THE DOE CSGF ANNUAL

2020

DEIXIS

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

SMALL DETAILS, HUGE RESULTS

Howes Award-winner approaches the nanoscale world with a big heart

PAGE 28

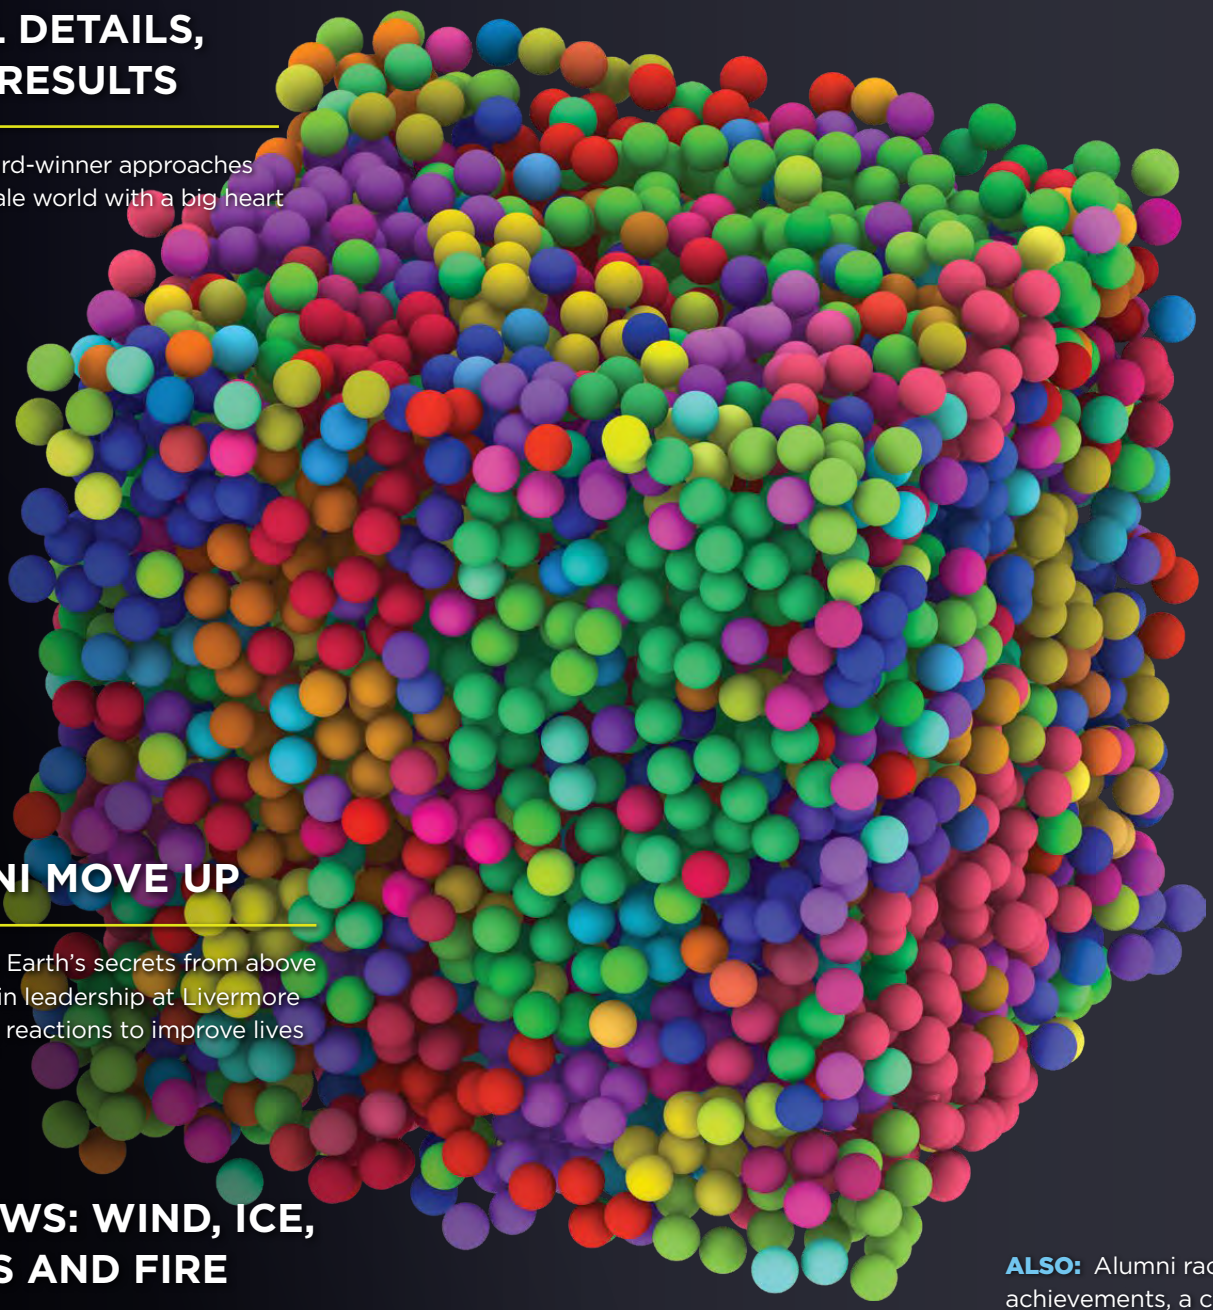

ALUMNI MOVE UP

- Exposing Earth's secrets from above
- A lesson in leadership at Livermore
- Modeling reactions to improve lives

PAGE 22

FELLOWS: WIND, ICE, FLUIDS AND FIRE

- Getting the drift on snow bedforms
- Substances that shake off freezing conditions
- How fluids and solids mingle
- Fuel-burning questions answered

PAGE 10

ALSO: Alumni rack up achievements, a climate modeling leader looks back and ahead, and our annual essay contest winner names materials' tunes

DEPARTMENT OF ENERGY

COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

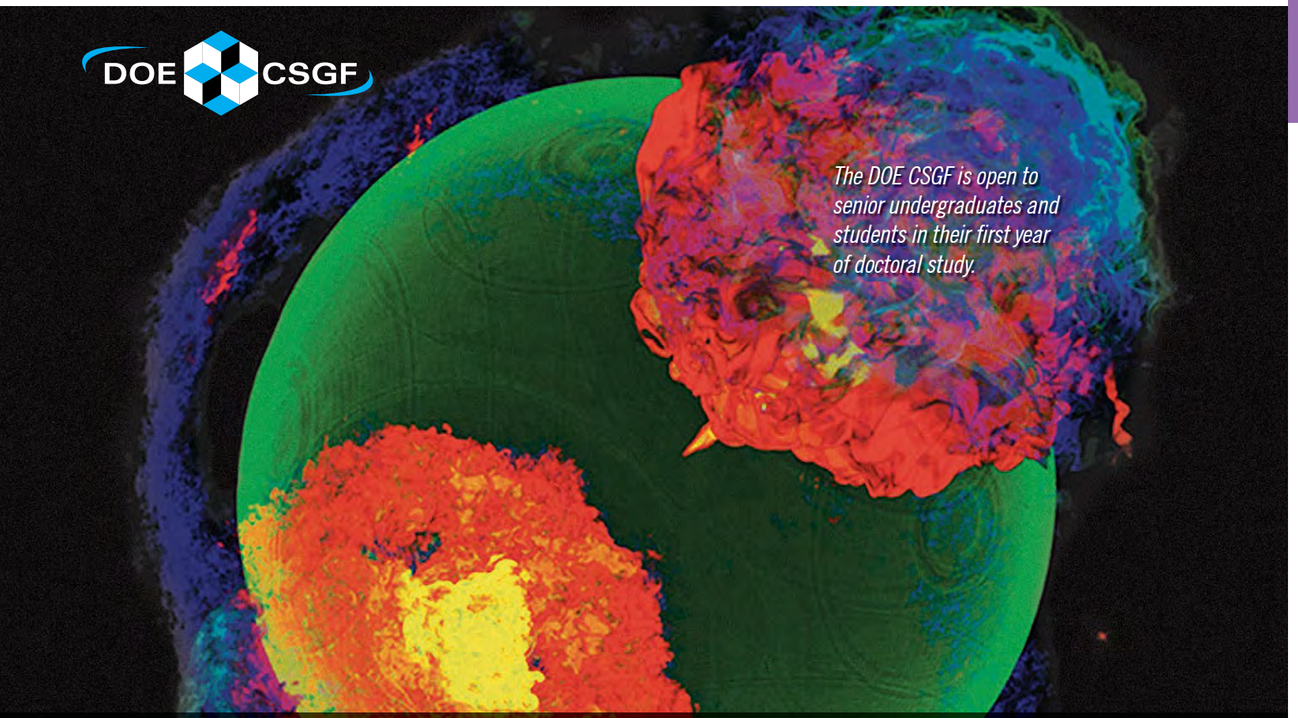

The DOE CSGF is open to senior undergraduates and students in their first year of doctoral study.

The Department of Energy Computational Science Graduate Fellowship (DOE CSGF) provides up to four years of financial support for students pursuing doctoral degrees in fields that use high-performance computing to solve complex problems in science and engineering.

The program also funds doctoral candidates in applied mathematics, statistics or computer science who are pursuing research that will contribute to more effective use of emerging high-performance systems. Complete details and a listing of applicable research areas can be found on the DOE CSGF website.

www.krellinst.org/csgf

BENEFITS

+ $38,000 yearly stipend
+ Payment of full tuition and required fees
+ Yearly program review participation
+ Annual travel allowance
+ 12-week research practicum experience
+ Renewable up to four years

A 3-D visualization of the Gravitationally Confined Detonation model, depicting a Type Ia supernova at the moment when a detonation wave sweeps through a white dwarf star. The simulation was performed on a high-performance computer at the Argonne Leadership Computing Facility, a Department of Energy user facility. Credit: Flash Center for Computational Science, University of Chicago.

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

U.S. DEPARTMENT OF ENERGY | Office of Science

National Nuclear Security Administration

KRELL institute

INCOMING DOE CSGF CLASS

The 2020 Department of Energy Computational Science Graduate Fellowship (DOE CSGF) class includes students applying computing power to subjects ranging from mechanical engineering and plasma physics to biology and the environment. More than 500 students now have entered the fellowship and more than 400 work at national laboratories, at universities and in industry.

Alexandra Baumgart
California Institute of Technology
Mechanical Engineering

Marc Davis
Massachusetts Institute of Technology
Computer Science

Emily de Jong
California Institute of Technology
Mechanical Engineering

Anthony Degleris
Stanford University
Electrical Engineering

Ian DesJardin
University of Maryland, College Park
Aerospace Engineering

Kiran Eiden
University of California, Berkeley
Astrophysics

Ethan Epperly
California Institute of Technology
Applied and Computational Mathematics

Margot Fitz Axen
University of Texas at Austin
Astronomy

Grant Johnson
Princeton University
Plasma Physics

Ariel Kellison
Cornell University
Computer Science

Nikita Kozak
Stanford University
Fluid Mechanics

Mary-Francis LaPorte
University of California, Davis
Plant Biology

Nishad Maskara
Harvard University
Physics

Kaishu Mason
University of Pennsylvania
Statistics

Albert Musaelian
Harvard University
Applied Mathematics

Laura Nichols
Vanderbilt University
Computational Solid State Physics

Graham Pash
University of Texas at Austin
Computational Science, Engineering and Mathematics

Danilo Perez Jr.
New York University
Neural Science

Justin Porter
Rice University
Mechanical Engineering

Luis Rangel DaCosta
University of California, Berkeley
Materials Science

Rachel Robey
University of Colorado Boulder
Applied Mathematics

David Rogers
Stanford University
Earth System Science

Benjamin Sepanski
University of Texas at Austin
Computer Science

Ellis Torrance
University of North Carolina, Greensboro
Environmental Health Science

Margaret Trautner
California Institute of Technology
Computing and Mathematical Sciences

Santiago Vargas
University of California, Los Angeles
Theoretical and Computational Chemistry

IN THIS ISSUE

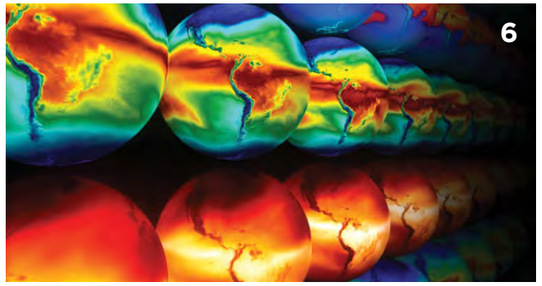

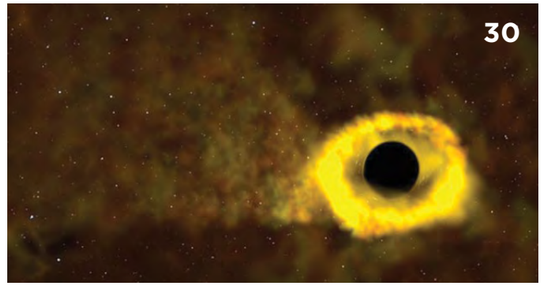

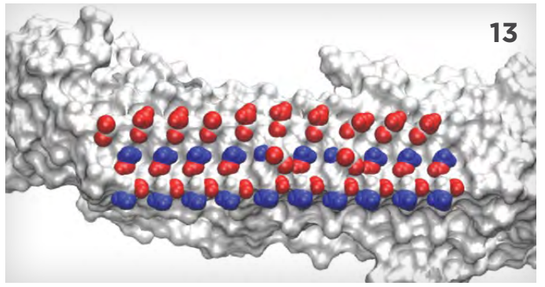

Departments

Features

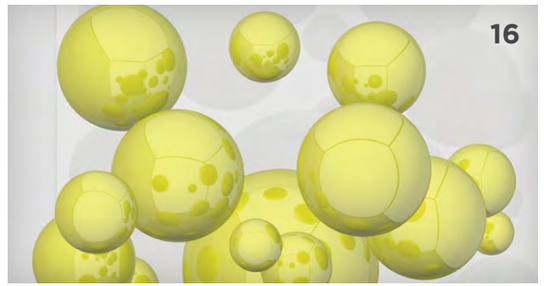

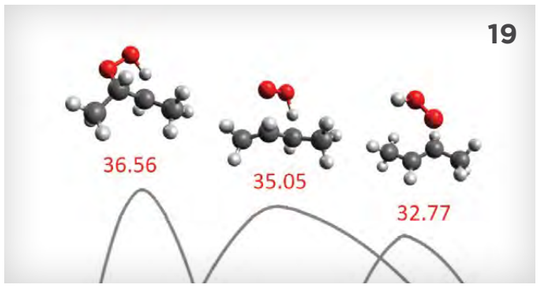

DEIXIS, The DOE CSGF Annual is published by the Krell Institute. Krell administers the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) program for the DOE under grant DE-SC0020347.

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

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
Senior Science Writer — Sarah Webb
Design — Stilt Studio, Inc.

ON THE COVER: Closely-packed particles in a colloidal suspension during friction, from a simulation by 2020 Howes Award-winner Gerald J. Wang. The colors differentiate particle groups whose motions are particularly interrelated, affecting the stickiness of the whole material. Read about Wang's research and community service on page 28. *Credit: Gerald J. Wang.*

CALCULATING CLIMATE

A computational science leader reflects on his career and on shaping the field's top fellowship program.

Credit: Jason Richards/ORNL.

James Hack is the former director of Oak Ridge National Laboratory's National Center for Computational Sciences. In 2007, he founded and directed the Oak Ridge Climate Change Science Institute, which integrates climate modeling, observation and experiments with high-performance computing (HPC). He previously worked for nearly 25 years at the National Science Foundation's National Center for Atmospheric Research (NCAR) in Boulder, Colorado, including as head of the Climate Modeling Section.

DEIXIS: HOW DID YOU GET STARTED IN CLIMATE MODELING?

Hack: In graduate school, I was interested in numerical weather prediction. I wanted to understand the dynamics of tropical storms and hurricanes. Much of this work needed to be done on a computer, and the simulation work got me into the HPC world. Later I worked at IBM Research on computational and engineering problems. But I missed working in the atmospheric sciences, and lo and behold there was a position available in NCAR's global modeling group in 1984. The NCAR group I eventually led built a global modeling and simulation tool, later known as the Community Climate System Model.

WHY IS CLIMATE MODELING SUCH A DIFFICULT PROBLEM?

Climate modeling has some of the same challenges as numerical weather prediction. Forecasters must be able to predict specific events – such as a tropical storm or a tornado outbreak – long before they occur.

With climate, you're studying the state of the Earth's ecosystem on timescales from decades to centuries. You want to run those simulations in a practical time window that keeps you actively engaged in the science. Like almost every computational science problem, we're driven by reducing the time to solution so that answering the underlying scientific questions is tractable.

Climate is also a highly multiscale science problem. Models must consider phenomena that occur from molecular to planetary scales and last from seconds to centuries. That's an up to 10^{12} difference. You can explicitly simulate a portion of the range of motions with a computer system, but there will always be components that you can't resolve and must treat in other ways.

WHAT PRIMARY ADVANCES IN CLIMATE MODELING HAVE YOU WITNESSED OVER YOUR CAREER?

When I first was doing this research, we didn't have the computing power to do more than a four- or five-month simulation for the whole globe. Those models could marginally resolve fundamental atmospheric processes – such as formation of a cold front – that are a thousand kilometers in horizontal scale. Thirty-five years later we can run a global climate configuration, including atmosphere, ocean, land and ice, to simulate decades at resolutions that allow us to begin capturing hundred-kilometer phenomena like tropical cyclones accurately.

But even the fastest computer in the world isn't enough to resolve all the details. And we need to keep making observations to build better models. I'm particularly interested in cloud processes and smaller-scale phenomena like the transport of aerosol particles and water vapor. Clouds absorb and reradiate energy in all directions, so they strongly modulate both parts of the radiation budget. You must get the clouds right if you're going to understand how other factors such as changes in atmospheric CO_2 affect the system.

Observations help us improve parameters that modelers construct for processes that we can't simulate using first principles. Starting in the 1980s, NASA's Earth Radiation Budget Experiment and later the Clouds and the Earth's Radiant Energy System project have measured both outgoing long-wave radiation from Earth and net downward-radiation from the sun. These global observations have been critical to our ability to better represent cloud processes in climate models.

The Defense Meteorological Satellite Program also had a passive microwave channel that let us measure global water-vapor distributions. These observations let us improve our treatments of all phases of water in our models.

HOW HAS CLIMATE SCIENCE DRIVEN PROGRESS IN COMPUTATIONAL SCIENCE?

Computer manufacturers have always been aware of the computational demands of certain disciplines, such as weather prediction and climate. Through the mid-'90s, Seymour Cray was a frequent visitor at NCAR, where we actually had the Cray-3 and Cray-4, basically laboratory prototype machines, on our floor. Why? Because we would break them. The modeling applications would eventually identify hard-to-find problems with both the hardware and the software. Climate was an important test application for Cray and is still important for today's vendors.

We're going through another transition with computer architecture as we come to the end of Moore's Law: We can't boost processing power simply by adding more transistors per square millimeter of silicon. That means HPC architectures will change and become much more complex. No one knows what this will look like, but navigating this change will require working with vendors and discussing the needs of demanding scientific disciplines like climate, among others. Science helps drive computing improvements, and technology improvements have enabled new science. You can't do one without the other.

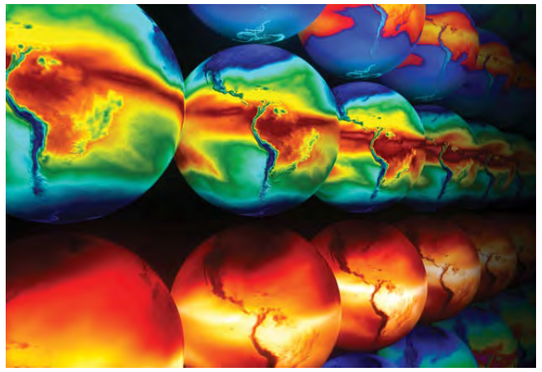

Planetary water-vapor distribution as displayed by the Community Climate System Model (CCSM). James Hack worked on CCSM's predecessor as part of the National Center for Atmospheric Research's global modeling group. He went on to lead computational science at Oak Ridge National Laboratory (ORNL). *Credit: ORNL.*

WHAT ARE TODAY'S KEY CHALLENGES IN CLIMATE MODELING?

A fundamental problem is uncertainty. There's a spread of simulation outcomes for metrics like average global temperature by the end of the century. All the models say global temperature is going up, but the question is how quickly. These changes drive other regional changes to Earth's ecosystems. Economists put numbers on these scenarios, and the costs of climate mitigation and changing energy strategies could cost trillions of dollars. If one model predicts an increase of 1.5 degrees Celsius and another 3 degrees, which one do you bet on?

The climate modeling community realizes that we need to understand why those differences in projections exist. How do we quantify that in a way that informs a policymaker? We must be able to talk about the likelihood of devastating outcomes and look at what-if simulations that could examine available intervention strategies when compared with no action. And the longer we wait the more difficult it will be to prevent potentially catastrophic outcomes. So reducing uncertainty gets back to the problem of doing a better job of representing small-scale processes in the hydrologic cycle – clouds, aerosols and water vapor.

YOU'VE BEEN ACTIVE IN THE DOE CSGF FROM THE START. DESCRIBE THE PROGRAM'S IMPACT.

The program (Department of Energy Computational Science Graduate Fellowship, or DOE CSGF) started with the recognition that many researchers who used computers to advance their science were not formally trained as computational scientists. The skills were self-learned, and there were no academic programs that combined math and computer science with a scientific discipline.

Jim Corones, founder of the Krell Institute, took on DOE CSGF management soon after the program started. He and the program's steering committee came up with a thorough, almost bulletproof selection process. Jim did a great job in guiding the selection committee to whittle down pools of strong applicants proposing computational science research at schools committed to supporting key aspects of the program, such as the national laboratory practicum.

The DOE CSGF has more than 400 alumni, people who are now leaders in computational science. That's the outcome we were hoping for. Some end up at universities, which is truly wonderful because they are training the next generation of scientists. And others are directly involved in research at national laboratories and industry.

Note: This space usually features a keynote speaker from the DOE CSGF Annual Program Review, the fellows' research conference. The meeting was held virtually in 2020; there were no keynotes.

SHAZAM FOR ATOMS

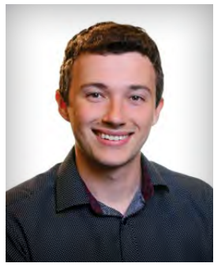

The DOE CSGF Communicate Your Science & Engineering Contest gives fellows and alumni the opportunity to write about computation and computational science for a broad, non-technical audience. The author of this year's winning essay is a third-year fellow studying materials physics at Harvard University.

By Steven Torrisi

One of my favorite phone apps is Shazam. Whenever I'm out in public, I've been known to embarrass my friends by leaping out of my seat, standing on my tiptoes, arms outstretched, holding my phone upside-down toward a ceiling-mounted speaker just for a chance to catch the song that's playing. It's remarkable that Shazam can identify a single tune from the millions of recorded tracks out there just by analyzing a few seconds of notes played through a tinny dive-bar speaker.

Similarly, my most recent research project focused on automatically analyzing scientific data. Like my phone app, the process I studied can use just a snippet of an experimental signal to tell us a lot about a material we're examining. When I explain this method to my friends, "Shazam for atoms" is the first phrase that comes to mind.

Instead of using the arrangement of musical notes to find a song's title, my research studies X-ray patterns to understand the properties of substances. Just as a doctor learns about your internal structure based on how much X-ray your body absorbs (a lot in the bones, less in the soft stuff), the patterns of how materials absorb X-rays give us clues about what's happening inside them on the scale of individual atoms. A lot of time and brainpower goes into detecting minute changes in arrangements of electrons and atoms in matter, because this knowledge is vital to understanding and improving the performance of useful materials. For example, when a battery charges and discharges, the stuff on the inside flexes and flows as electrons are brokered between its components. We can use X-ray absorption spectroscopy (XAS) to spy on these processes as they occur. In short, watching a battery's microscopic choreography can help us design better ones.

Now, Shazam's algorithm – its mathematical instructions for identifying musical recordings – doesn't show you exactly what it is about a song that uniquely identifies it, while your friend who recognizes the same song might say, "that voice is obviously Anderson .Paak." In the same way, researchers use black-box computer algorithms similar to Shazam that do just fine at analyzing experimental data to tell us what we want to know about a material, such as the charge state of a particular element. But it still takes an expert – like your music aficionado friend – to explain *which part* of the data matters and why.

While on an internship at Toyota Research Institute last summer, I set out to bridge this gap, collaborating with Lawrence Berkeley National Laboratory (LBNL) scientists and Matt Carbone, a dear college friend who, like me, is a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient. We set out to devise a procedure that could highlight *which part* of the X-ray absorption data made an analysis algorithm predict one property or another. Two benefits would flow from this: We would have an algorithm whose function was more easily interpretable – so we could gauge if it was working in a trustworthy way – and we could discover trends connecting absorption patterns to interesting properties.

When we shine X-rays on a material, they are absorbed by atoms inside it. And when an atom absorbs X-rays, it's actually the electrons enveloping it that take them in. The atoms' surroundings influence how those electrons behave, including their ability to soak up rays. In XAS, we sweep the energy of the incoming X-rays from low to high, across elements and materials, and see a universal pattern in the resulting curve. First, at low energy, there's little to no absorption. Then there's a sharp spike, and finally a gentle fall-off from the peak, with wiggles up and down as the energy increases and the absorption drops.

Watching a battery's microscopic choreography can help us design better ones.

Just as your friend may only need a snippet to identify a song, sometimes experts only need to examine a sliver of an XAS curve to determine what we want to know about a material. Machine learning (ML, sometimes known by its stage name, artificial intelligence) automates this process. Just as Shazam uses a brief recording to identify a song, ML lets us work backward from X-ray absorption patterns to learn facts about a material.

Most previous ML methods worked this way: Show an algorithm the entire XAS spectrum as well as a material property we care about that leaves a clue in the data, such as how widely the atoms are spaced. Repeat, showing the algorithm enough spectrum-property pairs until it learns to make predictions by itself. The algorithm can use whatever internal means it wants, whether it's easily understandable or not (hence the black box description), as long as it spots the patterns. One ML technique lets us rank which parts of the input matter most in making the prediction, but it can be hard to determine the *trends in the data* that inform the model's guess.

To extend this procedure, we added a pre-processing step, breaking the spectra down into trends instead of individual data points by fitting them to simple curves. This way, we would show the algorithm the *curve parameters*, telling us not only the spectra's important regions, but also which local absorption values (high or low), trends (up or down) or peaks (bump or no bump) counted for the prediction. Just as we hoped, we could work backward from the spectra to properties we cared about and also learn how the model made its call. This helped us discover spectrum-property trends that were either too subtle to attract expert notice or only emerged after integrating data from thousands of spectra.

Just as Shazam needs data from millions of songs to work, our algorithm requires thousands of curated spectra to learn how to accurately identify trends. As such, high-throughput, high-performance computing entirely enabled this work. The Materials Project (an LBNL collective where I had a delightful 2018 DOE CSGF practicum) made our project possible with data it freely shared and open-source software it developed to help us generate more.

With this undertaking, we moved one step closer to one of machine learning's most exciting promises: that it cannot only accelerate our design and discovery of materials, but also help us discover new trends along the way, deepening our understanding.

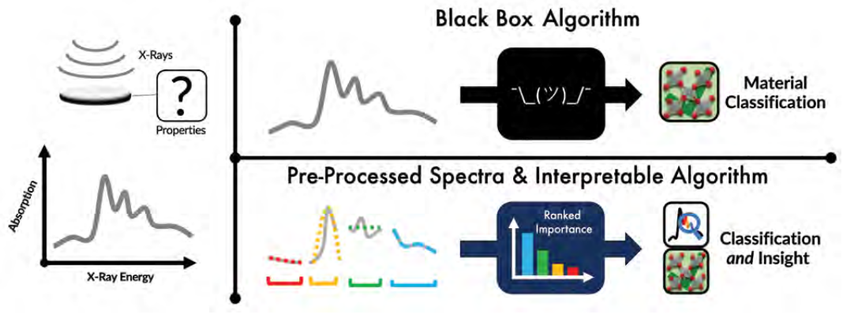

Shining X-rays of different energies on a material (left) yields a characteristic absorption pattern: a spiked edge with wiggles in the post-edge region and occasional bumps in the pre-edge regions. Some previous spectrum-property matching research (top right) used black box algorithms, which solve the problem but are difficult to understand. Fitting simple curves to different spectrum segments (lower right) and providing those to an algorithm that tells us the importance of different inputs lets us not only classify material properties but also better understand what parts of the spectra matter and how. *Credit: Steven Torrisi.*

THE SHAPE OF SNOW

For Kelly Kochanski, a frozen wilderness provides earth science and computational insights – and a source of inspiration.

By Jacob Berkowitz

When Kelly Kochanski dresses to collect her research data, she's all about high-performance layers. For her hands, she first pulls on skin-tight Spandex gloves, then cozy wool ones. Next come the wind-proof gloves, all topped with bulky, down-filled mitts.

"The problem with fieldwork on an exposed 12,500-foot-elevation mountain ridge is that you have to get your hands in and out of layers to manipulate things, and I've found that my hands get almost uselessly numb in less than a minute if the wind's really going," says Kochanski, a fourth-year Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient at the University of Colorado Boulder.

Such is the challenge of conducting pioneering computational snow science. For the past three winters at a site in the Colorado Front Range, Kochanski has collected unprecedented time-lapse images of snow surface topography evolution. She's used them to develop Rescal-snow, the first multiphysics computational model of wind-blown snow form evolution, a dynamic that shapes about 8 percent of the Earth's surface.

She's seeking "elegant ways to represent things I see out in the field that no one else has thought of modeling yet and put them into simulations that make useful predictions."

Picking up her family's penchant for math, Kochanski gravitated to her parents' alma mater, the Massachusetts Institute of Technology. While studying geophysics, she thrived on weekend trips into New Hampshire's White Mountains with the MIT Outing Club's winter school. There she fell in love with "all the wonderful, weird shapes that snow makes," says Kochanski, who spent her childhood in New Jersey and teenage years in Oxford, England.

Searching for Ph.D. programs in landscape evolution, she found ideal mentors in University of Colorado earth scientists Robert Anderson and Gregory Tucker.

"Bob brought up the idea of snow bedforms," Kochanski says, and in it she saw the perfect intersection of math, modeling, the outdoors and beauty.

Snow bedforms are the shifting shapes formed by dry, windblown snow, from arcing dunes to sastrugi, or parallel wave-like ridges. In the past several years, earth scientists have become increasingly aware that snow's topography plays a central role in its energy dynamics and is a missing detail in global climate models.

Although most of us are familiar with the shapes of windblown sand, Anderson says its "self-organization is a significantly simpler problem than with snow. If a snow grain slams into a snowy surface, it has the chance of actually welding to it to essentially become part of a solid. The problem that I posed to Kelly was how does that kind of surface evolve into bedforms?"

Kochanski knew there was only one way to find out. During her first winter in Boulder, she searched the surrounding mountains for an ideal field-research site. She knew she'd found it when she arrived at Niwot Ridge, a dry, tundra-like plateau two miles above sea level, where she watched consistent winds create an endless dance of snow bedforms. A nearby research station could provide weather data.

"It turns out it's really hard to do field work in the snow," says Kochanski, whose site was a 40-minute drive up a winding mountain road from Boulder, followed by a four-mile hike over snow. Wind, altitude and freezing temperatures blocked a third of her weekly visits.

It also was hard on her research equipment, financed through an independent campaign of grant applications. Before she found a waterproof camera and cold-hardy batteries, one set of cameras short-circuited when snow blew into the electronics and another set died when the batteries froze.

But with support from a team of enthusiastic undergraduate field assistants she recruited (and helped use the experience for their own research), Kochanski captured whole seasons in photos shot every 10 seconds, generating more than a thousand hours of unique video. Analyzing it, she produced the first quantitative prediction of formation and evolution under varying weather conditions, on scales from 10-centimeter-long ripples to 30-meter-long whaleback dunes.

"What I documented is that wind-swept snow is mostly not flat," she says of the research published in *The Cryosphere*. "This is critical in terms of the snow's energy balance because when you start roughening the snow you have more surface area that can absorb sunlight, and that makes it melt faster."

She also discovered that wind speed predicts bedform formation, "and that the cutoff is about 7 meters per second.

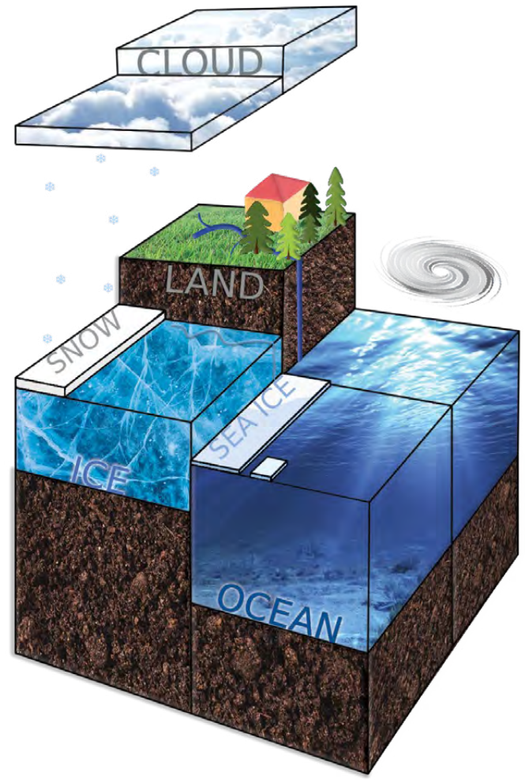

A schematic of Earth's water in its various forms. Fellow Kelly Kochanski and colleagues at Lawrence Livermore National Laboratory are interested in integrating machine learning and physics in Earth system models to refine climate predictions at low computational cost. *Credit: Stilt Studio, after Kelly Kochanski, Donald Lucas, Ghaleb Abdulla, Barry Rountree in the poster "Accelerating Earth Systems Models with Machine Learning" (2019).*

MARINE CORE

As a marine science undergraduate, **Riley Brady** expected to research whales or crabs but turned to ocean current and climate simulations instead. Now at the University of Colorado Boulder, he and advisor Nicole Lovenduski study the California Current, a band along the Pacific coast in which rising waters ferry nutrients to the surface, supporting phytoplankton that fish and other sea creatures eat. But the carbon cycle, in which the ocean absorbs atmospheric carbon dioxide, makes the current corrosive, harming corals and other animals. Brady has used a forecasting system to predict the water's acidity in these systems over a five-year period.

FLUID SITUATION

Fusion energy has intrigued **Brian Cornille** since middle school. At the University of Wisconsin-Madison, he now works with advisor Carl Sovinec to develop simulations of plasma's complex behavior in fusion experiments. The group simulates plasma as a conducting fluid over scales from nanoseconds to hundreds of milliseconds. For his Ph.D., Cornille explored whether an alternate numerical method – first-order system least squares – could help scientists better understand magnetic field evolution in fusion experiments. Next he'll simulate vertical displacement events, potentially catastrophic disruptions in the hot plasma that can damage reactors.

An example of snow bedforms in the Rocky Mountains. The structures are made from dry windblown snow and vary in appearance, from arching dunes and wave-like ridges to the type seen here, a barchan (or "two-horned") dune. Snow surfaces' arrangement and texture affects climate-relevant factors such as sunlight absorption and melting rate. *Credit: Kelly Kochanski.*

This is really useful for anybody who wants to incorporate the effects of bedforms in a larger climate model." She was inspired to think at a global level with her first practicum at Los Alamos National Laboratory, where she worked with DOE's leading sea ice code for prediction across scales, MPAS-Seaice.

As the computational bedrock for modeling snow bedforms, Kochanski uses the sand-dune code ReSCAL, developed by a French group.

"I took the sand model and started adding snow physics, things like the stickiness of snow and the fact that while sand conserves mass, snow can appear from out of the air and sublimate, or disappear into the air."

At SC17, an international computing conference in Denver, Kochanski arranged a coffee meeting with Lawrence Livermore National Laboratory computer scientist Barry Rountree. They'd both been top responders on the Q&A website Quora.

Rountree says that when he learned Kochanski was running ReSCAL on a single processing core in Colorado, he suggested she come to Livermore and its Center for Applied Scientific Computing "and we'll pay you to parallelize" the code so it simultaneously uses multiple processors, enabling bigger, faster simulations.

She came to the lab on a practicum the next summer and discovered that parallelizing Rescal-snow was even harder than slogging up to Niwot Ridge. "It's built on discrete event simulation, a particular class of algorithms that are known to be very hard to parallelize," she says.

Using Quartz, Livermore's 108,000-core supercomputer, she ran run hundreds of full or partial single-core serial runs, slightly varying the physics inputs. This let her, and the team of five undergraduate high-performance computing interns she mentored, debug and optimize the code, port it to different computing platforms and demonstrate that Rescal-snow works for initial conditions. She returned to California in the summer of 2019, working with Rountree and Ghaleb Abdulla in Livermore's Data Science Institute to begin developing a promising machine-learning approach to speed up the code.

By inputting her hard-won field data, Kochanski has used Rescal-snow to simulate for the first time the formation and movement of snow bedforms under numerous wind and snowfall conditions. She'd like to add the physics of snow bedforms to global climate models, but first, after graduation in summer 2020, she needs to land a job. She's hoping it'll be a research position that involves pulling on layers.

SURVIVAL OF THE COOLEST

Cold-hardy animals stunt ice-crystal growth in their tissues. Learning how they do it, Sean Marks and others think, will revolutionize airplane-wing deicing and donor-organ storage.

By Andy Boyles

At first, Sean Marks was frustrated with his model, which was meant to simulate water molecules as they link to form ice crystals on a solid surface. In his animations, the molecules danced, tugging back and forth between interacting with the substrate and their neighboring water molecules. But they refused to settle down.

"Crystal nucleation can be tricky, even when you attempt to utilize sophisticated simulation techniques," Marks says. "In my very first simulations, I got pretty good at making vapor."

Marks's ice-nucleation studies were a new direction for his advisor's team. As a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient, Marks joined Amish Patel's University of Pennsylvania lab in 2015.

Patel and his group had mainly studied how liquid water interacts with materials ranging from hydrophobic – water-repelling – coatings to proteins. They are experts in performing molecular dynamics (MD) calculations, which simulate chemical systems down to the level of atoms and molecules.

Patel was interested in ice nucleation and antifreeze proteins (AFPs). These naturally occurring molecules hold the key to winter survival for many insects, fish and other organisms. Although most animals seek warmth during winter months, these cold-hardy species weather freezing temperatures that can form large, deadly ice crystals inside cells throughout the body. AFPs in their systems stunt crystal growth, helping them survive.

DISEASE, DATA AND DEATH

Kelly Moran seemed destined to follow her father to a career in medicine before being drawn to the mathematics of epidemiology instead. At Duke University, she and doctoral advisor Amy Herring apply dimension reduction, which makes it easier to identify key aspects in data. For one project, she visited Tanzania to study verbal autopsy data – interviews with relatives about symptoms a person exhibited before death. Dimension reduction could help understand these data and predict new deaths. Moran also has applied dimension reduction to toxicology, analyzing information on a chemical's safety at varying doses to forecast outcomes for an untested compound.

INTRON SNIPPING

Kayla McCue, with MIT advisor Christopher Burge, focuses on pre-mRNA splicing, one way cells process genetic instructions after DNA is transcribed to messenger RNA (mRNA). Only exons, portions of these initial mRNA sequences, encode for the proteins that perform vital cellular tasks. Cells edit out the other parts, known as introns, using two different strategies. They can directly identify the introns to be removed or recognize the exons that will remain before snipping out the intron. McCue uses computational models to help understand differences between the strategies and the biological reasons for why they occur.

Scientists think AFPs could lead to a range of lifesaving and cost-reducing advances, such as ice-repelling aircraft wings and improved storage of transplant organs. Additive molecules in today's commercial antifreezes, such as ethylene glycol, prevent nucleation by running interference among water molecules. They work best in concentrations of 20 to 50 percent by weight. AFPs are effective at concentrations of less than 1 percent by – somehow – binding directly to tiny ice crystals and preventing their growth.

Marks's path to this work began with an appreciation of chemistry. His father is a research scientist and his mother works at a major pharmaceutical company. He grew up hearing stories of how science and medicine improve the human condition. In high school, he read about quantum mechanics and was captivated by the idea of a submicroscopic realm and the laws governing it.

He majored in chemical engineering at Northwestern University in Evanston, Illinois. "I heard about this idea of using a computer to directly access things at the molecular scale, and I was like, 'That's pretty neat,'" Marks says. He began undergraduate research in the laboratory of Randall Snurr, who has since become chair of Chemical and Biological Engineering. Snurr and a graduate student introduced Marks to computer simulations. He learned how to program in C++ and wrote a lot of code for modeling how gases interact with solid surfaces.

Marks had offers from several prestigious schools by the time he visited Patel's lab, but he liked the group and the AFP puzzle hooked him. Soon, he was in Philadelphia, tackling the problem.

"This solid phase of water was very new to us," Patel says. "And Sean had no fear, right? He took on the challenge. There were a number of hurdles that he had not anticipated, and he has just taken them down systematically, one at a time."

Other research groups had performed MD simulations of ice nucleation and some basics were known. Marks wanted to see molecules linking and use his calculations to determine how ice-friendly various surfaces are. He coined the terms *icephilic* and *icephobic*, borrowing suffixes from *hydrophilic* and *hydrophobic*, to describe surfaces that are, respectively, more hospitable or less hospitable to ice nucleation.

First, he had to figure out how his simulation could tell the difference between ice and liquid water. The two phases are chemically identical, and since individual water molecules interact strongly, even in the liquid phase two or more will temporarily form locally ice-like structures. He found that viewing a single molecule and its immediate neighbors plus those one step farther away reveals whether the structure is highly ordered (ice) or relatively disordered (liquid water).

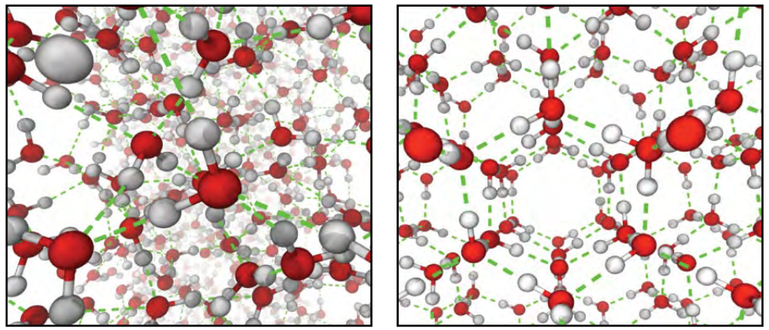

In bulk liquid water (left) and hexagonal ice (right), oxygen atoms are red and hydrogen atoms are white, with hydrogen bonds shown as dashed green lines. Ice exhibits a long-range order characteristic of solids. The hexagonal channels that give the phase its name also are visible. *Credit: Sean Marks.*

Next, he had to induce ice formation on the appropriate time scale. MD modeling yields details down to the level of atoms but typically can simulate activity lasting only nanoseconds at a time. Ice nuclei take much longer to form spontaneously. He introduced a bias into the system, loading the dice in favor of ice nucleation, and now the simulated ice forms within nanoseconds. "We completely sidestep this problem of the waiting game," Marks says.

The bias also leads to an important measure of icephilicity. "In the language of thermodynamics, this is expressed as the free-energy cost of forming a sufficiently large ice nucleus," he says. Surfaces that are more icephilic lower the cost and therefore exhibit shorter nucleation wait times. By measuring the system's response to the applied bias required for ice nucleation, he can compute the free-energy cost of growing the ice – a quantitative and robust measure of icephilicity.

By summer 2018, Marks could trigger ice nucleation at will. "To be able to watch movies that I create, of order forming from complete disorder, that's really cool," he says, no pun intended. Patel honored Marks's animation by showing it at the beginning of his presentations. Patel recalls enjoying it for months as a soothing screen saver.

Marks has used the simulation to study longstanding questions. For example, researchers have speculated that the key to ice nucleation is a surface whose hydrogen-binding sites have the same precise spacing as those of an ice crystal. To test this lattice-matching hypothesis, Marks simulated a family of model surfaces based on the most icephilic surface possible: ice itself.

"If lattice-matching were the end-all be-all, then I could do all kinds of things to this surface, and it would still nucleate ice very readily," he says. But it didn't. When the forces were equal to those of real ice, ice crystals formed. But when Marks reduced the forces, ice crystals formed slowly or not at all. When he cranked up the forces, water molecules piled onto the surface without adopting an ice structure.

Marks's contributions have taken on a life of their own, Patel says, sparking others in the group to take on related problems. "There's this whole class of research that our group is now engaged in, and we all have Sean to thank for it."

Patel also lauds Marks for his generosity. When researchers outside their group began approaching Patel to borrow their new methods, Marks jumped at the chance. "He's helping many of these other research groups go and now ask interesting questions in their own specialized areas that they could not before," Patel says. "I think that's really noteworthy."

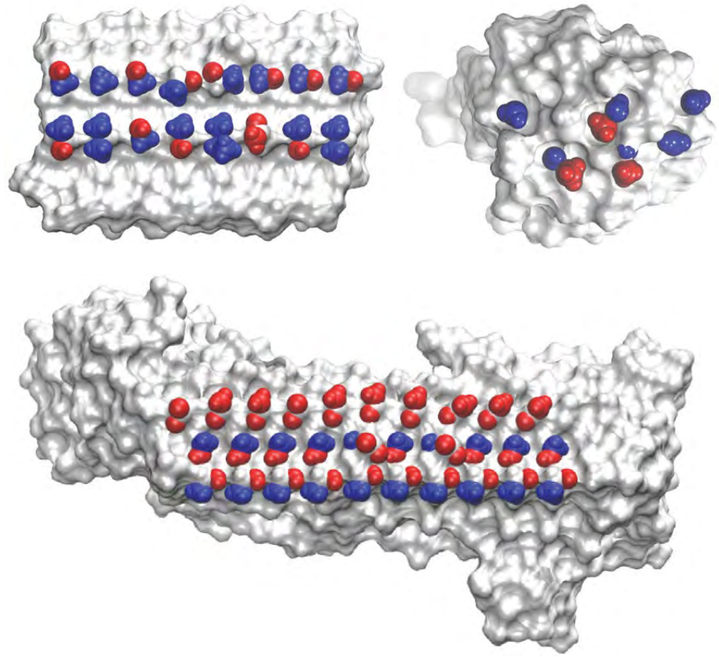

The putative ice-binding sites of three representative antifreeze proteins (AFPs), illustrating the diversity of ice-binding motifs found in nature. Protein side-chain groups are blue, indicating hydrophobic, and red for hydrophilic. Clockwise from top left: a weak AFP from winter rye grass; a moderately active protein from ocean pout; and a hyperactive AFP from the bacterium *M. primoryensis*. Credit: Sean Marks and Saeed Najafi.

Marks credits much of what he has learned to the fellowship, especially the community it has created. At the annual program review, he's shared ideas with other fellows and has learned about many computing approaches from masters at the DOE leadership computing facilities and from fellowship alumni at universities, startups and major tech companies such as Google.

For his DOE CSGF practicum, Marks went to Los Alamos National Laboratory, where he worked with Angel Garcia, director of the Center for Nonlinear Studies. "To be able to peek inside Angel's brain, as well as those of his postdocs, was a real treat," Marks says. As for Marks, "Sean was fantastic," Garcia says. "It was a great honor to have him."

Marks's more recent studies have focused more narrowly on how AFPs interact with water molecules. He is eager to discuss those results but frustrated that he can't, as they're yet to be made public. "Keep an eye on our group's publications, because cool things are coming soon."

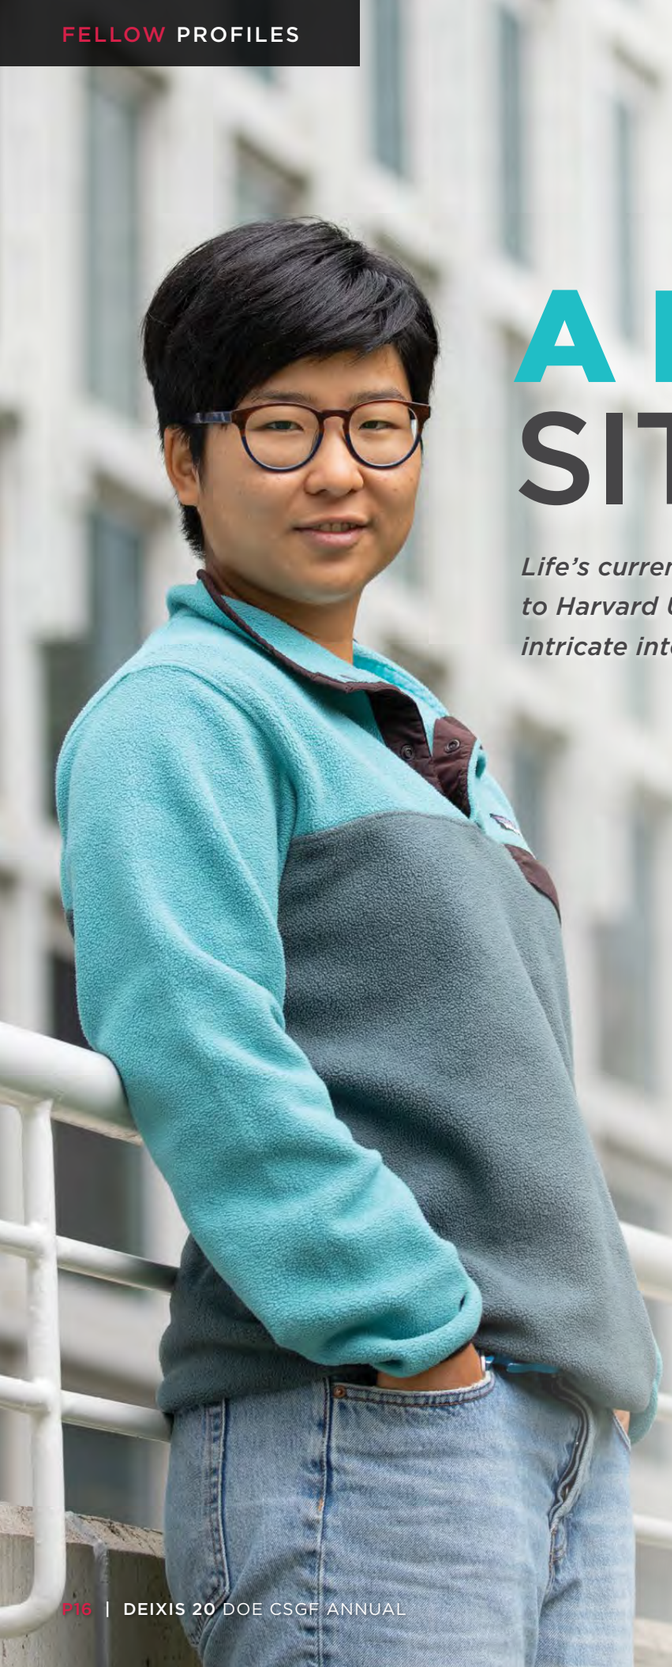

A FLUID SITUATION

Life's currents led Yuexia Lin from her native China to Harvard University, where her algorithms track the intricate interactions of fluids and solids.

By Thomas R. O'Donnell

As a child in her native China, Yuexia (pronounced "YOUsha") Lin would frequently dip into *Ten Thousand Whys*, a book addressing fundamental facts of nature – why the sky is blue and other "really basic, day-to-day phenomena in nature that we wonder about," says Lin, now a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient.

The book didn't actually cover as many questions as advertised, Lin says, but it ignited her interest in physics nonetheless. She's still exploring fundamental phenomena – now with algorithms and high-performance computing (HPC) – as a Harvard University applied mathematics doctoral student. Lin helps improve techniques that probe such things as the interactions between moving fluids and solid objects.

Lin, who uses the name Luna (the first character of her given name stands for moon), didn't set out to pursue doctoral research, and other factors besides books influenced her. She was 18 when she and her mother moved to the United States and rejoined her father, who had emigrated when Lin was four.

Although she had finished high school, a miscommunication forced Lin to repeat classes. For the first two years of school, she worked a low-wage job in New York City, reconditioning videocassettes for television stations to reuse. At night she attended a high school for young adult students, including recent immigrants and those who work during the day.

Meanwhile, Lin produced a documentary of her life with help from a nonprofit organization that helps youths tell their stories through film, leading her to consider a career in journalism and filmmaking.

A summer internship with public radio station WNYC, however, altered Lin's path. At a meeting with other interns, she asked talk-show host Brian Lehrer what to study in preparation for her career.

"He said journalism, good writing, is a craft you learn by doing," Lin recalls, and that young people should study something that would be difficult to master outside academia, such as a scientific discipline. So Lin decided to follow her love of physics, first to Barnard College in New York City, and perhaps work in journalism later.

Harvard applied mathematician Chris Rycroft quickly chose Lin for doctoral studies with his group, but she deferred entry. For a year, "I did a little soul searching, to put it poetically," she chuckles. Working through much of high school and college burned her out, and "I didn't know whether I wanted to spend another five years in school."

Lin worked on financial instrument models for a New York company, but "it didn't feel like I was creative enough or using my brain enough. So I decided grad school was the right call."

Rycroft, meanwhile, always was sure further study was Lin's best course. He invited her on a summer project at Lawrence Berkeley National Laboratory (LBNL), where he's a visiting faculty scientist. The task, which focused on modeling a dissolving solid, was perhaps less important than easing Lin into graduate studies and improving her math and programming skills. "Now Luna is an expert programmer – absolutely expert, masterful – and I think that really stems from those beginnings," Rycroft says.

Lin now largely focuses on algorithms to model fluid-structure interaction (FSI), a phenomenon found throughout nature. Scientists need better methods to address complex FSIs spanning huge differences in scale, such as nanotubes and other sub-microscopic structures floating in vats of fluid. "Without new algorithms that have a baked-in structure to exploit current HPC platforms, these problems will be really, really hard," Lin says.

FSI modeling is difficult because it unites two largely incompatible approaches. In fluid dynamics, mathematicians typically break the modeled region into a regularly spaced fixed grid or mesh of square (for a two-dimensional simulation) or cubic (for three dimensions) cells. At each node or cell, computer processors calculate properties such as the fluid's velocity.

Solid mechanics, however, uses a grid laid over the material. "That grid is now glued onto the solid matter," Lin says, and processors calculate the stretching or compression at each point.

Mathematicians have devised ways to bridge these contradictory approaches. Rycroft and a colleague, Ken Kamrin of the Massachusetts Institute of Technology, called theirs the reference map technique. It expresses the solid mechanics equations completely in the fixed grid normally applied to fluids. "The computational mesh for fluid and solids is united," Lin says.

A fixed mesh lets researchers use standard HPC techniques to solve problems more quickly and efficiently, Rycroft says. The reference map technique also is more effective than previous approaches for certain types of problems, especially simulating biological materials that have complex mechanical properties.

So far, the researchers have used the method only on a two-dimensional grid. They want to extend it to three dimensions

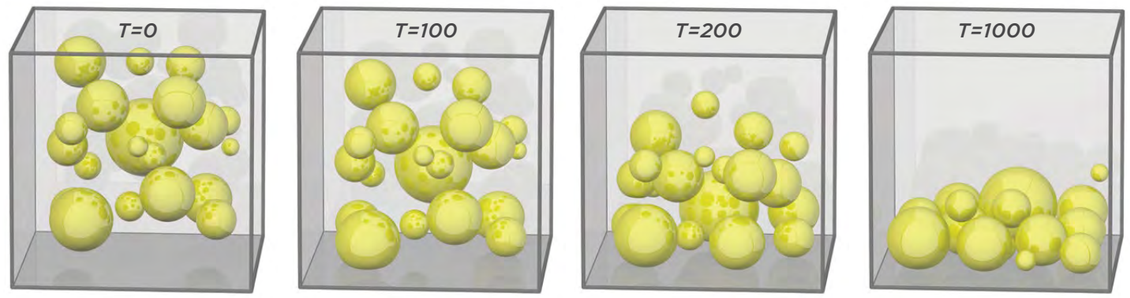

Snapshots from a fluid-structure-interaction animation of three simulated sedimentation stages of 20 variously sized spheres with soft consistencies like rubber and hydrogels, contained in a box with six solid sides. The animation proceeds through time from left to right. At T=200 some spheres are deformed due to contact with each other and the bottom of the box. At T=1000 the spheres have mostly reached equilibrium, although they could still slowly rearrange. The computational resolution is coarse (128 grid points in each dimension), but the simulation reasonably captures interactions between the spheres. *Credit: Y. Luna Lin.*

CHARGING AHEAD

Emily Crabb wants to build better lithium-air batteries, which pack more power into less space than standard lithium ion units. With Jeffrey Grossman at the Massachusetts Institute of Technology (MIT), she addresses lithium peroxide buildup on electrodes and considers properties of electrolytes, substances that conduct ions as a battery charges or discharges. The team has compared the simulations' mathematical methods, seeking the best combination of speed, accuracy and low demand for computer resources.

STAR GRAZER

At Harvard University, **Harshil Kamda**r and advisor Charlie Conroy model 4 billion stars as they form in clusters and dissipate, like a drop of ink spreading as smaller dots through water. By comparing their models with real data tracking a billion stars, Kamdar and colleagues identify weaknesses in the simulations and learn about physics governing star and galaxy formation.

LIGHT SUBJECT

Nicholas Rivera joined Marin Soljačić's Massachusetts Institute of Technology group as an undergraduate. He's still there, studying optical materials and how their atomic arrangements influence or control photon properties. He's shown that nanometer-sized structures could prompt optical materials to emit pairs of quantum-entangled photons rather than just one and has explored how researchers could build table-top X-ray sources to replace miles-long electron accelerators. Most recently Rivera used parallel computing to study whether a beam of free electrons moving through a nanostructure can amplify photons to produce a lightsaber-like coherent light source.

and run it on the newest HPC systems, which yoke millions of processing cores and accelerators to tackle problems previously thought impossible. But that creates multiple issues, particularly for handling communications between processors.

Lin is tackling some of those challenges. Rycroft says she's "made some big steps toward advancing our knowledge in this area."

Lin's quest advanced during 2018 and 2019 LBNL practicums. Her project was to extend the code and test it on a health problem: sleeping sickness. The disease swept parts of Africa in the 20th century, the World Health Organization says, but there were fewer than a thousand new cases in 2018. FSI models can help understand how trypanosomes, the disease-causing parasites, move through the bloodstream so quickly – 30 millionths of a meter, three times its body length, per second.

Lin first worked with Ann Almgren in LBNL's Center for Computational Sciences and Engineering and John Bell, chief scientist for the lab's Computational Research Division, to mate FSI methods with AMReX, an adaptive mesh refinement (AMR) framework. AMR allocates smaller cells to areas of the greatest interest, such as where fluids meet solids, to simulate them in detail. The mesh is less refined elsewhere, conserving computing resources. AMReX is designed for exascale computation, capable of achieving a billion billion (10^{18}) scientific calculations per second.

In her first summer, Lin learned AMReX and studied graphics processing units, computer chips used to accelerate calculations. She ran tests on Summitdev, an intermediate to Summit, Oak Ridge National Laboratory's newest HPC system.

When Lin returned in 2019, she corrected bugs and tested aspects of the code's solid mechanics, running calculations on Cori, a Cray XC40 system at LBNL's National Energy Research Scientific Computing Center. She continued the project after returning to Harvard, collaborating with LBNL postdoctoral researcher Johannes Blaschke, now working in industry, and expects it to become part of her doctoral thesis.

Lin's AMReX education has been ideal, Rycroft says. "We're hoping we'll be able to adapt this reference map technique so it can be done within the AMReX framework. That's what Luna's trying to do with this trypanosome example."

It's not easy, Bell says. Combining such an intricate algorithm with AMR is "a complex problem. She did a phenomenally good job."

Lin says the practicums exposed her to real-world HPC and "made me more ambitious for what kind of science I can do with this method." She had planned to test it on a local computing cluster. Now she can scale it for supercomputers.

Lin may tackle those goals before her projected 2021 graduation – or after. She expects to work as a postdoctoral researcher before taking a university or DOE laboratory post.

Almgren and Bell hope to recruit Lin. Graduate students often toil in isolation, Almgren says, and adjusting to a team setting can be difficult. Lin "was easily able to adapt to working in a collaborative environment and really thrived in it."

Rycroft also prizes Lin's methodical approach. "That's really an important skill for a graduate student – not only being smart but also being able to plan how you're going to approach a problem. Luna is fantastic at that."

CALCULATING COMBUSTION

Sarah Elliott pivoted from atmospheric chemistry to burning fuels, contributing to a new tool that engine designers can use to model combustion.

By Sarah Webb

When Sarah Elliott arrived at Argonne National Laboratory for a 2017 practicum, she envisioned a short summer detour modeling reactions in burning fuels before she returned to the University of Georgia to continue her graduate research in atmospheric chemistry.

But she arrived at the perfect time to embark on an ambitious project. A few months earlier, Argonne's Stephen Klippenstein had received Advanced Scientific Computing Research funding from the Department of Energy (DOE) Office of Science. The goal: develop a simulator into which engineers can plug information about an engine design, fluid flow patterns and a fuel. The code would then predict combustion products, reaction rates and heat capacity, a measure of how molecules absorb thermal energy.

Klippenstein had selected a practicum project: modeling how atoms in heated, combusting molecules writhe around their chemical bonds, then connecting those results to other parts of the calculations.

By summer's end, "I just fell in love with the project," says Elliott, a DOE Computational Science Graduate Fellowship (DOE CSGF) recipient. Combustion is such a direct, important application, she notes, and the field needs improved computational tools. She soon changed course, making combustion the central focus of her Ph.D. research.

Elliott made important contributions to the project, Klippenstein says. "And it was clear that I could use more manpower," he adds. "So we made the plan to have her keep coming out whenever she could manage."

The change was a natural evolution, Elliott says. “Atmospheric chemistry is just very-low-temperature combustion chemistry.” Working on kinetics – reaction rates – is a logical next step, says her Ph.D. advisor, Henry Schaefer, from the molecular electron movement and function modeling that his University of Georgia group performs.

So Elliott combined the Schaefer group’s quantum chemistry methods with Argonne-developed techniques for modeling thermochemistry – the heat energy connected with reactions. Though computational chemists have made great strides in both fields over the years, Elliott says, “they haven’t really communicated with each other that well.”

Last fall she returned to Argonne for her third practicum there. “Every time she comes out we carve out a new part of the project for her to work on,” Klippenstein says. Most recently, Elliott tied together the simulator’s higher-level structure, uniting components various team members developed so they interact seamlessly for users. The simulator follows several steps to explore the molecular energies of reactants, products and the transient structures, known as transition states, that form as fuels burn. The simulator also considers the rapid wiggling and spinning of molecules at high temperatures, up to 2,700 degrees Celsius. Combustion can involve 10,000 reactions and a thousand chemical species, Klippenstein notes. By hand, researchers can only generate complete information about a small fraction of those reactions.

Her practicum wasn’t the first time that a summer research project reoriented Elliott’s science perspective. As an undergraduate at Bethel University in Minnesota, she struggled to choose from among possible majors in math, physics and chemistry. She once considered dropping her chemistry major; the real-life applications fascinated her, but mixing compounds in labs felt formulaic, like following a recipe.

Considering her interests in math and physics, her academic advisor sent Elliott to talk with a quantum chemist, Rollin King. King had completed his Ph.D. in Schaefer’s lab and encouraged Elliott to apply for a summer research position there.

Soon Elliott was driving south, where she would discover an interest in both programming and computational quantum chemistry. She had just completed her sophomore year and was less experienced than Schaefer’s other summer students. Nonetheless, “she outran them all,” Schaefer says.

At Bethel, Elliott settled on a double major in chemistry and physics and a minor in math. She returned to Georgia for a second summer and started her Ph.D. in Schaefer’s lab just a week after completing her undergraduate degrees in 2015. In 2016, she ran the same summer research program where she got her start and mentored three undergraduate students.

Elliott applied for the DOE CSGF because it focuses on interdisciplinary training. “When you’re an undergrad and going on to grad school, everyone is telling you that you have to become more niche,” she says. As a self-taught programmer, the fellowship’s math and computer science coursework provided expertise and confidence to incorporate numerical methods into her code and new techniques and algorithms into her research.

Schaefer’s group mainly models the location and behavior of subatomic particles in atmospheric molecules. Before Elliott’s first practicum, she had described the behavior of a reactive molecule in the atmosphere. Since changing her focus, she has trained her University of Georgia colleagues in the thermochemistry methods she learned at Argonne. “She has become a mentor to other students who want to work on kinetics,” Schaefer says. “She’s really changed the direction of the group.”

After three years, Elliott and her Argonne collaborators are nearing their goal of a complete combustion simulator. She has been a key team member in a field where researchers tend to work independently, Klippenstein says. From her Python programming skills to the detailed work needed to produce milestone data on the project, “she’s been one of the best at communicating with all the different people on the team and at making all the different efforts work together. She’s very good at listening to what the problem is and attacking that problem.”

In March 2019, Elliott spoke about the research at an American Chemical Society national meeting, a notable achievement for a Ph.D. student, Schaefer says. Two months later she presented the work at the 17th International Conference on Numerical Combustion in Aachen, Germany. Engineers there were impressed with her team’s results, she says. Designers have been using approximations of combustion chemistry to devise engines, in part because they didn’t believe it was possible to simulate these reactions efficiently and accurately. “The fact that they were really excited about [these tools] just justifies what we’ve been doing all this time,” she says.

Elliott is interested in more than improving combustion engine efficiency. She also hopes to use these models to track byproducts, such as carbon monoxide, carbon dioxide, and sulfur and nitrogen oxides. These molecules can harm human health or redistribute the atmospheric mix of chemical species, affecting global climate. The methods also can predict unexpected hazardous byproducts before they are generated in a lab.

After graduation in 2020, Elliott will return to Argonne, this time as a postdoctoral scientist to continue her research with Klippenstein’s group.

“It’s very collaborative,” she says. “You’re doing interesting science but with long-term goals that are useful to people.”

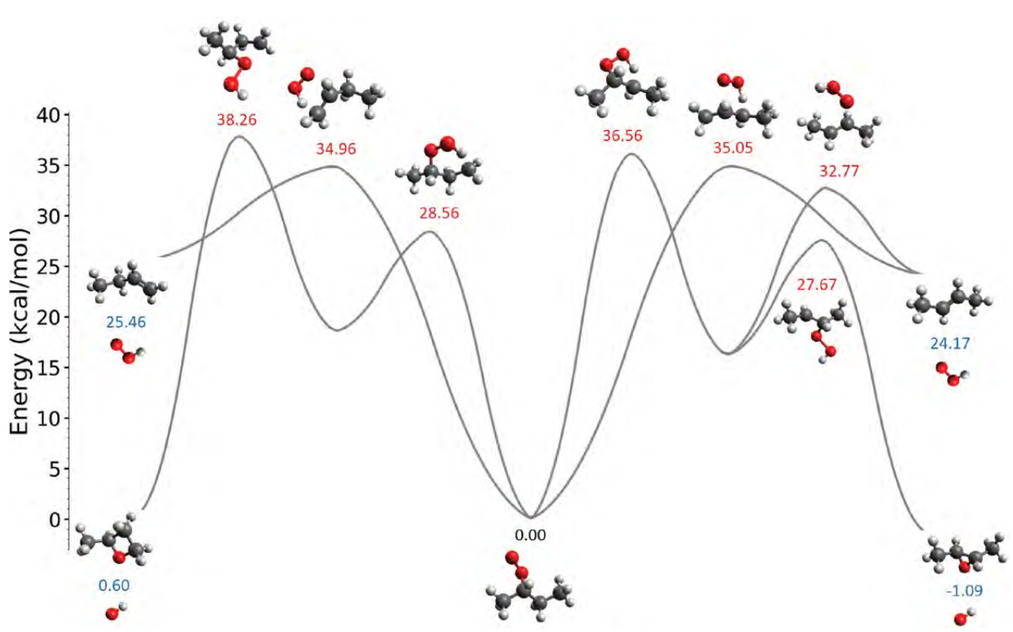

Sarah Elliott and her Argonne National Laboratory colleagues have developed a combustion simulator. Using information about fuel and engine conditions such as pressures and temperatures, their Autochem code computes key information about the combustion products and their energies (colored numbers). This diagram shows just one small slice of those combustion results as the 2-butyl radical (black and white spheres, center bottom), produced from the common fuel n-butane, reacts with oxygen (red spheres). This reaction can follow four different primary pathways (curved lines) with numerous transient species – transition states and reaction intermediates (structures with red numbers) – as it makes four products (structures with blue numbers). *Credit: Sarah Elliott.*

MEMBRANE MAVEN

Amaresh Sahu was a successful software engineer until a documentary about the Higgs boson changed his path. Now he’s at the University of California, Berkeley, working with advisor Kranthi Mandadapu to model cell-membrane deformation. Membranes are layers of lipid molecules that can behave as a 2-D viscous liquid and as a 3-D elastic shell. Most models don’t consider these behaviors simultaneously, but Sahu and Mandadapu found they intricately influence one another. Connecting the two conditions is difficult, so they’re combining known mathematical approaches in new ways using high-performance computing.

QUANTUM PATH

Sukin Sim thought her father’s processor design work was dull but today is deep into a similar field on the edge of the quantum computing revolution. With Alán Aspuru-Guzik at Harvard, she develops circuit designs and algorithms for near-term quantum machines. Sim described a metric called expressibility, a way of defining the number of potential pathways a quantum circuit can use to find a molecule’s low-energy ground state. Some circuits might have a restricted set of states, whereas others might be more expressive. Understanding these features could help scientists design and scale up experiments on coming quantum computers.

FLU FIGHTER

Laura Watkins, with Gregory Voth at the University of Chicago, uses computers to track protein channels and transporters – large, complex molecules that shuttle protons from one side of the cellular membrane to the other. Many important proteins, including M2 in the flu virus, have this function, but it’s unclear how these molecules achieve it. Using special computing methods, Watkins has tracked and examined how protons move and interact with M2 and alter the hydrogen bonding network of surrounding water molecules. Besides improving understanding of proton transport, her work could eventually help develop flu drugs.

Longer versions and additional profiles of fourth-year fellows are available at ***krellinst.org/csgf/fellows/fellow-reflections****.*

LOOKING INWARD

Trained in cosmology, Samuel Skillman gazed at the heavens. Now, with connections made at a national lab, he's using satellites and powerful analytical tools for new views of Earth.

By Thomas R. O'Donnell

As a student and postdoctoral researcher, Samuel Skillman studied distant galaxy clusters. Now the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) alumnus seeks ways to access and analyze information on our home planet.

"The joke was I used to look up toward the cosmos. Now I look the other way around," says Skillman, a fellow from 2009 to 2013.

Skillman is head of engineering at Descartes Labs, a Santa Fe company with origins in nearby Los Alamos National Laboratory (LANL). He joined in early 2015, a few months after the company launched.

In those early days, Skillman was "basically a firefighter. You're putting out fires in the system. You're trying to make things more fireproof. You're also trying to not start any fires." Since then, Descartes Labs has established itself as a platform for geospatial data from government and commercial sources – satellite imagery, weather observations and unconventional sensor information like maritime shipping transponders.

These data are spread across servers worldwide, with varying protocols to access and analyze them. Descartes Labs uses cloud computing to gather and preprocess the information and builds application program interfaces (APIs) for users to retrieve and analyze it. "It's having access to all that data through the same API instead of hundreds of different APIs, latencies and bandwidths," Skillman says.

For example, Skillman was principal investigator for a Geospatial Cloud Analytics (GCA) contract with the Defense Advanced Research Projects Agency. As the name suggests, the program explored how combining commercial cloud computing with public and private remote sensing data can enable new ways to identify global geographic trends. GCA teams used the Descartes Labs platform to access and study geospatial data, addressing such problems as assessing food security and tracking hydraulic fracturing – fracking – for hydrocarbon extraction.

"One of the main enabling technologies is having access to all of the different data sources and modalities," Skillman says, such as Automatic Identification System (AIS) information. AIS transponders report positions for maritime ships, helping avoid collisions, but "also can be used as ways to track nefarious activities," such as vessels going dark as they enter protected fishing areas.

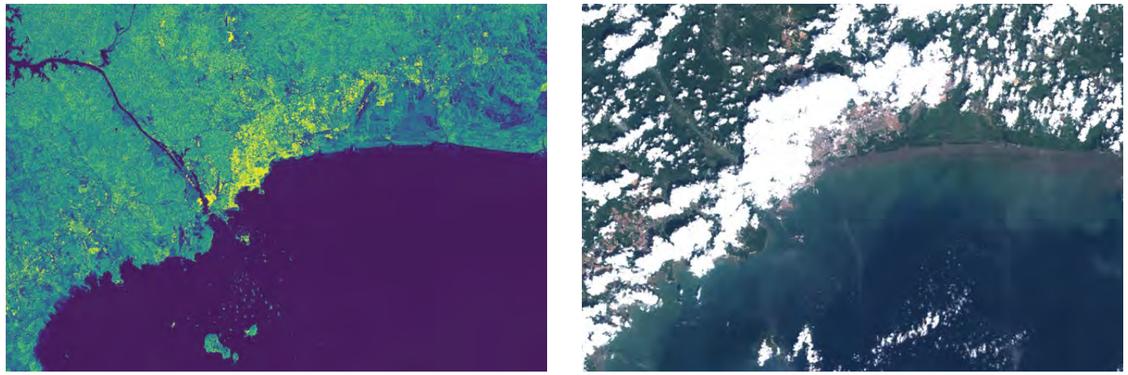

Left: Visible optical satellite imagery from the Sentinel-2 constellation over the southern entrance to the Panama Canal. Right: Synthetic aperture radar (SAR) from the Sentinel-1 constellation, showing polarization fraction over the same area. An example application built on the Descartes Labs platform lets users switch between the two views. While clouds can obscure optical imagery, SAR penetrates them and illuminates large metallic objects – including ships. Optical and SAR imagery can be used in tandem to help identify and track vessels' movements throughout the world. *Credit: Samuel Skillman, Descartes Labs.*

"With the right tools, you can start to identify patterns and develop predictive models," Skillman says. "You might be able to combine that with other data sources, perhaps satellite photos of ship locations."

Meanwhile, researchers also are creating frameworks that make working with geospatial data easier for subject-matter experts. "They may not know how to pull in satellite data and work with that to extract useful signals," Skillman adds, but have knowledge that directs their searches. "Can we build an abstraction that sits on top of our core APIs that lets them access those data more seamlessly and simply?"

The goal: cut the time it takes to develop a hypothesis, access and analyze data, and verify or discard an idea.

'The joke was I used to look up toward the cosmos. Now I look the other way around.'

Skillman is familiar with the discovery cycle. As a University of Colorado Boulder Ph.D. student and a postdoctoral researcher at Stanford University's Kavli Institute for Particle Astrophysics and Cosmology, he simulated galaxy cluster formation and interaction on supercomputers. Each produced huge data sets – similar in size to those Descartes Labs handles – that required new tools to analyze and visualize.

At Stanford, Skillman collaborated with LANL scientist Michael Warren, who had mentored him during a Science Undergraduate Laboratory Internship. Warren left the lab to co-found Descartes Labs and invited Skillman to meet the team.

Although Skillman thought he might work in industry someday, he didn't expect the opportunity to arise so soon. The idea of collaborating with a respected friend and joining a company focused on data and physical phenomena was appealing. "I didn't see that opportunity coming around any time soon after that, so I decided to go for it."

His job gives Skillman the chance to hire and work with smart, engaging scientists and software engineers. It also lets him delve into problems that combine large-scale data and computing with visualization and analysis – an intersection he calls his "happy place."

For example, soon after joining the company he and his colleagues ran a program that used 16 hours on 30,000 cloud-computing cores to take in and preprocess a petabyte of satellite data gathered over a 30-year span. As the number and variety of remote sensing satellites grows, the industry and research community must be ready to process as much as a petabyte every day.

"Thinking about data at that scale is an incredible task," Skillman says. How to help analyze them is "an interesting and challenging problem. I kind of like those."

SCALING HER CAREER

Teresa Bailey's path has grown from code physicist to up-and-coming leader.

By Jacob Berkowitz

It's early on a Friday in mid-December, and Teresa Bailey is enthusiastically describing her upcoming day's schedule at Lawrence Livermore National Laboratory (LLNL). Among a half-dozen major tasks, she's interviewing a prospective new hire for one of her teams and there's a strategic planning meeting for Livermore code development. She's also conferring with a Los Alamos National Laboratory colleague to compare their world-class codes for determining particle transport in physical systems governed by nuclear reactions.

"I enjoy it all," says Bailey, a 2002-06 Computational Science Graduate Fellowship (DOE CSGF) alumna who received her nuclear engineering Ph.D. from Texas A&M University.

This interest in the dynamics of people, projects and particles has shaped Bailey's LLNL career. Over the past dozen years, she's scaled from her initial job as a code physicist to lead teams on two tasks involving diverse groups of scientists and engineers – more than 30 people collaborating to push the limits of predictive science and high-performance computing (HPC).

"I run multidisciplinary projects, and the people on the projects do the research in support of national security," says Bailey, whose career has grown in parallel with LLNL's HPC resources. "It's exciting to take on problems nobody else can and really fulfilling to know that our work makes a real impact."

Bailey completed her Ph.D. in 2008 and was hired at LLNL that February as a code physicist on deterministic transport – simulating and modeling neutron interactions to guide the design of shielding for nuclear reactors, space-based electronics and other structures.

"It was a very small team," Bailey says: just her, mathematician Peter Brown and computer scientist Adam Kunen.

Their challenge was to scale up the lab's deterministic transport code – "when I started, the biggest computation anyone was doing was running on 8,000 processors," Bailey notes – so it would be ready to run in parallel on up to 1.5 million processors on Sequoia, then the lab's leadership-class supercomputer. "We basically decided that we had to switch codes." Five years of collaborative work produced a new LLNL neutron transport production code they named Ardra.

Developing Ardra, Bailey says, "was one of the best experiences of my career because I got to work very closely with two really great people, and we did a job that I think at that time we might have thought was impossible."

The experience also propelled her into a leadership role. Responding to her interest in and aptitude for seeing the big picture and imagining the future, LLNL made Bailey head of an eight-person, lab-wide deterministic transport team with a thermal radiative transfer code group besides the particle transport code. She received a year of on-the-job management training through LLNL's Leadership Institute, including a week on conflict resolution and mentorship.

"Managing science projects is really interesting," Bailey says. "You have to be courageous. The whole job is kind of high-risk because we are folding in everyday research with production. It's fun because everybody learns from each other, and I don't have to go any further than my teams to have a world-class expert view into areas from nuclear physics to computer hardware."

Immediately after developing the new deterministic transport software for Sequoia, her team was tasked with scaling it up for its successor, the 125-petaflop Sierra and its distinctive architecture based on graphics processing units. Because the deterministic project included two different code bases, both had to be ported to the new system.

"It was back to the drawing board," Bailey says in a tone suggesting she relished the challenge. Harnessing Sierra's increased computing power, Bailey's team has helped to make three-dimensional simulations routine. "This is a big deal since it reduces the uncertainties that we have with 2-D and 1-D approximations."

Sierra, the 125-petaflop supercomputer at Lawrence Livermore National Laboratory (LLNL). DOE CSGF alumna Teresa Bailey helped develop a deterministic transport code – used to simulate neutron interactions that guide shielding design for nuclear reactors, space-based electronics and other structures – on Sierra's predecessor, Sequoia. She later led the effort to scale up the code for Sierra. *Credit: LLNL.*

'You have to be courageous. The whole job is kind of high-risk because we are folding in everyday research with production.'

Bailey's ability to coordinate, inspire and facilitate a multidisciplinary team – and deliver results – led to a second management role: In 2016, barely nine months after taking on the deterministic transport project, she became group leader of LLNL's Nuclear Data Project. These two dozen experimentalists and physicists develop the underlying test data and nuclear theory codes that inform the deterministic transport program.

On that Friday morning, Bailey had recently returned from a meeting of leaders from Sandia, Lawrence Livermore and Los Alamos national laboratories on developing next-generation software ready for LLNL's first exascale machine, El Capitan, due for delivery in 2023. It's a big task – and one that she's now ready to help lead.

"Based on our experience with Sequoia and Sierra," she says, "we have continued to evolve and improve our process on how we're going to address the El Capitan challenge. Exascale is a brave new future for us. We must leave no lesson behind."

CATALYST FOR CHANGE

Brenda Rubenstein models energy storage materials and catalysts, designs molecular and quantum computers and works to level inequities in science education.

By Sarah Webb

Chemist Brenda Rubenstein won't settle for science problems that are hard and interesting alone. Though a theorist, she works on practical challenges such as energy storage, new catalysts and molecular computing. Rubenstein, now an assistant professor at Brown University, credits her experience at Department of Energy (DOE) national laboratories with steering her toward grand challenge questions. "Working at the labs showed me a different perspective," she says. "There are societally important problems that need to be solved."

Over the past two years she's gained traction, financial support and accolades for that research. In a recent *Nature Communications* paper, Rubenstein described how molecules can store data – up to gigabytes of images and audio clips. In early 2019, she received a prestigious Alfred P. Sloan Fellowship to support her work. *Chemical & Engineering News* recognized her as one of 2019's Talented 12 chemists, an honor reserved for early-career researchers tackling ambitious and risky questions.

The recognition is "a nice pat on the back," says Rubenstein, a DOE Computational Science Graduate Fellowship (DOE CSGF) recipient from 2008 to 2012, but her core commitment lies in doing the best possible science while helping humanity.

As a high school student, she was interested in history and biology and was a nationally ranked debater. As a Brown undergraduate, science took over. "I realized that if you actually want to make a faster impact and improve people's lives, technology is an excellent way of doing that."

That eagerness, even impatience, also shaped how Rubenstein chose to work. By her early college years, she'd completed several internships in biochemistry and biophysics, fields that can require lengthy experiments. On one job, she killed time while purifying molecules by reading *Crime and Punishment*, *East of Eden* and other classics.

"I moved toward theory, recognizing that it has a faster pace. So if you have a good idea and you can formulate it, you can solve it much more rapidly." Rubenstein wanted to combine equations with the strong scientific principles of chemistry and physics. She majored in chemical physics and applied mathematics and completed her Ph.D. in David Reichman's group at Columbia University.

The DOE CSGF launched her work on quantum Monte Carlo (QMC) methods, tools that help overcome a tradeoff between speed and accuracy. As the number of atoms modeled increases, computation times balloon exponentially. To compensate, chemists often use approximations, reducing computing time and accuracy. To ease this computational burden, QMC employs random sampling. Highly accurate QMC methods help researchers produce phase diagrams, detailed maps of molecules' energies and the behaviors they display under different temperatures and pressures. Building on postdoctoral research she started as a Lawrence Fellow at Lawrence Livermore National Laboratory, Rubenstein designs methods to accurately model molecules and materials for superconductors and other energy applications.

Such calculations remain computationally intensive and require DOE leadership computing resources at Oak Ridge National Laboratory and the National Energy Research Scientific Computing Center (NERSC). Rubenstein and her colleagues work to model transition states, the transient chemical species that occur as reactants convert to products. They want to connect these more-accurate results about species that chemists can't observe to existing experimental data.

Rubenstein and Brown engineer Jacob Rosenstein also are blazing a trail in a nascent field: molecular computing. Small molecules can store a wealth of data just as compactly as organisms do in DNA. Bits can be encoded in the presence or absence of molecules built from combinations of chemical fragments. They read out the bits by measuring the molecules' masses.

Chemistry also can serve as a basis for computation. For example, linear algebra and much of machine learning are built on multiply-accumulate operations. "Molecules can do that pretty much automatically," Rubenstein says. Weighting the amount of different molecule components or a reaction's speed is like multiplication, she says. Products then pool as a reaction progresses. "That's like accumulating," she explains.

'Working at the labs showed me a different perspective. There are societally important problems that need to be solved.'

Besides her research, Rubenstein wants to improve underserved students' access to science education. She attended struggling New Jersey schools in poor districts before landing a spot at a top science magnet high school. "I won the lottery there," she says, and she's determined to help others. Rubenstein now runs two science-education initiatives: the Rhode Island Advocate Program, supported by the Society for Science & the Public; and Rhode Island Project SEED, sponsored by the American Chemical Society. These programs provide paid internships to underserved high school students and help them produce science fair projects.

Improving the diversity of the U.S. science talent pool is important for research quality, Rubenstein adds. Fewer than 10 percent of chemistry Ph.D.s and fewer than 5 percent of physics Ph.D.s were awarded to African-American and Hispanic students, 2015 National Science Foundation statistics show. She notes that science lacks diversity, "and that's not sustainable."

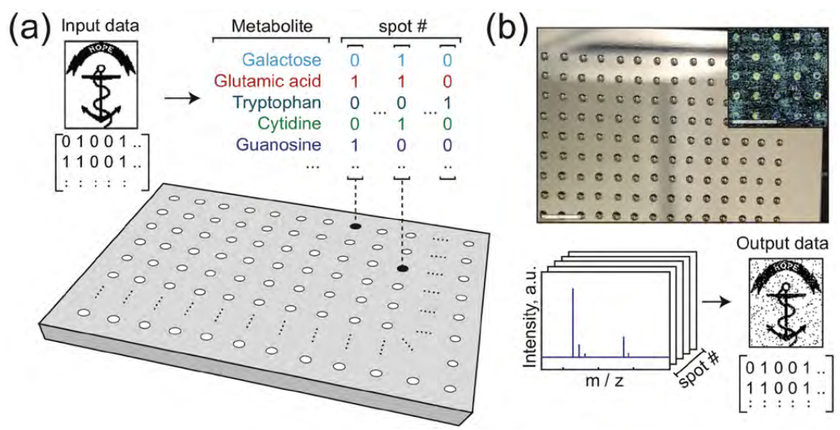

Brenda Rubenstein and Brown University engineer Jacob Rosenstein have been developing approaches for storing digital data using molecular mixtures. The 1s and 0s are encoded based on whether certain molecules are included or excluded from the mixtures. In a 2019 *PLOS ONE* study, they used mixtures (a) of common metabolites – sugars, amino acids, nucleic acids – to encode images as large as 2 kilobytes. These mixtures were dried in spots on a steel plate (b). With mass spectrometry, they could detect the chemical mix in the sample and retrieve the original image (lower right). They built on that work in a *Nature Communications* paper, in which they synthesized more complex chemical libraries from multicomponent reactions and encoded much larger images – including a Picasso drawing. They can now store up to gigabytes of images and audio clips. *Credit: E. Kennedy, C.E. Arcadia, J. Geiser, P.M. Weber, C. Rose, B.M. Rubenstein, et al. (2019) "Encoding Information in Synthetic Metabolomes."* PLOS ONE *14(7): e0217364.*

COMMUNITY BUILDER

Gerald Wang receives the award for his research applying particle-based simulations and for his leadership and empathy on the MIT campus and beyond.

Gerald Wang

By Sarah Webb

Gerald "Jerry" Wang sweats the small stuff. In both his research and in community outreach, he has shown how details can add up in big ways – whether they're interactions between tiny particles or individual acts of kindness and service.

The 2014 to 2018 Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient used particle-based simulations to study hard-to-measure interactions throughout his Massachusetts Institute of Technology doctoral research. By presiding over MIT's Edgerton House and co-creating the Stick With Me! project, Wang supported other graduate students and demonstrated how small acts of compassion and gratitude can make an outsized difference in people's lives. For these contributions to science and society, Wang, now a civil and environmental engineering professor at Carnegie Mellon University, has been named the 2020 Frederick A. Howes Scholar in Computational Science.

Although particle-based simulations and high-performance computing are central to all of Wang's research, he's gradually expanded how he applies them. During his Ph.D. with Nicolas Hadjiconstantinou, Wang focused on the nanoscale, studying questions such as how water molecules pack into layers and slip past each other when confined within spaces only marginally wider than the molecules themselves. Understanding these interactions and patterns can be useful for a range of practical problems, including improving water quality and desalination strategies. During his MIT postdoctoral fellowship with James Swan, Wang moved up in particle size to study the attraction and friction between solid particles suspended in liquids. This material type, called a colloidal suspension, appears in paints, pastes, gels and even foods such as yogurt and cheese.

At CMU, Wang explores problems spanning a range of sizes in which simulating components as individual particles can offer insights. Particles needn't be small, he notes. "They can be anything, as long as a particle is a good descriptor for one discrete entity."

In one project, Wang's team studies nanoscale heat transfer between solids and liquids, which could help researchers devise more efficient cooling strategies for high power-density electronics such as computer processors. He's also examining fluid transport in plants and plant-inspired materials. Biologists know little, he says, about the molecule-scale details of fluid transport through a plant's first-line defense, the cuticular membrane, which admits important nutrients but blocks harmful compounds.

Wang also is excited about taking particle simulations to the human scale in collaboration with his CMU civil engineering colleagues. Particles can represent individual human beings in models of large crowds. "We're studying things like flows of pedestrians in spaces of various geometries and under competing influences: desires for speed, desires for personal space, desires for social distancing."

But Wang doesn't just simulate human interactions. He's shown leadership and empathy in the MIT community and beyond. Wang served four years on the house association and student advisory board at Edgerton House, one of MIT's graduate residence halls, including two years as president.

Besides handling routine tasks, such as planning social events, Wang grappled with larger community challenges, such as improving building accessibility and supporting residents' mental health. "Going back to Edgerton House, it really feels like home to me and a lot of that is that sustained commitment and investment in a community. It matters a lot," he says.

The Stick With Me! project started on the door to Wang's Edgerton House living room, which he shared for three years with mechanical engineering graduate student Nick Demas. Before Wang's first midterm exam, Demas left him a Post-it note message: "Good Luck, Jerry! I think you'll do great today." Wang bombed the midterm, he recalls, but the note made him smile. Over time, the roommates posted a series of supportive notes – as Wang presented research at conferences or Demas pitched inventions to potential investors. After two years, Demas and Wang realized how much joy their sticky note collage had brought to their lives. "If every single day you see this little mural of positivity," Wang says, "you're going to feel a lot more positive over time."

Late one night the two engineers wondered if their mutual support strategy could alleviate some of the isolation and cumulative stress that all graduate students face. With support from MIT's MindHandHeart Innovation Fund, Wang and Demas launched Stick With Me! They created a postcard series that included pictures of encouraging sticky notes along with campus locations important to MIT graduate students. They organized programs in graduate residences. They challenged participants to take a specially designed Post-it pack and share encouragement and gratitude with others. Over a three-day period in spring 2017, MIT students posted more than 4,500 encouraging notes on the large three-story columns in the lobby of the main campus building.

"We were floored" by the enthusiasm and participation in the public art display, Wang says. "I've had some fun experiences scaling up code for research, but this was a whole other kind of scaling up." The original postcards have reached far-flung locations, including Portugal, South Africa and China.

In recognition of his service to the graduate student community, Wang received MIT's Edward L. Horton Fellowship Award in 2017.

Wang also is known for his enthusiasm and wit. As a Yale undergraduate he wrote for the campus humor magazine, The *Yale Record*. For several years he has participated in BAHFest, a science comedy event celebrating bad *ad hoc* hypotheses. In 2019, for example, Wang pitched the numerous benefits of SURF-N-TURF, a bio-inspired amphibious naval vehicle powered by chickens.

Wang praises the support the DOE CSGF has provided his career through human connections and collaboration. "What's so special about the DOE CSGF is the community," he says. The annual program review and other informal opportunities encourage fellows to engage with others, even those doing very different science, throughout the year. "That's the one aspect of the program, that – head and shoulders above the rest – made a difference in my life and that I'm endlessly grateful for."

ABOUT FRED HOWES

The Frederick A. Howes Scholar in Computational Science award, first presented in 2001, 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.

FROM *NATURE* AND *SCIENCE* TO EARLY-CAREER AWARDS

Former fellows advance their professions and computational science.

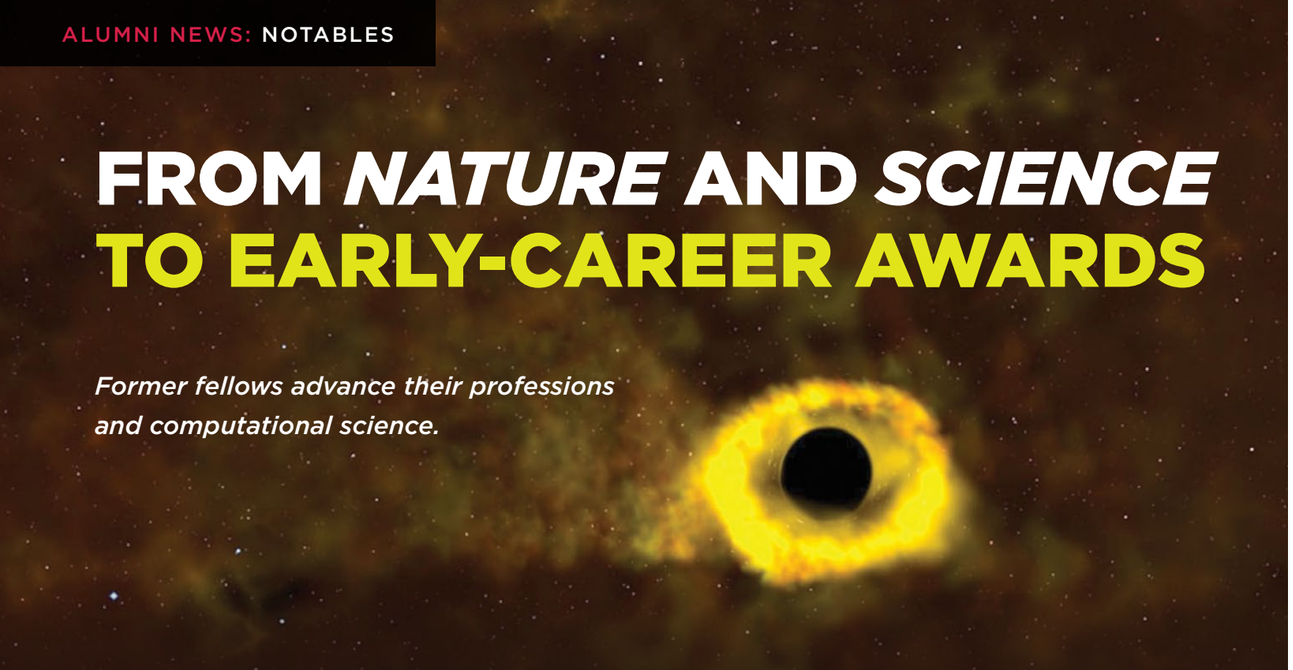

Kristen Grauman (2001-2005) was elected a fellow of the Association for the Advancement of Artificial Intelligence. She is a computer science professor at the University of Texas at Austin.

Oregon State University honored alumna **Teresa Bailey** (2002-2006), a code development team leader at Lawrence Livermore National Laboratory (LLNL), with the Council of Outstanding Early Career Engineers award. She also appeared in a video highlighting how the lab's upcoming El Capitan supercomputer will enable materials science research.

Bailey and **Jeffrey Hittinger** (1996-2000) played prominent roles in a video that debuted at the 2019 International Congress on Industrial and Applied Mathematics in Valencia, Spain. It described how LLNL's Center for Applied Scientific Computing, which Hittinger directs, uses high-performance computing for research in the interest of national security.

A Department of Energy (DOE) Office of Science web post highlighted work on Antarctic ice sheet modeling, one of the projects **Daniel Martin** (1993-1996) works on at Lawrence Berkeley National Laboratory (LBNL). Martin, head of LBNL's Applied Numerical Algorithms Group within its Computational Research Division, also was named head of the Exascale Computing Project's (ECP) Earth and Space Science Application Development focus area.

Julian Kates-Harbeck (2014-2018) was lead author of a *Nature* paper describing an artificial intelligence code that can help predict damaging disruptions in fusion energy reactors. The development could be key to enabling the production of clean, safe and plentiful energy. Oak Ridge National Laboratory (ORNL), where some of the calculations ran, and Harvard University (Kates-Harbeck's graduate institution) touted the research in web articles. Kates-Harbeck's deep-learning code will be among the first to run on Argonne National Laboratory's Aurora, DOE's first exascale machine, as part of its early science program.

Amanda Randles' (2010-2013) computations contributed to research featured on the cover of the journal *Science*. Randles and her Duke University group modeled structures that Rice University bioengineers developed via a three-dimensional printing process. The structures mimic capillaries surrounding air sacs in the lungs. SC19, the international supercomputing conference, also interviewed Randles for its website, covering her innovative biomedical computer models. She also was selected for the National Science Foundation's 2020 Faculty Early Career Development (CAREER) Award, which recognizes outstanding young faculty. It comes with a five-year, $500,000 grant that she and her team will use to build and improve supercomputer simulations of how fluids interact with various cellular structures in the human body.

A Carnegie Mellon University video touted alumnus **Zachary Ulissi** (2010-2014) and his research group's early processor time allocation on Perlmutter, the newest computer at the National Energy Research Scientific Computing Center (NERSC), when it begins operating this year. Ulissi, an assistant professor of chemical engineering, and his team will use its time to accelerate searches for new materials that catalyze critical chemical reactions for renewable energy.

In a *Nature* paper, **Anubhav Jain** (2008-2011) and colleagues demonstrated a machine-learning algorithm that can find and predict new materials with useful properties based on an analysis of millions of scientific publications. The *Vice* website covered the advance.

Bacteria that **Tal Danino** (2006-2010) and colleagues altered could help destroy cancer tumors from within. In a *Nature Medicine* paper, Danino, a Columbia University assistant professor of biomedical engineering, and his collaborators described how an engineered, non-disease-causing organism colonized tumors in mice and released tiny antibodies that made the cancer susceptible to immune system attack. The university also featured Danino's research and his artful bacterial colony images in its magazine.

Three alumni received the 2019 Presidential Early Career Award for Scientists and Engineers. The U.S. Intelligence Community, a federation of 17 agencies, nominated **David Markowitz** (2005-2009), a program manager at the Intelligence Advanced Research Projects Activity, for his "innovative research at the intersection of neuroscience, machine learning and high-performance computing." The National Science Foundation (NSF) nominated **Christina Payne** (2003-2007), director of its Molecular Separations Program in the Division of Chemical, Bioengineering, Environmental and Transport Systems. She also is an adjunct associate professor of chemical and materials engineering at the University of Kentucky. Her PECASE selection is for research examining how a class of enzymes decomposes biomass into sugars, for which she also received a 2016 NSF CAREER award. The NSF also nominated **Alejandro Rodriguez** (2006-2010), a Princeton University associate professor of electrical engineering, recognizing his theoretical and computational research in nanophotonics – the study of light in artificial materials with features on the scale of electromagnetic waves. Markowitz also was named manager of the Intelligence Advanced Research Projects Administration's FELIX (Finding Engineering-Linked Indicators) program. It focuses on detecting engineered biological systems that have been accidentally or deliberately released into the world.

The American Physical Society honored **Norman Yao** (2009-2013) with the George E. Valley Prize, recognizing his outstanding scientific contributions to physics as an early-career researcher. Yao, an assistant professor at the University of California, Berkeley, was chosen for his work elucidating non-equilibrium quantum phases of matter and for enabling the realization of these phases in quantum optical systems.

Bioprinting research from Rice University produced a scale model of a lung-mimicking air sac, with airways and blood vessels that never touch yet still provide oxygen to blood cells. DOE CSGF alumna Amanda Randles and her Duke University group modeled the structure. *Credit: Jordan Miller, Rice University.*

ORNL highlighted the research of **Ethan Coon** (2005-2009) in a web article. Coon, a computational hydrologist at the lab, models complicated ice and soil physics.

Thomas Holoien (2014-2017), a fellow at the Carnegie Observatories in Pasadena, California, led the observation of a rare astronomical phenomenon – a black hole destroying a star – and its progress. The team's results were published in *The Astrophysical Journal*. Holoien explained the discovery in a video for San Francisco television station KGO.

Jarrod McClean (2011-2015), a senior research scientist at Google, is part of a company research team that reported achieving quantum supremacy, using a quantum computer to solve a problem that is virtually impossible for a standard computer. He's one of 70 authors on a *Nature* paper, a list that includes scientists from ORNL and other institutions. The group reported that its quantum processor took about 3 minutes and 20 seconds to sample one instance of output from a pseudo-random quantum circuit a million times, a task that would have taken a state-of-the-art classical computer about 10,000 years.

In a paper presented at the Institute of Electrical and Electronics Engineers High Performance Extreme Computing

Above: When a star comes too near a black hole, intense tides break it apart into a stream of gas. The tail escapes the system, while the rest swings back around, surrounding the black hole with a disk of debris. This depicts a tidal disruption event called ASASSN-19bt. Alumnus Thomas Holoien was part of the team that spotted the disintegrating star in sky surveys. *Credit: NASA Goddard Space Flight Center.*

Conference, **Jeremy Kepner** (1993-1996) and colleagues from the Massachusetts Institute of Technology described a model that captures what internet web traffic looks like on a given day. Using a dataset of 50 billion data packets exchanged over the network for several years, the researchers trained a neural network of processors that captured relationships for the links it contained.

The National Institutes of Health awarded **Ashlee Ford Versypt** (2006-2010) a five-year, $1.8 million grant to develop computational tools for investigating how different diseases and pathogens affect the body's balance between constructing and decomposing fibrous structures. Conditions such as cancer, osteoporosis, arthritis and fibrosis of various organs can be traced to this imbalance. And the Chemical Engineering Division of the American Society for Engineering Education chose her for the 2020 Ray W. Fahien Award, recognizing her vision and contributions as a faculty member in her first 10 years.

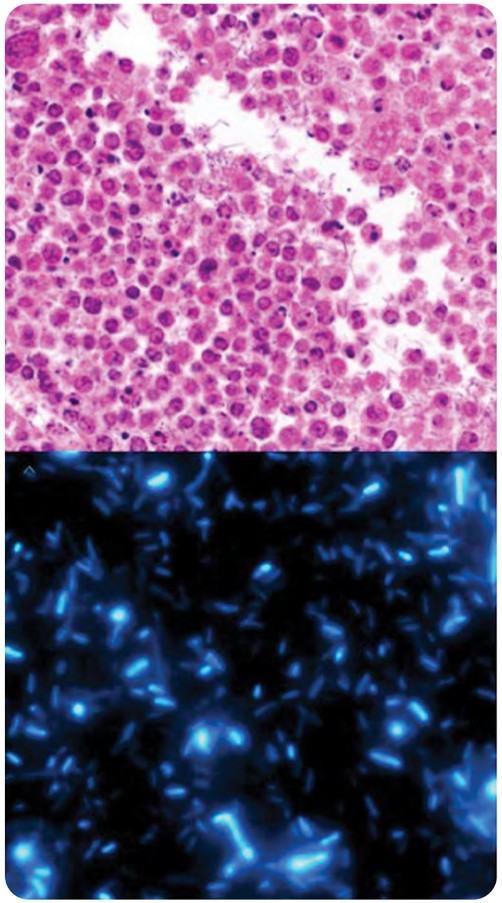

Microscope image (top) of bacteria growing within necrotic regions of lymphoma tumors. Alumnus Tal Danino is researching how to program bacteria to undergo waves of growth and self-destruction, leading to immunotherapeutic release (bottom). *Credit: Danino Lab, Columbia Engineering.*

The Society for Industrial and Applied Mathematics Activity Group on Supercomputing presented **Edgar Solomonik** (2010-2014) an award for outstanding early-career research. Solomonik, an assistant professor of computer science at the University of Illinois at Urbana-Champaign, received the honor at the 2020 SIAM Conference on Parallel Processing for Scientific Computing. The group cited his work on communication-avoiding algorithms for numerical linear algebra problems and on tensor contraction algorithms and software.

Carl Boettiger (2008-2012), assistant professor in environmental science, policy and management at the University of California, Berkeley, and colleagues used artificial intelligence to calculate the optimal choices to maintain sustainable fisheries. The work appeared in the *Proceedings of the National Academy of Sciences*. In April, the Ecological Society of America named Boettiger an Early Career Fellow, recognizing his "emerging leadership in the role of stochastic dynamics in both population ecology and open science."

Alex Perkins (2007-2011) of the University of Notre Dame coauthored a paper that uncovered a previously unknown 2017 outbreak of the mosquito-borne Zika virus in Cuba. The study, published in *Cell*, was covered in *The New York Times* and *Science*.

Alumni **Jack Deslippe** (2006-2010) and **Timothy Germann** (1992-1995) were among seven authors of a report assessing the progress of projects aimed at preparing software tools and algorithms for exascale computers. Deslippe is leader of the Application Performance Group at NERSC. Germann is part of the Physical and Chemical Materials Group at Los Alamos National Laboratory.

LBNL Computing Sciences chose **Aditi Krishnapriyan** (2014-2018) for its Luis W. Alvarez Fellowship in Computing Sciences. During the prestigious two-year postdoctoral appointment, Krishnapriyan will use applied mathematics methods, including machine learning, in a range of areas, such as condensed matter systems. She's at least the fourth program graduate to be named an Alvarez Fellow.

DOE's ECP interviewed **Hal Finkel** (2007-2011) of Argonne National Laboratory for its "Let's Talk Exascale" podcast in March. He leads the LLVM Project, a collaboration to create an open-source compiler infrastructure for HPC systems. Finkel later was a guest on the Code Together podcast, where he discussed programming portability and models.

Stefan Wild (2005-2008), an Argonne National Laboratory computational mathematician, earned a 2020 DOE Early Career Research award. He'll receive at least $500,000 per year for five years to support his project, "Structure-Exploiting, Adaptive, Zeroth-Order Optimization to Improve Efficiency," which will address complex design, decision and control problems.

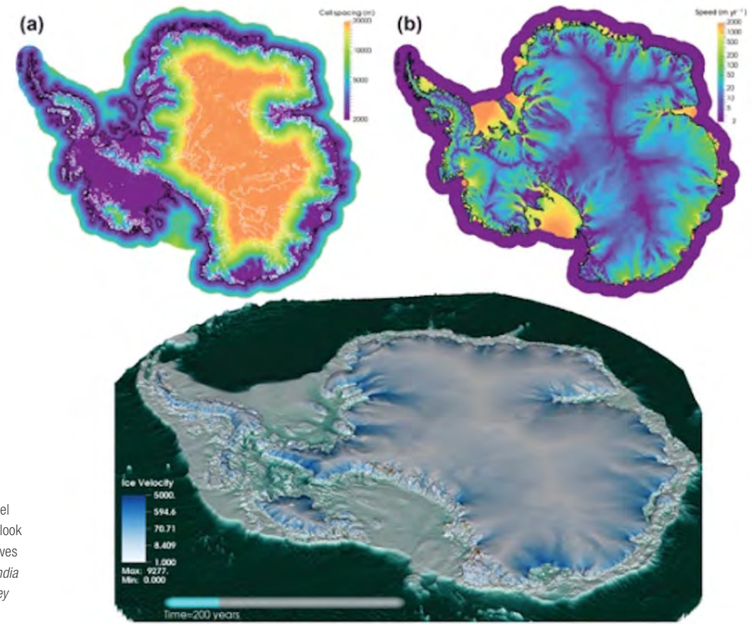

Screenshots from the MALI and BISICLES computer models alumnus Daniel Martin has helped develop, simulating what the Antarctic ice sheet would look like and how fast ice would move 200 years after the sheet's floating shelves have disintegrated. *Credit: MALI model, developed by Los Alamos and Sandia national laboratories and BISICLES model, developed by Lawrence Berkeley National Laboratory, University of Bristol and University of Swansea.*

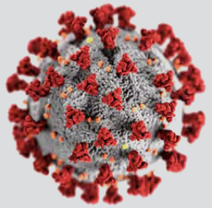

ALUMNI VS. THE PANDEMIC

When SARS-CoV-2, better known as coronavirus, spread to the United States, DOE CSGF alumni were among those lending their expertise to track its progress and address its effects.

Oklahoma State University's **Ashlee Ford Versypt** (2006-2010) was part of a group that – in fewer than 12 hours – created a model of viral dynamics in lung and intestinal tissue. The project's goal was to develop a comprehensive multiscale simulation framework to understand and test interventions in the coupled dynamics of COVID-19, the disease the virus causes. Such a model would let researchers investigate vulnerabilities in viral replication and infectious spread and find approaches to control the immune response that causes adverse reactions.

Meanwhile, other alumni used computational models to project the disease's spread. In March, **Alex Perkins** (2007-2011) of the University of Notre Dame and colleagues estimated the real number of American coronavirus infections at that time probably was at least double the number health officials reported. In April, Perkins, an expert in infectious disease epidemiology and population biology, led a study projecting the consequences of easing restrictions in May – the time when many states did exactly that. The paper's conclusion: Policymakers should have maintained a high level of control measures, such as social distancing and contact tracing, well into the summer to minimize transmissions.

Timothy Germann (1992-1995) normally models material properties, but when the outbreak struck he returned to a project that briefly made him famous 15 years ago: the first large-scale agent-based simulations of a disease spreading through a population. The Los Alamos National Laboratory scientist is adapting the smallpox-modeling code to take on COVID-19. Germann has run hundreds of jobs – remotely, from home – on the lab's Trinitite computer, testing disease scenarios and representative demographics for parts of the nation. He projected the number of cases that could arise in New Mexico if the lab's home state opened all versus some of its schools. "Pandemic spread has the additional complexity of unknowns such as human behavior, policy changes and medical interventions such as vaccines or treatments," he told the *LANL Today* newsletter in June. "By making assumptions about likely scenarios for each of those, we can project how the future disease spread will change."

Such a forecast, **David Ketcheson** (2006-2009) said in a podcast for King Abdullah University of Science and Technology (KAUST), "even when that future is frightening, is powerful and is reassuring." Ketcheson, a KAUST applied mathematics and computational science professor, discussed how researchers in his field view and calculate transmission of the infectious disease.

Credit: Alissa Eckert, MSMI, Dan Higgins, MAMS

CLASS OF 2020

Riley Brady
University of Colorado Boulder
Ocean Biogeochemistry
Advisor: Nicole Lovenduski
Practicums: Los Alamos National Laboratory
Oak Ridge National Laboratory

Harshil Kamdar
Harvard University
Computational Astrophysics
Advisor: Charlie Conroy
Practicum: National Renewable Energy Laboratory

Nicholas Rivera
Massachusetts Institute of Technology
Optical Physics
Advisor: Marin Soljačić
Practicum: Argonne National Laboratory

Brian Cornille
University of Wisconsin-Madison
Nuclear Engineering and Engineering Physics
Advisor: Carl Sovinec
Practicum: Lawrence Livermore National Laboratory

Kelly Kochanski
University of Colorado Boulder
Earth and Climate Science
Advisor: Robert Anderson and Gregory Tucker
Practicums: Los Alamos National Laboratory
Lawrence Livermore National Laboratory

Amaresh Sahu
University of California, Berkeley
Chemical Engineering
Advisor: Kranthi Mandadapu
Practicum: Argonne National Laboratory

Gabriela Correa
Cornell University
Materials Science
Advisor: David Muller
Practicum: Lawrence Berkeley National Laboratory

Brett Larsen
Stanford University
Physics
Advisor: Shaul Druckmann
Practicum: Sandia National Laboratories, California

Andres Salcedo
The Ohio State University
Astronomy
Advisor: David Weinberg
Practicum: Argonne National Laboratory

Emily Crabb
Massachusetts Institute of Technology
Computational Condensed Matter Theory
Advisor: Jeffrey Grossman
Practicum: Argonne National Laboratory

Yuexia Lin
Harvard University
Applied Mathematics
Advisor: Chris Rycroft
Practicums: Lawrence Berkeley National Laboratory
(2018, 2019)

Clay Sanders
Duke University
Civil Engineering/Computational Mechanics
Advisor: Wilkins Aquino
Practicum: Sandia National Laboratories, California

Julia Ebert
Harvard University
Computer Science and Bioengineering
Advisor: Radhika Nagpal
Practicum: Lawrence Livermore National Laboratory

Thomas Ludwig
Stanford University
Chemical Engineering
Advisor: Jens Nørskov
Practicum: Argonne National Laboratory

Sukin Sim
Harvard University
Chemical Physics
Advisor: Alán Aspuru-Guzik
Practicum: Lawrence Berkeley National Laboratory

Sarah Elliott
University of Georgia
Computational Chemistry
Advisor: Henry Schaefer
Practicums: Argonne National Laboratory
(2017, 2018, 2019)

Sean Marks
University of Pennsylvania
Chemical Engineering
Advisor: Amish Patel
Practicum: Los Alamos National Laboratory

Laura Watkins
University of Chicago
Theoretical Chemistry
Advisor: Gregory Voth
Practicum: Los Alamos National Laboratory

Jenelle Feather
Massachusetts Institute of Technology
Neuroscience
Advisor: Josh McDermott
Practicum: Lawrence Livermore National Laboratory

Kayla McCue
Massachusetts Institute of Technology
Computational and Systems Biology
Advisor: Christopher Burge
Practicum: Lawrence Berkeley National Laboratory

Blake Wetherton
University of Wisconsin-Madison
Plasma Physics
Advisor: Jan Egedal
Practicums: Fermi National Accelerator Laboratory
Los Alamos National Laboratory

Claire-Alice Hébert
Stanford University
Cosmology
Advisor: Patricia Burchat
Practicums: Los Alamos National Laboratory
(2018, 2019)

Kelly Moran
Duke University
Statistics
Advisor: Amy Herring
Practicums: Los Alamos National Laboratory
(2017, 2018, 2019)

Cristina White
Stanford University
Mechanical Engineering
Advisor: Charbel Farhat
Practicum: Lawrence Berkeley National Laboratory

Daniel Jacobson
California Institute of Technology
Computational Chemistry
Advisor: Tom Miller
Practicum: Lawrence Berkeley National Laboratory

Ian Ochs
Princeton University
Plasma Physics
Advisor: Nathaniel Fisch
Practicum: Lawrence Livermore National Laboratory

DOE CSGF: A RECORD OF ACHIEVEMENT

The Department of Energy Computational Science Graduate Fellowship (DOE CSGF) builds a community of leaders who apply high-performance computing to problems of national importance. Graduates go on to prominent positions at their respective organizations, as demonstrated in their professional accomplishments. Large majorities of alumni surveyed in 2017 had presented research at national and international meetings, lectured at host institutions or mentored others in the previous five years or since completing their fellowship (if less than five years out). *Source: https://www.krellinst.org/csgf/about-doe-csgf/2017-longitudinal-study.*

Percent of alumni reporting how often they have presented research by setting (N=211)

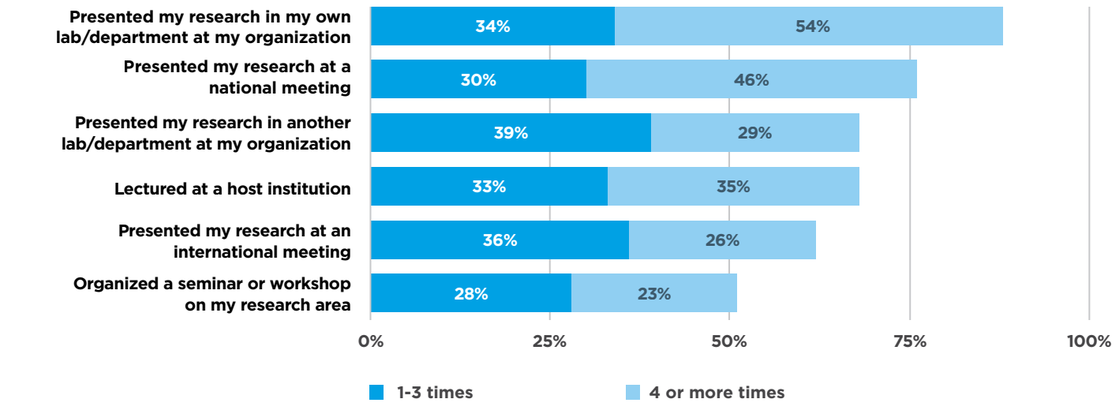

Percent of alumni reporting how often they had made other professional achievements (N=211)

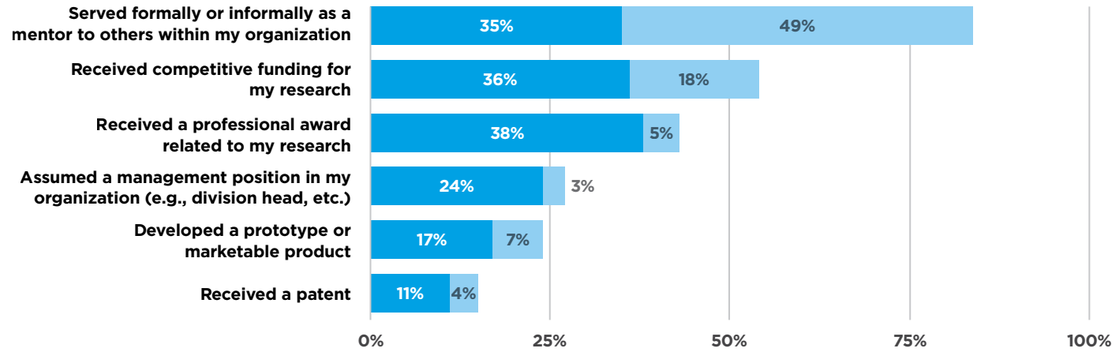

NOTE: Percentages may not sum to the totals referenced in the text because of rounding.

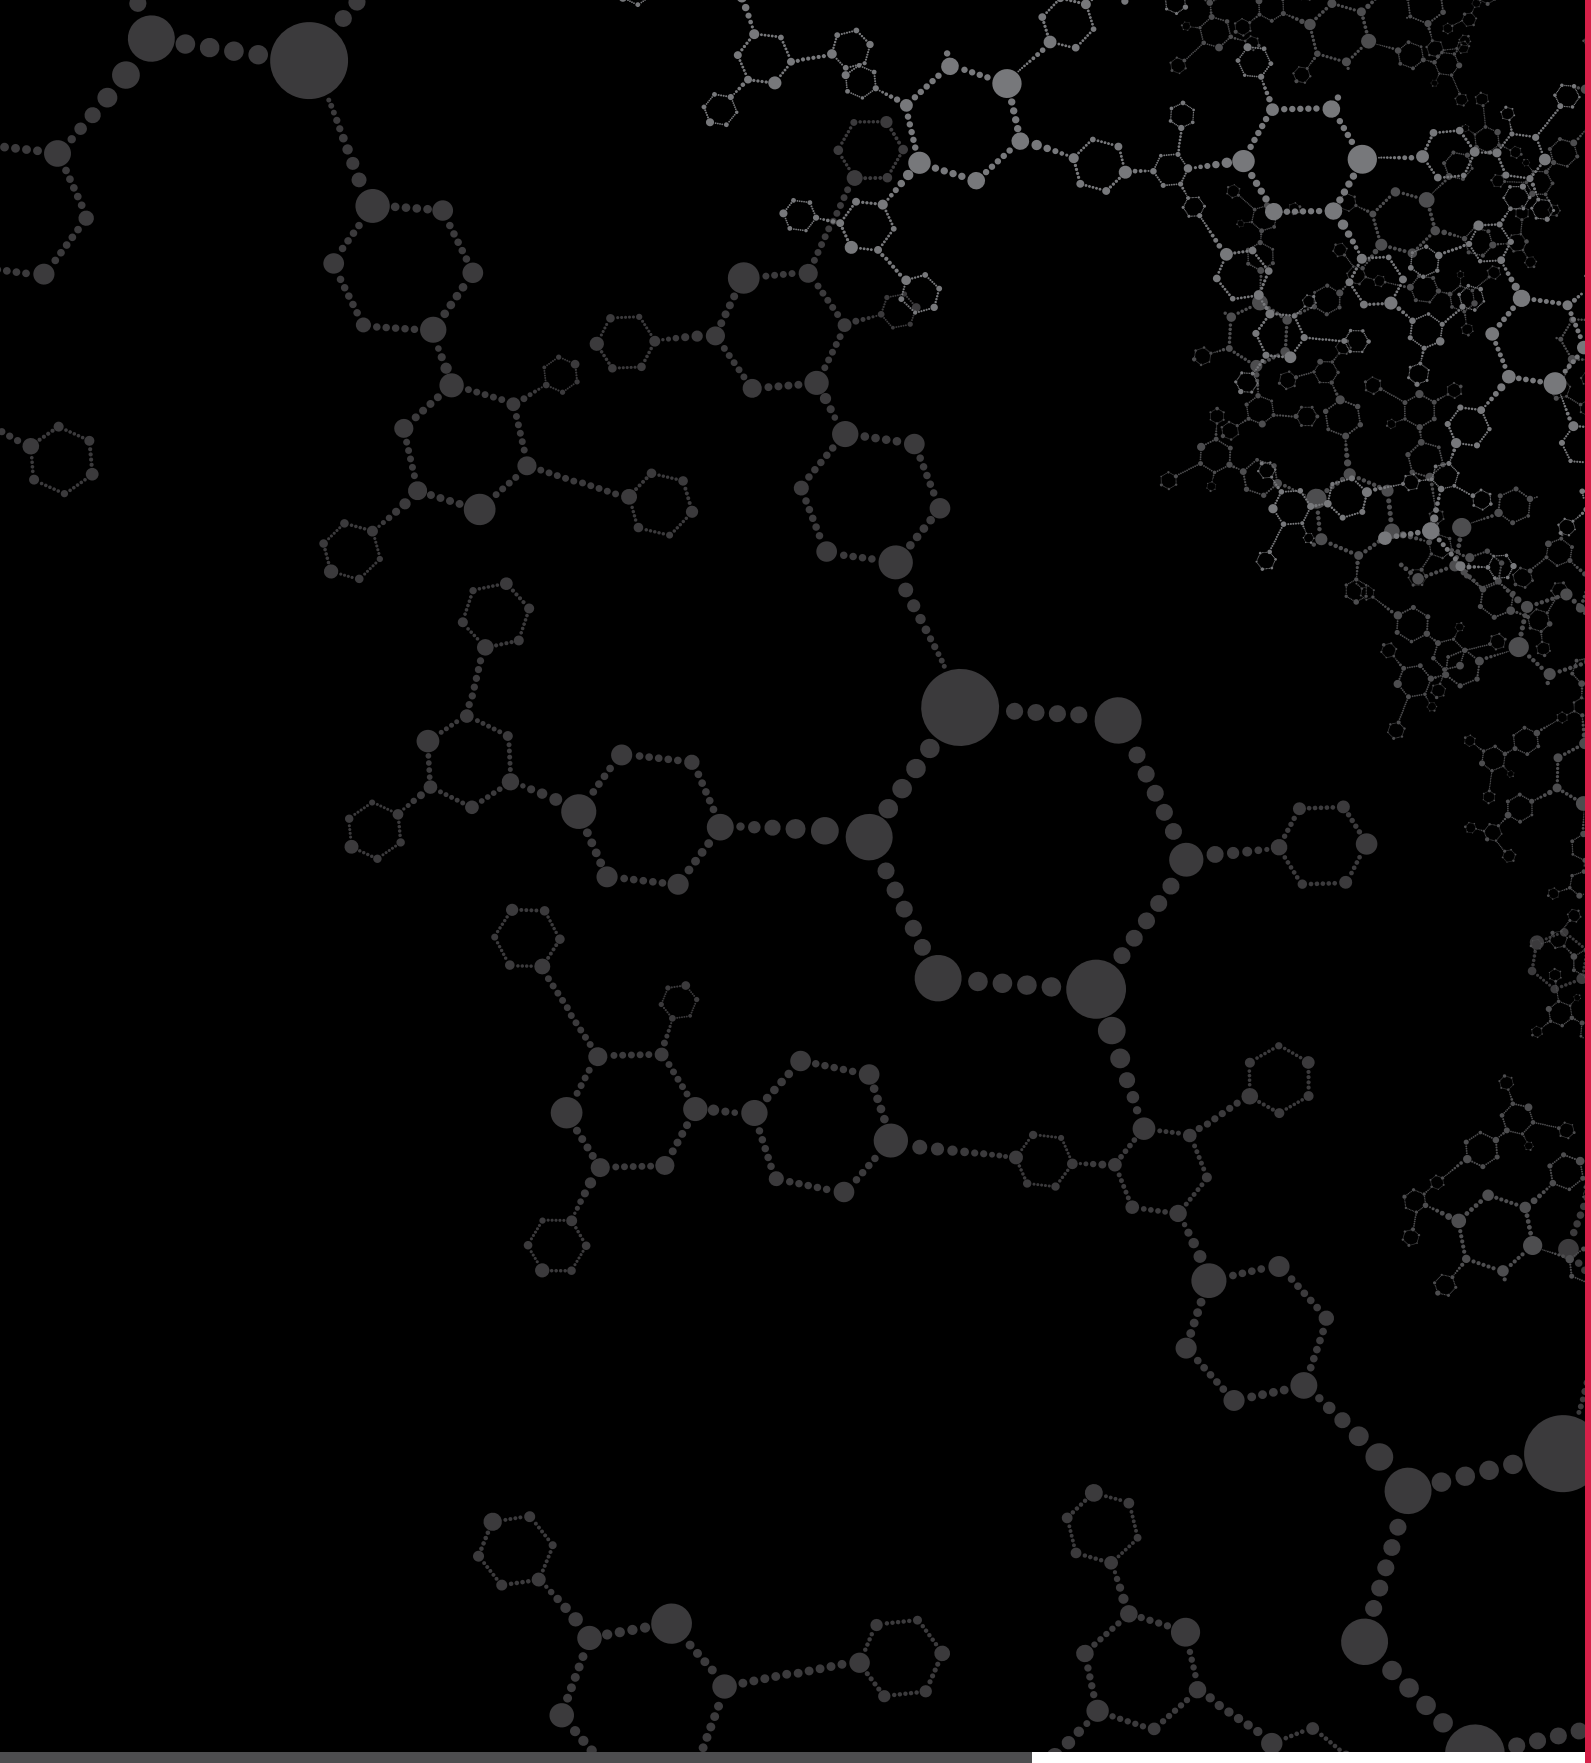

The Krell Institute
1609 Golden Aspen Drive, Suite 101
Ames, IA 50010
(515) 956-3696
www.krellinst.org/csgf

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