THE DOE CSGF ANNUAL

DEIXIS

2019

DEPARTMENT OF ENERGY COMPUTATIONAL SCIENCE GRADUATE FELLOWSHIP

ALUMNI GO BIG

Former fellows are among those leading the quest for exascale, the next high-performance computing horizon

PAGE 22

FELLOWS TAKE ON BOTS, HURRICANES AND MORE

- Casey Berger captures quantum powers
- Julia Ebert wrangles robot swarms
- Max Bremer speeds a surge model
- Carson Kent seeks an ultimate optimizer

ALSO: Alumnus Jarrod McClean on quantum computing, dual Howes Scholars and our essay winner herds computational cows

INCOMING DOE CSGF CLASS

The 29th Department of Energy Computational Science Graduate Fellowship (DOE CSGF) class includes students applying computing power to a range of subjects, such as materials science, neuroscience, climate dynamics and biological oceanography. Like all previous fellows, they'll receive yearly stipends, payment of full tuition and fees, and other benefits for up to four years.

Christopher Balzer
California Institute of Technology
Chemical Engineering

Thomas Blommel
University of Michigan
Physics

Kyle Bushick
University of Michigan
Materials Science

Lindsey Byrne
Northwestern University
Astronomy

Madelyn Cain
Harvard University
Condensed Matter Physics

Gabriel Casabona
Northwestern University
Computational Astrophysics

Scott Emmons
University of California, Berkeley
Computer Science

Nicholas Ezzell
University of Southern California
Physics

Koby Hayashi
Georgia Institute of Technology
Computer Science

Louis Jenkins
University of Rochester
Computer Science

Christopher Kane
University of Arizona
Physics

Arianna Krinos
Massachusetts Institute of Technology
Biological Oceanography

Peter Lalor
Massachusetts Institute of Technology
Nuclear Engineering

Margaret Lawson
University of Illinois at Urbana-Champaign
Computer Science

Kyle Lennon
Massachusetts Institute of Technology
Chemical Engineering

John Lindsey
Columbia University
Computational Neuroscience

Rebekah Loving
California Institute of Technology
Computational Biology

Cole Miles
Cornell University
Theoretical Condensed Matter

Guy Moore
University of California, Berkeley
Applied Mathematics

Jamin Rader
Colorado State University
Climate Dynamics

Michael Toriyama
Northwestern University
Computational Materials Science

Jason Turner
University of Wisconsin-Madison
Mathematics

Amalee Wilson
Stanford University
Computer Science

Boyan Xu
University of California, Berkeley
Computational topology

Claire Zarakas
University of Washington
Atmospheric Sciences

Lauren Zundel
University of New Mexico
Physics

IN THIS ISSUE

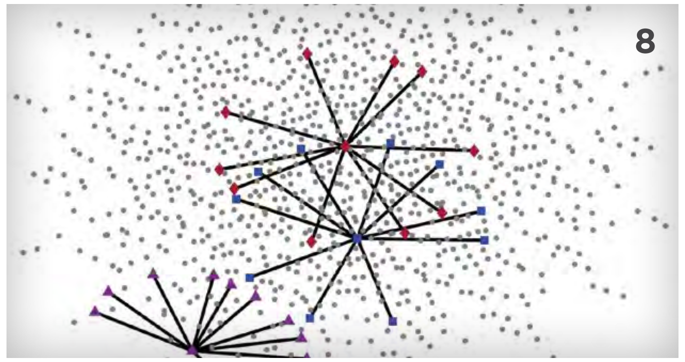

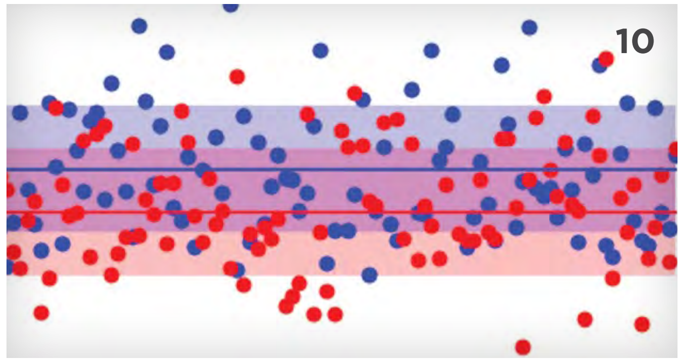

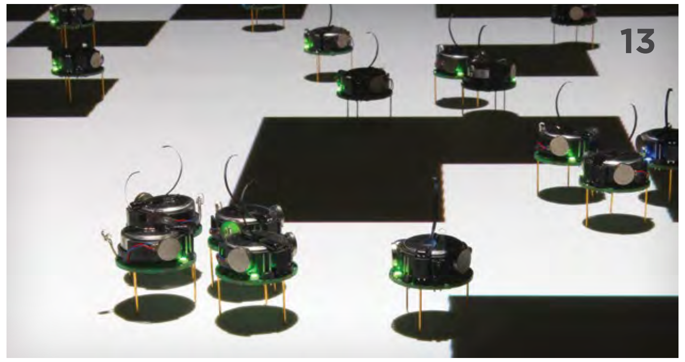

Departments

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-SC0019323.

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: Blood velocity flow profile shown in cerebral vessels as calculated with HARVEY, a whole-body model of the human circulatory system devised by DOE CSGF alumna Amanda Randles and colleagues. Turn to page 22 to read about how she and other former fellows are part of a Department of Energy push to break the exascale computing speed barrier. *Credit: Science Visuals.*

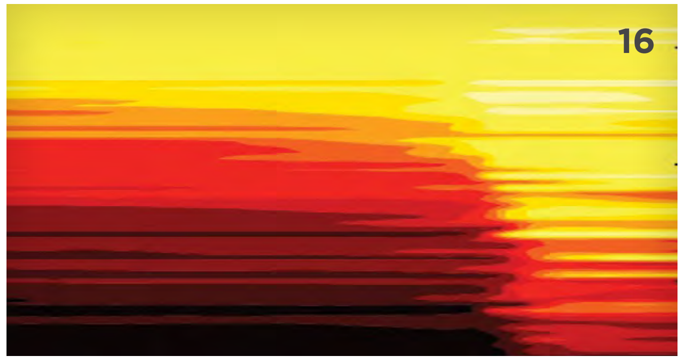

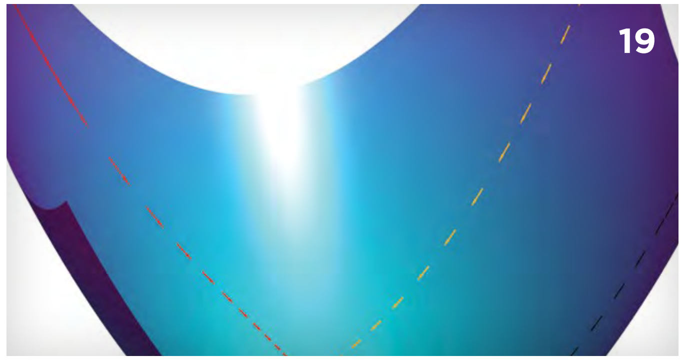

QUANTUM CALCULATIONS

He's out to make quantum computers useful.

Jarrod McClean is a senior research scientist at Google LLC, where he develops algorithms to simulate materials, chemistry and mathematics to advance the potential of quantum computers. The DOE CSGF alumnus previously was an Alvarez Postdoctoral Fellow at Lawrence Berkeley National Laboratory. He was an invited speaker at the 2019 DOE CSGF Annual Program Review.

DEIXIS: HOW DID YOU START WORKING ON QUANTUM COMPUTING?

Jarrod McClean: In my graduate work, I was focused on how to simulate electronic and chemical systems using conventional methods. My advisor [at Harvard University], Alán Aspuru-Guzik, had some projects focusing on this question: If you had a quantum computer, would it be useful for simulating chemical reactions? This was interesting, and I hadn't heard about quantum computers before. In some of my first projects, I was lucky enough to collaborate with groups that had early experimental apparatuses. While they were extremely small, they let you do proof-of-principle experiments. That got me excited about the fact that the technology was not entirely theoretical – that people were building these things and you could even do small-scale tests.

WHAT IS THE CURRENT STATE OF QUANTUM COMPUTING?

Today's quantum computers have been summarized with the acronym NISQ: noisy intermediate-scale quantum – devices before we have resources to do full fault-tolerant quantum error correction. The number of qubits has increased dramatically, but the quality matters just as much. For example, in ion traps, you could load in thousands of ions, which are, in principle, qubits. But today you would lack the control to do anything meaningful with them. If we don't have the ability to control at that level, it's about as good as only having a few qubits. It's been challenging to battle the noise before full quantum error correction, so there's a lot of innovation both in new applications and taking existing applications and making them more resilient to errors. What we would really like is some practical application that could take, say, a quantum device that has around 70-plus qubits with good but not perfect control and do a useful application.

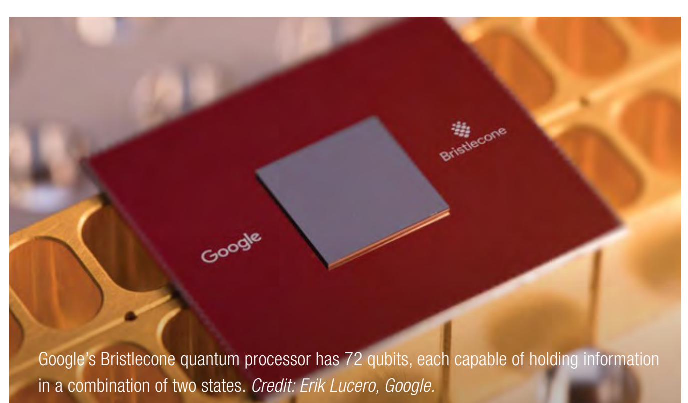

Google's Bristlecone quantum processor has 72 qubits, each capable of holding information in a combination of two states. *Credit: Erik Lucero, Google.*

IS IT NECESSARY TO SOLVE THE ERROR-CORRECTION PROBLEM TO MAKE QUANTUM DEVICES USEFUL FOR GENERAL-PURPOSE COMPUTING?

That's debated in the field. I want to be optimistic that we might reach useful applications, perhaps not general-purpose but in specific areas, before we have full quantum error correction. We might be able to increase resilience to a level that we can do something like a chemistry, optimization or machine-learning problem before we reach full fault tolerance. Other people believe the only way to reach practical applications will be full-scale quantum error correction and that the best course is to both improve the device and reduce the overhead requirements by improving our error-correction methods. I think these things aren't so different from what we're already trying to do with near-term devices. As we try to make things more robust for NISQ devices and run applications on them, we learn about the actual noise models. In principle, learning about that and testing codes on near-term devices will let us optimize quantum error-correction methods for real models of noise and bring the overhead down dramatically.

WHAT IS YOUR ROLE AT GOOGLE'S QUANTUM ARTIFICIAL INTELLIGENCE LABORATORY?

My focus more generically is how to make quantum computers useful. More specifically it's developing algorithms for both the long-term – an error-corrected device – and (for) the near-term, the next five years. That inevitably requires looking at methods to reduce errors, but I also try to think about unique ways to use quantum resources for big problems people haven't thought about yet. That can include different methods to store probability distributions for quantum machine learning or how to use optimization subroutines. A lot of the work is exploratory in the sense that we don't know that it will pan out, but it's still important because the more applications we have, the better the motivation for building quantum computers.

WHAT DO YOU MOST WANT TO DO WITH A QUANTUM COMPUTER?

If I had an ideal quantum computer, I would study some basic physical phenomena, including, say, strongly correlated chemical systems – biometallic enzymes or small catalysts responsible for either nitrogen fixation or processes in the body we have not been able to see or study before. If you go to an even cooler scale, there are ideas for doing mechanistic simulations of processes like nuclear fusion and fission. It would be amazing to see the real dynamics and answer questions there that perhaps haven't been addressed. Its use as a digital spectrometer for accessing parts of the world that we've not really been able to get good insights into is amazing and is one of the things that excites me most. And of course, I just want to know what quantum computers are good at. I want to just play with it and see. Can I come up with other algorithms? Can I iterate as people do on classical computers by exploring problems we didn't think quantum computers would be good for?

Quantum computers based on superconducting qubits, such as Google's Bristlecone, must be chilled to near absolute zero to both maintain minimal electrical resistance and delicate coherences in the device. That means the chips must be housed in specialized cryogenic equipment like that shown here. *Credit: Erik Lucero, Google.*

WHERE DO YOU SEE YOUR CAREER TAKING YOU?

At the moment, this is where I want to be. It's where some of the most exciting hardware and theoretical development is. Industry is starting to attract the best theorists, even out of academia. It's a hive of activity for quantum computing, and I imagine myself staying in this capacity for some time. When I was deciding between academia and industry, someone who had worked at Bell Labs shared their experience with me, which was it's almost always the right choice to go where the resources you'll have will be best. Right now, it seems like there's a ramp-up in industry where you can do your best work and have the best resources. It's a good strategy to keep going where I have the opportunities to do the best work that I can and I think that's here. If that changes, then I'll adapt.

SPHERICAL COWS

Using barnyard animals to understand quantum computing.

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

By Jacob Bringewatt

There's a joke in the physics community that pokes fun at theoretical physicists like me: A farmer's cows weren't producing enough milk, so he wrote to the local university asking for help. A week later a theoretical physicist calls him. "I have a solution," the professor says. "But it only applies to a spherical cow in a vacuum."

This highlights a common feature of how we do physics. The world is a complicated place; it would be impossible to describe it in exact detail, so a key part of a physicist's job is to simplify things in an intelligent way. If they do a good job of picking the correct abridged model of some complicated system like a cow, it should accurately describe the essential pieces of the real system.

It's not obvious that this should work, but remarkably it does. For example, Newton's theory of gravity tells us that the attraction between two objects depends only on their mass and the distance between them. Nothing else matters in any significant way – not the color of the objects, or their shapes, or their smells, or anything else.

That means this method of simplification is also one of unification. As far as gravity is concerned, an octahedral chicken or a conical goat is essentially the same thing as a spherical cow. It works the same whether we're discussing playing catch or planets orbiting the sun. The fact that we can figure out such unifying relationships is the amazing beauty of physics.

Understanding complicated systems this way doesn't just work for centuries-old physics. We can apply the same approach to a new and exciting area: quantum computation.

Many people know that the ordinary computers we're familiar with perform calculations on the basis of zeros and ones called bits. Physically, they're just tiny electrical components that can exist in either the state we call 1 or the state we call 0. In quantum computing, scientists and engineers hope to use quantum bits (qubits) instead. Under the rules of quantum physics, qubits have extra properties that regular bits don't have. For example, a qubit can exist in a state that is part 0 and part 1. Or they can be entangled, which means that the state of one qubit depends on another. So if one qubit is in state 0 and is entangled in the right way with a second, then that qubit also is in state 0. The idea is that quantum computers can use these extra properties to solve all sorts of interesting problems that standard machines can't.

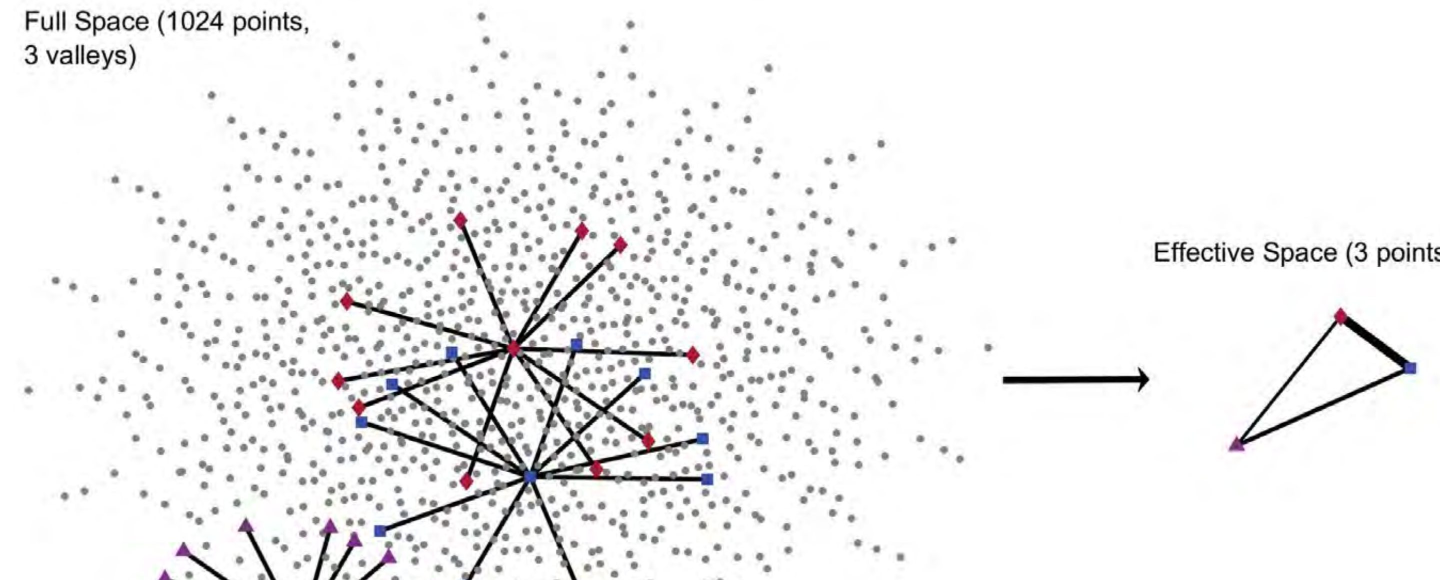

A representation of how Jacob Bringewatt's model reduces the quantum adiabatic problem. The figure at left depicts a pasture-like area with three valleys. Every point in the space is connected to 10 other points, but for visual clarity only connections within the valleys are labeled. The reduction process approximates this problem with just the three valleys with different probabilities for hopping between them. *Credit: Jacob Bringewatt.*

For my research, I look at a particular model called quantum adiabatic computation. This solves a computational problem by describing it as a landscape in which we're trying to find the lowest point. If we still use the barnyard analogy, we can think of it as finding the deepest valley in a given pasture.

It turns out many real-world problems can be described as versions of this. For example, Amazon deciding where to put its supply centers to most cheaply and efficiently deliver packages can be seen as the same pasture problem.

So how can we find the lowest point? The best standard method we know of is just wandering around and looking. Let's say we add cows to the pasture. They can roam (either randomly or in some more intelligent way) or even be airlifted from place to place, all seeking the lowest point. But in a really big pasture, no matter how smart the cows are it would take a long time to find it. Even if they find the bottom of a valley, it's difficult for them to know if it's the deepest one in the entire pasture.

Quantum adiabatic computation lets us do better. Instead of ordinary cows, imagine we have "quantum cows" who (like qubits) come equipped with special powers. One of these makes them aware of an extra-wide area around them. This should let quantum cows find out more effectively whether they are at the true lowest point or if there is some deeper valley over a hill. As a result, they can potentially find the lowest point quicker. What I want to know is how much faster these quantum cows actually are.

Unfortunately, this is still a hard problem. So, as physicists usually do, we devise a simpler model. In my research, I imagine that these high-dimensional pastures consist of a collection of perfectly symmetrical valleys of different depths and widths. Of course, like imagining that a cow is a sphere, this isn't true for a real landscape, but it picks out the key features.

I then calculate (using an actual standard computer) how quickly the quantum cows could find the lowest point of each valley, if that was the only one in the whole pasture. I add all these results in a particular way to approximate how long it would take to find the lowest point in the full collection of valleys. Once again, this process of breaking the problem into smaller, simpler chunks works remarkably well. This model of the quantum cows allows me to investigate when and how a quantum computer can do better than an ordinary one.

So maybe the theoretical physicist solved the farmer's problem after all.

TOUR DE FORCE

Casey Berger abandoned the movie world to pursue quantum questions.

By Thomas R. O'Donnell

Math was Casey Berger's first love. "That was my area" in elementary, middle and high school, she says. "But I was also involved with theater and sang in a choir, so I really got to explore both sides of my personality."

When she had to choose a major at Boston University, Berger picked what appeared to be the more creative direction: film and television. (She later added philosophy and earned dual bachelor's degrees.) "It seemed like the kind of thing that you have to give a shot once in your life," says Berger, who had more confidence in her artistic instincts than in her math and science abilities.

Berger graduated and pursued a film production career in the Los Angeles area. She met big-name stars and directors, but a what-if question from a show-business mentor prompted her to change course and pursue a doctoral degree in theoretical and computational physics.

Now a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient at the University of North Carolina, Chapel Hill, Berger realizes her dilemma was false. "There's an immense amount of creativity that goes into research, because at a certain point you're not solving a problem someone else has given to you," she says. "You're exploring the unknown and that requires being able to think creatively and flexibly."

Berger focuses on the quantum mechanical world, where energy and matter interact as both particles and waves. Under quantum rules, subatomic particles such as electrons absorb energy only in discrete amounts and can influence each other's behavior over great distances.

Quantum physics fascinated her, Berger says, but "it also blew my mind how little we actually understand about the way important elements of our universe work." High-performance computing (HPC) is necessary to find answers.

With advisor Joaquín Drut, Berger develops techniques to solve the many-body Schrödinger equation, the canonical – and gnarly – formula for describing a quantum system's state. They focus on rotating bosons – elementary or composite particles, including certain nuclei, that have spin (a form of angular momentum) characterized as a whole number. Rotating bosons are found in large-scale phenomena like superfluids, which flow without losing kinetic energy, and superconducting materials, which allow electrons to move without resistance. Spinning a quantum system radically alters its properties, entering complex regimes that computer power has only recently made accessible.

Suppose someone wants to calculate the properties of a population, such as the average income of United States residents. "That's a well-defined number," Drut says. "You have a finite number of people, and you can calculate it easily."

But for quantum systems, "you have a humongous number of elements that you want to average." To cope with that larger space, researchers typically sample a range of probabilities and average them. For Drut and many others, the standard tool for this is the quantum Monte Carlo method, named for the eponymous Monaco casino because it randomly tests probable values in a way that's similar to games of chance.

"In conventional quantum Monte Carlo, you have a well-defined probability, and you use that to choose the important parts to sample to get your averages," he explains. It's like choosing to compute the per capita income in Oklahoma: sampling would focus on that state rather than another.

But rotating quantum systems – and problems in many other physics areas – suffer from the sign problem. The probability is complex and "you find that the averages are all over the place," Drut says. "They keep fluctuating." This statistical noise, with averages oscillating between positive and negative and cancelling each other, obscures the final answer – an expected value, such as a quantum system's total energy, that an experiment might find.

Berger and Drut tackle the problem with stochastic quantization, using these complex probabilities to help choose the samples. "Stochastic quantization gives you a way to make sense of that probability and compute what you want to compute," Drut says. It handles the positive and negative pieces of the physics separately so they don't cancel. A final step averages the two, zeroing out the negative and producing an expected value.

To test the algorithm, Berger and Drut recreated a model of a relativistic Bose gas, one comprised of bosons moving at a significant fraction of light speed. It showed the same sudden, sharp increase in density that a 2009 computation found and that echoed experiments.

The rotating bosons Berger focuses on, however, are low-energy systems, with atoms and particles moving at nonrelativistic energies. "That changes what kind of equation you use to solve the quantum mechanics," she says. Berger rewrote the code and tested it on a problem with a precise answer: a nonrotating, noninteracting free boson gas. The algorithm produced statistical noise, prompting some adjustments to better balance its precision and efficiency.

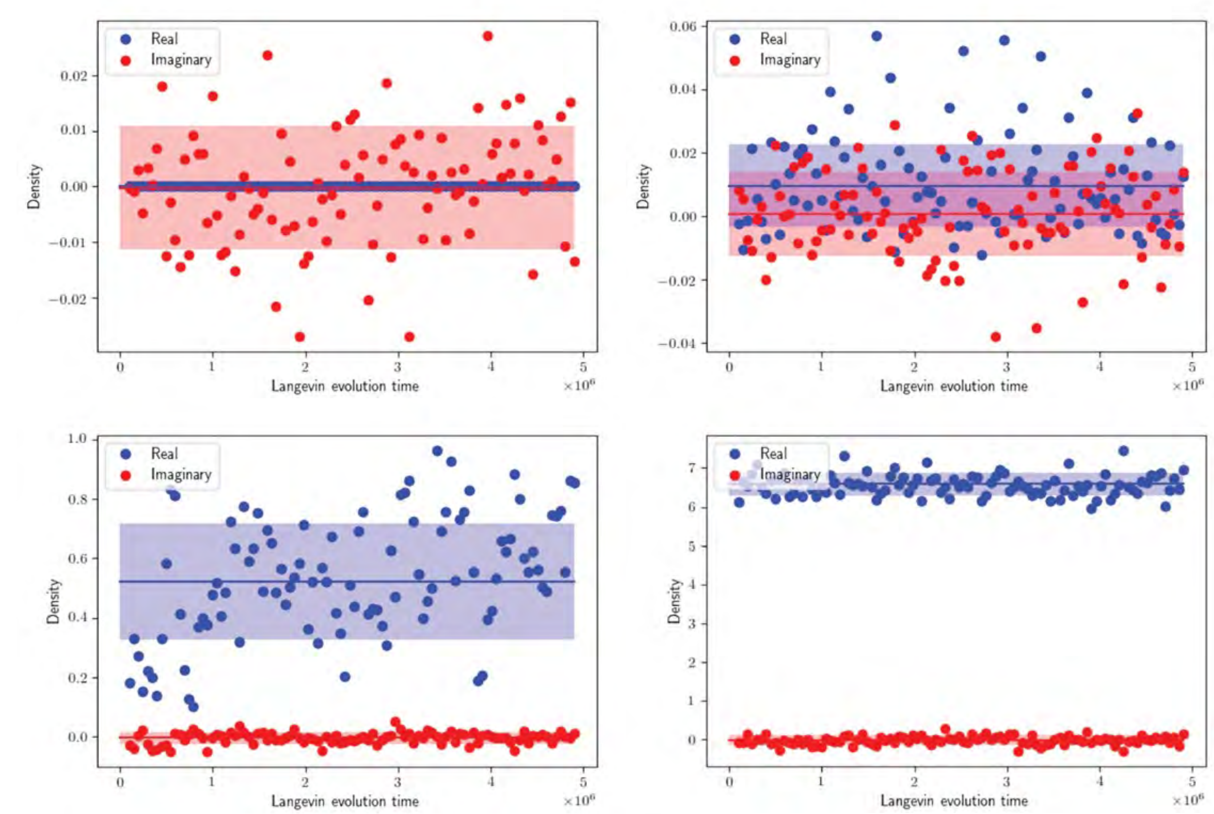

These show calculations of the density of a particle system, from least dense to most dense (top left to bottom right). The complex Langevin method lets researchers separate the real part (blue) from the imaginary part (red). As the system gets denser, the real values increase while the imaginary values remain at zero. *Credit: Casey Berger.*

ONE TOOL, MANY TASKS

Richard Barnes, a University of California, Berkeley, computational ecology and geoscience doctoral student, has studied perennial grain-producing plants, how landscapes change over time and election gerrymandering. In a recent collaboration with Lawrence Berkeley National Laboratory staff, Barnes examined whether GPUs (graphics processing units) could effectively reassemble disparate strands of DNA into a complete genome. He says his resulting code can assemble the whole human genome in around eight minutes. His Cal advisor is John Harte.

ON THE EDGE

As a University of Maryland undergraduate, **Noah Mandell** so impressed physics professor William Dorland that he took Mandell on annual trips to Britain, where Dorland collaborates with Oxford University faculty on fusion energy research. Mandell has continued pursuing that interest in his Princeton University doctoral studies, modeling the super-hot plasma that must be created and contained to trigger nuclear fusion. Gkeyll, the code that Mandell and advisor Greg Hammett are developing with researcher Ammar Hakim, tracks turbulence in the edge of plasmas contained in tokamaks, donut-shaped fusion reactors. With supercomputer power, Mandell is extending Gkeyll to include the effects of electromagnetic fluctuations.

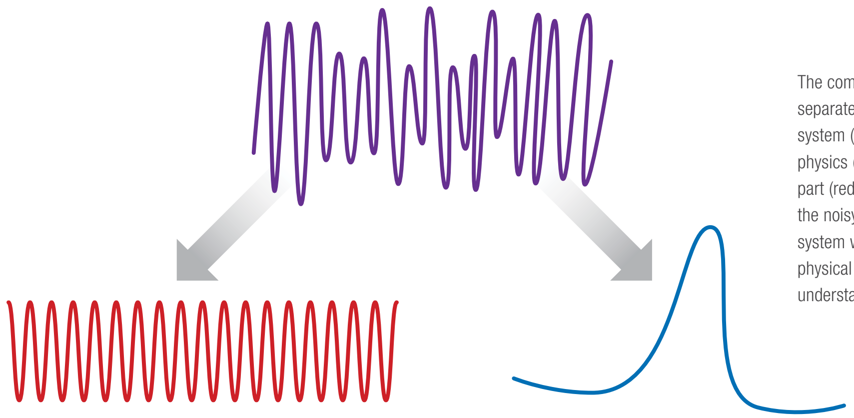

The complex Langevin method separates a complex physical system (purple) from its real physics (blue) and its imaginary part (red). Without this method, the noisy imaginary part of the system would drown out the real physical signal researchers want to understand. *Credit: Casey Berger.*

The next step – adding rotation and particle interaction – increases the difficulty. A preliminary test found the algorithm again generated noisy data, Berger says. "We saw things we were expecting with a rotating system, but to really see vortex formation and other more detailed elements, we needed a cleaner algorithm." She's since tweaked the code, producing results that agree with previous outcomes.

Such algorithms can run poorly on HPC systems, but Berger's method works well and is easy to parallelize, distributing calculations among many processors. She hopes to test it on supercomputers at DOE's Fermi National Accelerator Laboratory or Argonne National Laboratory.

It's a long way, physically and intellectually, from the talent management agency where Berger formerly worked. The company connects actors, writers and directors with projects and produces movies and TV shows.

The job was glamorous; Berger met major stars and directors but worked 12-hour days and read scripts on weekends. She also found a surprising lack of intellectual curiosity. Berger was thrilled to learn one morning that physicists had identified the Higgs boson, a crucial particle physics achievement, but at her office "no one really knew or cared what that was."

She missed the intellectual engagement that had been part of her life. When a mentor asked Berger what she'd do if money was no object, she automatically replied "go back to school to study astrophysics."

At Ohio State University in her hometown of Columbus, she joined Richard Furnstahl's nuclear theory group. She later did summer research in computational physics at North Carolina with Drut and naturally slid into his group as a doctoral student.

Berger's dual career reflects her character, Drut says. "She remains a person with multiple interests," something that distinguishes her from other graduate students. Berger also "sees her career in a very broad way," understanding how the low- and high-energy aspects of physics link. That's "not at all common, but it's something I always wish more people in our community were able to see," Drut says. After watching her speak, colleagues have told him that Berger is as skilled as an advanced postdoctoral researcher.

Berger hopes her research will help "chip away at a big gap in our knowledge of quantum many-body physics." She plans to research computational physics after graduation – in academia, she hopes – with a characteristically broad perspective. "Sometimes there are these wonderful secrets locked away in one area of physics that nobody has heard of in another area of physics," she says. Methods like hers "can be used to make huge progress somewhere that had been stuck for a long time."

ROBOT WHISPERER

Julia Ebert weaves together biology, technology and more to explore behavior, swarms and space.

By Sarah Webb

Julia Ebert creates, builds and learns. But her varied interests – in neuroscience, coding, baking, ice-rink sports, web design and robotics – make her hard to pin down, even for those who know her best. "I had no idea what you were going to do," Ebert's father once told her. "You were interested in everything."

It's no wonder then that Ebert merged a behavioral neuroscience major with computer programming at Northeastern University in Boston. Now the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient is combining coding and biology again, developing swarms of small robots that collaborate to perform tasks. Someday similar groups may do jobs that bore or endanger humans.

Ebert is part of Radhika Nagpal's Harvard University lab, where researchers take lessons from biology, such as termites building collectively or bacteria cooperating, to make groups of robots that work well together.

So far Ebert has focused on collective decision making using Harvard-developed Kilobots, half-dollar-sized robots with limited intelligence, sensing and communication capacities that are deployed in groups of up to a thousand.

Ebert initially examined how 100 Kilobots used environmental cues to make collective decisions. She based her experiments on the house-hunting problem – strategies that groups of bees and ants use as they search for a new home.

Ebert starts by using computers to simulate test runs of her experiments, letting her identify potential pitfalls and try out scenarios before setting up longer trials with the Kilobots. She has shown that the miniature robots can use simple sensing and communication to detect colors and make decisions about environmental features. “You’re taking advantage of the fact that even though you’re small and can’t see much, everyone else is working on the same problem,” Ebert says.

Someday robot swarms could handle a range of dirty, dull and dangerous tasks, such as setting up a habitat for humans on Mars, she says. But in that scenario, the devices must deal with complex situations in which one set of decisions leads to others. On Mars, for example, robots might have to spread out to locate an area with enough water for humans. “So one piece of this bigger puzzle that we’re trying to sort out is how you get more complex behavior out of groups of robots,” she says.

Besides herding a robot gaggle, Ebert explores a range of algorithmic approaches for solving problems. She wants to find general strategies groups of devices could use in many scenarios, not just in a single situation. “Swarm robotics is at a very early stage. There aren’t unifying theories about how to approach and how to solve problems when they’re spread out over this many robots.”

In fall 2017, Ebert designed and built her own devices, nicknamed LARVAbots. The machines mimic insects crawling over each other to see whether the group can stay together and move faster than an individual. Ebert completed the initial design in a single semester, producing six brightly-colored, hand-sized treaded robots. In experiments, the machines coordinated their actions and even enabled one of them to escape from a walled track.

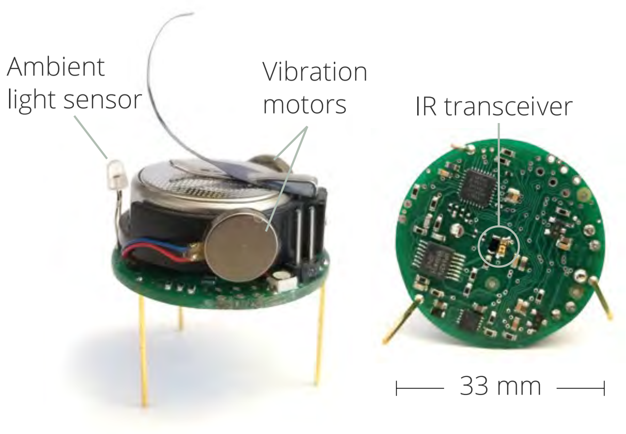

Harvard-developed Kilobots are half-dollar sized robots that can make simple measurements in their environment and communicate with their nearest neighbors using infrared signals. **Right:** In these experiments, modeled after the way insect communities search for a new home, a group of 100 Kilobots navigated a square environment patterned with black or patterns of other colors. Each robot can sense its immediate environment and signal to neighbors who are no more than three body-lengths away. *Credit: Julia Ebert.*

The LARVAbots have fed another passion: communicating science to diverse audiences. Ebert soon demonstrated the machines at the Cambridge Science Festival. She also has done numerous outreach events with the Kilobots, including in her Wisconsin hometown. It’s a chance to appreciate “how excited a kid gets about being able to play with a robot that’s being used for real research.”

Ebert’s many side projects often have programming and web design components. As an undergraduate, she produced a website that used data from Amazon to generate fake but realistic book blurbs. Other projects have more practical applications: While in the United Kingdom, Ebert and her teammates designed a website that helped medical students learn how to suture without a doctor present. It won a National Health Service contest.

Ebert is fearless, Nagpal says. “She has a raw talent at learning, and she crosses learning curves at speeds that are a little intimidating to the rest of us.”

Ebert’s polymathic tendencies appeared at an early age. Growing up, she built with LEGOs, wrote fiction, baked pizzas from scratch and went on plant-focused hikes with her science-teacher father. In high school, she became fascinated with psychology, fueled by discussions with her mother, who had a degree in the subject. Ebert loved learning, and it felt natural to study how the brain works, so she pursued a bachelor’s in behavioral neuroscience at Northeastern.

In an early course, Ebert learned to code in MATLAB, the ubiquitous scientific computing language. “It turns out I was pretty good at it, and I really enjoyed it,” she says. By her second year she had combined her biological and computational

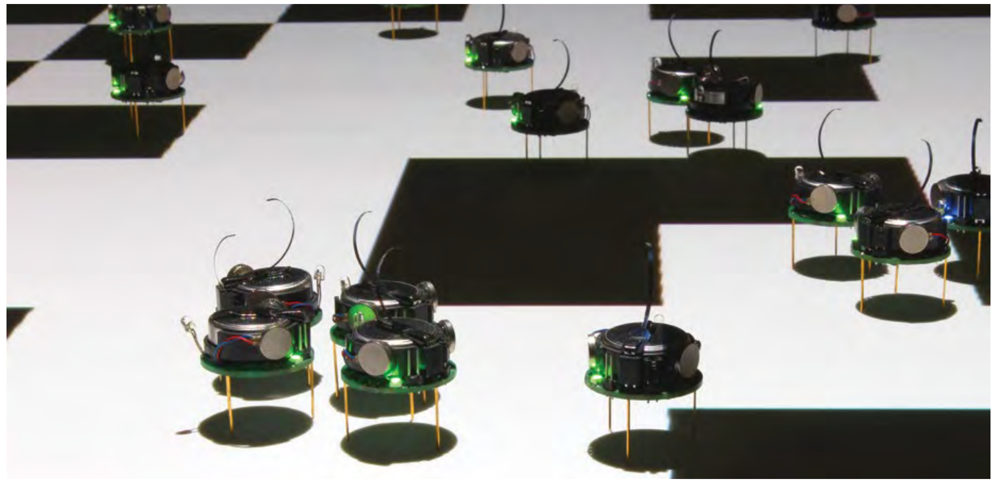

Kilobots move through a patterned environment, sensing as individuals and signaling to their nearest neighbors. Over time the group uses that information to make collective decisions about how to move while staying together. *Credit Julia Ebert.*

interests in Dagmar Sternad’s laboratory, picking apart how people learn motor tasks based on feedback cues and eventually programming robots for human experiments.

In 2015 Ebert earned a Marshall Scholarship and headed to Imperial College London, where she studied human-robot interactions and earned a master’s degree in bioengineering. Working with researchers Etienne Burdet and Ildar Farkhatdinov, Ebert helped devise the human-robot interface for the lower-limb powered exoskeleton (LOPES), a wearable robot designed to help injury or stroke victims maintain or recover their balance.

For her doctoral work, Ebert “wanted to flip it around a bit and look at the other side of how we could use biology to make better robotics.” Ebert also has turned her swarm robotics skills toward space. During her summer 2018 Lawrence Livermore National Laboratory practicum, she worked with physicist Michael Schneider to develop software enabling small satellites to communicate and make collective decisions about their observations. She started early by coding software for the satellite simulator as a course project before arriving at the lab.

Schneider was impressed with Ebert’s independence, creativity and knowledge. “She took off running very quickly,” he says, and produced useful programming. “Even understanding what the code was doing was a challenge, but she built some pretty impressive visualization tools so we could just show it.”

“Space is really cool,” Ebert notes, and the small satellites offered an interesting parallel to her work with Kilobots. She’s still collaborating with Schneider and considering new extraterrestrial directions for her graduate research, including working with the Space Exploration Initiative at the MIT Media Lab. Ebert hopes to finish her Ph.D. by 2021 and stay in academia.

When not programming robots, developing side projects or baking cinnamon rolls, Ebert continues a high school interest as a member of Harvard’s curling team. The sport involves sliding and accurately placing heavy stones, shuffleboard style, on ice. “It’s a lot of fun; it’s not too athletic for me.”

Ebert also is the life of the lab, Nagpal notes – a ringleader of activities from covering the group’s workspace in googly eyes to building an evolving LEGO model of the lab.

Like Ebert’s father, Nagpal says the longer she knows her, the less sure she is of where Ebert’s interests and abilities will lead. “Once I realized the breadth of her talents, I realized that there is no limit on the number of exciting things she may end up working on.”

BRAIN TEASER

With Stanford University’s Surya Ganguli, **Alexander Williams** creates techniques to analyze results from neuroscience experiments, helping sift data from hundreds or thousands of individual neurons. “There’s a big question in the field of what to do with all those data,” Williams says. The mathematical methods he develops will provide statistical descriptions that serve as foundations for computational models. Williams based his proposed method, tensor component analysis, on techniques he explored with Sandia National Laboratories mathematician Tamara Kolda during his DOE CSGF practicum.

CLOSE ENCOUNTERS

At the Massachusetts Institute of Technology, **Helena Qi** and advisor Heather Kulik use GPUs to examine atomic interactions in protein structures, especially when atoms are nearer together than scientists expect. Qi pinpointed the location of all atoms in more than 13,000 structures, then extracted cases in which atoms are closer than expected and applied quantum mechanical methods to understand them. Qi looked for whether certain amino acids, such as tyrosine, frequently occur in these situations. Tyrosines appear in these interactions more often than expected, suggesting that its tight interaction with other residues stabilizes protein structures.

EFFICIENCY SURGE

Max Bremer took on a hurricane simulation and blew away limits on its performance.

By Andy Boyles

When Max Bremer arrived at Lawrence Berkeley National Laboratory (LBNL) for his Department of Energy Computational Science Graduate Fellowship (DOE CSGF) practicum, he didn't realize he already had the seed of a perfect research project. Bremer had earned a bachelor's degree in aerospace engineering. He also had taken steps toward computationally modeling hurricane storm surges. At LBNL, he soon learned that he could make a unique contribution to computational science.

"I was going to do rockets or planes or something, and this hurricane simulation stuff just fell into my lap," he says. His work doesn't center on writing new modeling codes but on revising them to boost efficiency on future high-performance computing (HPC) architectures. "It's a really good problem, and I find it really interesting. So I just kind of stuck with it."

It's also a potentially life-saving project. Since 1980, seven of the 10 costliest weather disasters were hurricanes, the National Oceanic and Atmospheric Administration (NOAA) reports. Besides intense gales and torrential rains, a hurricane generates storm surge – widespread flooding as winds push ocean waters onto land. Storm surge is "often the greatest threat to life and property," NOAA's National Hurricane Center (NHC) says. It blames storm surge for most of the estimated 1,500 deaths hurricane Katrina caused directly in 2005.

When a hurricane threatens a coast, local officials must make rapid evacuation decisions based on NOAA's predictions. Those forecasts rely, in part, on simulations from an HPC code called Advanced Circulation (ADCIRC).

"They are life-or-death decisions," says Clint Dawson, Bremer's advisor at the University of Texas at Austin, where he heads its computational hydraulics group. "For example, if they decide to close a road, and the road didn't need to be closed, then that could slow down evacuations by several days."

Bremer's move into Dawson's group happened naturally. He was an engineering undergraduate student, also at UT Austin, when a course in numerical methods lured him into the field. The instructor, Troy Butler, showed Bremer that the techniques can yield approximations with real-world relevance to heat- and fluid-flow models and other physics.

At the time, Butler was a postdoctoral researcher in Dawson's group, which helped develop ADCIRC over the past two decades. (Butler is now at the University of Colorado Denver.) Still an undergraduate, Bremer soon was working in Dawson's group on DGSWEM, an experimental code used to test concepts for later incorporation into ADCIRC.

Bremer still didn't see storm surge simulation as his career's next chapter. After graduating, he went to the University of Cambridge intending to delve deeper into pure mathematics. As he worked through the university's prestigious Part III of the Mathematical Tripos, he blew off steam playing Ultimate Frisbee – and, unrelated to the sport, realized he was mainly interested in applied mathematics.

Bremer was chosen for the DOE CSGF and was on track to begin doctoral studies with Dawson. He expected to work on any one of several projects – rain-induced flooding, fluid flow through porous media, hurricane storm surge.

When he began his 2016 summer practicum with Cy Chan in LBNL's Computer Architecture Group, Bremer was particularly interested in the unit's emphasis on HPC optimization.

Bremer says he expected this work would differ from his doctoral thesis research. "My attempt to branch out was to do something computer science-y like 'let's look at one of these really out-there solutions to problems in high-performance computing.'"

He was especially interested in learning more about task-based parallelism – a way to improve simulations' efficiency. "There's only so much smaller you can make these computer chips," Bremer says. So to speed up computers, "you just have more computers" – additional processors.

More processors and greater parallelism breaks big problems into discrete tasks and doles them out to individual processors, which then solve the separate pieces simultaneously. This approach still entails inefficiencies. MPI, Open-MP and similar parallelism methods run tasks in lockstep. Because some jobs are bigger than others, processors can sit idle, waiting as all the tasks are completed before they all move on to the next set. In a big simulation, these brief idle times add up to losses. They also prevent researchers from drawing on a machine's full capability for more sophisticated and accurate simulations. In task-based parallelism, each task runs on its own time step, reducing the wait between jobs.

After Bremer arrived at Berkeley Lab, he and Chan saw that Bremer already had the kernel of an excellent project. He set out to learn task-based parallelism and load balancing, another optimization method. Bremer planned to take the methods to Dawson's group and apply them to DGSWEM and, perhaps someday, to ADCIRC and other codes.

ADCIRC is a mature code, Bremer says. "It's probably hard to move the bar on what's been done in the mathematical or algorithmic space," he says. "Rather than come up with a new mathematical model, what if we can just use the machines better?"

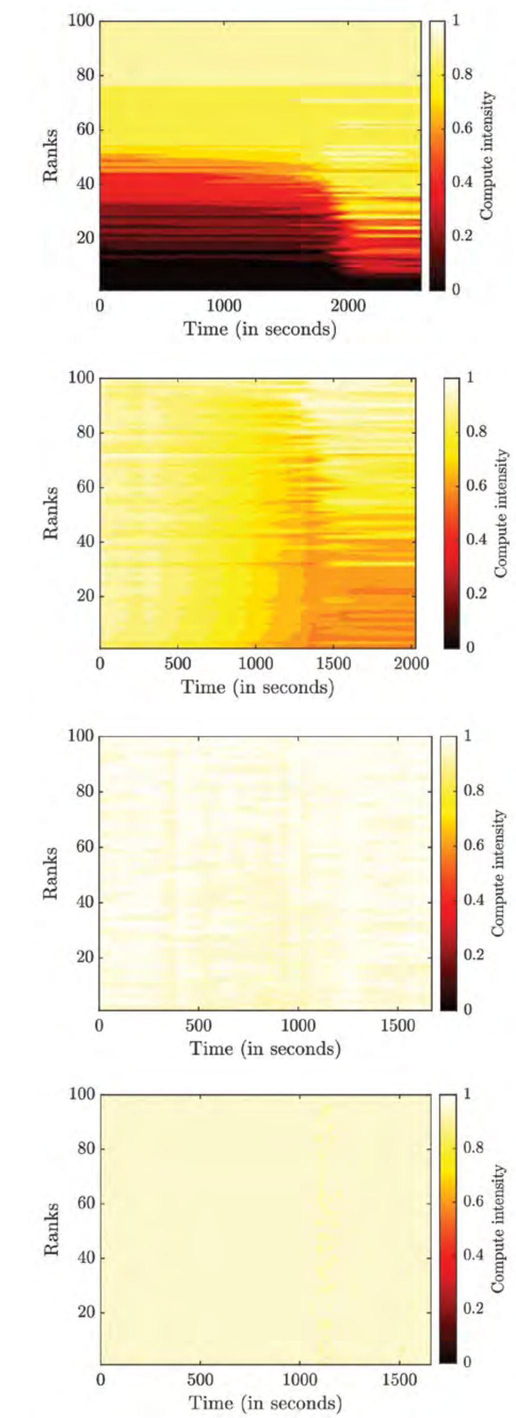

Predicted use of computational ranks throughout a hurricane simulation. In the standard static load-balancing approach (top), work is concentrated in the higher computational ranks (light colors along top) during runtime while lower ranks mostly sit idle. A multi-constraint-static technique (second from top) distributes work more evenly across processors, but still concentrates work in the higher ranks late in the computation. Asynchronous diffusion (third from top) and semi-static load-balancing (bottom) strategies redistributed work nearly uniformly. The sustained high computational intensity for these two approaches produced a speed-up of 1.5 times compared to static load-balancing. *Credit: Maximilian H. Bremer, John D. Bachan, Cy P. Chan, Semi-Static and Dynamic Load Balancing for Asynchronous Hurricane Storm Surge Simulations,* 2018 Parallel Applications Workshop, Alternatives to MPI (PAW-ATM), *November 16, 2018.*

PROPER PARAMETERS

Emmet Cleary uses computing to understand fluid flows. As a California Institute of Technology mechanical engineering doctoral candidate, he works with Tapio Schneider on mathematical techniques to select simulation parameters that produce the most accurate results and to calculate the inherent uncertainty in those choices. They focus on climate models, whose often-unknown parameters can force a researcher to set them manually. "That begs the question," Cleary says. "Did you get the best parameters? Did you pick the right ones? That leads to a lot of uncertainty." He tests algorithmic tools to determine the best choices for real-world problems.

STELLAR SCHOLAR

Hannah Klion explores high-energy phenomena, such as black holes and neutron stars, in which complicated physical properties interact in unusual ways. She's particularly interested in understanding astronomical changes that occur on a human timescale. With Eliot Quataert at the University of California, Berkeley, Klion has simulated neutron star mergers that produce gravitational waves, eject and attract matter around them and launch gamma-ray jets. Klion uses this model to simulate how light moves through the resulting stellar material, aiming to predict the light signatures Earth-based astronomers should observe from these cataclysmic celestial events.

Chan and his co-workers helped Bremer tackle another inefficiency in hurricane storm surge simulations. At the outset, dry areas demand no computational work. But some soon require significant processing as they become inundated. Because no one can predict precisely which dry areas will suddenly demand more computing resources, the computer's workload becomes imbalanced. "To achieve efficient utilization of the machine, you need to move these patches around on the fly," Bremer says. "That's load balancing."

During the practicum, Bremer learned the C++ programming language, then used it to implement task-based parallelism and load balancing. Before incorporating the methods directly into DGSWEM, Bremer created DGSim, a skeletonized version of the program. "We wrote a simulator for our simulator, which is bizarre to explain," Bremer says. "But it allowed us to use a lot less of the machine. So I can run the simulation on my laptop whereas normally I would need thousands of cores to do it."

The group then validated DGSim on Edison, a Cray XC30 supercomputer at LBNL's NERSC (the National Energy Research Scientific Computing Center), Chan says. "We were able to reduce the number of time steps calculated but still capture the same overall dynamic load profile of the hurricane." DGSim saved more than 5,000 core-hours compared to running a DGSWEM simulation, and the new algorithms improved hurricane-simulation performance by more than 50 percent. Bremer, Chan and their colleague John Bachan presented the results at SC18, the international supercomputing conference in Dallas.

Dawson recalls Bremer's new mastery with DGSWEM after his studies at Cambridge and LBNL. "When he came back, he sort of took over as the lead of that code development," Dawson says. "He and another student (Kazbek Kazhyken) basically rewrote the code from scratch, and they put in a lot of much more modern coding paradigms."

Today, Bremer finds himself "sitting between these two camps. One is this high-performance computing community, where you have all these people who are trying to figure out how we can get these algorithms to run efficiently on the new computers, and then Clint and his collaborators, who have a very concrete idea of a problem they'd like to solve."

Dawson and Chan say they admire Bremer's ability to serve as that bridge. In the summer of 2019, he'll begin another practicum with Chan's group, exploring how to make time-stepping not only asynchronous but also locally determined, with each element in a simulation taking its cues from neighboring elements. "We're excited to have him come back," Chan says. "I think it will be another good collaboration of him coming with expertise in the domain and us providing expertise in the computer science side."

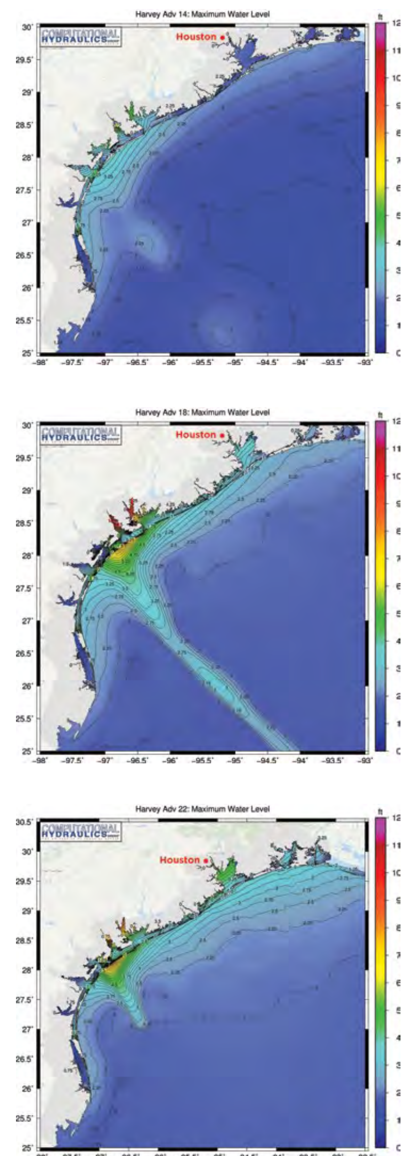

This sequence shows forecasted maximum water levels over the span of two days for 2017's Hurricane Harvey on the Middle Texas Coast, based on National Hurricane Center advisories and generated using the ADCIRC+SWAN Surge Guidance System (ASGS). The water level color scale is on the right; latitude and longitude are along the left and bottom. The city of Houston is at the top. The computational model uses the advisories as inputs to simulate how the storm affects the ocean and to predict flooding. Changes between advisories 14 and 18 show the hurricane's rapid intensification and advisory 22 shows it making landfall. The lighter-colored, lower water-level areas jutting into the gulf are the effect of the hurricane's low-pressure eye. For more ASGS results, visit the CERA website: http://cera.cct.lsu.edu. *Credit: UT Austin Computational Hydraulics Research Group.*

EARLY BLOOMER

Stanford's Carson Kent began working on DOE supercomputers in high school and never stopped.

By Monte Basgall

Carson Kent's tenure as a Department of Energy researcher long predates his Ph.D. work at Stanford University as a Department of Energy Computational Science Graduate Fellowship (DOE CSGF) recipient. He's been DOE-supported since his Albuquerque High School days, thanks to the New Mexico Supercomputing Challenge.

The state's DOE national laboratories are major sponsors of this competitive program to boost interest in high-performance computing (HPC) among high school or younger students. Encouraged by an instructor, Kent met with potential mentors at Sandia National Laboratories who threw out a couple of ideas for the contest. He ended up creating a hydrocode – a program that could simulate fluids flowing fast enough to create shock waves, he says, to evaluate whether hyper-powered water flows could be shaped to break up an improvised explosive device.

"Of course the one involving explosives is the one any high school student chooses," recalls Kent, though noting that they didn't actually blow up anything.

At Stanford, Kent works in the high-demand field of optimization, "the mathematics of efficiency," as he puts it. "What we're really interested in are algorithms for computing solutions to problems. That means you define some measure of cost and find a method to reduce that cost as much as possible."

Kent focuses on optimal transport, which his Stanford advisor Jose Blanchet describes as the cheapest way to move mass from one place to another – sand, for instance, from point A to point B to cover a sinkhole. Optimal transport is a problem that has been around for 250 years, adds Blanchet, an associate professor of management science and engineering and an affiliate of Stanford's Institute of Computational and Mathematical Engineering.

Kent's an expert, Blanchet says, at "developing and analyzing algorithms for a wide range of problems – complex, large-scale, high-dimensional computational problems." The math can be applied to a variety of other puzzles, such as matching donor kidneys with recipients, linking commercial products with customers in the marketplace, or using machine learning – which feeds known data to an algorithm so it can identify similar characteristics in unknown data – to generate human-looking faces for online advertising.

Kent entered the fellowship after starting at Stanford in 2015. His DOE CSGF practicum work with Argonne National Laboratory's Sven Leyffer exposed Kent to the field that would become the setting for his thesis work in optimization. He also collaborated with Leyffer on a related subject, robust optimization, which Kent describes as "taking optimization and then adding uncertainty." During his practicum "we worked on faster, better methods for solving those types of problems under very difficult constraints and conditions that the Department of Energy cares about."

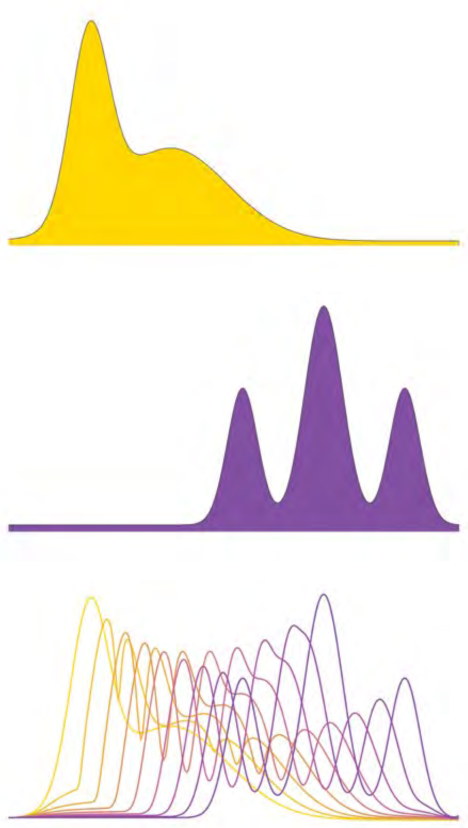

Optimal transport is a method for finding the most efficient way to move mass. Here, two different univariate densities – probability distributions based on single variables – illustrate the method. The gold and purple shaded plots display the respective densities of two different distributions. The third plot shows displacement interpolation – a geodesic, or length-minimizing, curve in the abstract space of probability distributions – between these densities that's induced by optimally transporting the mass of one distribution into the mass of the other. *Credit: Carson Kent.*

A poster presented as part of his fellowship described another related problem-solving technique used thousands of times every day in optimization. It works on linear programs, a simple way of expressing a bunch of costs associated with a bunch of decisions. For instance, a linear program can model how Amazon would route packages to customers. The National Football League uses a small modification of a linear program to schedule games each season.

As Kent's CV suggests, it wouldn't be far-fetched to conclude he's among the best-prepared students to enter the DOE CSGF program.

After high school, in 2010, his supercomputing challenge experience led to work with other Sandia teams on methods to detect malware in certain Windows files and on using geographic information system web applications to simulate the Western U.S. power grid for ways to detect and prevent failures.

More Sandia collaborations on other real-world projects continued while he worked on a bachelor's of science in mathematics and statistics at the Colorado School of Mines, which he selected for its intensive emphasis on engineering, math and science. Each summer he interned at Sandia, and for his last three undergraduate years also telecommuted with the lab during the winter while attending classes in Golden, Colorado.

The fruits of those interactions included leading development of a tool called Cyber Shopper. The program sought to model an adversary's actions when attempting gain access to a computer system, he says. This led to a U.S. patent, issued in July 2018, on a method and apparatus for managing such an attack.

Another project applied machine learning to automate something called cognitive radio, which Kent describes as a method for listening in on and deciphering wireless transmissions.

While at Mines, Kent also picked up a research interest in uncertainty quantification, which aims to calculate how much trust researchers can put in their computational models of real-world conditions.

That interest began in the summer of his junior year, when Paul Constantine, then a Mines professor, invited him to take a graduate-level short course on the subject at Stanford.

When they returned to Colorado, Kent began a research project with Constantine and a University of Texas at Austin professor that led to a 2016 paper in a Society for Industrial and Applied Mathematics journal on a computational workhorse called the Markov chain Monte Carlo (MCMC) method. MCMC is a

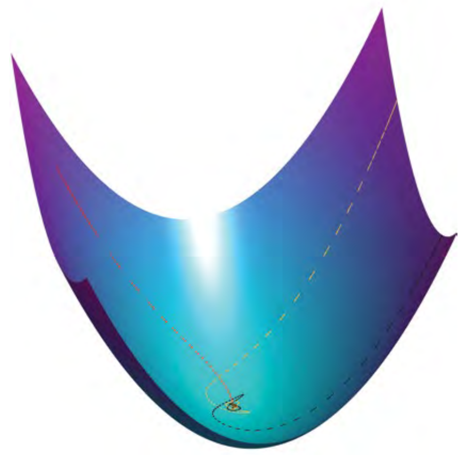

Smoothed versions of the paths taken by various first-order optimization algorithms –mathematical techniques for finding a problem's most efficient solution – testing various candidate solutions. These optimization methods minimize the loss/ objective function (the property the algorithm seeks to minimize or maximize) pictured here. The gold and black trajectories arise from accelerated first-order methods – ones that maximize efficiency by considering past solutions the algorithm has tested – and converge to the minimum much faster than the standard, gradient descent algorithm (red), even though the paths they take to the solution appear to vary much more. *Credit: Carson Kent.*

class of algorithms for sampling among probability distributions. These distributions, for instance, can represent physical quantities of interest in a simulation problem. Some of the authors' calculations ran on Mio, an HPC cluster at Mines.

While an undergraduate, Kent – a veteran of many an elementary and middle school science fair – volunteered as a judge at similar local events. He also taught programming at a Montessori school. Kent says those activities have been rewarding and have helped him recognize how valuable mentors have been in his own career.

One such guide was Sandia's Uzoma Onunkwo, who Kent worked with his senior year at Mines and "was a main cause of my desire to go to grad school rather than industry after undergrad." Onunkwo wrote one of Kent's supporting letters for attending Stanford, as well as one backing his DOE CSGF application.

All that happened after Onunkwo and another researcher enlisted Kent to help estimate the effectiveness of error-correcting codes in quantum computation. "He was phenomenal and delivered beyond the tasks we assigned to him," Onunkwo recalls of Kent's performance as a student intern. "We needed a C++ programmer with expert-level experience. He delivered remarkable results for us in that role."

Kent chose Stanford for his doctorate because its programs provide "the right combination of exposure to all the different areas you need if you're going to do computational work." For his Ph.D. work to come, Kent is gearing up to use Stanford's Sherlock HPC cluster to solve some large-scale optimal transport problems.

Kent calls the fellowship "amazing. I've definitely spent a while around the Department of Energy and there are few other programs in it that I value as highly as the DOE CSGF in terms of its ability to build the workforce that the department constantly needs." The same is true if fellows choose to go into industry, he adds.

Stanford's Blanchet calls Kent's attitude, personality and work ethic an impeccable match for such a program. "He's independent. He's very creative and approachable. It's just a joy to interact with him."

AMORPHOUS GOAL

Nicholas Boffi studies soft-matter physics with Harvard University's Chris Rycroft, modeling bulk metallic glasses (BMGs), moldable amorphous metals. BMGs are promising materials for many uses, but they can fail when subjected to certain forces. Boffi has extended Rycroft's computing tools from two dimensions to three dimensions to test a new theory describing how and why BMGs fail. With Jean-Jacques Slotine at the Massachusetts Institute of Technology, Boffi also uses mathematical tools to explore how artificial intelligence algorithms work. He hopes to use algorithms he's developed to extend the BMG models.

A MODEL DESIGN

Michigan State University's **Zane Crawford** studied finite element algorithms for electromagnetics, but his Sandia National Laboratories practicum – and a discussion with advisor Shanker Balasubramaniam – changed that. Now Crawford researches topology optimization – finding the best design for an object to perform a specific task. He seeks novel designs for electrical devices that "best get me from some input signal to some desired output signal." His algorithms must choose the best orientation and properties to meet that goal. After applying a finite element approach, Crawford uses a topology optimization algorithm to fine-tune selected parameters affecting the system to get the best possible device design.

REALITY: IT'S COMPLICATED

As a Harvard University undergraduate, **Ian Dunn** realized his chemistry interests focused on physical theory, differential equations and mathematics. He wanted to know not just how molecules react with each other but also why. With David Reichman at Columbia University, Dunn wrestles with many-body quantum physics problems, attempting to incorporate more realistic – but complicating – factors into his models. These simulations often require approximations or clever computational schemes to make calculations manageable. His work examines fundamental physics, but someday such models could be used to design and optimize the properties of new materials such as solar cells.

EXASCALING THE HEIGHTS

DOE CSGF alumni are everywhere in the drive toward the next big computing milestone.

By Thomas R. O'Donnell

The simulations **Amanda Randles** runs on high-performance computing (HPC) systems are huge – large enough to contend for the Gordon Bell Prize, the premier award recognizing outstanding supercomputing achievement. Her codes portraying the human circulatory system in unprecedented detail could help doctors improve treatments of aneurysms and other conditions.

But understanding some diseases, such as how cancer spreads via veins and arteries, "requires solving problems we literally can't do on today's supercomputers," the Duke University biomedical engineering professor says. For example, her models portray bulk fluid flow, "but we can't capture the movement and mechanical properties of the cells. It's really a coarse-grained view."

To achieve cellular-level detail, Randles needs the next HPC generation: Exascale machines capable of a billion billion – 10^{18}, or a 1 with 18 zeroes after it – scientific calculations per second. That's about five times the theoretical top speed of Summit, an IBM AC922 at the Department of Energy's Oak Ridge Leadership Computing Facility (OLCF), rated as the world's most powerful supercomputer in November 2018.

Randles is one of many Department of Energy Computational Science Graduate Fellowship (DOE CSGF) alumni working toward that goal. Former fellows are involved with nearly every aspect of the exascale push, from hardware to enabling technologies to applications. This article highlights contributions from a few of them.

It's not surprising that DOE CSGF alumni lead the exascale push, says **Aric Hagberg**, deputy leader of the Computer, Computational and Statistical Sciences Division at Los Alamos National Laboratory (LANL). "The fellowship was developed to train people for large-scale computational modeling," the 1994 alumnus says. "Doing that for more than 27 years, it's obviously going to generate people who are interested in leading these kinds of projects."

Aurora, slated to launch at the Argonne Leadership Computing Facility (ALCF) in 2021, will be DOE's first exascale computer. The second, Frontier, will arrive at the OLCF the same year. In 2020, a third DOE computing facility, the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory, will launch Perlmutter, a pre-exascale machine that will help pave the way for the others. The DOE Office of Science's Advanced Scientific Computing Research (ASCR) program oversees the three centers.

Petascale computers like Summit help scientists understand such physical phenomena as quantum chemical reactions, plasma interactions in fusion energy and astrophysics, and the origins of the universe. Exascale, they say, will let them model bigger systems for longer times and with greater precision to probe these processes more deeply.

"Every time we deploy a new supercomputer generation where we increase the performance by a factor of five, I believe we actually get more than a factor of five increase in productivity" for science, says **Judith Hill**, a 2003 DOE CSGF alumna who leads the OLCF Scientific Computing Group. She expects exascale computing will further integrate simulations with observational data, more closely connect HPC to experimental facilities like powerful DOE X-ray and particle beam generators, and boost statistical confidence in model results.

ABOUT THE APPLICATIONS

Randles is counting on it. The 2013 alumna leads a project titled "Extreme-Scale *In-Situ* Visualization and Analysis of Fluid-Structure-Interaction Simulations," one of the first chosen to run on Aurora via the ALCF's Early Science Program (ESP). Of the 10 ESP projects, DOE CSGF alumni are involved with three. Each initiative was picked for its focus on handling and analyzing large data files and on machine learning, in which algorithms sift data to identify or classify unknown information.

Randles' models already generate around a petabyte of data for each step taken through time. Exascale machines, capable of modeling individual blood or cancer cells in detail, will produce even more.

"Trying to download a petabyte or several petabytes of data for every time step and then visualize and analyze it to find what's going on is just not tractable," Randles says. Her Aurora ESP project will develop methods to scrutinize data *in situ* – as the simulation runs. "We want to see something like a high-resolution video of how a cell is moving and interacting" without slowing the calculations.

Randles hopes to run detailed simulations that help doctors understand how cancer cells move, what makes them adhere to or gather in certain body locations, and how the cells' shape or stiffness affect their propensity to grow and spread.

Recent DOE CSGF graduates **Julian Kates-Harbeck** and **Kyle Felker** track a different kind of circulation – hot charged particles swirling through nuclear fusion reactors – as part of the "Accelerated Deep Learning Discovery in Fusion Energy Science" Aurora ESP project. It focuses machine-learning algorithms on tokamaks, donut-shaped chambers that house plasmas – intensely hot gases of atomic nuclei and free electrons. If the plasma is hot enough, the nuclei fuse, releasing tremendous energy, but it also can escape the magnetic fields containing it, damaging the reactor walls.

"We want to see something like a high-resolution video of how a cell is moving and interacting."

– Amanda Randles

The project, led by William Tang of DOE's Princeton Plasma Physics Laboratory, employs deep-learning models, a form of machine learning, to spot signs of damaging plasma instabilities and then act to stop or minimize them.

Kates-Harbeck and Felker joined the project during DOE CSGF practicums. Felker contributes knowledge gained from modeling astrophysical plasmas and Kates-Harbeck is the code's chief architect.

"The ESP is an opportunity to scale this deep-learning application to leadership-class supercomputing facilities," Kates-Harbeck says. He'll continue pushing the code to run well on a growing number of processors while incorporating additional and more complex training data.

Nicholas Frontiere, a 2017 DOE CSGF alumnus, is part of "Dark Sky Mining," an Aurora ESP project headed by Argonne National Laboratory's Salman Habib. It will connect meticulous cosmological simulations with massive data from the Large Synoptic Survey Telescope (LSST), which will rapidly examine the sky in detail when it begins observing in 2020.

Frontiere began collaborating with Habib's group as an undergraduate, working on HACC, a cosmology code that simulates the evolution of large parts of the universe. (Argonne computer scientist **Hal Finkel**, a 2011 DOE CSGF alumnus, also is on the HACC team, helping improve the code's data storage and retrieval performance and other factors.)

Observatories such as the LSST will produce data of unprecedented precision, so simulations that attempt to explain observations must be similarly accurate, Frontiere says. "If the data are good to one percent, you'd better be able to simulate to one percent and understand everything to that point. That's the goal" of exascale.

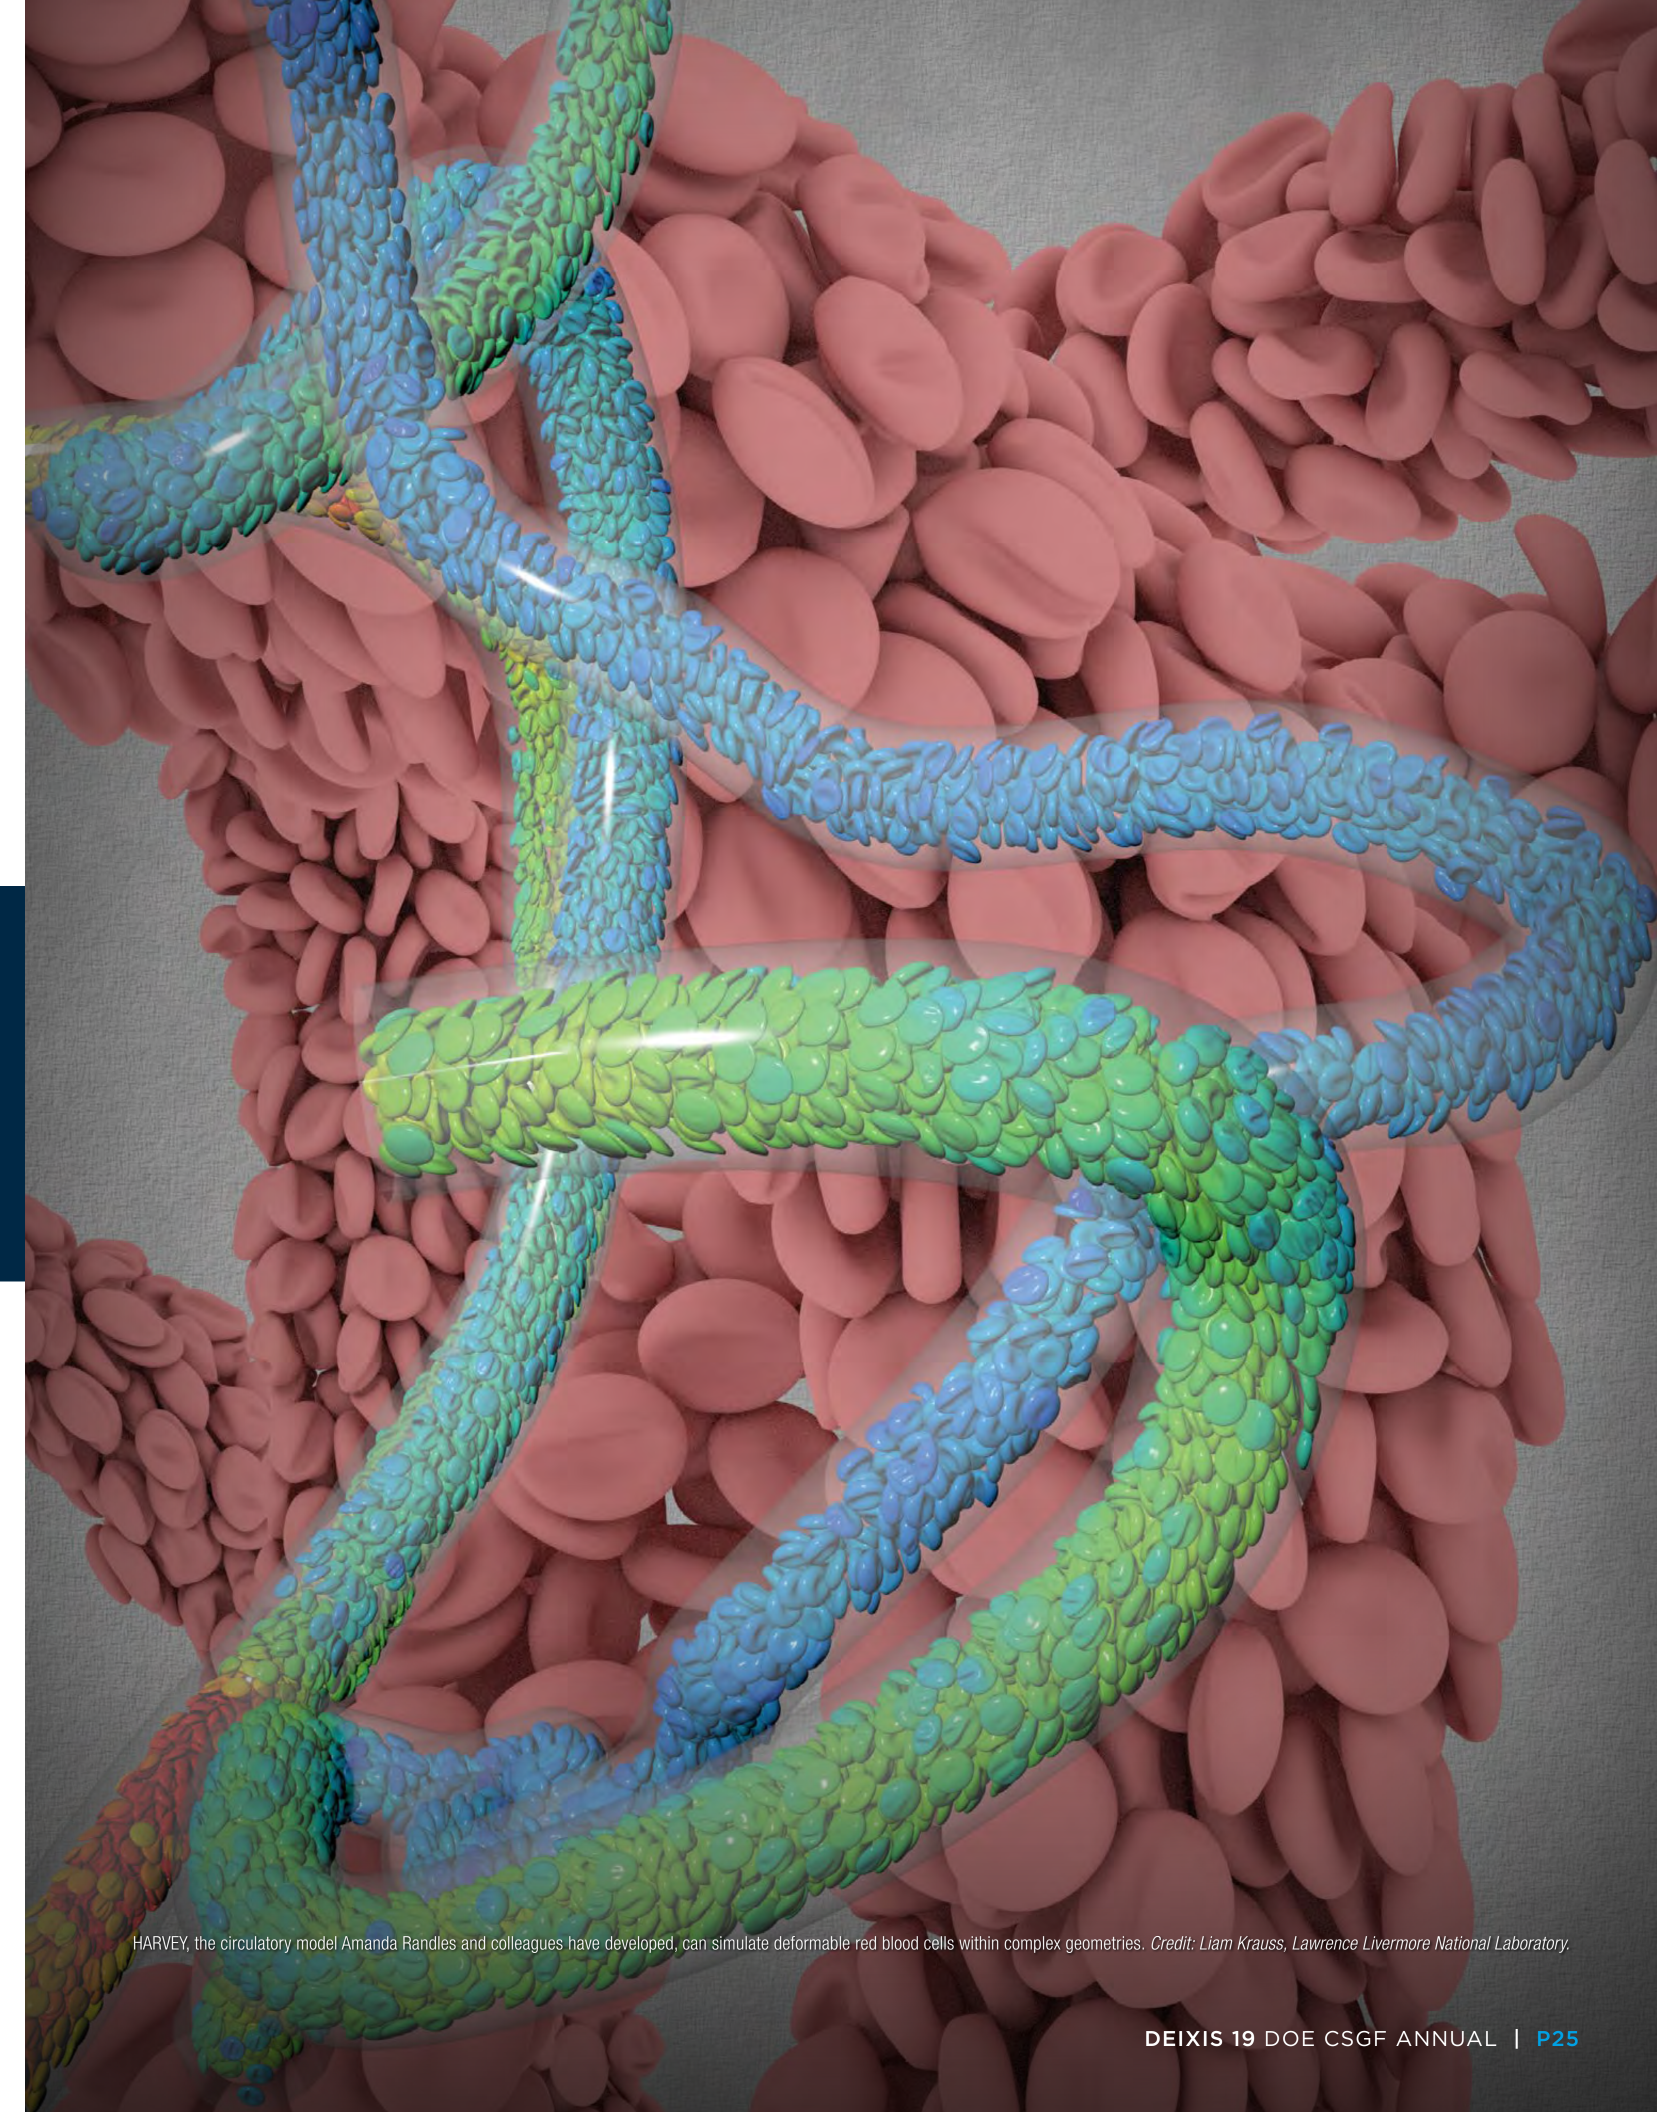

HARVEY, the circulatory model Amanda Randles and colleagues have developed, can simulate deformable red blood cells within complex geometries. *Credit: Liam Krauss, Lawrence Livermore National Laboratory.*

A worker inside General Atomics' DIII-D tokamak during a maintenance period in 2017. Princeton Plasma Physics Laboratory researchers use data from the fusion device to train a machine-learning algorithm to identify and head off plasma disruptions in similar devices. *Credit: Rswilcox via Wikimedia Commons; republished under a Creative Commons attribution license.*

22 TO GET READY

Significant hurdles, however, block the path to these science objectives. The new systems must be resilient, overcoming failures in their hundreds of thousands of processors and parts. The power barrier also looms: An exascale machine could devour enough electricity to fuel a city. To minimize consumption, engineers are pairing standard processors with energy-efficient accelerators such as graphics processing units (GPUs), but that complicates programming. And if processors must constantly communicate, it could impede calculations.

Perhaps most important, researchers must develop new or improved algorithms and enabling software to maximize exascale's potential. Existing programs must be revamped to operate on computer architectures that will drastically differ from ones used today.

"As these machines become more complex, we in the (HPC) facilities have a tough job getting ready to deploy" them, Hill says.

To surmount these challenges, DOE launched the multiyear Exascale Computing Project (ECP), connecting national laboratory researchers, academics and hardware and software vendors in collaborations encompassing every exascale issue.

Hill leads the Application Integration at Facilities activity, part the ECP's Hardware and Integration focus area. She oversees efforts ensuring that application developers understand the coming exascale architectures so their programs are suited to run on DOE machines. System designers, in turn, will learn about hardware properties the applications will need.

"We have performance engineers and computational scientists at the ASCR computing facilities who are experts at this type of work," Hill says. Application Integration at Facilities is "an opportunity to bring some of that knowledge into the ECP." The application readiness efforts at each ASCR center will connect specialists with developers working on codes from both ECP and the broader community to prepare their programs for the labs' newest supercomputers.

Hill will oversee the OLCF's program to ready codes for Frontier. The program is ramping up, she says, and should choose applications once the machine's design is announced.

Jack Deslippe, a 2010 DOE CSGF alumnus and leader of NERSC's Application Performance Group, oversees a similar effort to prepare programs for Perlmutter. He leads the NERSC Exascale Science Application Program (NESAP), now entering its second round of helping DOE developers adapt their science codes.

"Parallelism is becoming harder for application teams to exploit."

– Jack Deslippe

The first NESAP targeted Cori, NERSC's Cray XC40 that came on line in 2017. Cori used Intel Corporation's Xeon Phi many-core processors. Each core can concurrently compute on chunks of data twice as large as ones previous Xeons could handle. These energy-efficient chips mark a trend likely to continue in exascale architectures, toward "more and more parallelism," Deslippe says. "In some ways that parallelism is becoming harder for application teams to exploit. That's the big challenge" NESAP is to overcome.

Perlmutter presents a next step toward exascale, since it will contain both the latest multicore processors in standard CPU-only nodes as well as GPU-accelerated nodes. At least five ECP applications will be among the 25 chosen for the second NESAP.

ENABLING TECHNOLOGY

Several DOE CSGF graduates are leaders in another ECP program: co-design centers, multi-institution collaborations that develop applications and hardware simultaneously, each influencing the other to make the systems relatively easy to program. The teams create enabling technology such as code libraries and toolkits that applications focused on similar tasks can use.

Timothy Germann, a 1995 DOE CSGF alumnus at LANL, leads ECP's Co-Design Center for Particle Applications (COPA), targeting particle-based programs used in plasma physics, cosmology and other disciplines. COPA has produced libraries that particle application developers can tap for codes performing common functions, such as fast Fourier transforms to convert wave data over time into frequency data.

COPA and other ECP co-design centers also build proxy applications, small versions of big programs to test prototype HPC hardware, often on computers and programs that emulate full-scale supercomputer designs. Proxy apps use hundreds of lines of code rather than thousands in the full-scale versions.

COPA researchers work directly with a handful of ECP-chosen applications, helping "to bridge between them and the vendors and computer science folks on new algorithms and how they impact some of the (HPC) architectural decisions" for particle codes, Germann says.

Frontiere and his mentor, Habib, work with a COPA code project, ExaSky, that will alter HACC to run at exascale. One of its contributions to COPA is SWFFT, a three-dimensional fast Fourier transform code library Frontiere helped write for HACC.

ExaLearn, another ECP co-design center, focuses on exascale machine learning. It connects eight laboratory-based teams, two of them led by DOE CSGF alumni: Hagberg and **Michael Wolf** of Sandia National Laboratories.

"You should be able to decrease the turnaround time on finding out if you learned something from the experiment," Hagberg says. "Instead of going away for a week and processing your data, you may be able to do it off-shift" or as the beam runs. ExaLearn methods might even help adjust beam parameters during experiments to achieve optimal results, he adds.

Argonne's Finkel also helps applications prepare for and run well at exascale as part of the ECP's Software Technology focus, concentrating on compilers – programs that translate code into machine-readable directions. "The quality of the compiler has an impact on both the performance of the application and on the productivity of developers and users," says Finkel, who leads the ALCF's compiler and programming languages research.

Much of Finkel's ECP work revolves around LLVM, an open-source, widely used compiler technology he has helped advance. Collaborating with vendors and other researchers, he's adapting LLVM to ensure it works well at exascale with programming models such as OpenMP and languages such as Fortran.

Finkel also works with exascale hardware development as one of two lead technical representatives in a project involving Intel Corp.'s federal government business arm. He provides feedback on the company's technical direction as its processor, memory and input/output equipment evolves for next-generation computing.

"Instead of going away for a week and processing your data, you may be able to do it off-shift."

– Aric Hagberg

The machine-learning tasks DOE researchers face are far more complex than those confronting industrial HPC users, says Wolf, scalable algorithms manager for Sandia's Center for Computing Research. DOE teams also have far less of the information needed to train their algorithms.

ExaLearn's products will have many applications, but among the ones researchers will attack first are DOE light sources – powerful X-rays beams that examine submicroscopic structures and fast-moving reactions. Machine learning, already in testing at these facilities, could help analyze the growing mountains of experimental data they produce.

The work is important, Finkel says, but in the end it's discovery that matters. "We're enhancing scientific computing" to tackle unsolved problems in materials science, climate, aerodynamics, biology and other areas. "The range of scales you have to simulate" to address these tasks "you can really only do with an exascale-class computer." Surmounting that challenge is "truly what excites me about exascale."

That DOE CSGF alumni are entrenched in the effort, he says, is "a reflection on two things: The program is quite important; and it does a really good job of preparing graduate students for productive careers in science."

PROBLEM SOLVERS

Howes Award recognizes winners for mentoring and communication.

By Sarah Webb

Chelsea Harris

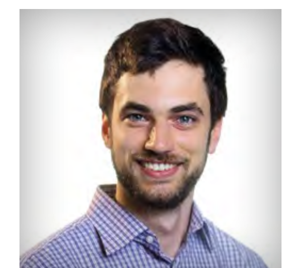

Adam Riesselman

Chelsea Harris and Adam Riesselman solve problems. Working in disparate disciplines – astronomy and biology – they both use computing's power to untangle complex mysteries. Meanwhile, the 2019 Frederick A. Howes Scholar in Computational Science award winners also are making their marks outside their fields as they communicate science to broad audiences and mentor researchers at a range of career levels.

"When I see a problem, I think about how can it be fixed and in what I think of as the best way," says Harris, now a postdoctoral researcher at Michigan State University. As an undergraduate, she dove into general relativity at the University of California, Santa Barbara's College of Creative Studies. Thinking about how small she is as an individual on the Earth's surface, in the solar system and beyond, spurred her interest in astronomical distances. Scientists measure these vast expanses by studying Type IA supernovae, the remains of sunlike stars that exploded while interacting with another stellar body. In her doctoral work at the University of California, Berkeley, Harris used simulations to examine how these supernovae interact with the gas surrounding them to better understand the original stars that spawned them. She pursued the Department of Energy Computational Science Graduate Fellowship (DOE CSGF), she says, because computing is the best way forward in astronomy. "You can't go out and observe this because *Star Trek* isn't real." Harris was a fellow from 2013 to 2017.

For her postdoc, she's digging deeper into computational work and massive star explosions, an area in which exascale computers many times faster than today's most powerful machines will play a key role. She's implementing state-of-the-art codes for high-order calculations of supernovae from massive stars with rotating cores. In these systems the star's magnetic fields and rotation interact, generating more powerful explosions. She continues to run interaction simulations on the workstation, named Sparky, that she built with an allowance from the DOE CSGF.

During his Ph.D. at Harvard University, supported by the DOE CSGF from 2014 to 2018, Riesselman realized that a lot of computational biology's power is linked to experiments. With the ability to test millions or billions of hypotheses through DNA engineering, experimental teams need software tools to guide and prioritize their research.

One critical biological conundrum: matching DNA sequences with the physical traits or maladies they produce. If a diagnosed patient has 20 potentially disease-causing mutations, Riesselman notes, researchers would like to zero in on the culprit. He's used machine learning models to group and categorize families of biological data to understand which mutations are beneficial or neutral versus those that are harmful. This approach can also help teams prioritize experimental research, narrowing billions of testable hypotheses down to the top 200,000 that are most likely to prove fruitful in understanding genetic diversity.

Since completing his Ph.D. last year, Riesselman has taken a position as a machine-learning engineer at Insitro, a drug-discovery startup company in the San Francisco Bay Area. He works closely with wet-lab researchers to design high-quality experiments. The synergy between experimental biology and machine learning is where both fields can have the strongest impact, Riesselman says. "You can make as many predictions as you want, but if they don't work in the real world, then they're worthless."

Riesselman has also focused on communicating science to broad audiences. "Generally, we need to make sure that taxpayers understand where their money is going and see how

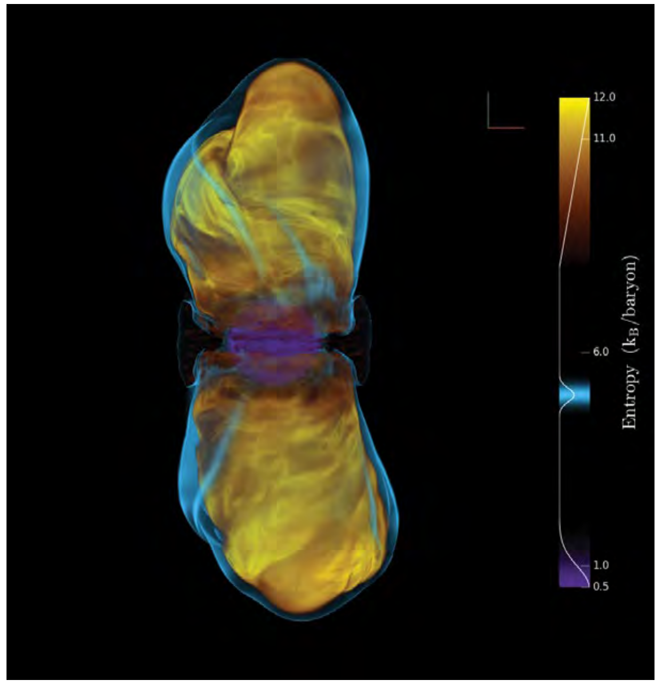

A FLASH code-generated simulation of a 25-solar-mass star's rotating core that, through magnetic effects, is developing a jet. The jet is thought to rip through the star's outer layers and create a long gamma-ray burst. *Credit: Chelsea Harris.*

university science research helps them," he says. During his Ph.D. studies, he worked with Science in the News, a Harvard graduate student organization focused on communication. Riesselman wrote and edited for the group's online publication and participated in a Reddit "Ask Me Anything" discussion about artificial intelligence. For several years, he also represented computational biology at career fairs for the Health Professions Recruitment and Exposure Program (HPREP), an outreach organization focused on underprivileged high school students in the Boston area. He'd give short presentations to groups of seven students, describing research and job opportunities in computational biology.

Harris has also done outreach, including speaking about the science of *Star Wars*, often in costume. But she's also distinguishing herself as a mentor and mental health advocate, drawing on her experience overcoming such issues as a doctoral student to help others. In 2016 she attended the UC Berkeley Mental Health Conference and gave the astronomy department's first-ever talk about mental health to destigmatize those concerns and inform others about campus resources.

"Talking about mental health within the framework of graduate school is hard," Harris says. People who don't fit traditional science stereotypes face additional coping challenges. They can sense that others doubt them and think they need to project strength. "I'm never surprised when someone tells me that they're struggling," she says, but she encourages them to seek help and a path forward. At Michigan State, she is one of three coordinators for the Stellar Mentorship program, an initiative within the astronomy department that offers broad-based career support from the undergraduate level to faculty.

Among the DOE CSGF's many benefits, both Riesselman and Harris highlight the importance of their practicum experiences. At the DOE's Joint Genome Institute, Riesselman worked with wet-lab researchers to design experiments and understand high-dimensional genomic data. That helped him focus on building a common language and understanding between computational and experimental cultures.

Harris's practicum at the National Renewable Energy Laboratory allowed her to study radiation transport in renewable energy, examining how mirrors can degrade in concentrated solar power installations. Though she's aiming for an academic career in astronomy, the experience showed her she could apply her expertise beyond supernovae."If I wanted to pursue renewable energy work," she says, "I would have the confidence and the skill set."

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.

HOT PAPERS, AWARDS AND MORE

Achievements accumulate for former fellows.

The Association for Computing Machinery presented **Amanda Randles** (2010-2013) with the 2018 Grace Murray Hopper Award, recognizing Randles' work in building detailed models of the human circulatory system. The Duke University assistant professor of biomedical engineering developed HARVEY, a massively parallel fluid dynamics simulation of red blood cells moving through the human arterial system. The ACM citation says it's the first time a researcher has effectively modeled blood flow at the cellular level.

Seth Davidovits (2010-2014) received the 2018 Marshall N. Rosenbluth Outstanding Doctoral Thesis Award from the American Physical Society. The prize recognizes "exceptional young scientists who have performed original thesis work of outstanding scientific quality and achievement in the area of plasma physics." Davidovits' dissertation focused on the theory and simulation of turbulence in compressing fluids, with an emphasis on effects unique to plasma, such as a novel sudden viscous dissipation mechanism. Davidovits, now a postdoctoral fellow at DOE's Princeton Plasma Physics Laboratory, earned his doctorate in 2017 from Princeton University.

Joshua Vermaas (2011-2016) is lead author on a *Proceedings of the National Academy of Sciences* paper describing a key mechanism behind antibiotic resistance in bacteria. Vermaas, now a postdoctoral researcher at DOE's National Renewable Energy Laboratory, worked on the project while still a graduate student at the University of Illinois at Urbana-Champaign. The project computationally modeled a biological pump in the bacterium *Escherichia coli*. In antibiotic-resistant microbes, these pumps eject unfamiliar small molecules, such as germ-killing chemicals, before they can damage the organism. If researchers can block the pumps, they may be able to overcome the bacterium's resistance.

Tal Danino (2006-2010) and his Columbia University collaborators received a $500,000 award for young investigators from the Bonnie J. Addario Lung Cancer Foundation and the Van Auken Private Foundation. The grant supports their research into engineering bacteria for lung cancer immunotherapy treatments. *The New York Times* also quoted Danino in a September article about synthetic biology, the quest to genetically modify bacteria to treat disease.

Carnegie Mellon University highlighted research from **Zachary Ulissi** (2010-2014) into a method to identify prospective energy storage materials. Ulissi, a CMU assistant professor, and his colleagues say their automated screening method uses machine learning and optimization to guide molecular models that predict the performance of new catalysts used to reduce carbon dioxide or to split water into hydrogen and oxygen. Ulissi and alumna **Brenda Rubenstein** (2008-2012), a Brown University chemistry professor, also are part of a project to develop advanced software capable of designing chemicals and processes for energy production and other potential applications.

Jordan Hoffmann (2014-2018) and his Harvard University colleagues developed a computer code that can build realistic models of the complex patterns seen in insect wings. The research, published in the *Proceedings of the National Academy of Sciences*, could help study the evolution of wing structure and other patterned shapes.

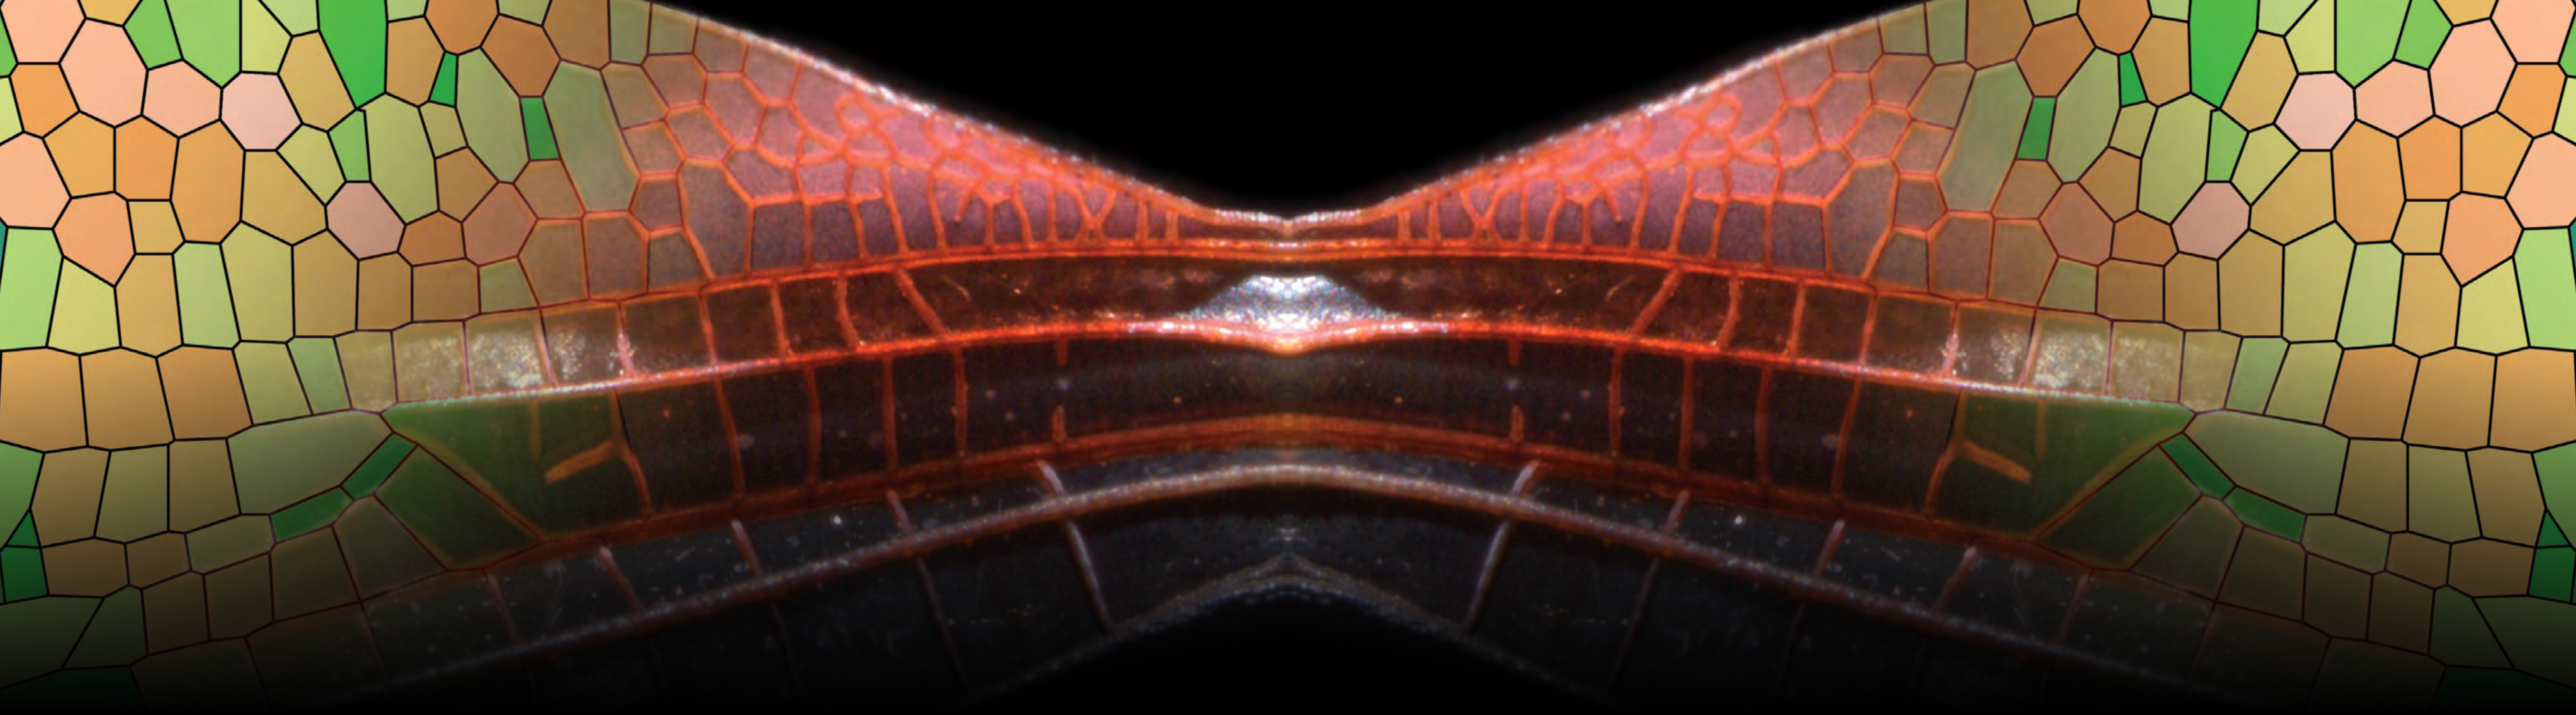

Alumnus Jordan Hoffmann helped develop a computer code that can build realistic models of the complex patterns seen in insect wings, as in this damselfly, *Hetaerina americana*, with polygonized vein domains colored by their circularity. The photo has been altered to appear in mirror image. *Credit: Seth Donoughe and Hoffmann, Harvard University.*

The Institute of Electrical and Electronics Engineers Computer Society named **Edgar Solomonik** (2010-2014) winner of the 2018 Technical Consortium on High Performance Computing Early Career Researchers Award. Solomonik, assistant professor in the scientific computing group of the Department of Computer Science at the University of Illinois at Urbana-Champaign, was chosen for his work on numerical algorithms and high-performance computing libraries, particularly the Cyclops library for tensor computations.

Norman Yao (2009-2013) was chosen for the 2018 Packard Fellowship in Science and Engineering. He'll receive $875,000 per year for five years to support his condensed matter physics research at the University of California, Berkeley.

Teresa Bailey (2002-2006), a code physicist at Lawrence Livermore National Laboratory, appeared in a lab video introducing its Sierra supercomputer.

A team that included **Jack Deslippe** (2006-2010) from DOE's National Energy Research Scientific Computing Center (NERSC) won the 2018 Association for Computing Machinery Gordon Bell Prize. The team, including personnel from NVIDIA Corp. and Oak Ridge National Laboratory, exceeded an exaop (a billion billion calculations per second) as it used deep learning methods to extract detailed information from climate data produced at NERSC. It ran on Summit, the world-leading system at DOE's Oak Ridge Leadership Computing Facility.

The *National Geographic* website highlighted research from **Brenhin Keller** (2012-2016) and colleagues that offered a possible explanation for why layers of Earth's crust representing millions of years of history are missing. The paper, published in the *Proceedings of the National Academy of Sciences*, suggests that glaciers that covered the planet beginning about 715 million years ago pushed the rocks into the oceans, where subducting tectonic plates sucked them into the mantle.

In November 2018 Prometheus Books published *Your Place in the Universe: Understanding Our Big, Messy Existence* by **Paul Sutter** (2007-2011), an astrophysicist at Ohio State University with an avid YouTube and podcasting following. He also consulted on (and appeared briefly in) *UFO*, a science fiction movie released in fall 2018 by Sony Pictures Home Entertainment.

The Economist magazine featured **Alexander Turner's** (2013-2017) research at the University of California, Berkeley, to explain why levels of methane, a powerful greenhouse gas, have shot up in the last decade. His take: There are fewer hydroxyl radicals, volatile compounds that react with methane to break it down into water and carbon dioxide.

Alex Perkins (2007-2011) is part of a Notre Dame University team to receive a $33.7 million, five-year award from Unitaid, an international preventive health organization, to research spatial repellents to battle mosquito-borne disease. Perkins is Eck Family Assistant Professor in Notre Dame's Department of Biological Sciences. The award is the largest research grant in the university's history.

Danino and **Ashlee Ford Versypt** (2006-2010) each received 2019 National Science Foundation Faculty Early Career Development Program (CAREER) grants to support their research. Ford Versypt, an assistant professor of chemical engineering at Oklahoma State University, will develop multiscale computational models to understand the progression of diabetic kidney disease. Danino, an assistant professor of biomedical engineering at Columbia University, plans to engineer bacteria's swarming behavior to provide an inexpensive macroscopic detector to identify and avoid diseases and other toxic agents. Each grant is about $500,000 spread over five years. In 2018, **Eric Chi** (2008-2011), an assistant statistics professor at North Carolina State University, received a CAREER grant to help develop a new framework for identifying patterns in multiway arrays, which will help analyze high-resolution data collected in bioinformatics, neuroscience and other fields.

Aurora Pribram-Jones (2011-2015), an assistant professor at the University of California, Merced, is part of the Consortium for High-Energy Density Science, a National Nuclear Security Administration-supported collaboration of multiple universities and Lawrence Livermore National Laboratory. It aims to expand and diversify the pipeline of students who pursue careers in the field, including using lasers to explore how matter and energy behave under extreme pressures and temperatures.

CLASS OF 2019

Richard Barnes
University of California, Berkeley
Computational Ecology/Geoscience
Advisor: John Harte
Practicum:
National Renewable Energy Laboratory

Zane Crawford
Michigan State University
Electromagnetics
Advisor: Shanker Balasubramaniam
Practicum:
Sandia National Laboratories, New Mexico

Noah Mandell
Princeton University
Plasma Physics
Advisor: Greg Hammett
Practicum:
Lawrence Livermore National Laboratory

Casey Berger
University of North Carolina, Chapel Hill
Theoretical and Computational Physics
Advisor: Joaquín Drut
Practicum:
Lawrence Berkeley National Laboratory

Ian Dunn
Columbia University
Chemical Physics
Advisor: David Reichman
Practicum:
Lawrence Livermore National Laboratory

Mario Ortega
University of California, Berkeley
Nuclear Engineering
Advisor: Rachel Slaybaugh
Practicums:
Sandia National Laboratories, New Mexico
Lawrence Livermore National Laboratory

Nicholas Boffi
Harvard University
Soft Condensed Matter Physics
Advisor: Chris Rycroft
Practicum:
Lawrence Berkeley National Laboratory

Sarah Gady
Michigan State University
Computational Mathematics, Science and Engineering
Advisor: Shiv Karunakaran
Practicum: Argonne National Laboratory

Helena Qi
Massachusetts Institute of Technology
Chemistry
Advisor: Heather Kulik
Practicum:
SLAC National Accelerator Laboratory

Maximilian Bremer
University of Texas at Austin
Computational Science, Engineering and Mathematics
Advisor: Clint Dawson
Practicums:
Lawrence Berkeley National Laboratory

Carson Kent
Stanford University
Computational and Mathematical Engineering
Advisor: Jose Blanchet
Practicum: Argonne National Laboratory

Alexander Williams
Stanford University
Theoretical Neuroscience
Advisor: Surya Ganguli
Practicum:
Sandia National Laboratories, California

Emmet Cleary
California Institute of Technology
Mechanical Engineering
Advisor: Tapio Schneider
Practicum:
Sandia National Laboratories, California

Hannah Klion
University of California, Berkeley
Astrophysics
Advisor: Eliot Quataert
Practicum:
Oak Ridge National Laboratory

DOE CSGF: HPC LEADERS

The Department of Energy Computational Science Graduate Fellowship (DOE CSGF) equips alumni with training to lead scientific advancement. A recent survey and review of curriculum vitae found nearly two-thirds of fellows used DOE supercomputers while in the program and almost half have tapped these resources since leaving. The fellowship also led most graduates to apply their skills to developing scientific codes. *Source: https://www.krellinst.org/csgf/about-doe-csgf/2017-longitudinal-study.*

Percent of alumni who have contributed to and/or led development of each type of product (N=211)

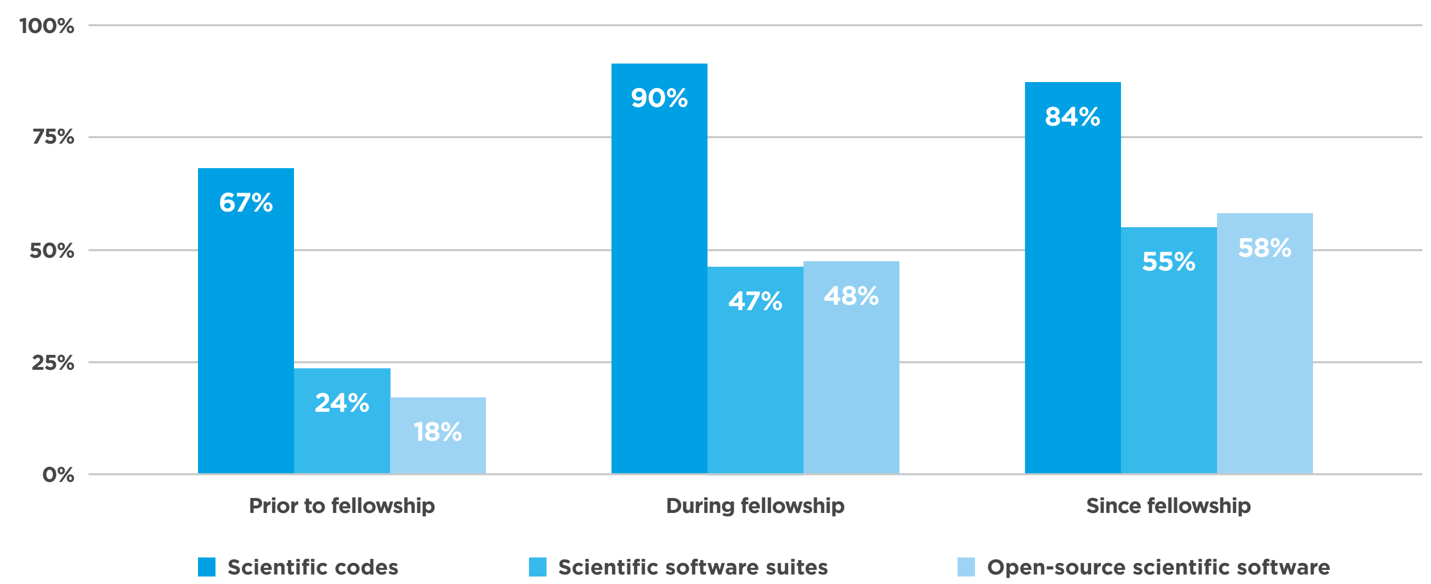

Percent of alumni who have taken advantage of dedicated computing time on DOE supercomputers (N=211)

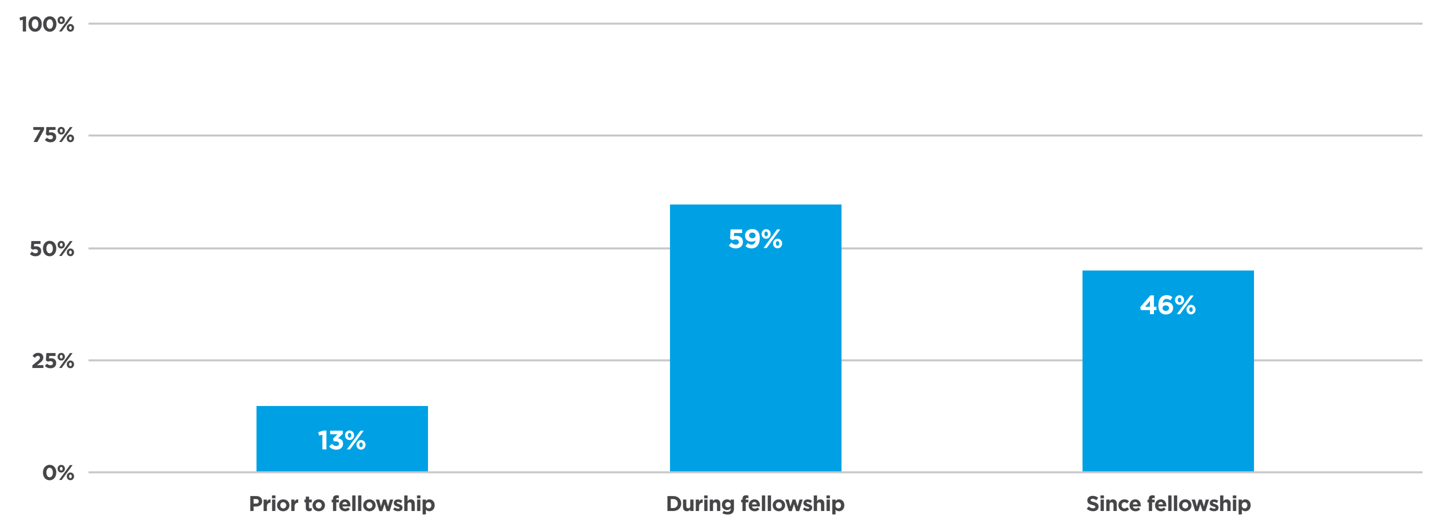

KRELL institute

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.