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."

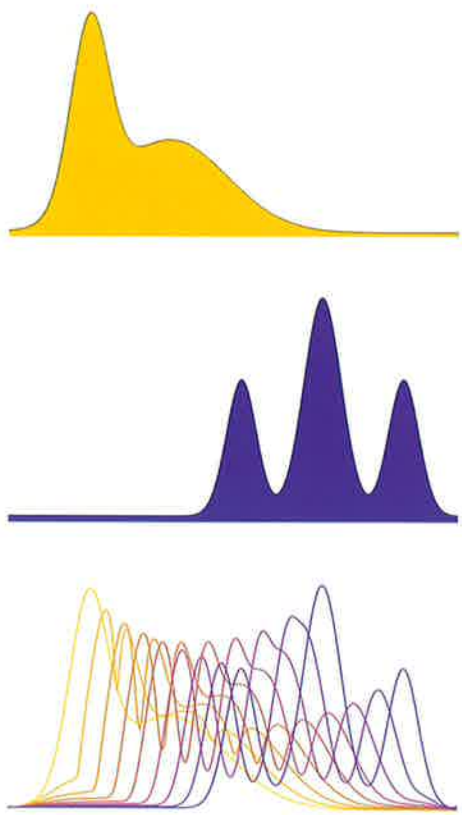

Optimal transport is a method for finding the most efficient way to move mass. Here, two different univariate densities – probability distributions based on single variables – illustrate the method. The gold and purple shaded plots display the respective densities of two different distributions. The 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

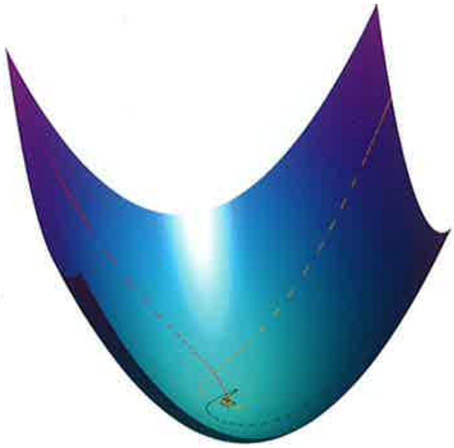

Smoothed versions of the paths taken by various first-order optimization algorithms –mathematical techniques for finding a problem's most efficient solution – testing various candidate solutions. These optimization methods minimize the loss/ objective function (the property the algorithm seeks to minimize 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.