THE DOE CSGF ANNUAL

2016

DEIXIS

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

ROCK STAR

PAGE 13

Fellow Brenhin Keller computes conditions of Earth's early geology and hikes the heights for hard evidence.

MORE FELLOWS' RESEARCH

- Pushing cosmic boundaries
- Probing a river's metabolism
- Cracking a fracturing problem

Alumna Bree Aldridge probes a tuberculosis puzzle, dual Howes Award winners, computing on the brain – and cancer – with Argonne's Rick Stevens, and an atomic boogie.

DEPARTMENT OF ENERGY

Computational Science Graduate Fellowship

DOE CSGF

The Department of Energy Computational Science Graduate Fellowship (DOE CSGF) program provides outstanding benefits and opportunities to students pursuing doctoral degrees in fields of study that utilize high-performance computing to solve complex problems in science and engineering.

APPLY ONLINE

www.krellinst.org/csgf

> $36,000 yearly stipend
> Payment of full tuition and required fees
> Attend yearly program review
> $5,000 academic allowance in first year
> 12-week research practicum
> $1,000 academic allowance each renewed year
> Renewable up to four years

Image courtesy of Oak Ridge National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

National Nuclear Security Administration

KRELL institute

This equal opportunity program is open to all qualified persons without regard to race, gender, religion, age, physical disability or national origin.

WELCOME TO THE 2016 INCOMING DOE CSGF CLASS

In fall 2016, 27 doctoral students – the biggest group yet – enter the Department of Energy Computational Science Graduate Fellowship (DOE CSGF). These fellows were chosen in a rigorous screening process from more than 350 applicants. Each will receive yearly stipends, full tuition and fees and other benefits for up to four years, freeing them to pursue research and training in applied mathematics, computer science and an application area of their choice.

Riley Brady
University of Colorado
Atmospheric and Oceanic Dynamics

Brian Cornille
University of Wisconsin-Madison
Nuclear Engineering & Engineering Physics

Emily Crabb
Massachusetts Institute of Technology
Condensed Matter Theory

Julia Ebert
Harvard University
Bioengineering

Sarah Elliott
University of Georgia
Theoretical Chemistry

Jenelle Feather
Massachusetts Institute of Technology
Neuroscience

Sarah Gady
Princeton University
Mechanical and Aerospace Engineering

Claire-Alice Hebert
Stanford University
Applied Physics

Daniel Jacobson
California Institute of Technology
Computational Chemistry

Harshil Kamdar
Harvard University
Computational Astrophysics

Kelly Kochanski
University of Colorado
Earth Surface Processes

Brett Larsen
Stanford University
Physics

Yuexia Lin
Harvard University
Applied Mathematics

Thomas Ludwig
Stanford University
Chemical Engineering

Sean Marks
University of Pennsylvania
Chemical Engineering

Kayla McCue
Massachusetts Institute of Technology
Computational and Systems Biology

Kelly Moran
Duke University
Statistics

Ian Ochs
Princeton University
Plasma Physics

Mario Ortega
University of California, Berkeley
Nuclear Engineering

Nicholas Rivera
Massachusetts Institute of Technology
Optical Physics

Amaresh Sahu
University of California, Berkeley
Chemical Engineering

Andrés Salcedo
Ohio State University
Astronomy

Clay Sanders
Duke University
Civil Engineering/Computational Mechanics

Sukin Sim
Harvard University
Chemical Physics

Laura Watkins
University of Chicago
Computational Chemistry

Blake Wetherton
University of Wisconsin-Madison
Plasma Physics

Cristina White
Stanford University
Mechanical Engineering

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 grant DE-FG02-97ER25308.

For additional information about the DOE CSGF program, the Krell Institute or topics covered in this publication, please go to: www.krellinst.org/csgf

Or contact:
Editor, DEIXIS
Krell Institute
1609 Golden Aspen Drive, Suite 101
Ames, IA 50010
(515) 956-3696

Copyright 2016 Krell Institute. All rights reserved.

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.

DOE CSGF funding is provided by the DOE Office of Advanced Scientific Computing Research (ASCR) within the Office of Science and the Advanced Simulation and Computing (ASC) program within the National Nuclear Security Administration.

Manager, Science Media — Bill Cannon
Science Media Editor — Thomas R. O'Donnell
Creative Project Coordinator — Buffy Clatt
Design — Stilt Studio, Inc.

ON THE COVER: Fellow Brenhin Keller collects a sample for analysis from the Bergell Pluton, a large rock formation, above Lago da l'Albigna (Albigna Lake) in Switzerland's central Alps. Read about Keller's research starting on page 13. *Credit: Kyle Samperton.*

IN THIS ISSUE

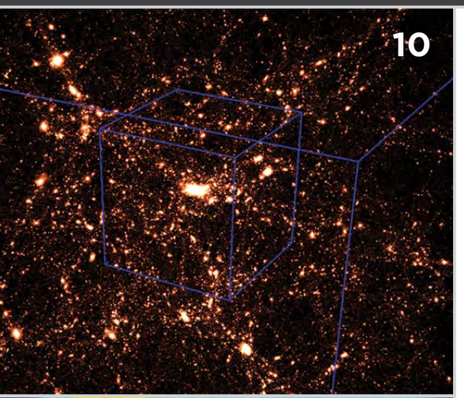

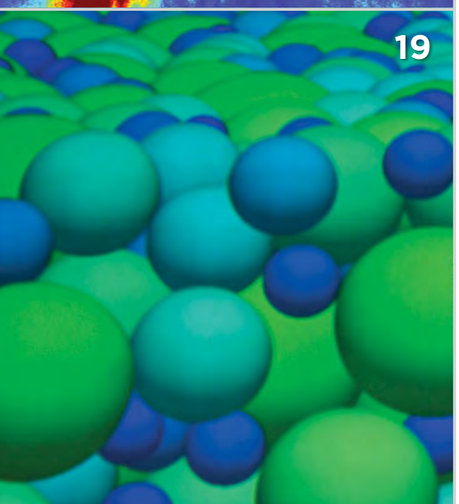

A QUARTER-CENTURY OF COMPUTATIONAL SCIENCE

The Department of Energy Computational Science Graduate Fellowship (DOE CSGF) could claim to be the internet's littermate.

Like the worldwide computing network, the fellowship arose from the High-Performance Computing and Communications Act of 1991, which directed DOE to "support basic research, education and human resources in computational science." That first year DOE CSGF enrolled 22 fellows. The program marks its silver anniversary in fall 2016.

Since that inaugural class, the fellowship has built a community of scientists dedicated to applying high-performance computing (HPC) to knotty problems. Its more than 350 alumni work at DOE national laboratories, in academia or in industry.

At the fellowship's program review each summer, one or two of these graduates receive the Frederick A. Howes Scholar in Computational Science award, recognizing their exemplary research excellence and leadership. This year's honorees are Aurora Pribram-Jones, a postdoctoral fellow at Lawrence Livermore National Laboratory and the University of California, Berkeley, and Alexander Rattner, a Pennsylvania State University assistant professor of mechanical and nuclear engineering. Learn about their contrasting, yet similar, research approaches on **page 6**.

We also talk with Argonne National Laboratory's Rick Stevens **(8)**, one of the 2016 program review's keynote speakers, about new programs focusing high-performance computing on cancer and brain research.

The program review is a showcase for fellows finishing the program, as each delivers a talk summarizing their research. We feature three graduating fellows here, plus a third-year student.

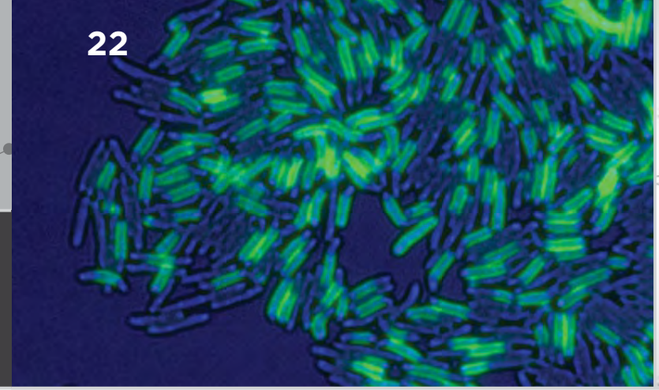

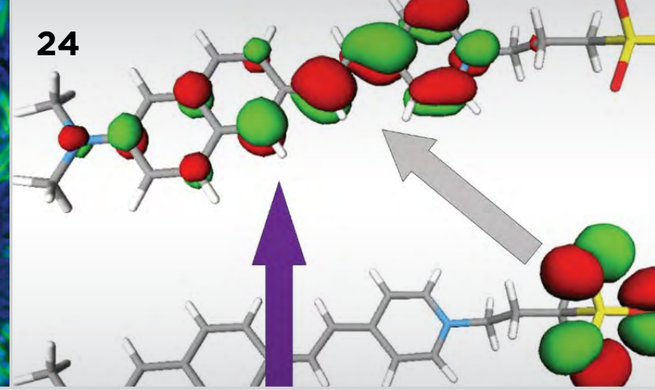

Nicholas Frontiere turns computer power toward the stars. The University of Chicago third-year fellow helps devise massive simulations of the universe's evolution. In "Cosmic Encounter" **(10)**, Science Media Editor Thomas R. O'Donnell tells how Frontiere's research could improve such models.

Brenhin Keller, at Cornell University, uses HPC to probe chemistry data, seeking answers to fundamental questions about Earth's early history. Keller also has sought rock samples on spectacular European mountainsides. Science writer Sarah Webb, who trained as a chemist, provides details in "Going Deep" **(13)**.

Meanwhile, Jesse López's science flows from the Columbia River. López, an Oregon Health & Science University doctoral student, uses computational models to analyze bioreactors – estuary regions where microbes digest organic material, enriching the river and supporting aquatic life. As O'Donnell notes in "Ebb and Flow" **(16)**, López's results could help preserve the Columbia's health.

Andrew Stershic's research studies materials' details: how fractures and cracks begin at the atomic scale. The Duke University fellow combines mathematical methods to better simulate shattering, as longtime DEIXIS contributor Karyn Hede explains in "Fracture Tracker" **(19)**.

These fellows' range demonstrates HPC's broad applications. Alumni featured in this issue work in similarly diverse areas.

Bree Aldridge, a fellow from 2002 to 2006, began her career simulating cancer-cell signaling. Now at Tufts University, she studies how tuberculosis bacteria disarm immune cells. Freelance science writer Andy Boyles, a former *Highlights* magazine science editor who has written extensively about biomedical research, tells the story in "The Single-cell Solution" **(22)**.

Since joining Lawrence Berkeley National Laboratory, alumnus Anubhav Jain has helped use HPC to predict the properties of thousands of compounds, creating a database for researchers developing useful substances. With support from a DOE Office of Science Early Career Research Program grant, he now probes the database himself. As O'Donnell tells it in "Materials Miner" **(24)**, Jain seeks substances that produce electricity as they warm up.

Alumnus Jeff Hammond, in contrast, is more interested in making HPC calculations run better. At Intel Corp.'s Parallel Computing Laboratory, he explores solutions where applications and hardware intersect, and he looks ahead to increasingly complex computer architectures, O'Donnell reports in "The Non-job Job" **(26)**.

Whether developing algorithms or applying them to problems, fellows and alumni must be able to explain their work, especially to those whose taxes support it. The annual Communicate Your Science & Engineering contest promotes those skills by recognizing a current or former fellow whose essay best describes their work to a lay audience. This year's winner, Eric Isaacs, waltzes readers through the microscopic world of phonons in "Atoms on the Dance Floor" **(28)**.

As Isaacs and other fourth-year fellows leave the program, 27 new students enter in fall 2016 as the 26th class.

Their talents are in demand – a demand the DOE CSGF addresses, Secretary of Energy Ernest Moniz says. "The needs are dire in terms of providing the people power for our future high-performance computing and other IT and cyberactivities." With exascale computers – about a thousand times faster than today's best – and other new technology on the horizon, "the workforce needs to be greatly expanded," he says. "We think this is a critical thing."

Happy 25th birthday, DOE CSGF.

MASTERS OF MATTER AND ENERGY

Pribram-Jones and Rattner are Joint 2016 Howes Award Winners

ABOUT FRED HOWES

Since it was first awarded in 2001, the Frederick A. Howes Scholar in Computational Science award has come to stand for 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 and advocated for the fellowship and for computational science as manager of the Applied Mathematical Sciences program.

Howes died unexpectedly in 1999 at age 51. Colleagues formed an informal committee to honor him and chose the DOE CSGF as the vehicle. With donations, including a generous contribution from Howes' family, they endowed an award in his name.

By Thomas R. O'Donnell

On the surface, it seems the recipients of the 2016 Frederick A. Howes Scholar in Computational Science award take their research in divergent directions.

Aurora Pribram-Jones, a Lawrence postdoctoral fellow at Lawrence Livermore National Laboratory, largely focuses on refining mathematical techniques to understand matter's fundamental properties. Alexander Rattner, an assistant professor of mechanical and nuclear engineering at Pennsylvania State University, creates simulations and experiments directed at improving energy systems like power plants.

But classifying them isn't that easy. Pribram-Jones, the theorist, is preparing to apply her work to materials. Rattner, the engineer, hopes researchers and energy-system designers adopt his group's computational modeling method.

Both are excellent researchers and leaders – qualities recognized by the Howes award, given to exceptional alumni of the Department of Energy Computational Science Graduate Fellowship (DOE CSGF).

Pribram-Jones and Rattner will receive honorariums and awards in July at the fellowship's annual program review, where they'll also lecture.

The 2015 graduates say they're humbled to be chosen. "The caliber of the work done by other (Howes) recipients is really impressive," Rattner says. It's "overwhelming to be put in that category."

The honor is significant, Pribram-Jones says, because she sometimes felt like an outsider among her DOE CSGF peers. "Not because I've ever been treated that way but because my research was different" from most doctoral projects. "It means a lot to be recognized by a group that you weren't ever sure you fit in with."

Pribram-Jones studies density functional theory (DFT), a method to calculate electron interactions in molecules and predict a material's properties. Pribram-Jones usually sticks to theory, seeking ways to extend DFT to new problems and improve its accuracy. Her latest work, however, could see her apply the method to real materials.

Besides her Lawrence fellowship, Pribram-Jones also has a Presidential Postdoctoral Fellowship at the University of California, Berkeley. She's using the dual appointment to develop a collaboration between her Berkeley mentor, Martin Head-Gordon, and Livermore scientists (led by DOE CSGF alumnus Brandon Wood) to investigate metal organic frameworks – microporous molecular structures – for hydrogen storage.

Head-Gordon's group excels at DFT quantum chemistry calculations, Pribram-Jones says, complementing the lab's high-performance computing and materials research. The challenge is connecting DFT's small-scale accuracy to the energetics of how the materials store and release hydrogen. "It's a big multiscale problem, so it's an ideal one to bridge" the research groups.

The problems Rattner tackles also span time and space. His Multiscale Thermal Fluids and Energy Lab studies two-phase flows, like those found in power plants, refrigeration cycles and many industrial processes.

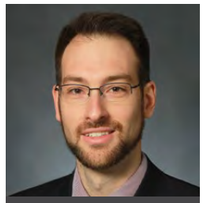

Alexander Rattner

"These processes are highly multiscale, meaning that you've got big-picture stuff going on" up to the size of power-plant steam generators, "but most of the heat transfer occurs in micron-thick regions underneath vapor bubbles." He also investigates multicomponent flows – those containing more than one chemical – found in petrochemical processing or refrigeration. High-performance computing is invaluable for handling the associated modeling challenges.

Rattner and his colleagues conduct experiments to validate their models, but he's most proud of a recently released open-source solver for two-phase flows with phase change. He hopes other researchers adopt and extend it.

Rattner's service includes organizing sessions on heat transfer for energy systems at American Society of Mechanical Engineers meetings. He's also helped the Lunar Lions, Penn State students who are developing a space mission.

But Rattner believes his biggest service may be as an educator. His courses emphasize engaging, hands-on experiments and activities and he wants to expose undergraduates to topics their core courses usually don't cover – such as computing.

"What I found to be some of the most important parts of my undergraduate education were with either engineering club activities outside the classroom or classes that gave the opportunity to develop new projects and get hands-on experiences. I've been trying to recreate some of those opportunities here," Rattner says.

'Helping scientists self-evaluate and self-reflect is a pretty powerful thing.'

Pribram-Jones also aspires to teaching and mentoring, building on a record stretching to her undergraduate days at Harvey Mudd College. For example, she's assisting Harvey Mudd students on a project to compute the properties of specialized alloys.

She's also helped incoming graduate students at the University of California, Irvine, write narratives defining their goals and values, often to submit as fellowship application essays. "Helping scientists self-evaluate and self-reflect is a pretty powerful thing," especially early in their careers. That has "a greater impact in a different way than any of my other outreach work."

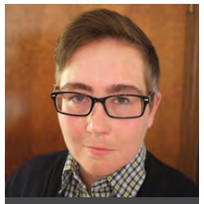

Aurora Pribram-Jones

Pribram-Jones' mentorship instinct partly arises from her difficult background. She left school as a teenager to help her family cope with poverty, mental illness and addiction that led to her father's death. Understanding professors helped her overcome her circumstances and later health problems to finish college. A national laboratory is probably the best place to continue her research, but "as soon as I'm teaching or mentoring again I know that's what I'm supposed" to do. She wants a faculty post that would let her collaborate with DOE lab staff.

Rattner, meanwhile, considers himself fortunate to find a tenure-track position at a prestigious school. It's a bonus that it's near family in Maryland, his home state.

It could not have happened without the DOE CSGF and the cross-disciplinary collaboration it encourages, Rattner says. "There are lots of opportunities in national labs and academia for people who have a range of experiences and capabilities in addition to their specialties."

But Pribram-Jones says fellows must extend themselves if they're to maximize the experience. She advises them to "pick new and exciting things. If something's exciting to you, (the DOE CSGF) is one of the only programs out there that encourages you to do it."

HEALTHY COMPUTING

Rick L. Stevens is associate laboratory director of the Computing, Environment and Life Sciences Directorate at Argonne National Laboratory. Among his many duties, he is a lead investigator for Department of Energy (DOE) collaborations with the National Cancer Institute (NCI) and White House BRAIN (Brain Research through Advancing Innovative Neurotechnologies) project. He was an invited speaker at the 2016 DOE CSGF Annual Program Review.

DEIXIS: DESCRIBE THE DOE-NCI COLLABORATION.

Rick Stevens: We work on three pilots. Fred Streitz from Lawrence Livermore National Laboratory leads the molecular-scale pilot, focusing on RAS, a gene involved in about 30 percent of cancers. That pilot is doing multiscale molecular modeling to get at not just the RAS molecule but the pathway it interacts with.

Gina Tourassi at Oak Ridge National Laboratory leads a pilot dealing with population-level data, focusing on a national cancer surveillance database with millions of records, including detailed pathology reports. The pilot will build a deep-learning text comprehension system that can read these reports and translate them into structured data we can compute on to build population-level simulations and optimize treatments.

I lead the cellular-scale pilot, building predictive models for drug treatment based on biological experiment and drug-response data. Machine-learning models will integrate those data, including additional drug structure information, and predict a given drug's effect on a given tumor. NCI also builds PDX – patient-derived xenograft – models, transplanting human tumor material into immune-suppressed mice. Those mice can become proxies for drug trials, letting us find the best drug for a specific tumor. With those data we'll build models to predict the best drugs to try. The results will help us create a model to identify the best treatments for patients with similar tumors.

Frank Alexander, the Los Alamos National Laboratory lead, has expertise in uncertainty quantification and optimal experimental design. We're applying that to all three areas. All labs work on all pilots, but we have primary responsibility for each.

WHAT ROLE DOES DOE AND ITS HIGH-PERFORMANCE COMPUTING (HPC) CAPABILITY HAVE IN THE BRAIN INITIATIVE?

The collaboration is broader than computing, involving DOE's Biology and Environmental Research, Advanced Scientific Computing Research and Basic Energy Sciences programs. But from a computing standpoint, there are big challenges.

One is the image-segmentation reverse-engineering problem. There are two ways to map a brain: First is a static reconstruction of the wiring diagram, called the connectome. Researchers slice a dead brain, usually from a mouse, into tens of thousands of pieces each about 10 nanometers thick and image them with scanning electron microscopes. We must recognize the 3-D-connected parts in the images and rebuild the structure. The data are large – for a mouse it's about 500 petabytes – so you have huge data-management and machine-learning problems. With a static map, you can model the neurons and topology and study the network's behavior. Again, it's a huge computational and mathematics problem.

Second is the functional problem: Instead of slicing dead brains, researchers stick probes into live brains and record signals. The problem is taking those data and producing a functional behavior map. That has visualization, large-scale computing, signal-processing and other challenges. Ultimately, we want to use the functional data to constrain behavior of the model in the connectome and get an integrated model.

WHAT CAPABILITIES DO DOE'S HPC RESOURCES BRING TO THESE PROJECTS?

We have the biggest, most capable infrastructure. We have more than a thousand computer scientists and mathematicians that not only are world-class scientists and mathematicians but also have experience solving real problems on these computers. Third, the labs know how to build software and do engineering at scale. For both projects, we need all three things. That's what's unique about the labs that neither NCI nor the BRAIN project had access to.

HOW WILL THESE COLLABORATIONS ADVANCE HPC?

They're a forcing function to make sure our computer architectures are more balanced. Historically, HPC has been optimized for particular kinds of simulations. These projects require us to do simulations but also support large-scale data analytics and machine learning. The vision we have for these systems is that they're equally good at those three domains. The collaborations also are examples where we need access to experiments or data. Prioritizing experiments or data acquisition to fill in missing theory of how these models will work is a good proxy for many DOE problems – the weapons program, climate or other areas where we can't easily do experiments.

'The data are large – for a mouse it's about 500 petabytes – so you have huge data-management and machine-learning problems.'

WHAT'S AHEAD FOR INTEGRATING HPC AND HEALTH RESEARCH?

We're seeing more opportunities to unite different kinds of data, like environmental data, that directly affect health. That will improve the ability to predict outcomes. And it's not just technical data; people are basically sensorizing themselves with Fitbits and Apple watches and so forth. Those data streams will find their way into complex models. We're also seeing more opportunity to apply machine learning to either drive, integrate with or complement simulations.

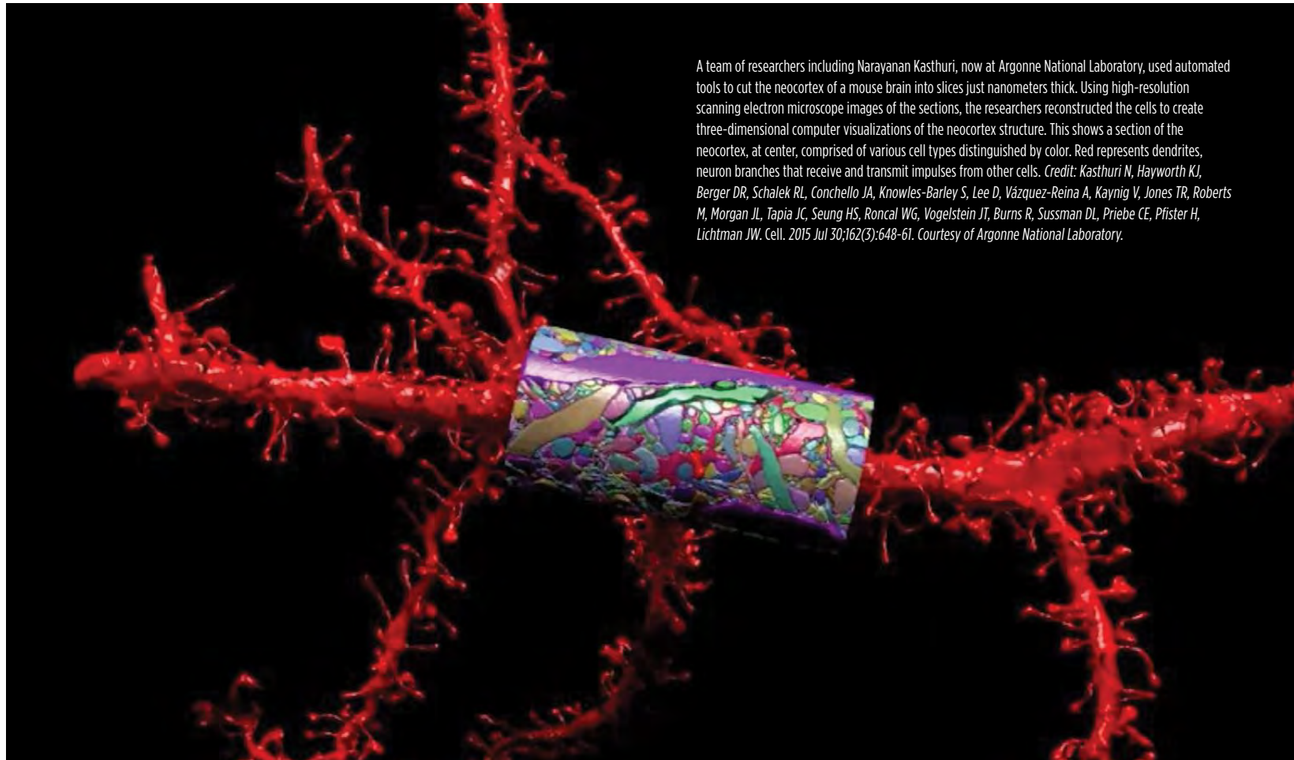

A team of researchers including Narayanan Kasthuri, now at Argonne National Laboratory, used automated tools to cut the neocortex of a mouse brain into slices just nanometers thick. Using high-resolution scanning electron microscope images of the sections, the researchers reconstructed the cells to create three-dimensional computer visualizations of the neocortex structure. This shows a section of the neocortex, at center, comprised of various cell types distinguished by color. Red represents dendrites, neuron branches that receive and transmit impulses from other cells. *Credit: Kasthuri N, Hayworth KJ, Berger DR, Schalek RL, Conchello JA, Knowles-Barley S, Lee D, Vázquez-Reina A, Kaynig V, Jones TR, Roberts M, Morgan JL, Tapia JC, Seung HS, Roncal WG, Vogelstein JT, Burns R, Sussman DL, Priebe CE, Pfister H, Lichtman JW.* Cell. *2015 Jul 30;162(3):648-61. Courtesy of Argonne National Laboratory.*

COSMIC ENCOUNTER

Nicholas Frontiere started national laboratory research at age 16. Now he's helping make some of the biggest and most precise simulations ever of the universe's growth.

By Thomas R. O'Donnell

It's common advice: Breakfast is the most important meal of the day. For Nicholas Frontiere, however, it may have been the most important meal of his career.

Frontiere, a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient, was a home-schooled teenager living near Los Alamos, New Mexico, home of a DOE national laboratory. He'd quickly completed his mother's courses and was flying through ones available at a local community college.

Frontiere's father, a retired composer, ate at a nearby café every morning and fell into an informal breakfast club. One regular, a former Los Alamos National Laboratory (LANL) scientist, introduced the younger Frontiere to lab astrophysicist Ed Fenimore.

Fenimore wanted to hire the youth, but he was only 14, too young under government regulations. Fenimore instead hired Frontiere's sister, Emily, also a precocious student, and told him to return in two years. (Emily Frontiere now is pursuing a master's degree in Medieval studies.)

At 16, Nicholas Frontiere (pronounced "FRONterry") finally began working at the birthplace of the atomic bomb. Under Fenimore's tutelage, he gained a security clearance and honed his computer science skills on classified projects.

The collaboration continued as Frontiere attended the University of California, Los Angeles, but the work's sensitive nature prohibited him from publishing scientific papers. Fenimore asked LANL physicists Katrin Heitmann and Salman Habib about finding Frontiere a group doing open research.

"We said, 'Well, he should work with us,'" Habib says. "That's how it started."

The team Habib and Heitmann lead builds simulations of the universe's evolution and runs them on the world's fastest supercomputers. The calculations, some of the biggest and most detailed yet, track trillions of tracer particles representing all the matter in the observable universe. As the particles evolve, gravity pulls them into clumps and filaments of visible

With Frontiere's high school and college experience, "you can't really compare him to your typical graduate student," Habib says. In terms of ability, "Nick is off the scale."

Frontiere also is the second author of one of the group's latest papers, published last year in the *Astrophysical Journal Supplement Series*, that describes a cosmological simulation called the Q Continuum. (Heitmann, a "Star Trek" fan, named it for the alternate dimension a race of super beings occupies in

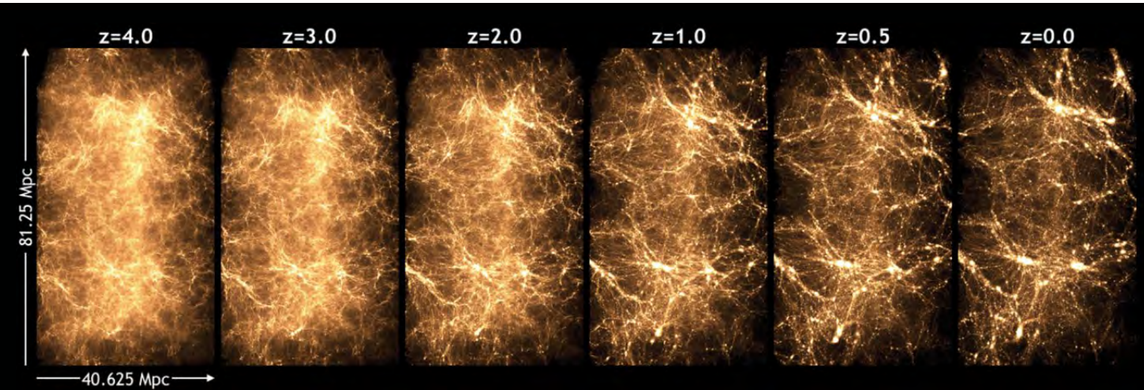

Opposite: Nicholas Frontiere and Salman Habib with Mira, Argonne National Laboratory's Blue Gene/Q. *Credit: Wes Agresta, Argonne National Laboratory.* **Above:** The evolution of dark matter distribution over time, from a redshift of 4 (about 12 billion light years) to today over just a piece of the Q Continuum simulation, about 81 million parsecs by 81 million parsecs by 41 million parsecs (around 264 million light years by 264 million light years by 134 million light years). It shows the detail the simulation was able to resolve in the dark matter web. *Credit: Katrin Heitmann, Nicholas Frontiere, Chris Sewell, Salman Habib, Adrian Pope, Hal Finkel, Silvio Rizzi, Joe Insley, Suman Bhattacharya. The Q Continuum simulation: Harnessing the power of GPU accelerated supercomputers.* The Astrophysical Journal Supplement Series, *2015; 219 (2): 34 DOI: 10.1088/0067-0049/219/2/34.*

galaxies and halos of dark matter over billions of years, from near the Big Bang to today.

Dark matter is a physics mystery: It makes up most of the universe's stuff, but is invisible because it doesn't interact with light. Physicists know it's there because, without its gravitational influence, galaxies and galaxy clusters would fly apart. The group's simulations can help cosmologists understand phenomena like the universe's accelerating expansion.

"They opened my eyes to the high-performance (computing) world, and I've been loving it and stuck in it ever since," Frontiere says. "I want to use this tool for whatever physics I come across." He kept working with the group even as it moved to Argonne National Laboratory and the University of Chicago.

The result: While still an undergraduate, Frontiere contributed to cosmology models in 2012 and 2013 that were finalists for the prestigious Gordon Bell Prize recognizing outstanding HPC achievements. After earning bachelor's degrees in physics and mathematics, Frontiere naturally enrolled as a Chicago graduate student to continue working with the group.

the science fiction franchise.) The model ran on Titan, Oak Ridge National Laboratory's Cray XK7 supercomputer, and tracks cosmological evolution from a mere 50 million years after the Big Bang to today.

The group's workhorse code, the Hardware/Hybrid Accelerated Cosmology Code (HACC), was designed for efficiency and adaptability. It runs well even on HPC systems comprised of both standard processor cores and graphics processing units (GPUs), cousins to video game chips. HACC has hit near 100 percent efficiency on Titan, using almost 90 percent of the machine's 299,008 processor cores and 18,688 GPUs. On Sequoia, Lawrence Livermore National Laboratory's IBM Blue Gene/Q, HACC has tracked the movement of 3.6 trillion particles on more than 1.5 million processor cores, hitting a speed of almost 14 petaflops (quadrillion calculations per second).

It's impressive, but like many cosmological codes, HACC calculates only gravity's effects on matter. It omits other physics governing baryons, the particles that comprise all visible matter: us, the stars and planets, and everything we see. Unlike dark matter, baryons are subject to forces beyond gravity.

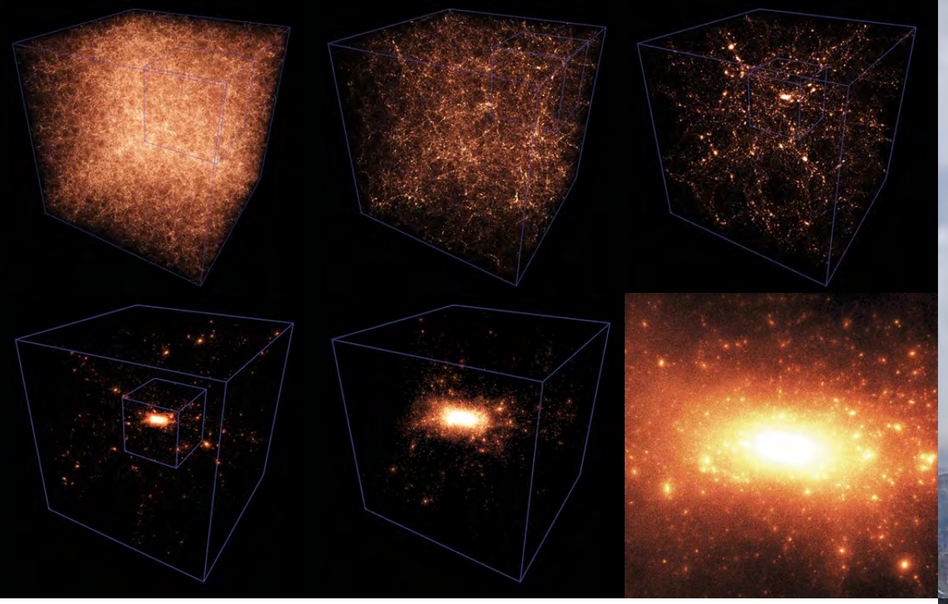

This visualization of the Q Continuum simulation shows dark matter particle halos distributed through a large part of the universe as they would be today after evolving from soon after the Big Bang. To highlight the cosmic web structure seen in the simulation, the visualization shows just 1 percent of the particles (each representing mass millions of times that of the sun) residing in the halos. Each image shows a zoomed-in section of the previous image, ending with a view of one of the most massive dark matter clusters in the simulation. *Credit: Katrin Heitmann, Nicholas Frontiere, Chris Sewell, Salman Habib, Adrian Pope, Hal Finkel, Silvio Rizzi, Joe Insley, Suman Bhattacharya. The Q Continuum simulation: Harnessing the power of GPU accelerated supercomputers.* The Astrophysical Journal Supplement Series, *2015; 219 (2): 34 DOI: 10.1088/0067-0049/219/2/34.*

Baryonic physics has little impact on structure formation when simulating huge pieces of the universe, and including it greatly increases the demand for computational power, Habib says. "So typically, you just turn the baryonic effects off and you run pure gravity."

But data from the latest astronomical missions, like the Dark Energy Spectroscopic Instrument and the Large Synoptic Sky Telescope, demand more precise models. As these devices look deeper and wider into the universe, "the statistical error bars are going down, to the extent that they're almost not there," Habib says. "That means our ability to model becomes really important" to interpreting the data. Researchers can ignore baryonic physics if it influences, say, 10 percent of structure formation and instrument error is 10 percent or more. "But when the measurement error is 1 percent, then you say, oops, I have to model the baryons."

Yet, Frontiere says, including baryonic physics is "quite a dreaded thing because (researchers) think it's too computationally expensive. And that's what we're hoping we can really bite a chunk out of."

Frontiere's quest to add baryonic physics to HACC began with his summer 2014 practicum at California's Lawrence Livermore National Laboratory. Working with physicist J. Michael Owen and postdoctoral researcher Cody Raskin, he began modifying a fluid dynamics method, Smoothed Particle Hydrodynamics (SPH), for use in HACC.

To understand why fluid dynamics applies, think of the universe as a tank of liquid. Early in its evolution, it's "almost homogeneous everywhere and everything's fine," Frontiere says. But as gravity squeezes the fluid, "things are going to move around and repel and push and mix." In cosmology, baryonic matter – largely ionized gas – is the fluid. "Once it collapses to the point where it starts pushing back, exciting things happen," such as shocks, vortices and flows.

SPH calculates fluid forces on baryonic particles in the simulation and interpolates those values onto neighboring particles to capture overall behavior. "It's almost as if you're smearing out the particles or smoothing them," Frontiere says, thus prompting the method's name.

Traditional SPH algorithms typically aren't accurate enough for cosmology simulations. With Owen and Raskin, Frontiere tweaked the method to create Conservative Reproducing Kernel SPH.

The name is a clue to the researchers' improvements: The reproducing kernel calculates forces even as the phenomena it simulates become more complex. Such kernels usually are poor at conserving, or maintaining, quantities like energy or momentum through the solution. This version corrects that. The technique is scalable and easily matches to GPUs, so it should run well on HPC systems, Frontiere says.

The team will test HACC's baryonic physics implementation on Titan in fall 2016. The code also is one of the first chosen to run on the newest DOE systems: Cori, at the National Energy Research Scientific Computing Center; Theta, to arrive at Argonne in 2016; and Summit, expected at Oak Ridge in 2018.

"We'll be supplying simulations that people currently think are not possible," Frontiere says.

GOING DEEP

Brenhin Keller crunches data from geologic samples to understand how our planet's rocky features formed.

By Sarah Webb

When he started his graduate studies in 2011, Brenhin Keller had planned on a laboratory-based research career focusing on the chemistry and features of ancient rocks. But at the time, his Princeton University Ph.D. advisor, Blair Schoene, was still setting up equipment in his new laboratory. He and Keller instead chose a computational project Keller could work on immediately, until the lab was ready.

That research – analyzing geochemical databases to probe fundamental questions about the Earth's early geology – soon led to a paper for the prestigious journal *Nature*. By the time it came out, Keller also had recognized that computation let him address questions other geochemists hadn't been able to examine quantitatively. With the support of a Department of Energy Computational Science Graduate Fellowship (DOE CSGF), this side project grew into Keller's primary research focus.

As a Cornell University undergraduate, Keller noticed there were areas of geology and geochemistry where computation could be useful but hadn't been widely applied. "Even though I didn't get into computation much as an undergrad, it was in the back of my mind," he says. He arrived at Princeton with little programming experience but got up to speed in courses and in consultations with colleagues whose work often used computational techniques.

Brenhin Keller's work focuses largely on computation, but he's participated in several research expeditions. In 2011, he was part of a team that mapped and sampled the Bergell pluton, a large rock formation in the central Swiss Alps. Here Keller looks across a glacial valley toward what was once the roof of the now-tilted formation. *Credit: Kyle Samperton.*

Schoene had experience working with rocks from the Archaean Eon, between 4 billion and 2.5 billion years ago. At the end of the eon oxygen first became abundant in the atmosphere, dramatically altering how chemicals like sulfur and iron cycle from soil to air to organisms and back again and leading to one of Earth's earliest ice ages.

"There are a lot of qualitative observations about the Archaean to suggest that there may have been a quantitative difference in the way things worked back then," Keller says. But those earlier studies were based on the analysis of only a few hundred to a couple of thousand rocks. Within the last 20 years, geochemists have compiled data about far more rock samples into online databases such as EarthChem.

Keller and Schoene realized computational tools could mine this chemical information so they could model magma formation during the early Archaean. "Instead of having 1,000 or 2,000 samples, which is what you'd typically see in a compilation paper at the time, we had 70,000," Keller says.

A program like Microsoft Excel often was the tool of choice when geochemists analyzed smaller data sets. But working with the mass of information Keller and Schoene were examining required greater computational resources and more sophisticated statistics. To scrutinize these data sets, Keller used weighted bootstrap resampling, a statistical technique that ensured the analysis accurately represented the entire globe and the uncertainties within the field measurements.

For example, the researchers had to account for samples' proximity to each other. "If there's one area where we sampled a whole bunch of rocks right next to each other," Keller says, "each one of them doesn't contain as much new information as a (sample) from someplace far away that's the only one of its kind." There's also some uncertainty inherent in analyses of a rock's age and chemistry. "If you look at rocks of any given age, the variability of their composition at any one time is much larger than their variability over time." Keller made that uncertainty part of the resampling process.

The researchers' models, as detailed in their 2012 paper, showed long-term continuous cooling of the Earth's mantle over the Archaean Eon. But around 2.5 billion years ago, the cooling patterns abruptly steepened. At that same time the continental rock record shows rapid changes in the abundances of a range of trace elements. Taken together, these suggest a modification in the process of crustal differentiation, in which mafic (high in magnesium, low in silicon) magma derived from Earth's mantle evolves into less-dense felsic (high-silicon) magma.

That timing also is consistent with when an abundance of oxygen first appeared in Earth's atmosphere, a critical point in life's evolution. Although this oxygen ultimately comes from photosynthesis, some of the details of this record don't align. Fossil evidence suggests that oxygen-producing photosynthesizers evolved well before 2.5 billion years ago, but the gas didn't accumulate then, indicating that it was consumed faster than it was produced.

Keller's analysis suggests an explanation based on chemical changes in rocks dating from the Archaean. They indicate that crustal differentiation occurred at high pressure deep within

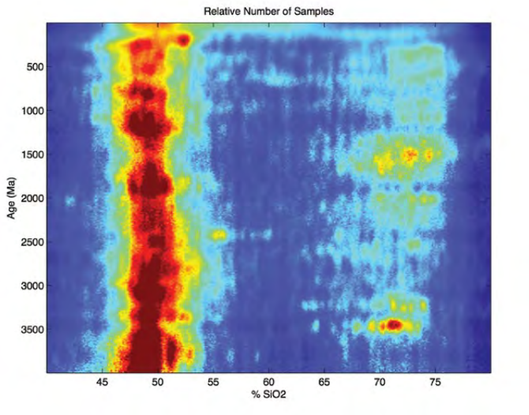

An estimated sample density map reveals the persistence through geologic time of silicate rock composition with abundant basaltic (about 50 percent silica) and granitic (about 70 percent silica) magmas but little in between. *Credit: Brenhin Keller.*

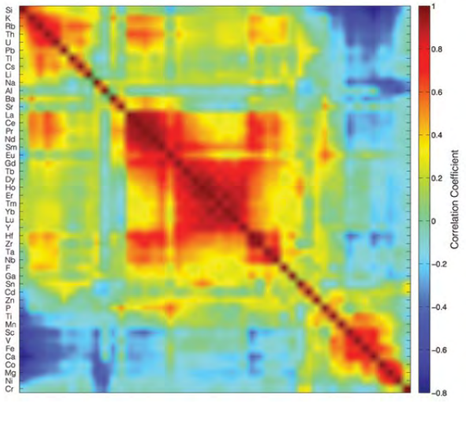

A many-dimensional geochemical data set reveals element groups with similar (red) and dissimilar (blue) geochemical behavior in silicate magmas. The horizontal axis is the same set of elements in the same order left to right. *Credit: Brenhin Keller.*

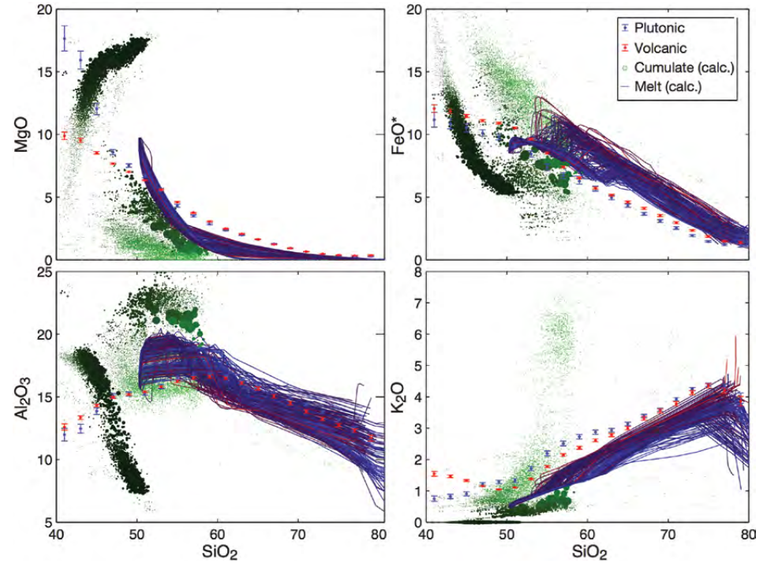

A comparison of observed average crustal compositions (error bars), calculated magma compositions (blue lines), and calculated crystal cumulate residue compositions (green dots) for 200 out of 1.3 million magma fractional crystallization simulations Brenhin Keller and his advisor, Blair Schoene, conducted for a 2015 study. In each plot, the horizontal axis shows the weight percent silica (SiO2); the vertical axis shows the weight percent abundance of some of the most important major elements (clockwise from top left, magnesium, iron, potassium and aluminum) in oxide form. In nature, high-silica rocks like granite form by fractional crystallization of low-silica magmas. Changes in the other major elements with increasing silica reflect the influence of different fractionating minerals. *Credit: Brenhin Keller*/Nature.

Earth, which is one way to produce magmas (and associated volcanic gases) that have more electrons available for bonding to other elements. Such magmas and volcanic gases would consume oxygen from the atmosphere, reducing it until the end of the Archaean.

The continental crust is crucial to life on Earth, both as a source of nutrients and as a part of silicate weathering feedback, the process that regulates carbon dioxide content, keeping temperatures on the planet suitable for liquid water on billion-year timescales. This weathering feedback requires the presence of both oceans and continents – easier said than done when the continents are constantly eroding. It works because the high-silica crust within continents floats like an iceberg on Earth's solid, convecting mantle.

Keller and Schoene have gone on to study (and publish in another *Nature* paper) how this lighter continental crust forms from basaltic magma. Basaltic magma is derived from the melting of Earth's mantle, the ultimate source of most of silicate rocks. Basalts, however, are heavy and contain high concentrations of iron and magnesium. Rocks in the continental crust lack many of these denser, higher melting-point minerals and instead have more silica-rich minerals that contain alkali metals such as sodium and potassium.

Scientists have long debated how these rocks transformed. Did these heavier, hard-to-melt materials crystallize out of molten basalts? Or were there processes that warmed an already-solid basalt and allowed lower melting-point minerals to flow away?

This time Keller and Schoene looked at 300,000 whole-rock analyses and considered silica content with other chemical components. Those patterns help them establish whether rocks formed from magma at the surface or crystallized deep underground. The tectonic environment also influences these processes, so they compared activity at rifts, where magma upwells and melts as pressure decreases, and at arcs, subduction zones where water flows into the mantle, liquefying rock much like salt water melts ice.

Their models showed that at least 60 percent of the crustal formation appears to occur through crystallization processes rather than partial melting. The researchers also weighed in on a longstanding debate about whether granites originate directly from magma. Their analysis supports direct formation from magma rather than weathering and water acting on the material.

Keller plans to incorporate lab and field research into future projects, but computational work will remain his core focus. "This project and the DOE CSGF have turned me into more of a computational scientist than I ever thought I would be," he says.

Keller has worked tirelessly to gain computational experience and shares his deep knowledge of physics, chemistry and electronics with other researchers, Schoene says.

Like other disciplines, geochemistry is moving toward increasing analytical and computational rigor. Keller's experience positions him well for those opportunities, Schoene says. "He bridges many gaps that many people can't. He's special in that way."

EBB & FLOW

Jesse López meanders far from his New Mexico roots, modeling biogeochemical activity in the lush Columbia River estuary.

By Thomas R. O'Donnell

The Columbia River drains more than 250,000 square miles of Canada and the United States and meanders more than 1,200 miles through Washington and Oregon. Its estuary – the final, broad stretch of more than 100 miles as it flows to the Pacific Ocean – hosts an abundance of creatures, from microscopic plants to salmon and seals.

This is Jesse López's laboratory. The Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient is part of a team researching the estuary's biogeochemistry – its physical, chemical, biological and geological processes and reactions. What they learn will help preserve its health and productivity.

The estuary's green valley is almost nothing like where López grew up: the high desert of Albuquerque, New Mexico. That may be why he's studying a river: "Definitely, I think I was drawn to something having to do with water."

Like a rambling river, López took some time reaching the coast. He studied history as an undergraduate, first at the University of New Mexico, and later at the University of Washington, where he transferred after falling for Seattle while on vacation.

After graduation, López joined the Americorps program and taught mathematics at a middle school in Brooklyn, New York. "That was a bit of a change of pace," he chuckles. The job was challenging and exhausting, but so rewarding that he stayed for three years. Many teachers leave after one.

López based his lessons on themes, including the environment and computing, leading him to study climate modeling. It interested him so much he decided to enter the field and searched for a graduate program in environmental modeling and high-performance computing (HPC). He found the Oregon Health & Science University – back in the Pacific Northwest – and joined a group led by António Baptista, who oversees the Center for Coastal Margin Observation & Prediction (CMOP), which studies the Columbia River estuary (CRE).

As the name suggests, the ETM is a cloudy mix of sediment, plant and animal detritus and microbes that digest organic material. The ETM moves with tides and river flow but usually follows the point where salty ocean water meets river water.

CMOP models suggest that residence time through most of the estuary ranges from hours to a few days. In the ETM, it's longer, but just how much is something López hopes to learn.

"What we need from Jesse is a quantitative description of these estuarine turbidity maxima from the perspective of their genesis, their dynamics, so that others then can look at the biogeochemistry inside," Baptista says. The team wants models that can predict the effects of such factors as climate change, upstream dam management and coastal earthquakes.

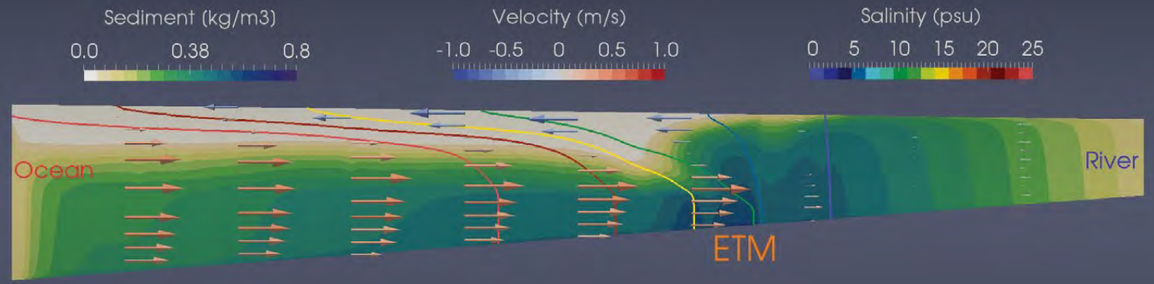

Opposite: Jesse López, on shipboard during a 2011 Pacific Ocean research cruise to measure biogeochemical profiles, including effects on the plume of water coming from the Columbia River. **Above:** This simulation of estuarine dynamics shows the intrusion of dense salt water (colored contours) from the ocean on the left creating upstream velocity near the bed (in red) and downstream velocity elsewhere (in blue). The opposing velocity fields converge near where salinity is low. Particles are trapped there and suspended sediments concentrate (filled contours) in the estuarine turbidity maximum (ETM). *Credit: Jesse López.*

The group, Baptista says, wants to "understand estuaries as bioreactors – places where significant biogeochemical transformations occur" with help from microbes that decompose dead plants and animals. These reactions supply fundamental nutrients for plant and animal life.

But the CRE, unlike most other estuaries, has "very fast water-flushing, so there's not a lot of time for transformations to occur," Baptista says. This short water residence time inhibits reactions, López says, yet "there's definitely a lot of fish and there's a lot of biogeochemical activity in the system, so where does it happen?"

The answer, the researchers hypothesize, is biological hotspots, where microbial communities overcome the limitations of short water residence. López focuses on one candidate: the estuarine turbidity maxima (ETM).

It's complicated, López says. "You have the ocean water that's trying to go upstream and generally you have the river water that's going downstream, so you have this convergence. It's at this point where you have high levels of turbidity" and mixing. The computer models must capture it all, including sediment, detritus and microbes.

One of CMOP's tools is SELFE (pronounced "self"), a code that models circulation in layered fluids. Although originally designed to simulate the CRE, researchers around the world use SELFE on similar problems.

The code uses tracer fields, applying the rules of physics and solving fluid dynamics equations to track qualities like salinity and temperature over time. López's research introduced tracers for sediment – a more complex input because sediment sinks. His algorithms add a term for settling velocity and deal with

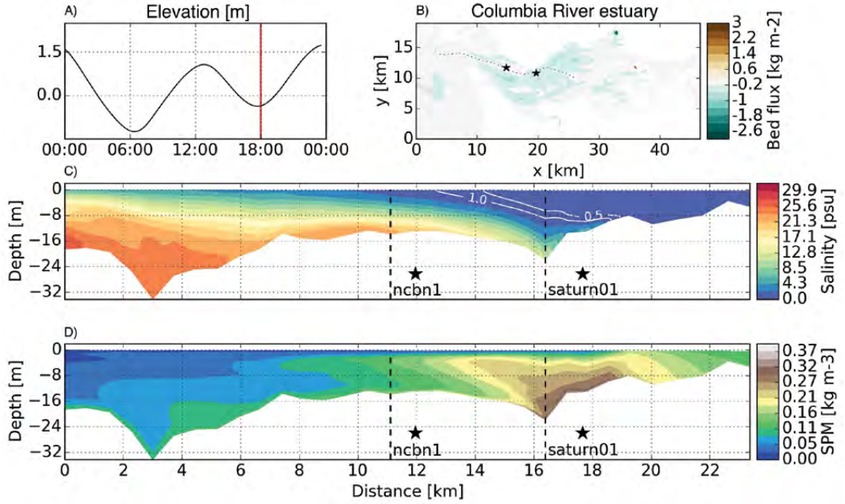

This simulation of the Columbia River estuary demonstrates: B) deposition of suspended sediment, C) intrusion of salt water from the ocean into the estuary, and D) suspended sediment concentrations during a slack tide. (See water elevation over time in Panel A.) The estuarine turbidity maximum (ETM) is observed near the limit of salinity intrusion as the brown color in Panel D. The stars indicate the locations of data monitoring stations designated ncbn1 and saturn01. *Credit: Jesse López.*

other factors, like how sediment moves along the bed of the river or ocean. "I'm able to follow sediment through water and understand where it moves and how the concentration changes."

With the models' help, "we have a much better understanding of the physics and the dynamics that are causing the ETM to exist," López says. He and his colleagues are estimating how long water stays in the ETM compared to elsewhere in the estuary and how much the ETM contributes to productivity.

But López was frustrated because SELFE's performance stalled when running on more than 128 processor cores. "I wasn't impressed with its strong scaling," in which a problem reaches a solution faster in proportion to the number of processors it uses.

López used his 2013 Argonne National Laboratory practicum to address the problems. Working with Jed Brown, an assistant computational mathematician, he analyzed SELFE's performance and improved its scalability and workflow. He linked it to code libraries that SELFE can summon for routines that perform functions like solving equations. He also addressed input and output bottlenecks.

The improved code drastically cut the time to solution for most simulations, Baptista says. Researchers now can more easily simulate estuary circulation over multiple years, helping them understand the ecosystem's variability. "With the type of performance we got from the code prior to Jesse, that was very, very difficult to do."

With the improved SELFE, Baptista's group went from using its own small clusters to running on supercomputers like Stampede, the University of Texas's Dell Linux machine. López also has used HPC systems at the National Energy Research Scientific Computing Center, including Hopper, a Cray XE6, and Oak Ridge National Laboratory's Titan, a Cray XK7.

Using modeling and other tools, the research group hopes to clarify some of the CRE's basic properties, including whether it releases carbon or absorbs it from the environment. Researchers also want to understand when, which and why parts of the estuary generate their own energy while others depend on energy from elsewhere.

López is thrilled that HPC can help answer some of these questions. He's also happy to see his efforts bear fruit. "I work directly with folks engaged with policy and management" of the CRE. "Our model results are actually used in the system and have implications immediately. That's incredible motivation" to continue his research. He hopes to find a postdoctoral post at a national laboratory after graduating in 2016.

To inform and validate its models, CMOP relies on sensors that gather data on salinity, temperature, chemistry and other conditions throughout the CRE. López has gone on several cruises to check the units and gather CRE data.

"I love it," he says. His computational science work is as important as gathering data, but "the feel of it when you're actually out in the field is just incredible."

And it's a long way from the desert.

FRACTURE TRACKER

Andrew Stershic applies new models to study how materials crack at the atomic scale.

By Karyn Hede

The forces that lead to fractures and failures, the bane of civil engineers everywhere, are pretty well understood on the macro level, but modeling and predicting cracks and shattering behavior at the atomic scale remains a demanding problem in computational physics.

With an undergraduate degree in civil engineering from the University of Maryland and a proclivity for taking on complex computational challenges, Andrew Stershic is combining two mathematical methods to advance modeling of fragmentation, a fundamental engineering problem and one of the most numerically complex to simulate.

Portraying a realistic fragmentation problem (picture a hammer hitting glass or a bullet striking an armored car) can require millions of computational degrees of freedom – enough to keep the most powerful computers busy for days, says Stershic, a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient.

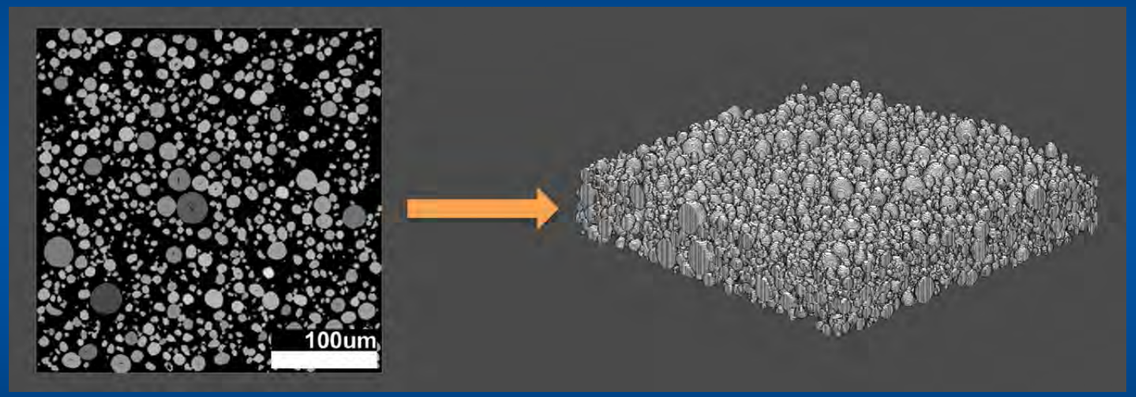

X-ray computed tomography slices, like the one on the left, of lithium-ion cathode microstructure are stacked to form a three-dimensional image of the microstructure. Individual particles from this stack can be identified and analyzed. *Credit: Andrew Stershic.*

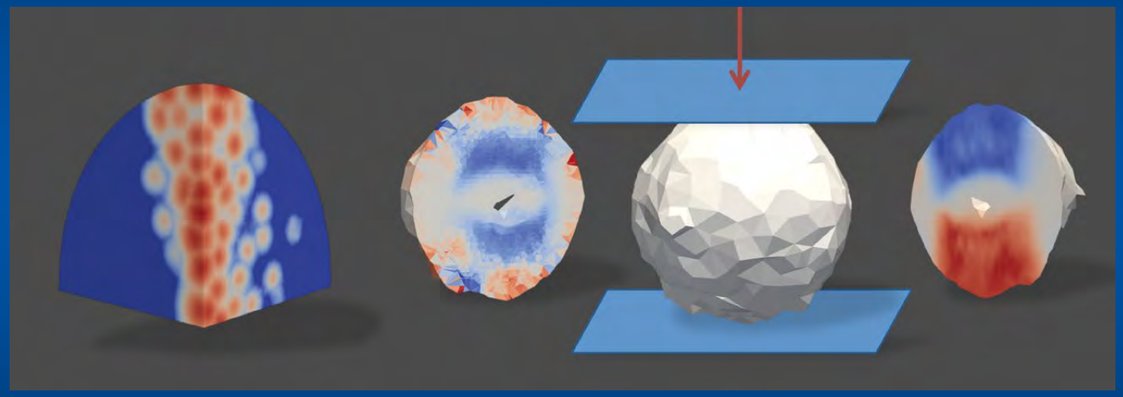

An individual cathode particle taken from the imagery stack to be simulated in compression. The applied model computes critical quantities such as damage, stress, and velocity that can be used to predict the particle's failure. *Credit: Andrew Stershic.*

While on his DOE CSGF practicum at Oak Ridge National Laboratory, the Duke University doctoral candidate wrestled with just such a problem: simulating damage to a lithium-ion battery cathode when particles split during manufacturing. (He wrote about it in an essay published in the 2015 *DEIXIS*.) Drawing on his graduate research, Stershic reasoned that he might have a solution for weaknesses in the lab's atomistic damage simulation model.

Getting that answer turned the summer project into what was essentially a second dissertation's worth of work and a "seminal contribution" to computational physics, says Stershic's graduate advisor, John Dolbow, a DOE CSGF alumnus and Duke professor of civil and environmental engineering.

Stershic describes the project more prosaically as an idea to combine his Ph.D. research with what he learned at Oak Ridge.

At Duke, he's studied modeling cracks and other damage with the finite element method (FEM), a mathematical technique to divide an object or area to be modeled with a mesh of elements so a computer can calculate the physical processes happening in each. Like the pixels in a digital picture, the elements taken together portray the entire object or area. At Oak Ridge, Stershic learned a technique more efficient than the FEM: the atomistic discrete element method (DEM), which divides a computational domain by treating it as a collection of separate, discrete particles.

"I thought 'Let's put them together and see what we can come up with,'" Stershic says. His goal is to maintain the DEM's efficiency while minimizing its shortcomings – chiefly the inability to represent significant microstructural changes.

"One weakness with the DEM model is that we know that real cathode particles break when enough load is applied, but in DEM it's really hard to have particles split apart in a scientifically meaningful way," Stershic says. "This problem provided a really natural connection to what I was doing at Duke."

He recognized a potential solution in a method first devised in France by Nicolas Moës of the Institut de Recherche en Génie Civil et Méchanique, École Centrale de Nantes. Instead of identifying individual points within the mesh of elements as either totally damaged or undamaged, Moës' method – the thick level-set approach, or TLS – models the crack as damage spreads along a field. It can track the movement of multiple merging cracks, making it an attractive method to capture the shattering phenomena Stershic was studying.

"Cracks can come together or branch apart just by changing the value of damage at all those points," Stershic says. "With our (TLS) method, if there is a crack, damage by definition must exist within a certain region or thickness, hence the name thick level-set. The `level-set' part refers to the algorithm to determine the exact crack location from the damage field."

He spent several months in France learning the technique and how others were using it, then brought that knowledge – and the code – back to Dolbow's laboratory to model the cathode particles and simulate shattering phenomena. Once experiments validate its accuracy, he hopes it can be used across a variety of fracturing problems.

'Research at a national laboratory provides a distinct experience from working at a university.'

The resulting method eliminates some of the shortfalls found in the most widely used technique, the cohesive zone method, in which fracture damage can only grow between the vertices of the modeling mesh. It can't split open a modeling element, so it can represent cracks only in certain places. Stershic's solution allows cracks to grow in a way that more closely mimics physical experiments.

Stershic has already shown that the method works to model the shattering of a brittle one-dimensional bar and now is working on a three-dimensional version. He expects to finish at about the same time he receives his doctorate in 2016.

"Andy is the first person to demonstrate that such gradient-based damage models can reproduce theoretical estimates for the scaling of fragment sizes with strain rates," Dolbow says. "He has also shown that such scaling requires a damage model that exhibits the proper energy dissipation."

Stershic's first research experience with Dolbow began during an undergraduate summer internship at Duke, where Dolbow encouraged him to pursue his budding interest in complex engineering problems and to apply for the DOE CSGF.

Stershic admits he really didn't know much about DOE research or the national laboratory system before going to Oak Ridge. "Research at a national laboratory provides a distinct experience from working at a university," he says. "I learned that the DOE places an emphasis on developing interdisciplinary computational scientists who can put its massive high-performance computing resources to good use."

Now Stershic wants to combine his childhood interest in bridges and buildings with all he has learned about modeling and apply computational firepower to a practical engineering problem. He's thinking of working either at a national laboratory or an engineering firm. He's particularly interested in forensic engineering, the study of why structures fail.

"My real passion is for civil engineering-type problems," he says. "Computers are getting more powerful by the day, so the amount you can do with finite element modeling and other numerical models is growing just as fast. In forensic engineering, you can use numerical models to figure out why a building collapses or a bridge fails. It's hands-on, looking at the evidence in the field, building a numerical model, putting in a flaw and asking, 'does that cause the model to fail?'"

With his civil engineer's mind and proven computational modeling chops, Stershic figures that's a problem he can tackle head on.

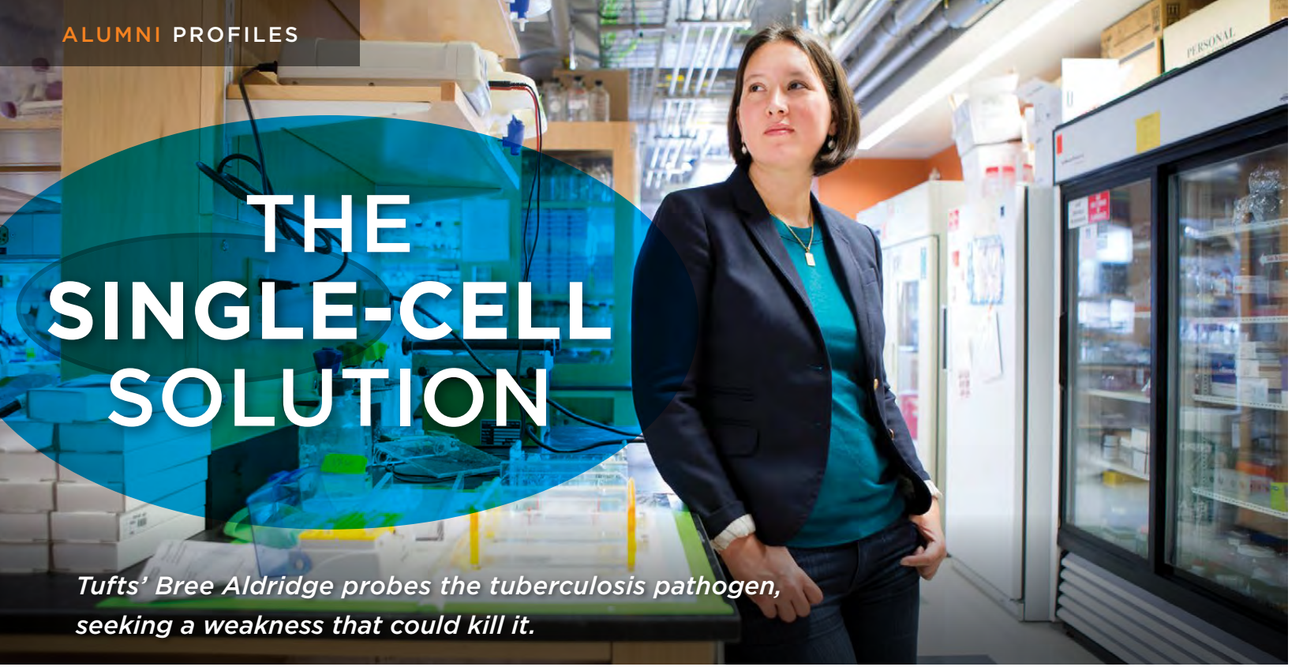

THE SINGLE-CELL SOLUTION

Tufts' Bree Aldridge probes the tuberculosis pathogen, seeking a weakness that could kill it.

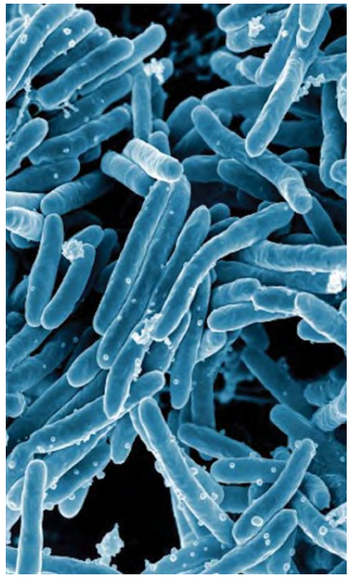

Above: A scanning electron micrograph of *Mycobacterium tuberculosis*. *Credit: National Institute of Allergy and Infectious Diseases.* ***Credit for Bree Aldridge photo, top of page:*** *John Soares for Tufts University. Copyright 2015, Trustees of Tufts College.*

By Andy Boyles

Bree Aldridge's career researching a deadly bacterium flows, figuratively, from what she learned as a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient from 2002 to 2006.

In her doctoral research at the Massachusetts Institute of Technology, Aldridge developed computer models of key chemical signal pathways inside cancer cells. She and her advisors, Douglas Lauffenburger and Peter Sorger (now at Harvard University), looked at the enzyme reactions that make a tumor cell self-destruct. This process, known as programmed cell death, or apoptosis, may provide targets for future drug therapies.

The fellowship's required plan of study directed Aldridge into courses she wouldn't have taken otherwise, including one called Nonlinear Dynamics and Chaos. "That altered my research project," she says. "What I learned in that class was how other researchers use high-performance computing to model how fluids flow."

The calculations of flowing, interacting currents reminded her of seemingly unrelated biological interactions. She knew that a signaling pathway inside a cell is a cascade of reactions that ebbs and flows in dynamic interplay with other processes. Such complex exchanges can defy the traditional experimental approach of altering one variable at a time – one gene or one protein – and watching what happens. "I saw that I could borrow and adapt methods from fluid-flow modeling and apply them to cell biology," Aldridge says.

That's what she's done. Now an assistant professor of molecular biology & microbiology at Tufts University School of Medicine, Aldridge uses novel microbiological techniques and mathematical models to answer questions about the pathogen that causes tuberculosis. A third of the world's population carries the disease and in 2014 it sickened nearly 10 million people and killed about 1.5 million. The World Health Organization ranks it with AIDS as one of the two deadliest diseases.

Aldridge and her colleagues combine laboratory cell studies and computer modeling to reveal new information about how the pathogen, *Mycobacterium tuberculosis* (Mtb), resists drug therapies, offering clues to more effective treatments. In 2013, her work received a prestigious two-year fellowship from the Alfred P. Sloan Foundation and a National Institutes of Health Director's New Innovator Award.

Aldridge didn't always see computer science in her future, although it might have been inevitable. "I come from a family of programmers. Everyone in my family worked at IBM at one time or another."

She resisted the tradition for as long as possible but in high school gave in to the lure of mathematical problems. "I couldn't shake it anymore," she says, laughing. But Aldridge also felt pulled toward biology and worked in the University of Arizona laboratory of Jesse Martinez, where apoptosis was the focus.

Aldridge couldn't decide which direction to take in college: computer science or biology. Martinez told her she didn't have to choose because a new field was emerging, one in which researchers applied computer science to biological questions. The idea sustained Aldridge through her undergraduate years; in 2002 she received University of Arizona bachelor's degrees in computer engineering and molecular and cellular biology.

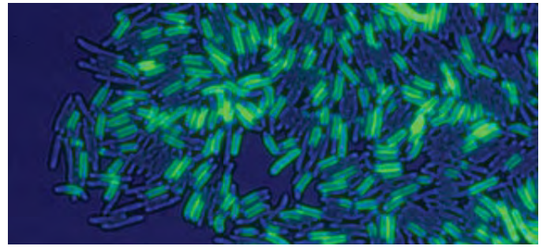

Microscopic image of *Mycobacterium smegmatis*. Researcher Bree Aldridge and her colleagues use this nonpathogenic mycobacterium as an experimental stand-in for *Mycobacterium tuberculosis*. A green fluorescent dye stains old cell walls in the growing bacteria. The image is rendered with a blue pseudo coloring to make unlabeled portions of the bacteria visible. *Credit: Bree Aldridge.*

She noticed similarities between cancer biology and tuberculosis biology, including that in both not all cells respond equally to a single drug. "In the case of tuberculosis, we know that the cells are often genetically similar, but they respond differently" to the same stressor.

By isolating and observing single living Mtb cells, Aldridge discovered one of the mycobacterium's prime defenses. Instead of growing and dividing symmetrically, the cells grow only at one pole and then divide into cells of different sizes, which also grow at different rates. The genetically similar population spins off subpopulations of varying cell sizes that mysteriously respond differently to the same drugs. The discovery earned Aldridge's report a spot in the prestigious journal *Science*. It's important because the pathogen's varied drug response forces patients to remain on cocktails of strong antibiotics for months.

'I wondered why the cutting-edge technologies in quantitative biology were being applied to cancer and rarely to other conditions.'

While focusing on cancer, she abandoned previous models' all-or-nothing approach and applied fuzzy logic to rules governing when and how various biological interactions take place inside a cell. The simulation employed weighted rules that responded in nuanced ways to the passage of time and to the presence of various signaling molecules. As a result, it matched laboratory outcomes using real tumor cells and replicated much of the cell-death cascade occurring inside a tumor cell, including the cross-talk among signaling pathways. Using methods borrowed from fluid-flow modeling, Aldridge discovered that the relative concentrations of signaling molecules can determine whether and how a cell self-destructs.

Aldridge wanted a new direction as she began postdoctoral research at the Harvard School of Public Health. "I wondered why the cutting-edge technologies in quantitative biology were being applied to cancer and rarely to other conditions," she says. "I began looking for another biological question."

At Tufts, Aldridge and her co-workers are exploring how Mtb cells disarm macrophages, the immune cells that become their hosts during infection. Some of these germ-eating defense cells can restrain Mtb while others lose the struggle for dominance and become a haven for the infectious agent.

To discover what causes one macrophage to dominate and another to submit, the team is examining changes in protein levels after Mtb infection. The researchers see some of the same responses that occur in cancer cells. They believe their unique applications of single-cell studies, computing and modeling will lead to a step forward in treating one of the world's deadliest diseases.

THE NON-JOB JOB

Intel's Jeff Hammond plays with codes to boost hardware's performance.

By Thomas R. O'Donnell

Jeff Hammond's task is to tinker: with computer code, new ideas in algorithms and virtually any interesting high-performance computing (HPC) problem.

"I don't feel like I have 'job job,'" says Hammond, a research scientist at Intel Corp.'s Parallel Computing Laboratory in Portland, Oregon. "I don't actually get told by anybody, 'Hey, do this thing.'" Instead, "it's more like, 'Hey, here are some questions we want to look at,' or 'Here are some topical areas. Go tinker around. Play with stuff. See what works. See what doesn't. Try to come up with something nobody's ever looked at before.' That, to me, is the most fun. That's research in a nutshell: just being able to tinker."

Besides the freedom to explore, Hammond also likes the impact his work has. He posts as much programming as possible to GitHub, a repository for open-source code. "Every time I find out that somebody's using my code, I get a lot of pleasure out of that," says Hammond, a Department of Energy Computational Science Graduate Fellowship recipient from 2005 to 2009.

After graduating from the University of Chicago, the Seattle native stayed in the area as a postdoctoral fellow and, later, assistant computational scientist at Argonne National Laboratory. While there, he mentored more than 20 students – many of them DOE CSGF recipients. He counts his work with them as one of his best achievements.

Mentoring was pivotal to Hammond's own career. On his DOE CSGF practicum at Pacific Northwest National Laboratory, he worked with researchers Karol Kowalski and Bert de Jong on his first attempts at programming for parallel HPC systems. He grasped the concepts quickly and added significant capabilities to NWChem, the lab's open-source quantum chemistry code. Hammond changed his doctoral research to focus on NWChem, and he still toys with improvements to the code.

As Hammond's career progressed, he identified himself less as a computational chemist and more as a computer scientist. At Argonne, he delved into HPC's details, working with others to run codes and improve performance on machines like Mira, the Blue Gene/Q supercomputer at the Argonne Leadership Computing Facility.

In 2014, Intel pulled Hammond away with a job offer that returned him to the Pacific Northwest. The company is focusing more on programming that contributes to the performance of its well-known computer processors, and it wanted to tap his expertise in optimizing codes.

"I try to come up with software that makes the Intel hardware work better, particularly in the high-performance computing domain," Hammond says. That often means better models for parallel processing, which divides problems among many processors to reach a solution more quickly than on a single processor. Hammond considers scientific applications to solve complex equations, perform detailed chemistry simulations and linear algebra and to do other jobs and how those tasks map onto Intel's processors.

He places his research in the "middle ground between applications and hardware," including open standards like MPI (message passing interface) and OpenMP (the "MP" here stands for "multiprocessing") that handle the basic functions of many parallel systems. He and several colleagues also focus on adapting standards and finding new ones to run on future exascale computers a hundred to a thousand times faster than today's biggest machines.

Codesign, when scientists who create and use computer applications collaborate with processor engineers and HPC designers, also is part of Hammond's research. "The way we come up with new features for hardware is by looking at the requirements" set out in software. "Computer architects are brilliant people, but they rely on other people like me in the software world to explain what application programmers are doing."

'What happens if we actually build a quantum computer? How do we use it?'

Hammond says his "perspective actually hasn't changed all that much" since leaving Argonne. "I still care a great deal about the end user and what features are there" in processor architecture and in software.

Meanwhile, Hammond eagerly anticipates what's to come in HPC research. Over the past several decades, computers have advanced according to Moore's Law, which says the number of transistors on a processor – and thus its speed – should roughly double every one-and-a-half to two years. Some experts see that growth tailing off, but Hammond believes it will continue, albeit with new manufacturing processes and computer architectures.

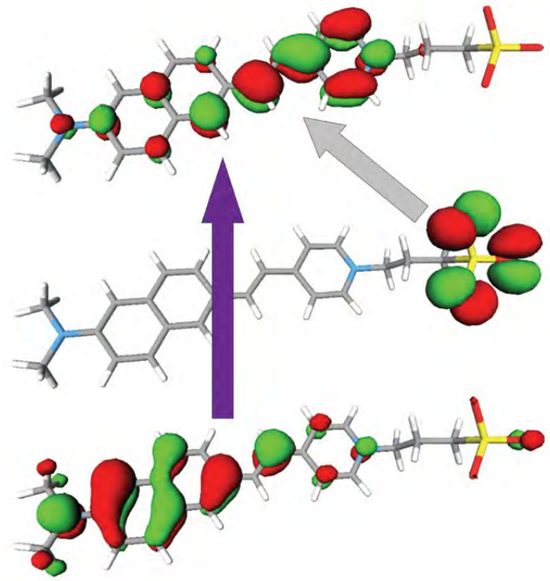

A comparison of two quantum chemistry simulation methods to calculate charge-transfer excited states of a dye used to monitor cells' physiological processes. The colored blobs represent electron orbitals. Density functional theory, a quantum chemistry method, incorrectly predicts the electron transition designated by the gray arrow. Coupled cluster methods, in which DOE CSGF alumnus Jeff Hammond has specialized, correctly predicts the transition marked by the purple arrow. *Credit: Jeff Hammond.*

But if processor speeds do stall, that doesn't necessarily mean application speeds also will, Hammond says. "Most of the codes DOE people run today are running at less than 10 percent" of a supercomputer's projected peak performance. "That's a huge opportunity." Even if HPC systems' rated speeds remain flat, "you can still take advantage of getting a higher fraction of peak with better software, better compilers (which translate programming into machine instructions), better hardware design that eliminates" bottlenecks in communication and other non-computing functions.

Hammond and his colleagues also face even more drastic technological changes – "the fun stuff" – like quantum computing, which relies on strange subatomic physics to drastically accelerate calculations. "What happens if we actually build a quantum computer? How do we use it? What's the software model? It doesn't matter what we can build unless people can do new and interesting things with it."

He'll continue tinkering to find answers.

MATERIALS MINER

Anubhav Jain tunnels ever deeper into materials' specific properties.

By Thomas R. O'Donnell

Anubhav Jain has been in the prediction business since starting graduate school at the Massachusetts Institute of Technology almost 10 years ago. Now, with help from a prestigious grant he earned in 2015, he's digging into those predictions to make better substances.

The Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient worked with his advisor, Gerbrand Ceder, to help start the Materials Genome Project. Using detailed simulations on high-performance computing systems, they predicted the practical properties of thousands of real and proposed compounds and compiled them into a database. Materials scientists used the information to identify good candidate substances for better batteries, solar cells and other applications before going to the expense of actually making and testing them.

The effort now is the Materials Project, based at Lawrence Berkeley National Laboratory, where the lead investigator (and Ceder's collaborator), Kristin Persson, is a staff scientist. Jain joined her at Berkeley Lab after graduation. Ceder brought his group there and to the University of California, Berkeley, in 2015.

Scientists, working from experiments, have compiled information about compounds' magnetic, electrical and other properties for decades. But progress was excruciatingly slow, inhibiting the flow of new materials to the market.

The Materials Project turns this drip of data into a torrent, predicting properties by combining massive computing power with advances in density functional theory, a technique that calculates the interaction of atoms and their electrons at the quantum level.

"Instead of working at the pace of experiments, where you get a certain amount of new data per year, you can start leveraging and piggybacking on supercomputers and generate lots and lots more data (on materials) than was ever available," Jain says. As lead developer, he oversaw simulations on machines at Berkeley Lab and its National Energy Research Scientific Computing Center, raising the number of computed compounds to more than 66,000.

Now Jain is moving to the next step: mining that information to understand how materials work, bringing "data science into materials science in a way that hasn't really been done before."

A DOE Office of Science Early Career Research Program award fuels Jain's investigation. Over the course of five years, his group will receive $2.5 million to apply machine-learning techniques to Materials Project data and to compute new data sets.

Jain targets thermoelectric materials, compounds that produce electricity as they warm up. When under an electrical charge, the materials also can move heat, opening possibilities for exceptionally quiet refrigerators with no moving parts. Thermoelectrics also are stable and reliable, leading aerospace engineers to choose them for satellites.

But thermoelectrics are expensive, made from rare elements like tellurium, and inefficient enough that they're a poor economic choice in all but the most extreme uses (like the aforementioned satellites). "So far, thermoelectrics have not been efficient enough or cost-effective enough to really be useful for generating electricity in a way that you would want to pay for it."

Jain and others want to find more effective, less expensive thermoelectric materials, but "it's very difficult to know in advance whether a material will be a good thermoelectric." For one thing, an effective thermoelectric material must balance conflicting properties, like thermal conductivity, electric conductivity and the Seebeck coefficient – a measure of the voltage produced relative to the change in temperature.

"For a lot of these properties, as you make one better, another one tends to become worse," Jain says. "It's really difficult to know exactly how things will play out for a proposed new material. That's where the simulations come in."

Materials Project calculations predict a compound's performance for each property, allowing scientists to search for likely thermoelectrics. But Jain wants to go further and use data mining to discover what it is about materials that leads to that performance.

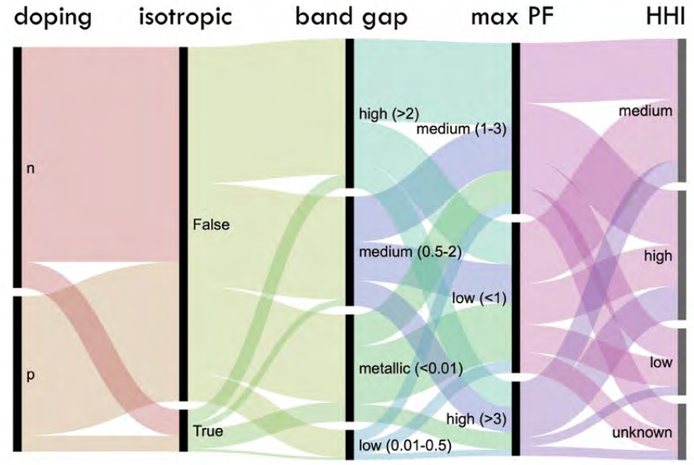

Anubhav Jain and colleagues diagramed a calculated data set of 48,770 materials. From left, doping is the type of elemental impurity added to improve performance that calculations show maximizes the power factor, a measure of a material's highest possible thermoelectric-power output. Isotropic compounds have power factors within 10 percent in all directions. Band gap is electron volts (eV) and max PF (power factor) indicates milliwatts for device applications. HHI stands for Herfindahl-Hirschman Index of elemental reserves for the compound, a measure of a material's availability in world supplies. High means the material's availability is confined to small geographic areas (making it undesirable) and low means resources are geographically dispersed (making it more desirable). *Credit: Anubhav Jain.*

'For a lot of these properties, as you make one better, another one tends to become worse.'

For instance, many thermoelectric material properties, like electrical conductivity and the Seebeck coefficient, relate to the electron band structure – the energy levels at which electrons reside in the substance. "Materials scientists would love to be able to say, 'I would like a band structure with these sorts of features'" to create desirable properties in materials.

That's not yet a reality, but with the Materials Project "we're computing all these band structures for tens of thousands of materials," Jain says. Using data science, "we'll be able to more rationally, maybe not quite draw band structure, but be able to understand how changing the material will affect band structure."

Scientists have used data science before to find what factors influence material properties, usually for small molecules. It's tougher to try the same thing on the materials Jain studies. The repeating arrangements of their crystalline structures mean that, in theory, the compounds contain interactions involving an infinite number of atoms. That's one of the problems he hopes to solve.

Jain also seems compelled to share and discuss his research with the materials and computational science research communities – and the world. He has a Twitter feed (@jainpapers) to summarize his research publications and a blog (hackingmaterials.com) on computational materials science. The blog provides a forum for discussing issues in the field that aren't necessarily fodder for a refereed science journal.

"There's a lot of important things to be said or important discussions to be had, and this was just a nice way that someone can do it." At scientific conferences, he often meets people who have read his work – although "usually they're a bit bashful about it," he says, laughing. "I'm not completely sure why, because my blog involves stick figure cartoons."

Jain laughs again. "For them to be the embarrassed one is probably weird."

ATOMS ON THE DANCE FLOOR

The DOE CSGF stages the Communicate Your Science & Engineering contest to give fellows and alumni an opportunity to write about computation and computational science and engineering for a broad, non-technical audience. The author of this year's winning essay is a fourth-year fellow studying applied physics at Columbia University.

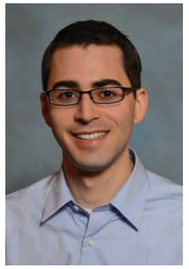

By Eric Isaacs

I love to play music while I cook. It's usually Motown or funk, which makes it impossible for my roommates and me not to start dancing. Before long, we're slicing eggplant to the beat and shimmying over to the oven with soon-to-be lasagna. It's probably my favorite way to relax on a snowy New York evening.

The scientist in me can't help but notice that we're not the only ones dancing – there's another, nearly imperceptible tango happening behind the scenes. It's in the metal foil covering our lasagna. Although you can't zoom in enough to see them, the aluminum atoms also are in motion, wiggling around in complex patterns like a choreographed dance.

These vibrations, called phonons, are present in most substances we encounter every day. From the salt in your shaker to the silicon chip in your smartphone, the atoms don't sit still but instead rapidly move back and forth.

And that's a great thing, since phonons are responsible for heat, sound and many other parts of daily life.

The microscopic dance party astonishes me. Learning about phonons in a physics class changed how I look at the world by highlighting that subtle, complex factors can have enormous impact.

What gets the atoms grooving is not a funk beat but rather the temperature. There's more dancing in a sizzling-hot oven than in an ice-cold freezer. Strangely, even at the lowest temperatures possible, the atoms still wiggle a bit due to quantum mechanics, the bizarre rules governing the microscopic world.

Phonons help determine whether a material is hard, like steel, or soft, like rubber. They tell us whether the material conducts heat efficiently, like diamond, or poorly, like ice. And they dictate how fast sound travels through the material.

But phonons are more than just a scientific curiosity. They have the potential to revolutionize how energy is generated and transported.

For instance, a material called lead telluride can convert exhaust heat from a car or power plant into usable electricity. Its secret is that collisions between its phonons disrupt heat flow. It's as if some atoms are moving to disco while others are rocking out to heavy metal. Their phonon dance steps clash, converting heat into electric current.

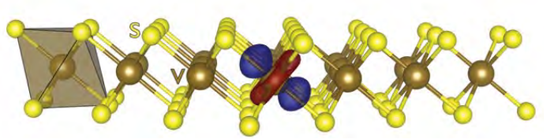

The atomic structure of monolayer vanadium disulfide (VS2) containing vanadium (gold) atoms coordinated octahedrally by sulfur (yellow) atoms. The unpaired d orbital, shown as the red and blue surfaces, causes the material's magnetism. At temperatures below 89 degrees Fahrenheit, a charge density wave modifies the structure and magnetic properties. *Credit: Eric Isaacs.*

Another example: In certain materials called superconductors, electrons flow freely without losing energy. Because of how atoms dance in the quantum world, electrons join in pairs to make this strange behavior possible. (I guess it must be a slow song.) These phonon-enabled superconductors can save energy by efficiently transmitting electricity in our power grid.

If you could squint enough to make out atoms vibrating in a sheet of aluminum foil or a crystal of table salt, what would you see? Using the laws of quantum mechanics and some of the world's most powerful computers, my colleagues and I are trying to answer this question in our Columbia University research group.

To determine how fast and in which directions atoms in a material wiggle, we need to know all the forces affecting them. For example, how does nudging a sodium atom in table salt (sodium chloride) away from its regular position in the cubic crystal structure influence a neighboring chlorine atom?

For most materials, a quantum mechanical approach called density functional theory (DFT) answers these atomic-scale questions. Since it's derived from fundamental physical laws, DFT uses only the composition and structure of the material to give us this information. Of course, it's vital to validate the theory by comparing it to known experiments so we can trust the answers it provides.

It turns out the DFT equations are far too difficult to solve directly with a pen and a pad of paper. Instead, we use what's called a numerical approach: Rather than trying to find a general solution to the equations, we plug in numbers and find a solution for the problem at hand, the way you punch numbers into your calculator.

But we don't employ any ordinary calculator for this number crunching. We use gigantic supercomputers such as Edison, housed at the Department of Energy National Energy Research Scientific Computing Center in Berkeley, California. This machine contains more than 130,000 individual computer processors, holds more than 1.5 million DVDs' worth of data,

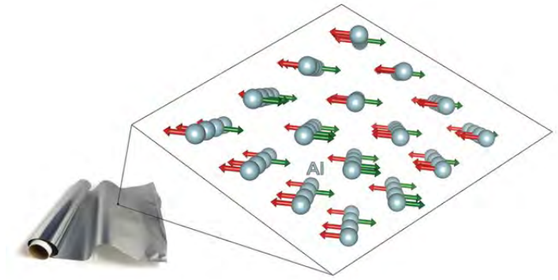

A schematic of an atomic vibration, or phonon, in aluminum. The aluminum atoms (light blue spheres) are in a repeating arrangement known as a crystal, with atoms only 11 billionths of an inch apart. They rapidly move back and forth in the directions indicated by the red and green arrows. *Credit: Eric Isaacs.*

and can perform 2.5 trillion math operations (like adding two numbers) each millisecond. Since it's difficult to directly measure the precise motion of atoms, studying phonons with a computer is quite valuable.

How do we do it? On the computer, we can virtually create a material we'd like to zoom in on and use DFT to calculate the forces between the atoms. After performing some arithmetic on data containing these forces, we know the direction and speed of the atomic vibrations. In other words, using a supercomputer like Edison to compute all the atomic forces lets us see the material's dance moves.

We create animations that depict the phonons as moving atoms. The beautiful and complex vibration patterns never cease to mesmerize me; it feels like watching the ripples in a puddle on a rainy day. But beyond this, the phonons also reveal new insights into our understanding of materials.

One example: What ultimately limits a material's strength? As the name indicates, monolayer materials are as thin as a single layer of atoms, yet are some of the planet's strongest materials. In an unexpected result, our calculations found that a particular phonon, in which groups of six atoms come together, breaks monolayer materials under extreme stress.

A magnetic monolayer material called vanadium disulfide has applications for electronics, biological sensors and energy storage. Here our calculations revealed a special atomic dance move known as a charge density wave, in which the atoms rearrange as the material cools and substantially modify the magnetic field's strength.

Phonons – the microscopic dance parties hidden in materials – have the capacity to enable technologies of the future. And with the help of massive supercomputers, we're uncovering mysteries of the atomic dance floor today.

CLASS OF 2016

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

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

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

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

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

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

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

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

Daniel Strouse
Princeton University
Theoretical Neuroscience
Advisor: William Bialek
Practicum: Lawrence Berkeley National Laboratory
Contact: danieljstrouse@gmail.com

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

Eileen Martin
Stanford University
Computational and Mathematical Engineering
Advisor: Biondo Biondi
Practicum: Lawrence Livermore National Laboratory
Contact: ermartin@stanford.edu

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

Brenhin Keller
Princeton University
Geochemistry and Geochronology
Advisor: Blair Schoene
Practicum: Lawrence Berkeley National Laboratory
Contact: cbkeller@princeton.edu

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

Dragos Velicanu
Massachusetts Institute of Technology
High Energy Physics
Advisor: Gunther Roland
Practicum: Lawrence Berkeley National Laboratory
Contact: velicanu@mit.edu

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

Victor Minden
Stanford University
Computational and Mathematical Engineering
Advisor: Lexing Ying
Practicum: Lawrence Berkeley National Laboratory
Contact: victorminden@gmail.com

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

FELLOWS AND ALUMNI: MAJOR DISCIPLINES

As it enters its 26th year, the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) now counts 436 current recipients and alumni. They pursue unique research in their doctoral studies, but as the accompanying graphic shows, the subjects can be grouped into broad interest areas. Regardless of discipline, however, all fellows and graduates comprise a community that leads the nation in employing computing to solve important problems.

The Office of Advanced Scientific Computing Research within the DOE Office of Science co-sponsors the fellowship with the National Nuclear Security Administration's Advanced Simulation and Computing program.

FELLOW AND ALUMNI DOCTORAL FIELDS*

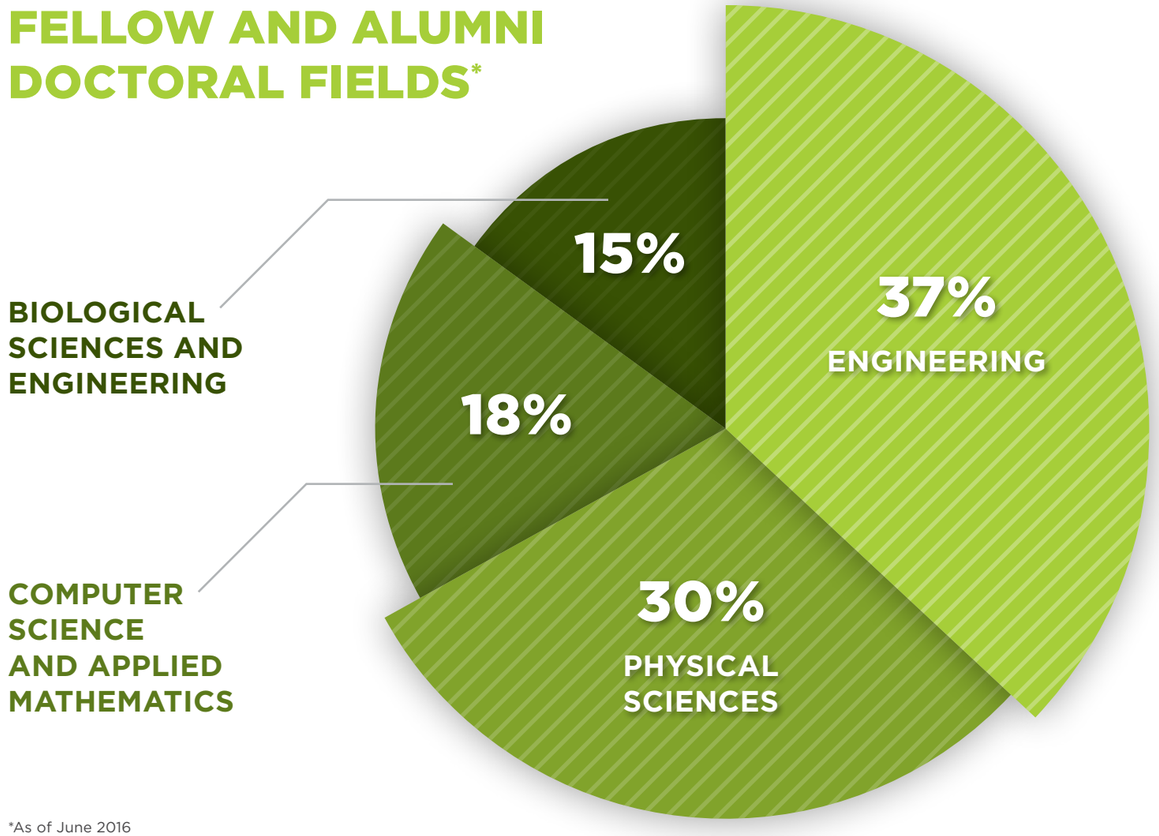

*As of June 2016

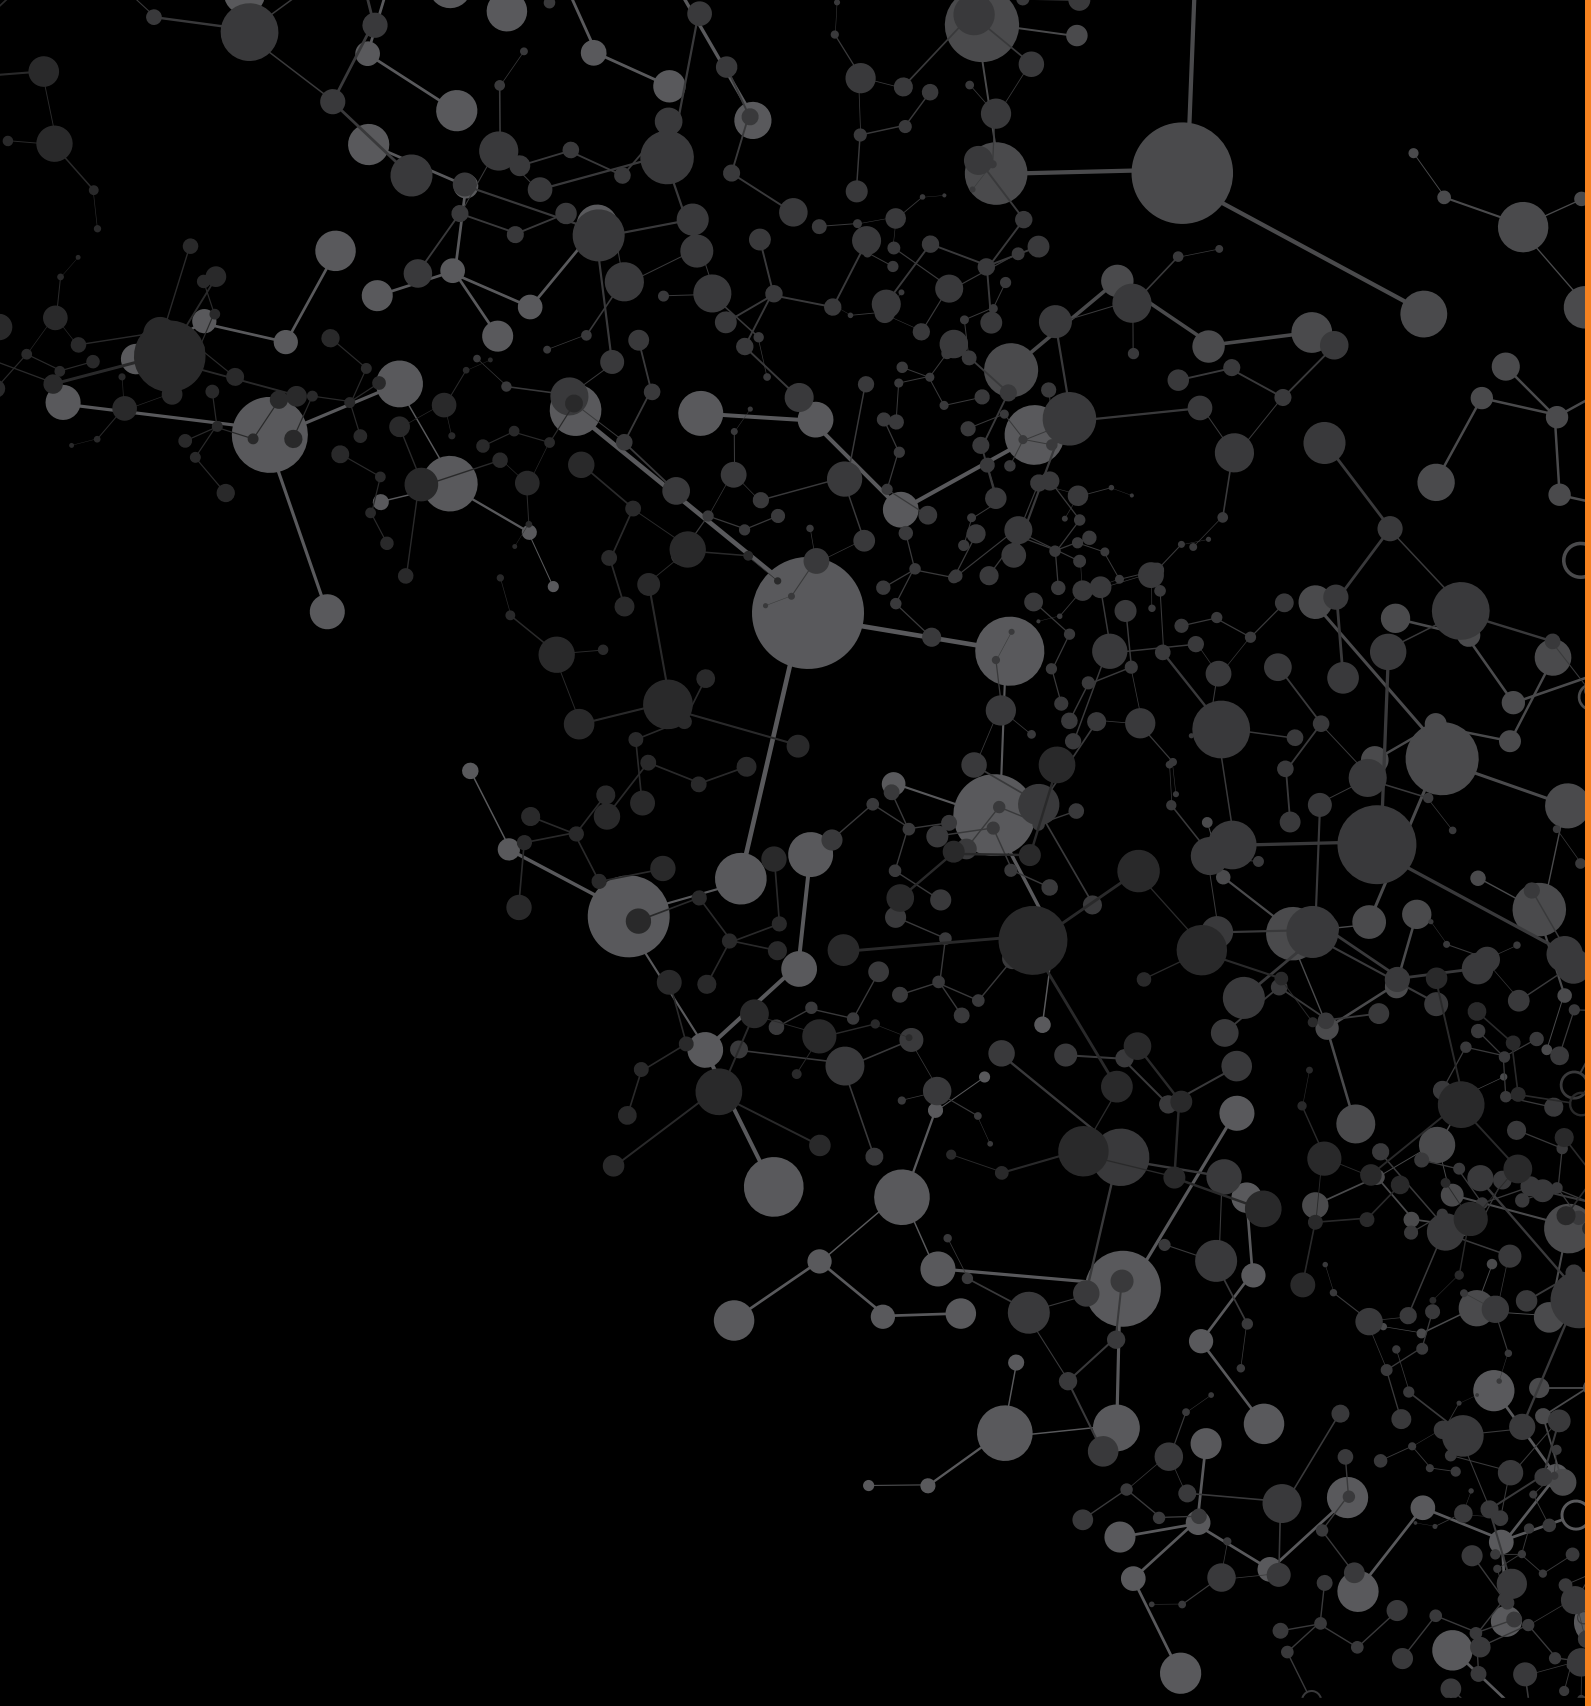

The Krell Institute
1609 Golden Aspen Drive, Suite 101
Ames, IA 50010
(515) 956-3696
www.krellinst.org/csgf

MIX
Paper from responsible sources
FSC® C015782

Funded by the Department of Energy Office of Science
and the National Nuclear Security Administration.