamos National Laboratory

TIMOTHY GERMANN lives in an extreme world, where copper, iron and other materials face crushing pressures, powerful shocks and intense radiation.

Germann, a Department of Energy Computational Science Graduate Fellowship recipient from 1992 to 1995, uses high-performance computers to simulate how atomic bonds break, molecules separate and materials disintegrate under the extreme impact and radiation environments found in outer space, the military and nuclear reactors.

But in lay circles Germann, a scientist in the Physics and Chemistry of Materials Group at Los Alamos National Laboratory (LANL), may be best known for a side project: the first large-scale agent-based simulations of disease spreading through a population. Seven years after Germann appeared on network television to discuss the research, he and his colleagues still work with the U.S. Centers for Disease Control and Prevention (CDC), most recently using their models to update pandemic influenza control strategies. "We're using simulations to do some of the 'what if' scenario evaluations," Germann adds.

The Illinois native has played with computers since the Commodore VIC-20 days. Still, he was unsure that computer science was a viable career on its own, so he double-majored with chemistry at the University of Illinois at Urbana-Champaign. "It shows how shortsighted I was," Germann says, but the combination led to a doctorate in chemical physics at Harvard.

At LANL, Germann has helped refine codes like SPaSM (Scalable Parallel Short-range Molecular dynamics) that calculate the movements and interactions of atoms and molecules. With increasing computer power and improved algorithms, he and his fellow researchers can simulate microns of matter and shocks stretching over nanoseconds, far beyond the capabilities of a few years ago.

X-RAY VISION

Historically, computer simulation was the only way to study how materials behave at these minuscule spatial and temporal scales, Germann says. Gas gun, flyer plate and laser-driven shock experiments provide information only about surface motion and the material's final condition. That's changing, Germann says, with powerful new X-ray facilities like the Linac Coherent Light Source at the SLAC Linear Accelerator Laboratory and Argonne National Laboratory's Advanced Photon Source. Both capture high-speed snapshots, generating both temporal and spatial data.

"It's really exciting," Germann says. At both facilities "they're able to make diffraction measurements with picosecond time resolution, which matches the time scale of our simulations." The thin foils used in experiments also are similar in size to computer simulations. This improved experimental capability "makes us more honest," Germann says, by providing better data to compare with models. The interchange helps computational scientists improve their codes and helps researchers design better experiments.

Germann now heads a project designed to push simulation capability even further. The Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx) is one of three interdisciplinary centers DOE's Advanced Scientific Computing Research program commissioned to surmount obstacles to creating applications that run well on exascale computers – ones about a thousand times more powerful than today's top machines. With scientists at Lawrence Livermore, Oak Ridge and Sandia national laboratories, Stanford University and the California Institute of Technology, LANL researchers will probe ways to connect software and hardware to better portray materials.

OPENING A TWO-WAY STREET

The trend toward computer designs that mix millions of processors of two or more types drives the project, Germann says. "With this very different heterogeneous architecture, we had to face the fact that some of the decisions we made in how the algorithms were implemented and how the data were laid out need to be rethought," Germann says. In the past, "it's been a one-way response of the codes to changing architecture trends … but the hope of co-design is that it can be a two-way street," with application scientists influencing computer hardware and system software plans.

The team is developing a multiphysics framework to better simulate materials under shock compression and high strain-rate loads, and pursuing algorithms that bridge spatial scales by adaptively refining models to create precise direct numerical simulations where needed. If all goes well, the team could have codes ready when exascale machines are available.

Getting computer scientists and materials science researchers to communicate effectively has been tricky, Germann says, but he enjoys the challenge. "It's this multidimensional jigsaw puzzle," he adds, "with the algorithms, the middleware and the architectures. It's neat to have a role in all three pieces."

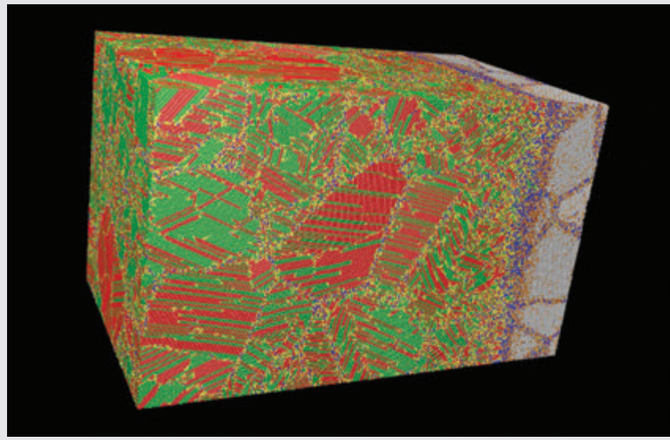

This visualization shows a simulation of nanocrystalline iron 14.6 picoseconds after a shock hit it at a velocity of 906 meters per second. The initial temperature is 50 degrees kelvin (K), reaching a temperature behind the shock front of 296 K and a pressure of about 39 gigapascals with a volume compression of 19 percent. The sample consists of 32 grains and about 30 million atoms confined in a 57.4 nanometer (nm) by 57.4 nm by 109.9 nm box. Color-coding denotes the local neighborhood crystal system of each atom. Gray represents body-centered cubic (bcc); blue is uniaxially compressed bcc; yellow is grain-boundary; red is hexagonal close packing; green is face-centered cubic.

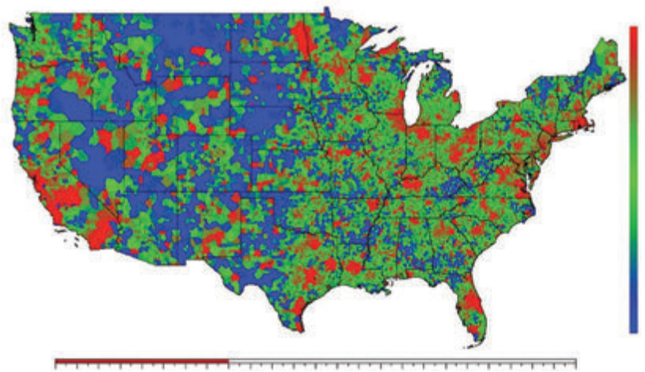

This map shows the predicted prevalence of pandemic influenza 80 days after the arrival of 10 infected individuals into the United States through Los Angeles International Airport. In this scenario, a highly virulent strain, comparable to the 1918 flu pandemic, is spread without any effective countermeasures (such as vaccines or antiviral drugs), resulting in a cumulative infection attack rate of 54 percent. Prevalence of symptomatic cases at any location is indicated on a logarithmic color scale, from 0.03 percent (blue) to 3 percent (red) of the population.

Christina Payne
University of Kentucky

Flexible Field Guides Alumna Down Many Paths

CHRISTINA PAYNE has sampled every sector open to Department of Energy Computational Science Graduate Fellowship (DOE CSGF) alumni.

First there was industry, working for a Denver engineering firm to model large-scale processes at oil and gas facilities and for nuclear waste reclamation.

Next was DOE's National Renewable Energy Laboratory (NREL), also in Colorado. Payne's attention turned from barrels-per-day processing to molecular-scale systems as she used computers to understand the action of cellulases – enzymes that break cellulose into sugars for conversion to ethanol.

Now Payne, a DOE CSGF fellow from 2003 to 2007, is in academia as an assistant professor of chemical and materials engineering at the University of Kentucky. As at NREL, her main tools are molecular dynamics (MD) codes, which calculate the physical movements and interactions of atoms and molecules.

The switch from sector to sector, Payne says, illustrates what attracted her to her field. "There's a lot of flexibility to do whatever you want in chemical engineering. You can do things at the molecular scale, like I do now, or you can do things at the process scale, like I used to."

But there are definite differences. At NREL, Payne worked under a cooperative research and development agreement (CRADA) with a corporation, working to improve an enzyme for industrial ethanol production. The company's data made it relatively easy to verify her simulations, but the work's proprietary nature means publication is unlikely in the near future.

PROPRIETARY VS. PUBLISHED

"As an academic, I like to publish, and I like to move forward and let the community know what I'm working on and learn from them as well," Payne says. "That's essentially what I'm trying to do now: start to develop my own experimental collaborations."

The first is with Norwegian University of Life Sciences researchers, who use X-ray crystallography to determine the structures of proteins Payne models. They've targeted a family of glycoside hydrolases (GHs), common enzymes that act on cellulose, hemicellulose and chitin, the tough material in insect shells and fungi cell walls. GHs catalyze a reaction that breaks polysaccharide molecules into smaller sugars. They're important in the natural carbon cycle and in biomass conversion for ethanol production.

Payne and her colleagues characterized a group of processive and nonprocessive GH chitinases the bacterium *Serratia marcescens* produces. Processive enzymes move along polysaccharide chains, successively cleaving chemical bonds to break sugar chains. Nonprocessive enzymes cleave a bond and detach in search of another surface, leaving new attachment sites for processive enzymes.

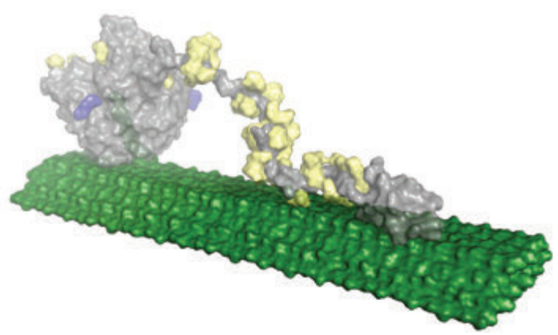

This shows a multi-modular glycoside hydrolase from the fungus *Hypocrea jecorina*, Family 6 cellobiohydrolase (Cel6A), bound to a cellulose microfibril (in green). The enzyme is decorated with glycosylation (yellow and blue), which, as findings published in the August 2013 PNAS Early Edition show, plays an essential role in substrate binding and processivity.

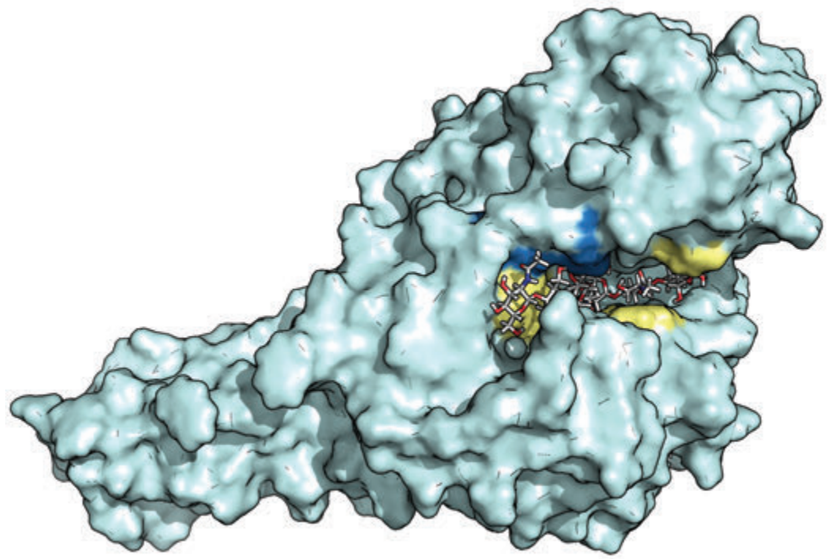

The processive glycoside hydrolase chitinase A (ChiA) from the *Serratia marcescens* bacterium is shown here bound to a chito-oligosaccharide (in gray). The aromatic residues lining the active site, characteristic of glycoside hydrolases, are shown in yellow. Polar residues, shown in dark blue, may play a role in processivity similar to the aromatic residues. The *S. marcescens* chitinolytic machinery includes two processive enzymes, ChiA and ChiB. These are responsible for the majority of hydrolytic cleavage in the degradation of biomass to soluble sugars, a process with ramifications for use in biomass-based liquid fuel production.

In a *Journal of Biological Chemistry* paper published last fall, Payne, NREL colleagues and the Norwegian group studied the structure of ChiC, a nonprocessive chitinase. Crystallography and computer simulations found three characteristics that appear to be hallmarks for determining processivity in GHs. Next, they'll study how the characteristics appear in other GHs and relate to processivity through thermodynamics.

PROBING BIOMASS BINDING

Payne is part of a team that last fall earned a DOE INCITE (Innovative and Novel Computational Impact on Theory and Experiment) award of 70 million processor hours on Argonne National Laboratory supercomputers to study the thermodynamics of binding biomass to cellulases for renewable fuel. Michael Crowley and Lintao Bu from NREL and Wei Jiang from Argonne are project investigators. "Ultimately, we will have an understanding of how these proteins work with their substrates in such a detailed fashion that we would be able to predict ways" to improve their potency.

Another project studies YKL-40, a protein that appears with diseases like inflammatory bowel disease and with many cancers. Its purpose is unknown. Payne's models will calculate the energy with which the protein binds to tissues. "If you understand what it binds to most tightly you can identify a therapeutic molecule that binds with greater affinity, so you can inhibit YKL-40's action."

For Payne, who defines herself first as a chemical engineer, computation is a tool. She'd rather focus on using MD programs others develop than devise her own algorithms. As good as those models are, however, they're still limited in predicting many of the ways atoms and molecules interact.

That's changing, Payne adds. "We're getting to a point where we can use computers to do what experiments can" to advance science more rapidly and inexpensively. "Getting to that point is a long way off, but being a part of that is really exciting."

Stephen Fink
IBM Thomas J. Watson Research Center

Research Helps Plug Computer-Programmer Gap

STEPHEN FINK is kind of a go-between for computers and programmers. As a researcher in programming languages at IBM's Thomas J. Watson Research Center in Yorktown Heights, north of New York City, he tries to connect the two with a minimum of problems.

"There's a big gap between the way people design things and the kind of instructions you give a computer" so it works as desired, says Fink, a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient from 1994 to 1998. "I work on better languages to make it easier to bridge that gap."

Many of Fink's projects are open source – free to modify and use. And used they are, by academics to research their own languages, compilers (which convert programs into machine-executable instructions), run-time systems (which handle the interface between software and operating systems) and other code.

Perhaps the most popular tool Fink helped develop is Jikes Research Virtual Machine (RVM). Virtual machines are software simulations of abstract computers and provide a portable interface that can run on many different types of physical computers. They provide an environment to run managed languages, like the popular Java.

VIRTUALLY UNPRECEDENTED

Jikes RVM is a Java virtual machine, but unlike most others it's also implemented in Java, making it self-hosted: It runs on itself without a second virtual machine. Jikes RVM also has a sophisticated adaptive optimization system, Fink says. "As you load programs and run programs, the system optimizes" them to run well by gathering performance information and using it for further optimization in a constant loop.

At the time Jikes RVM was released, there was nothing comparable, Fink says. "It was a very high-quality, state-of-the-art Java virtual machine that we made available open source." Building such an infrastructure is beyond most academic institutions, and researchers jumped on it. Jikes RVM has been cited in dozens of papers, dissertations and courses.

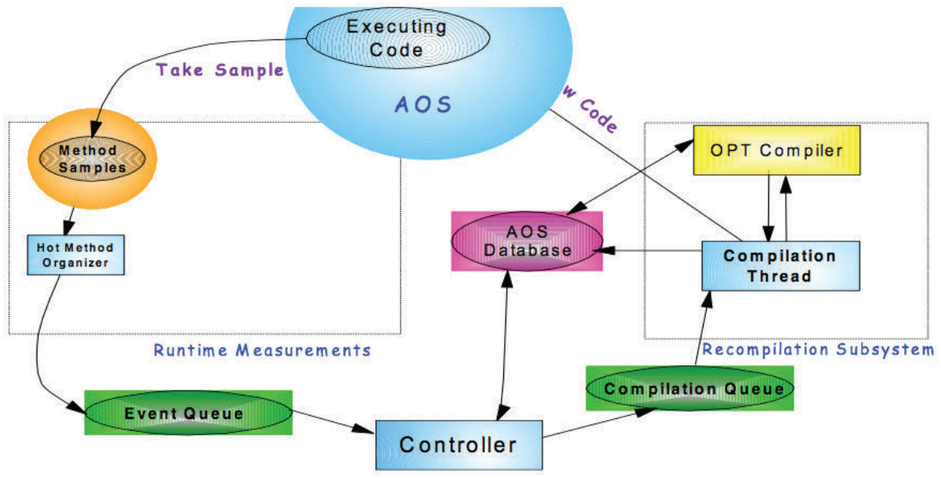

This diagram shows the architecture of the Jikes RVM adaptive optimization system. Jikes RVM is a leading infrastructure for virtual machine research.

WALA – T.J. Watson Libraries for Analysis – also has found a home in many academic research toolboxes. Unlike Jikes RVM, WALA performs static analysis, examining programs without running them. "You're given a computer program and you want to automatically analyze it to learn something about the program – some information about what the program may or may not do at run-time," Fink says. "WALA implements many algorithms to help extract information from program text to understand what the text is doing." Researchers tailor the library to search the code for specific aspects, such as pointers – variables referring to other objects or data – for analysis.

A SLICE OF LIME

Fink's latest project looks to computing's future: heterogeneous architectures, which combine standard multicore processors with other varieties, like graphics processing units (GPUs) and field-programmable gate arrays (FPGAs). The new structures offer more computing power while burning less energy, but each processor type has its own languages and rules. It's a bit like getting a mix of people speaking two or three different languages to collaborate on a big project.

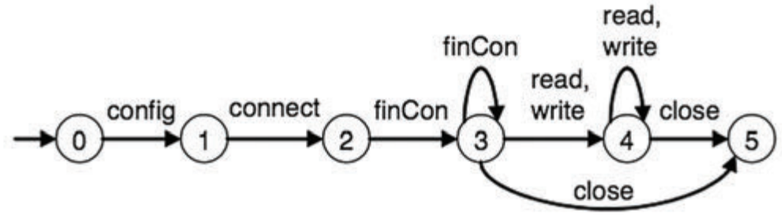

This illustration shows a typestate specification for a safety property for the Java Socket application program interface. WALA (T.J. Watson Libraries for Analysis) performs static analysis, improving the state of the art in verifying such safety properties in Java programs.

"Programming these systems is incredibly difficult," Fink says. IBM's Liquid Metal project seeks to "enable a single system so a person can learn one language and yet still be able to target these different kinds of devices." Lime, the team's new language, resembles Java but adds extensions that constrain parts of the code to run on specific types of devices. The extensions express qualities like parallelism so it's easier for the compiler to map the program down to particular kinds of devices, whether they're CPUs, GPUs or FPGAs. When it's released, Lime will include tools, a run-time system and other programming aids.

Such projects are important to maintain the growth in computer size and speed, Fink says, as physical constraints slow the decades-long acceleration in processor performance. "I think we'll see more and more exotic ideas in computer architecture start to flourish because that seems to be the most promising way forward."

It's a challenging path, but one the New Jersey native relishes. "I've always been excited by technology and the ability of technology to change people's lives," Fink adds. With programming languages, he hopes to make it "easier for people to use the technology and express what they want computers to do, leading to more and more sophisticated applications."

Krzysztof Fidkowski
University of Michigan

Former Fellow's Career Takes Flight

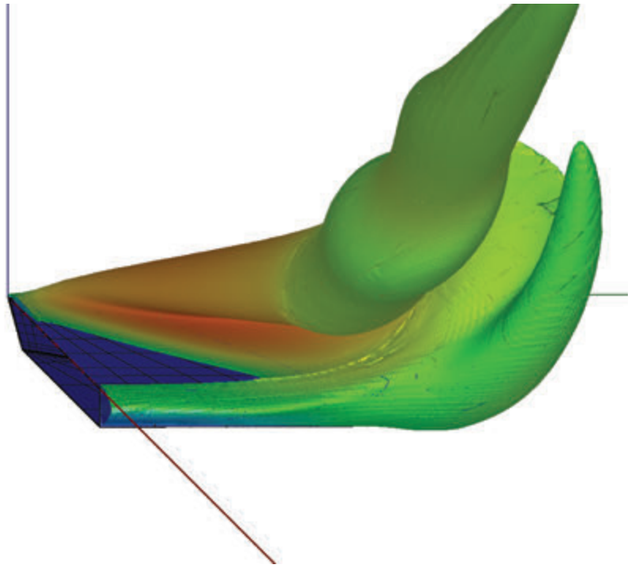

Entropy iso-surface colored by Mach number for an adaptive simulation of laminar compressible flow over a delta wing. The view is from the aft of the wing, showing the roll-up of the leading-edge vortex.

Airplane travel is different for KRZYSZTOF "CHRIS" FIDKOWSKI than for most people. As an assistant aerospace engineering professor at the University of Michigan, he uses high-performance computers to simulate airfoils and aircraft, and to seek better ways to calculate bigger fluid dynamics problems.

"I fly a lot and I know planes are safe, so I try to shelve the engineering side of things because every little thing you see outside might make you wonder," he chuckles. Yet Fidkowski, a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient from 2003 to 2007, likes watching ailerons and flaps deploy to control the aircraft. "Sometimes if you look very closely and the conditions are right, you might see a shock in the air, right over the wing," he adds.

Shocks and other flow phenomena are just what Fidkowski simulates. Yet it's impossible to capture each air molecule's movements, "so we're always making approximations to the fluid dynamics and that gives us errors" – the "plus-minus" that defines the output's accuracy. A native of Poland who dreamt of becoming an astronaut, Fidkowski focuses on cutting error by adapting the computational meshes researchers use to analyze things like airflow over wings. "We have the capability to reduce that error bar – that plus-minus – in fairly straightforward ways," usually by putting more mesh points in areas important to accurately predict an output.

This spring, Fidkowski's work on error estimation and mesh resolution earned him an Early Career Research Program award from the Department of Energy's Office of Science. The grant supports scientists in their early years, when they do their most formative work.

ADAPTATION FRAMEWORK

In a paper with graduate student Marco Ceze, Fidkowski describes an anisotropic *hp*-adaptation framework to cut output error. Like many fluid dynamics simulations, it focuses on the boundary layer – the critical, super-thin blanket of air nearest to the aircraft wing or body. Simulating airflow over the entire aircraft with an isotropic mesh small enough to resolve the boundary layer is far too demanding for even the best supercomputer. Instead, the algorithm makes boundary layer mesh elements anisotropic – flattened in one direction – with the *hp*-adaptation method flagging and refining those generating the most error.

As a rule, areas of nearly discontinuous airflow, like shocks or shears, are best represented with *h* refinement, so named because *h* denotes the length of a mesh element's side. It bisects each error-prone element to provide greater precision. In areas with smoother flows, error-causing elements are mathematically refined with higher-degree polynomial (*p*) approximations. "Increasing polynomial order is usually the best way to resolve smooth regions of the flow, where stuff happens relatively slowly," Fidkowski says.

Researchers can hand-design meshes to apply the appropriate method based on expected flow characteristics. But "we don't leave it up to the user to decide that," Fidkowski says. "We developed an algorithm that chooses the best option automatically" based on output error. Users specify a single output, such as drag. The algorithm picks the best refinement option, *h* or *p*, that most reduces drag error with the least computational effort. Sometimes the algorithm detects flow discontinuities where none were suspected and automatically bisects mesh elements. "That's not something we would have designed by hand," Fidkowski adds.

GOING NUCLEAR

Although Fidkowski focuses on aerospace, his methods could model other flows. He's part of a Michigan group participating in the Consortium for Advanced Simulation of Light Water Reactors, one of five DOE Energy Innovation Hubs. "There might be a few outputs you really care about – critical outputs for the safety of the reactor that you want to get just right. You don't want numerical errors and discretization errors polluting your results," he adds.

Regardless of the application, Fidkowski wants to design more efficient algorithms that solve bigger problems on the exotic computer architectures to come. "We want to do real-world problems," he adds. "We're just skimming the surface of that." Fidkowski also wants to make error estimation a routine part of uncertainty quantification, which puts a number on the degree to which a simulation's results can be trusted, and of optimization to improve designs.

"The most exciting part is getting an algorithm to work," he adds – something that rarely happens on the first try. It's just as important, he says, to disseminate what he learns to students. The flexibility to research and to teach thrills him.

"Every day there's not enough hours in the day," Fidkowski says. "That's when you know you're doing things you like."

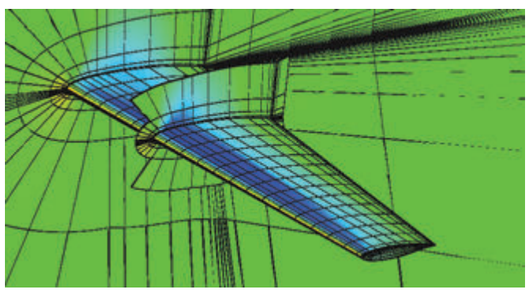

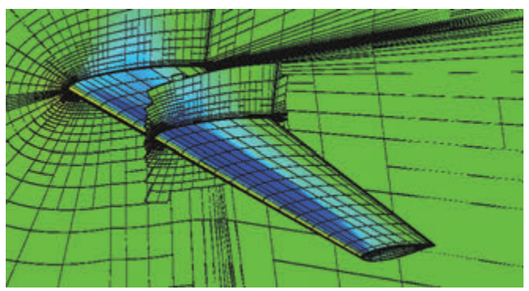

The far left image shows an initial computational mesh around the wing modeled in the American Institute of Aeronautics and Astronautics (AIAA) drag prediction workshop. This served as the input to the mesh-adaptive solver, which eventually produced the mesh shown in the near left image. The adapted mesh has significantly more elements in the leading-edge region, near the shock, and in the boundary layer near the surface of the wing. In addition, the mesh elements inside the boundary layer are efficiently adapted because they are "sliced," yielding improved resolution only along the direction perpendicular to the wing surface.

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.

HOWES AWARD

ALUMNA HONORED FOR RESEARCH AND OUTREACH

Ashlee Ford Versypt receives the 2013 Howes award from David Brown, director of Lawrence Berkeley National Laboratory's Computational Research Division.

Ashlee Ford Versypt learned she was the 2012 Frederick A. Howes Scholar in Computational Science just as she was preparing a presentation for teen-aged girls.

"The juxtaposition of those two things was very nice," says Ford Versypt, a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient from 2006 to 2010. Later that day, when the girls arrived at the Massachusetts Institute of Technology (MIT) laboratory where she is a postdoctoral researcher, Ford Versypt urged them to seize opportunities to make an impact. She talked about excellence and integrity. The Howes Award, she says, feels like an acknowledgement of those qualities in her own life.

The award recognizes a recent DOE CSGF graduate. Ford Versypt has heard her peers speak at the fellowship's annual conference, so "I knew the caliber of the graduates was very high. I was pleasantly surprised to be honored from among that cohort." She lectured on her work and received an honorarium and engraved award at the 2013 program review.

The program for teen-agers demonstrates a major reason a selection committee chose Ford Versypt. While the Howes Award recognizes outstanding technical achievements, it also honors exceptional leadership, integrity and character – qualities that reflect its namesake. Ford Versypt, who earned her doctorate in chemical engineering from the University of Illinois at Urbana-Champaign (UIUC) in 2012, "has a truly exceptional record" in service and outreach, the committee wrote.

It started, Ford Versypt says, with Girl Scouts, in which she has participated since the age of 5. Her small Oklahoma hometown had no engineering role models and "I didn't know what engineering really was." Scout camps, regional science camps and other experiences opened her eyes. "That's part of the motivation – that I had a lot of opportunities myself and I wanted to give back and help other people to have those."

TUTOR AND MENTOR

At UIUC, Ford Versypt organized a symposium on modeling and control in biomedical engineering and organized the department's graduate research symposium, the Howes Award citation says. She tutored engineering and computer science students, supervised undergraduate students doing research and prepared educational information about biomaterials. At MIT, she served on a committee planning a faculty workshop for women graduate students and postdocs. She also mentors high school, undergraduate and graduate students.

Off campus, Ford Versypt has reached out to girls and young women, working with the Society of Women Engineers, the Boys & Girls Club and, of course, Girl Scouts, beginning with her younger sister's troop back in Snyder, Okla.

Ford Versypt's research in Richard Braatz's group models the behavior of biodegradable polymers that encapsulate drugs, then slowly release them in the body. Such medications are a boon for patients, but designing them is tricky.

Ford Versypt's models portray drug-containing polymer microspheres and the chemical processes that degrade them. "What really fascinated me and motivated me was the two competing ways to get the drug out of the microsphere," she says. The first is diffusion through the polymer. The second is chemical degradation that creates pores, allowing the drugs to pass. How those processes cooperate or compete translates into different release rates for different medications.

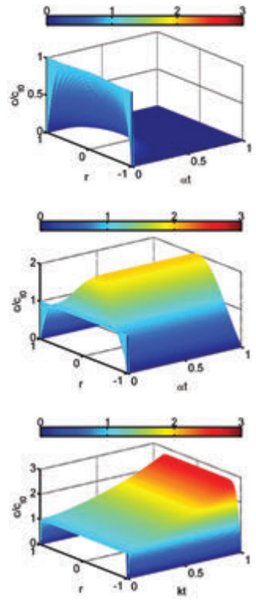

Predictions of polymer catalyst concentration inside microspheres of increasing size. Blue indicates diffusion through polymer dominates drug release, and red indicates transport through pores formed by chemical reaction dominates.

THE RIGHT COMBINATION

Until Ford Versypt's work, "no models addressed all of these important features," the selection committee noted in its citation. With her models, scientists could more quickly and inexpensively find the chemical combinations to release just the right amount of medication at the right time.

Ford Versypt also has worked with the Center for Continuous Manufacturing, a joint project between MIT and drug manufacturer Novartis. The partnership aims to bring a less expensive "assembly line" approach to the drug industry, which traditionally has produced medicines in batches. To help with design and quality control, Ford Versypt developed a computational model of the time-varying process involved in drying a chemical film spread on a conveyor belt. The system monitors the process and compensates for changes.

Computational science's range excites her, Ford Versypt says. As she told the recent young lab visitors, many fields play a role in her research, from mathematics to chemistry, and she can affect many of them. "I like that there are lots of tools from different disciplines that I can apply to these problems, and I look at them from an interdisciplinary perspective," she says.

But Ford Versypt's bigger message that day was about embracing change and pursuing opportunities for excellence with integrity. "This idea of not being just a scholar, but a scholar-citizen, has really been an important principle for me."

ABOUT FRED HOWES

In the 12 years since it was first conferred, the Frederick A. Howes Scholar in Computational Science award has become emblematic of research excellence and outstanding leadership. It's a fitting tribute to Howes, who was known for his scholarship, intelligence and humor.

Howes earned his bachelor's and doctoral degrees in mathematics at the University of Southern California. He held teaching posts at the universities of Wisconsin and Minnesota before joining the faculty of the University of California, Davis, in 1979. Ten years later Howes served a two-year rotation with the National Science Foundation's Division of Mathematical Sciences. He joined DOE in 1991.

In 2000, colleagues formed an informal committee to honor Howes. They chose the DOE CSGF as the vehicle and gathered donations, including a generous contribution from Howes' family, to endow an award in his name.

PAST HOWES SCHOLARS

2012	Carolyn Phillips and Matthew Reuter
2011	Alejandro Rodriguez
2010	Julianne Chung
2009	David Potere
2008	Mala Radhakrishnan
2007	Jaydeep Bardhan and Kristen Grauman
2006	Kevin Chu and Matthew Wolinsky
2005	Ryan Elliott and Judith Hill
2004	Collin Wick
2003	Oliver Fringer and Jon Wilkening
2001	Jeffrey Hittinger and Mayya Tokman

ALUMNI: WHERE ARE THEY NOW?

With the addition of 16 new graduates in 2013, alumni of the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) now number more than 300. A partnership between the DOE and National Nuclear Security Administration that began in 1991, the DOE CSGF has successfully generated top scientists for the national laboratory system. It has also seeded industry and academia with graduates who lead the way in employing computing to maintain and regain the nation's edge in discovery, generating jobs and income.

During the most recent 10-year period compiled, most DOE CSGF alumni have secured employment in highly specialized positions as well as in leadership positions in U.S. industry and university research (see chart), creating a multiplier effect that spreads the influence of the program well beyond the government agencies that sponsor the fellowship.

A complete listing of alumni (by last name, Ph.D. institution, fellowship start year, practicum location, current location and area of study) can be found at: **www.krellinst.org/csgf**

DOE CSGF ALUMNI: EMPLOYMENT BY CATEGORY

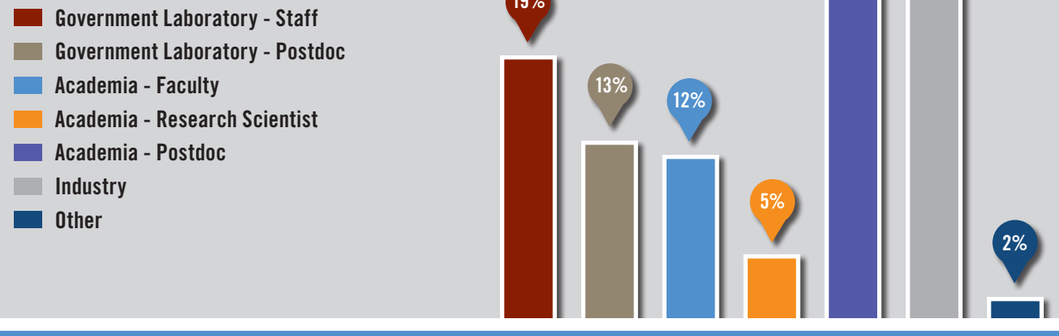

CLASS OF 2013

Edward Baskerville
University of Michigan
Ecology and Scientific Computing
Advisor: Mercedes Pascual
Practicum: Lawrence Berkeley National Laboratory
Contact: ed@edbaskerville.com

Sanjeeb Bose
Stanford University
Computational Fluid Dynamics
Advisor: Parviz Moin
Practicum: Lawrence Livermore National Laboratory
Contact: stbose@stanford.edu

Kurt Brorsen
Iowa State University
Physical Chemistry
Advisor: Mark Gordon
Practicum: Argonne National Laboratory
Contact: kurtbrorsen@gmail.com

Leslie Dewan
Massachusetts Institute of Technology
Nuclear Waste Materials
Advisor: Linn Hobbs
Practicum: Lawrence Berkeley National Laboratory
Contact: ldewan@mit.edu

Jeffrey Donatelli
University of California, Berkeley
Applied Mathematics
Advisor: James Sethian
Practicum: Oak Ridge National Laboratory
Contact: jdonatel@math.berkeley.edu

Virgil Griffith
California Institute of Technology
Theoretical Neuroscience
Advisor: Christof Koch
Practicum: Lawrence Berkeley National Laboratory
Contact: virgil@caltech.edu

Tobin Isaac
University of Texas
Computational and Applied Mathematics
Advisor: Omar Ghattas
Practicum: Los Alamos National Laboratory
Contact: tisaac@ices.utexas.edu

Mark Maienschein-Cline
University of Chicago
Physical Chemistry
Advisor: Aaron Dinner
Practicum: Los Alamos National Laboratory
Contact: mmaiensc@uchicago.edu

Amanda Peters Randles
Harvard University
Applied Physics
Advisor: Efthimios Kaxiras
Practicum: Lawrence Livermore National Laboratory
Contact: apeters@fas.harvard.edu

Noah Reddell
University of Washington
Computational Plasma Modeling for Fusion Energy
Advisor: Uri Shumlak
Practicum: Princeton Plasma Physics Laboratory
Contact: noah.reddell@gmail.com

Troy Ruths
Rice University
Bioinformatics
Advisor: Luay Nakhleh
Practicum: Lawrence Berkeley National Laboratory
Contact: troy.ruths@rice.edu

Samuel Skillman
University of Colorado at Boulder
Astrophysics
Advisor: Jack Burns
Practicum: Oak Ridge National Laboratory
Contact: samuel.skillman@colorado.edu

Hayes Stripling
Texas A&M University
Nuclear Engineering/ Uncertainty Quantification
Advisor: Marvin Adams
Practicum: Lawrence Livermore National Laboratory
Contact: h.stripling@tamu.edu

Travis Trahan
University of Michigan
Nuclear Engineering
Advisor: Edward Larsen
Practicum: Argonne National Laboratory
Contact: tjtrahan@umich.edu

Sean Vitousek
Stanford University
Environmental Fluid Mechanics and Hydrology
Advisor: Oliver Fringer
Practicum: Lawrence Berkeley National Laboratory
Contact: seanv@stanford.edu

Norman Yao
Harvard University
Condensed Matter Physics
Advisor: Mikhail Lukin
Practicum: Los Alamos National Laboratory
Contact: nyao@fas.harvard.edu

4TH YEAR FELLOWS

Front, left to right: Miles Lopes, Zachary Ulissi, Edgar Solomonik, Mary Benage, Kenley Pelzer, Leslie Dewan*, Thomas Fai and Christopher Quinn; **Middle, left to right:** Devin Matthews, Charles Frogner, Irene Kaplow, Aleah Caulin, Christopher Ivey and Evan Gawlik; **Back, left to right:** Amanda Peters Randles*, Carmeline Dsilva, Aaron Sisto, Christopher Eldred, Seth Davidovits, Scot Miller and Peter Maginot

*Finished with Class of 2013

Mary Benage
Georgia Institute of Technology
Geophysics
Advisor: Josef Dufek
Practicum: Lawrence Livermore National Laboratory
Contact: mary.benage@eas.gatech.edu

Aleah Caulin
University of Pennsylvania
Genomics and Computational Biology
Advisor: Carlo Maley
Practicum: Lawrence Berkeley National Laboratory
Contact: alefox@mail.med.upenn.edu

Seth Davidovits
Princeton University
Plasma Physics
Advisor: Greg Hammett
Practicum: Lawrence Berkeley National Laboratory
Contact: sdavidov@princeton.edu

Carmeline Dsilva
Princeton University
Chemical Engineering
Advisor: Ioannis Kevrekidis
Practicum: Argonne National Laboratory
Contact: cdsilva@princeton.edu

Christopher Eldred
Colorado State University
Climate Modeling
Advisor: David Randall
Practicum: Lawrence Berkeley National Laboratory
Contact: chris.eldred@gmail.com

Thomas Fai
New York University
Applied Mathematics
Advisor: Charles Peskin
Practicum: Lawrence Berkeley National Laboratory
Contact: tfai@cims.nyu.edu

Charles Frogner
Massachusetts Institute of Technology
Computational Biology
Advisor: Tomaso Poggio
Practicum: Lawrence Berkeley National Laboratory
Contact: frogner@mit.edu

Evan Gawlik
Stanford University
Applied Mathematics
Advisor: Adrian Lew
Practicum: Argonne National Laboratory
Contact: egawlik@stanford.edu

Christopher Ivey
Stanford University
Computational Fluid Dynamics
Advisor: Parviz Moin
Practicum: Argonne National Laboratory
Contact: civey@stanford.edu

Irene Kaplow
Stanford University
Computational Biology
Advisor: Hunter Fraser
Practicum: Lawrence Berkeley National Laboratory
Contact: ikaplow@stanford.edu

Miles Lopes
University of California, Berkeley
Machine Learning
Advisor: Peter Bickel
Practicum: Sandia National Laboratories, California
Contact: mlopes@stat.berkeley.edu

Peter Maginot
Texas A&M University
Nuclear Engineering
Advisor: Jim Morel
Practicum: Knolls Atomic Power Laboratory
Contact: pmaginot@neo.tamu.edu

Devin Matthews
University of Texas
Chemistry
Advisor: John Stanton
Practicum: Argonne National Laboratory
Contact: dmatthews@utexas.edu

Scot Miller
Harvard University
Atmospheric Sciences
Advisor: Steven Wofsy
Practicum: Lawrence Berkeley National Laboratory
Contact: scot.m.miller@gmail.com

Kenley Pelzer
University of Chicago
Theoretical Physical Chemistry
Advisor: Greg Engel
Practicum: Argonne National Laboratory
Contact: kpelzer@uchicago.edu

Christopher Quinn
University of Illinois at Urbana-Champaign
Communications
Advisor: Negar Kiyavash
Practicum: Sandia National Laboratories, California
Contact: quinn7@illinois.edu

Aaron Sisto
Stanford University
Computational Chemistry
Advisor: Todd Martinez
Practicum: Sandia National Laboratories, California
Contact: asisto@stanford.edu

Edgar Solomonik
University of California, Berkeley
Computer Science
Advisor: James Demmel
Practicum: Argonne National Laboratory
Contact: solomonik@berkeley.edu

Zachary Ulissi
Massachusetts Institute of Technology
Chemical Engineering
Advisor: Michael Strano
Practicum: Argonne National Laboratory
Contact: zulissi@gmail.com

3RD YEAR FELLOWS

Front, left to right: Alexander Rattner, Heather Mayes, Michael Rosario, Maxwell Hutchinson and Rogelio Cardona-Rivera; **Middle, left to right:** Aurora Pribram-Jones, Curtis Lee, Hansi Singh, Jason Bender, Robert Parrish, Matthew Zahr and Chris Smillie; **Back, left to right:** Phoebe Robinson, Sarah Loos, Jarrod McClean, Joshua Vermaas, Daniel Dandurand* and Omar Hafez

*Withdrew in 2013

Jason Bender
University of Minnesota
Hypersonic Computational Fluid Dynamics
Advisor: Graham Candler
Practicum: Argonne National Laboratory
Contact: jbender73@gmail.com

Rogelio Cardona-Rivera
North Carolina State University
Artificial Intelligence
Advisor: R. Michael Young
Practicum: Sandia National Laboratories, New Mexico
Contact: recardon@ncsu.edu

Omar Hafez
University of California, Davis
Computational Solid Mechanics
Advisor: Mark Rashid
Practicum: Lawrence Livermore National Laboratory
Contact: omhafez@ucdavis.edu

Maxwell Hutchinson
University of Chicago
Physics
Advisor: Robert Rosner
Practicum: Lawrence Berkeley National Laboratory
Contact: maxhutch@gmail.com

Curtis Lee
Duke University
Computational Mechanics
Advisor: John Dolbow
Practicum: Lawrence Berkeley National Laboratory
Contact: calee181@gmail.com

Sarah Loos
Carnegie Mellon University
Verification of Hybrid Systems
Advisor: Andre Platzer
Practicum: Oak Ridge National Laboratory
Contact: sloos@cs.cmu.edu

Heather Mayes
Northwestern University
Chemical Engineering
Advisor: Linda Broadbelt
Practicum: National Renewable Energy Laboratory
Contact: hmayes@u.northwestern.edu

Jarrod McClean
Harvard University
Chemical Physics
Advisor: Alan Aspuru-Guzik
Practicum: Los Alamos National Laboratory
Contact: jmcclean@fas.harvard.edu

Robert Parrish
Georgia Institute of Technology
Theoretical Chemistry
Advisor: David Sherrill
Practicum: Lawrence Livermore National Laboratory
Contact: robparrish@gatech.edu

Aurora Pribram-Jones
University of California, Irvine
Theoretical Chemistry
Advisor: Kieron Burke
Practicum: Sandia National Laboratories, New Mexico
Contact: apribram@uci.edu

Alexander Rattner
Georgia Institute of Technology
Mechanical Engineering
Advisor: Srinivas Garimella
Practicum: Idaho National Laboratory
Contact: Alex.Rattner@gatech.edu

Phoebe Robinson
Harvard University
Earth Science
Advisor: Brendan Meade
Practicum: Lawrence Berkeley National Laboratory
Contact: phoebemaherrobinson@gmail.com

Michael Rosario
Duke University
Organismic and Evolutionary Biology
Advisor: Sheila Patek
Practicum: Sandia National Laboratories, California
Contact: mrosario@bio.umass.edu

Hansi Singh
University of Washington
Atmosphere-Ocean Physics
Advisor: Cecilia Bitz
Contact: hansi@atmos.washington.edu

Chris Smillie
Massachusetts Institute of Technology
Biology, Computer Science and Bioengineering
Advisor: Eric Alm
Practicum: Lawrence Berkeley National Laboratory
Contact: csmillie@mit.edu

Joshua Vermaas
University of Illinois at Urbana-Champaign
Biophysics
Advisor: Emad Tajkhorshid
Practicum: National Renewable Energy Laboratory
Contact: vermaas2@illinois.edu

Matthew Zahr
Stanford University
Computational and Mathematical Engineering
Advisor: Charbel Farhat
Practicum: Lawrence Berkeley National Laboratory
Contact: mzahr@stanford.edu

2ND YEAR FELLOWS

Front, left to right: Miles Lubin, Sherwood Richers, Brian Powell, Dragos Velicanu, Victor Minden and Andrew Till;
Middle, left to right: Melissa Yeung, Jamie Smedsmo, Daniel Strouse, Andrew Stine, Brenhin Keller, Sarah Middleton, Justin Lee and Britni Crocker; **Back, left to right:** Eileen Martin, Eric Isaacs, Thomas Catanach, Samuel Blau, Derek Macklin, Andrew Stershic and Jesse Lopez

Samuel Blau
Harvard University
Chemical Physics
Advisor: Alan Aspuru-Guzik
Contact: sblau@fas.harvard.edu

Thomas Catanach
California Institute of Technology
Applied and Computational Mathematics
Advisor: Jim Beck
Contact: picatanach@gmail.com

Britni Crocker
Massachusetts Institute of Technology
Computational Neuroscience
Advisor: Richard Cohen
Practicum: Los Alamos National Laboratory
Contact: intirb@hotmail.com

Eric Isaacs
Columbia University
Applied Physics
Advisor: Chris Marianetti
Practicum: Brookhaven National Laboratory
Contact: ebi2104@columbia.edu

Brenhin Keller
Princeton University
Geochemistry and Geochronology
Advisor: Blair Schoene
Contact: cbkeller@princeton.edu

Justin Lee
Massachusetts Institute of Technology
Computational Imaging/Biomedical Optics
Advisor: George Barbastathis
Practicum: Lawrence Berkeley National Laboratory
Contact: jlee08@gmail.com

Jesse Lopez
Oregon Health and Science University
Environmental Science and Engineering
Advisor: Antonio Baptista
Practicum: Argonne National Laboratory
Contact: lopezj@stccmop.org

Miles Lubin
Massachusetts Institute of Technology
Operations Research
Advisor: Juan Pablo Vielma
Contact: miles.lubin@gmail.com

Derek Macklin
Stanford University
Computational and Systems Biology
Advisor: Markus Covert
Contact: derek.krellinst.org@nrm.com

Eileen Martin
Stanford University
Computational and Mathematical Engineering
Advisor: Margot Gerritsen
Contact: ermartin@stanford.edu

Sarah Middleton
University of Pennsylvania
Genomics and Computational Biology
Advisor: Junhyong Kim
Contact: sarahmid@mail.med.upenn.edu

Victor Minden
Stanford University
Computational and Mathematical Engineering
Advisor: Margot Gerritsen
Contact: victorminden@gmail.com

Brian Powell
North Carolina State University
Nuclear Engineering
Advisor: Yousry Azmy
Practicum: Oak Ridge National Laboratory
Contact: bppowell@ncsu.edu

Sherwood Richers
California Institute of Technology
Astrophysics
Advisor: Christian Ott
Practicum: Lawrence Berkeley National Laboratory
Contact: srichers@caltech.edu

Jamie Smedsmo
University of North Carolina
Environmental Modeling
Advisor: Marc Serre
Practicum: Oak Ridge National Laboratory
Contact: jsmedsmo@live.unc.edu

Andrew Stershic
Duke University
Civil Engineering/Computational Mechanics
Advisor: John Dolbow
Practicum: Oak Ridge National Laboratory
Contact: ajs84@duke.edu

Andrew Stine
Northwestern University
Chemical and Biological Engineering
Advisor: Linda Broadbelt
Contact: andrewstine2015@u.northwestern.edu

Daniel Strouse
Princeton University
Theoretical Neuroscience
Advisor: William Bialek
Contact: danieljstrouse@gmail.com

Andrew Till
Texas A&M University
Multiphysics Scientific Computational Nuclear Engineering
Advisor: Marvin Adams
Practicum: Los Alamos National Laboratory
Contact: attom@tamu.edu

Dragos Velicanu
Massachusetts Institute of Technology
High Energy Physics
Advisor: Gunther Roland
Contact: velicanu@mit.edu

Melissa Yeung
California Institute of Technology
Mathematics
Advisor: Mathieu Desbrun
Practicum: Lawrence Berkeley National Laboratory
Contact: myeung@caltech.edu

1ST YEAR FELLOWS

Front, left to right: Alexander Turner, Nicholas Frontiere, Adam Richie-Halford and Daniel Rey;
Back, left to right: Isha Jain, David Plotkin, David Ozog, Chelsea Axen and Kathleen Alexander;
Not pictured: Will Fletcher

Kathleen Alexander
Massachusetts Institute of Technology
Microstructure Modeling
Advisor: Chris Schuh
Contact: katcalex@mit.edu

Chelsea Axen
University of California, Berkeley
Astrophysics
Advisor: Peter Nugent
Contact: c.axen@berkeley.edu

Will Fletcher
Stanford University
Biophysics
Advisor: Malkiat Johal
Contact: will.r.fletcher@gmail.com

Nicholas Frontiere
University of Chicago
Physics
Advisor: David Reid
Contact: nfrontiere@gmail.com

Isha Jain
Harvard University/Massachusetts Institute of Technology
Computer Science and Systems Biology
Advisor: Vamsi Mootha
Contact: ijain@mit.edu

David Ozog
University of Oregon
Computational Science
Advisor: Allen Malony
Contact: ozog@cs.uoregon.edu

David Plotkin
University of Chicago
Earth Sciences
Advisor: Dorian Abbot
Contact: dplotkin@uchicago.edu

Daniel Rey
University of California, San Diego
Biophysics
Advisor: Henry Abarbanel
Contact: nadrey@gmail.com

Adam Richie-Halford
University of Washington
Physics
Advisor: Aurel Bulgac
Contact: richiehalford@gmail.com

Alexander Turner
Harvard University
Atmospheric Science
Advisor: Daniel Jacob
Contact: aturner@fas.harvard.edu

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

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