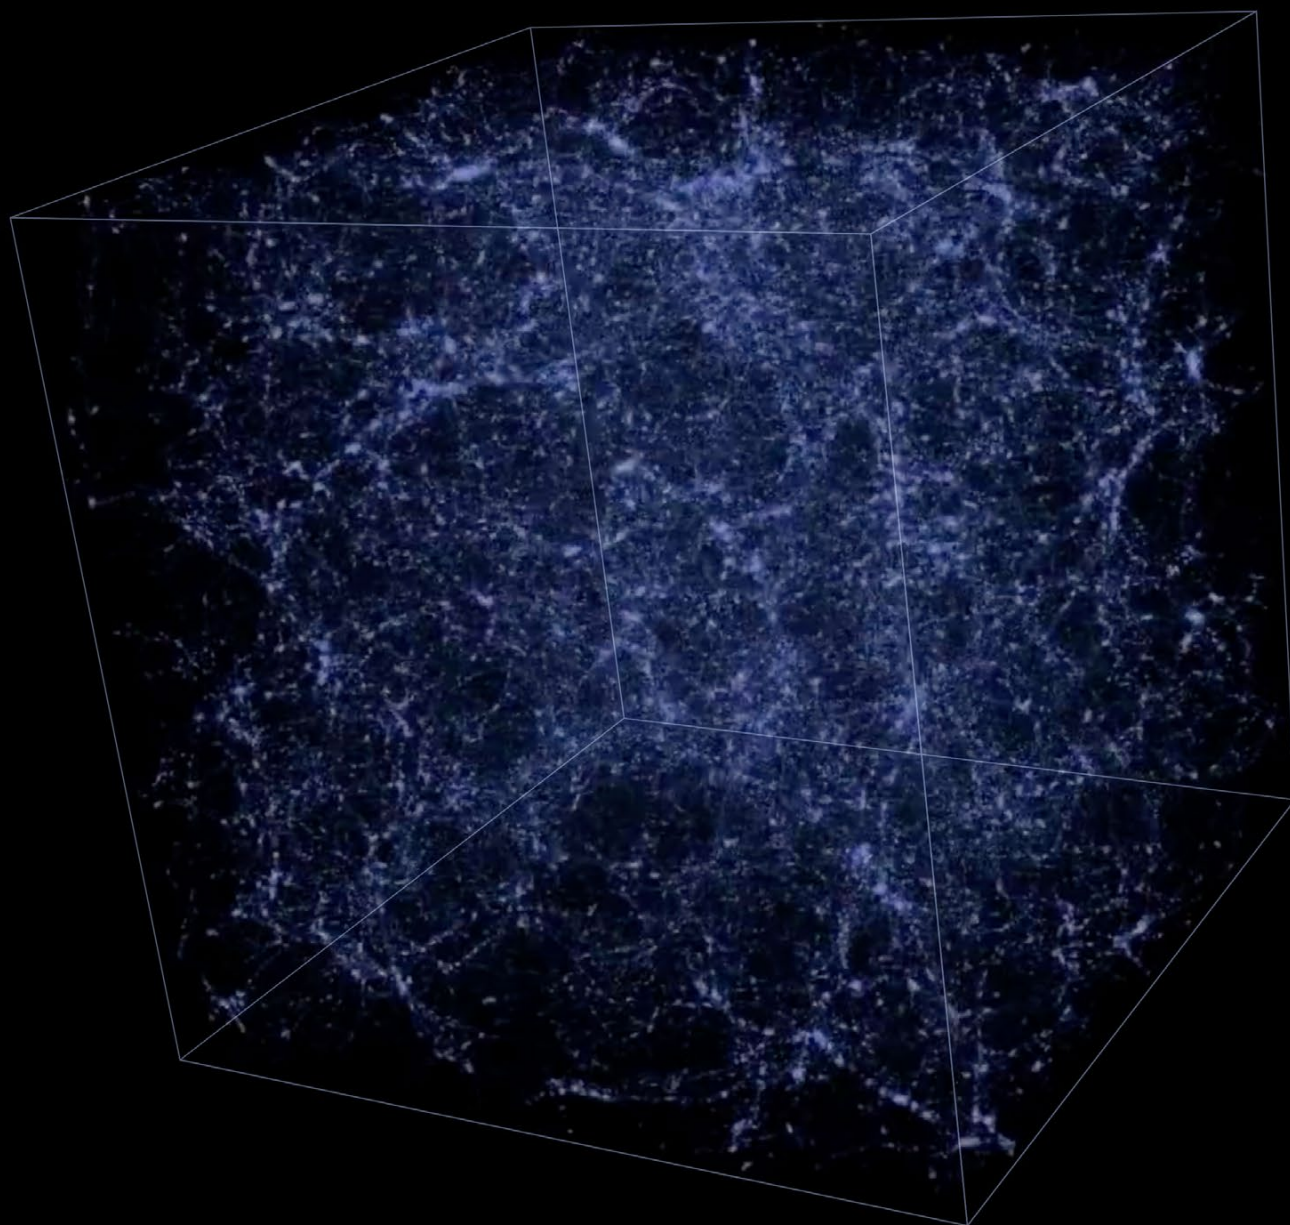


ASCR@40

FOUR DECADES OF DEPARTMENT OF ENERGY LEADERSHIP IN ADVANCED SCIENTIFIC COMPUTING RESEARCH



ASCR@40

ASCR'S MISSION

The Advanced Scientific Computing Research (ASCR) program's stated mission: to discover, develop and deploy computational science and high-performance networking tools and services that enable researchers in scientific disciplines to analyze, model, simulate and predict complex phenomena in energy, the environment and security that are important to the Department of Energy (DOE). To accomplish this mission, ASCR funds research at public and private institutions and at DOE laboratories to foster and support fundamental research in applied mathematics, computer science and high-performance network research. In addition, ASCR supports multidisciplinary scientific research through computational science partnerships with other Office of Science programs and other elements of DOE.

ASCR also operates high-performance computing (HPC) centers and maintains a high-speed network infrastructure (ESnet). The HPC facilities include leadership computing facilities at Oak Ridge National Laboratory, Argonne National Laboratory and the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory.

ABOUT ASCR@40

This document showcases high points in computational science and HPC that are a direct result of four decades of ASCR support and leadership. It is a sister publication to a technical report that covers ASCR's 40-year history, *Highlights and Impacts of ASCR's Programs*. Both documents were produced under the direction of the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee on the 40-year history of ASCR.

ON THE COVER

A frame from the animation "ESnet: Visualizing the Universe at 100 Gbps," which featured a simulation that tracked the evolution of the universe from its homogenous dawn 13.7 billion years ago to today. The simulation was performed on 4,096 cores of NERSC's Cray XE6 Hopper system with the Nyx code. The project produced more than 5 terabytes of data that were transferred to the SC11 supercomputing conference exhibit floor in near real-time on ESnet's then-new 100 gigabit-per-second (Gbps) network, part of ESnet's 25th anniversary celebration. ESnet's network has enabled remotely located scientists to make discoveries from large data volumes located at NERSC and other DOE leadership computing facilities. For more about ESnet, see "Big-Data Networking" on page 24.

ASCAC Subcommittee on the 40-year History of ASCR

Chairs: Bruce Hendrickson and Paul Messina

Members: Buddy Bland, Jackie Chen, Phil Colella, Eli Dart, Jack Dongarra, Thom Dunning, Ian Foster, Richard Gerber, Rachel Harken, Wendy Huntoon, Bill Johnston, John Sarrao, Jeff Vetter

ASCR@40: *Four Decades of Department of Energy Leadership in Advanced Scientific Computing Research*

Editors: Bruce Hendrickson and Bill Cannon


Design: Stilt Studio Inc.

Copy editor: Thomas R. O'Donnell

Contributing writers: Jacob Berkowitz, Steve Koppes, Mike May, Chris Palmer, Sarah Webb

ASCR@40

FOUR DECADES OF DEPARTMENT OF ENERGY LEADERSHIP IN ADVANCED SCIENTIFIC COMPUTING RESEARCH



11

To Take on Turbulence, DOE Applied Mathematicians

DOE's most advanced computers depend on constant algorithm improvements to solve fluid-flow equations—and depend on great mathematicians. Alexandre Chorin's career offers a case study.

16

Supercomputing to Compete

The rise of DOE leadership computing facilities is helping drive American commercial innovation and global competitiveness.

20

Petaflops for the People

ASCR opens its user facilities and expertise to researchers working on some of the world's most difficult scientific computing challenges.

24

Big-Data Networking

DOE's path-breaking ESnet moves science at the speed of light.

29

Discovery from Data

ASCR research produces detailed science simulations and the tools to sift through the massive data files they generate.

33

Embracing Uncertainty

DOE has pioneered the fields of uncertainty quantification, improving simulations' reliability on problems with large economic, environmental and security implications.

4

Executive Summary: Forty Years of Computing Leadership

6

Laying the Foundation for Computational Science

The Advanced Scientific Computing Research program helped create a third pillar of science and continues to shape its future.

37

Rules of the Road

Building a landscape for the high-performance computing community.

41

The Human Element

The rapid growth in high-performance computing and workforce training to fuel it go hand in hand.

EXECUTIVE SUMMARY:

Forty Years of Computing Leadership

Throughout its long history, the Office of Advanced Scientific Computing Research (ASCR) has built the critical technologies to ensure U.S. leadership in energy science and national security. It has made its parent agency, the Department of Energy (DOE) and its Office of Science, the world's recognized leader in computational science.

ASCR's stated mission is "to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the DOE." To accomplish this goal, ASCR oversees a large complex of computing and networking facilities and is responsible for procuring, deploying and operating high-performance computing (HPC), networking and storage resources; conducting basic research in mathematics and computer science; developing and sustaining a large body of software; and collaborating with other Office of Science programs, academia and industry.

ASCR's computational science leadership has a long history, predating even DOE's inception. Applied mathematics and advanced computing were both elements of the Manhattan Project's Theoretical Division. In the 1950s, DOE's predecessor, the Atomic Energy Commission, created a mathematics program to develop and apply digital computing by supporting researchers at universities and AEC laboratories. Several organizational and name changes later, this program would grow and become ASCR.

A sister report, *Highlights and Impacts*, provides a detailed technical history of ASCR and its predecessors. Like this publication, it was produced by the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee for ASCR's 40-year History.

In contrast, this report is primarily intended for a nontechnical audience and aims to showcase the ASCR program's living legacy, stemming from its role building the foundations of computational science and the applied mathematics that make the entire enterprise possible. Most of what you will read over the next 40 pages shines a light on research and related activities since DOE's advanced computing office adopted the ASCR name in 1999, including the highly influential early-and-ongoing SciDAC (Scientific Discovery through Advanced Computing) program and ASCR's more recent partnership with the DOE National Nuclear Security Administration on the forward-looking Exascale Computing Project. Other prominent ASCR programs noted include INCITE (Innovative and Novel Computational Impact on Theory and Experiment) and the ASCR Leadership Computing Challenge, which both make extensive use of ASCR's leadership computing facilities based at Oak Ridge, Argonne and Lawrence Berkeley national laboratories.

Specifically, this document places those programs into context, demonstrating how ASCR programs have increased the United States' computing and industrial competitiveness; pushed the limits of computing power and made leading-edge resources available to national laboratory researchers, universities and the commercial sector; built the premier network for scientific research, extending its reach to scientists overseas; led the push to convert massive amounts of data into scientific discoveries; pioneered uncertainty quantification and other computational science disciplines; established many of the ground rules in use worldwide for HPC architectures, applications and algorithms; and embarked on a decades-long commitment to training the next generation of computational science leaders.

Together with its sister report, a larger picture, built from these many specific contributions, will emerge: The disciplines and activities whose creation ASCR led have produced a whole that is greater than the sum of its parts. ASCR has played a pivotal role in envisioning, delivering and sustaining this rich ecosystem for innovation and discovery. Fundamental mathematical advances applied to software libraries enable unprecedented simulation capabilities. Enormous simulation and experimental data sets are exchanged via information services and collaboration frameworks, managed by input/output tools and analyzed by visualization and analysis algorithms and tools, all developed and deployed with ASCR support. Developments in programming models underpin message-passing libraries and portable performance abstractions to ensure that applications make the best use of ASCR leadership computing facilities.

In the chapter that begins on the next page, "Laying the Foundation for Computational Science," we characterize the leadership-facility supercomputers as "breakthrough machines." They reflect ASCR's accumulated decades of investment in applied mathematics, computer science and scientific-software tools. DOE's Exascale Computing Project has accelerated simulation-code development on an unprecedented scale. Exascale is right around the corner. When it arrives, its scientific and technical advances will owe an enormous debt to ASCR's visionary leadership and sustained investment.

LAYING THE FOUNDATION FOR COMPUTATIONAL SCIENCE

The Advanced Scientific Computing Research program helped create a third pillar of science and continues to shape its future.

A virtual prototype of a fusion reactor's swirling fuel enables scientists and engineers to gain crucial insights that guide reactor design. A model of a human heart's erratic beating suggests new ways to heal cardiac diseases. A multicolored digital simulation of water as it's squeezed from pulp saves hundreds of millions of dollars in energy costs, making U.S. paper companies more competitive and improving the environment.

Half a century ago, all of these real-world examples were considered science fiction. But for more than four decades, the Advanced Scientific Computing Research (ASCR) program and its Department of Energy predecessors have led the world in developing a powerful new tool of discovery: computational science, one of the most significant scientific advances of the 20th century and an indispensable ingredient of discovery in the 21st.

Computational science is the combination of mathematics, software and computer science with high-performance computers (HPC), applied to the nation's most pressing scientific and technical challenges. It lets researchers model, simulate, visualize, prototype and analyze physical phenomena that would be too expensive, too experimentally difficult or simply impossible to do otherwise—from global climate modeling to new energy technologies.

Through a deliberate vision, ASCR has pioneered the creation of a dynamic computational science community and talent pipeline bridging national labs, academia and industry; fostered the development of a mathematics, algorithm and software toolkit for tackling leading-edge challenges; and enabled this research through world-leading supercomputing centers and

Computational modeling is particularly valuable when experiments are prohibitively expensive, dangerous, time-consuming or impossible.

the fastest science network providing access to those facilities. This extensive human and technological infrastructure supports ASCR's Exascale Computing Project, the most ambitious HPC program in U.S. history.

FROM MATH TO COMPUTATIONAL MODELS

In the early 1950s, pioneering computer scientist John von Neumann at Princeton's Institute of Advanced Studies made a suggestion that would change the face of science. He foresaw that computers would be crucial to solving many of the complex technical challenges facing DOE's precursor, the Atomic Energy Commission (AEC). To enable these solutions, von Neumann said, researchers had to understand the heart of the problem: the mathematics underlying computation.

In response, the AEC created a mathematics program. Von Neumann inspired the pursuit of what until then was a fanciful dream: using equations to model and simulate real-world events *ab initio*, from first principles, in computers. Scientists, particularly physicists and engineers, describe the world in mathematical expressions, from Schrödinger's famous wave equation encapsulating an electron's behavior, to the formulae of fluid dynamics. Ever more powerful computers offered the possibility of doing the otherwise impossible: use mathematics to test a car's crash-worthiness and predict airflows over aircraft wings.

This is what ASCR has led the world in accomplishing. Until the advent of computational science, researchers could tackle questions about how the world works with either experiments (for example, lab testing how materials behave in an engine) or with theory (as in the development of quantum physics). Computational science has added a third pillar to researchers' toolkits: the ability to develop detailed, predictive models *in silico*.

Computational modeling is particularly valuable when experiments are prohibitively expensive, dangerous, time-consuming or impossible. For example, ASCR-supported research in data-intensive computational chemistry has been used to discover promising new drug designs and develop accurate, long-range global climate models, feats made possible only through the advent of computational science.

ASCR is unique among the six DOE Office of Science programs—it focuses on providing computational science resources to advance—and inspire—research in all DOE-mandated areas and beyond. As a result, ASCR has facilitated an idea-to-discovery computational science ecosystem that leads from equations to algorithms for HPC codes and programs in modeling and simulation that provide unparalleled discovery in fields from astrophysics to biology.

A SUPERCOMPUTER NEXT DOOR FOR EVERY SCIENTIST

In the 1980s, DOE leadership recognized trends that would define ASCR's evolution: growing demand among a wide group of DOE researchers for computational science resources to tackle the toughest technical problems and a blossoming of supercomputing technologies that paralleled the rise of personal computers.

The challenge was clear: While supercomputers required centralized, specialized facilities and expertise to operate, DOE researchers and experimental facilities were distributed across the United States.

DOE's solution turned out to be transformative. The agency devised a first-of-its-kind nationwide infrastructure of supercomputing centers connecting DOE scientists and thousands of university researchers via the world's fastest and most advanced scientific network, ESnet (Energy Sciences

Network). (See “Big-Data Networking,” page 24.) In essence, this vision has delivered a supercomputer as if it were in the room next door to all DOE and thousands of the nation’s university researchers.

The first ASCR supercomputing user facility was the National Energy Research Scientific Computing Center (NERSC), now at Lawrence Berkeley National Laboratory (LBNL). In the 2000s, DOE created two more, the Argonne Leadership Computing Facility (ALCF) and the Oak Ridge Leadership Computing Facility (OLCF). (See “Petaflops for the People,” page 20.)

Computational science advances are driven, in large part, by faster, more powerful computers, and ASCR’s supercomputer user facilities have led the world in providing researchers in

PROVIDING THE HPC COMPUTATIONAL SCIENCE TOOLKIT

Everyone knows that a laptop is only as useful as its software—the computer instructions that enable us to perform tasks. The same is true for supercomputers. For decades, mathematicians and computer scientists from ASCR and its predecessor programs have created a unique and growing HPC toolkit that provides the intellectual bedrock for computational science.

This HPC math and software repository enables computational science researchers to plug-and-play with ready-to-use packages of advanced algorithms and software, enabling research results in days or weeks that would otherwise take months to years to develop independently—if it could be done at all.

For the past 40 years, DOE centers have dominated the industry-tracking TOP500 list of the world’s most powerful supercomputers.

national laboratories, universities and industry with ever more powerful leadership-class machines.


For the past 40 years, DOE centers have dominated the industry-tracking TOP500 list of the world’s most powerful supercomputers. Notably, in the past decade, DOE-ASCR facilities have developed three of world’s No. 1 most powerful machines. In 2009, OLCF launched Jaguar, only the second computer to break the petaflops barrier (after LANL’s Roadrunner), eclipsed in 2012 by the center’s Titan and in 2018 by Summit, also at OLCF. In just the past 15 years, the computing power these three national centers provide has increased more than a hundred-thousand-fold, leading to extraordinary advances in physics, chemistry, materials science, climate prediction, biology and medicine.

To provide remote access to HPC user facilities, in 1986 ESnet grew from the merger of two specialized DOE networks, operating at 56K modem speed, that already served scientists in fusion and high-energy physics research. Today, ESnet5, the LBNL-managed network’s fifth iteration, is a many-hundreds of gigabits per second optical network connecting more than 50 DOE, government and academic facilities and integrated into other international research and education systems. ESnet5 moves data about 15,000 times faster than the average home network.

As a case in point, beginning in World War II, one of the central physics challenges was numerically calculating and predicting the dynamics of combustion and a contained explosion, such as in an engine’s piston. These complex events involve multiscale physics: from billionth-of-a-second atomic interactions up to visible-scale turbulent dynamics that happen over a period a billion times longer. To efficiently model these problems on supercomputers, ASCR researchers and colleagues developed a mathematical modeling technique called Adaptive Mesh Refinement (AMR) that lets researchers change the resolution of the computational mesh automatically as a function of space and time. This has led to extensive advances in fields from astrophysics to combustion.

ASCR researchers have packaged mathematical insights into solvers, software that includes libraries of tools for cracking computational science-related problems. LAPACK, one of the most successful ASCR-produced software packages, has profoundly influenced the development of computational science. It’s a cornerstone of software libraries provided by hardware vendors (including IBM, Intel, Cray and NVIDIA), and by software companies from MathWorks to Red Hat.

Similarly, the ASCR-supported SuperLU package, designed to more efficiently solve numerical systems coming from partial



Radial velocity field (red is positive; blue, negative) on the equatorial plane of a numerical simulation of Earth’s core dynamo performed at the Argonne Leadership Computing Facility. These small-scale convective flows generate a strong planetary-scale magnetic field. Credit: Rakesh Yadav/Harvard University.

differential equations, has become a computational science workhorse. First released in 1997, it is downloaded more than 30,000 times a year and dozens of companies have adapted it into their simulation software, including AMD for circuit modeling, Boeing for virtual aircraft prototyping and Walt Disney Animation for creating entire digital worlds on the silver screen.

PEOPLE, TEAMS, PARTNERSHIPS

Computational science is inherently interdisciplinary and requires access to and training on high-performance computers. Thus, from its beginnings, ASCR has fostered the creation of a nationwide computational science community and talent-development pipeline that bridges academia, industry, national laboratories and international partnerships.

Predating many university computational science departments, the DOE Computational Science Graduate Fellowship (DOE CSGF) has played a key role in seeding and catalyzing the national computational science community. (See “The Human Element,” page 41.) Established in 1991, fellows are supported in doctoral degree research in either mathematics, computer science or science or engineering disciplines. Notably, DOE CSGF participants must take courses in all three domains, serve a computational science research practicum at a DOE lab and get DOE HPC access and mentoring.

To spur collaborative research between professional computational scientists, computer scientists and mathematicians, in 2001 ASCR launched the Scientific Discovery through Advanced Computing (SciDAC) program, now an international model for driving the

ASCR researchers have packaged mathematical insights into solvers, software that includes libraries of tools for cracking computational science-related problems.

development of computational science. SciDAC has fostered hundreds of team-based collaborations between scientists exploiting computation to achieve breakthroughs across all Office of Science programs as well as at the National Nuclear Security Administration and Office of Nuclear Energy.

In 2004, the DOE High-End Computing Revitalization Act extended ASCR resources to researchers in U.S. industry on a competitive, merit-reviewed basis. In response, ASCR created three mechanisms to give government, academic and industry investigators access to DOE's best open-science HPC facilities: INCITE (Innovative and Novel Computational Impact on Theory and Experiment); the ASCR Leadership Computing Challenge (ALCC); and the Director's Discretionary Program. The three pathways enable researchers to compete for time allocations at NERSC, OLCF, and ALCF.


Since its creation, ASCR's laboratory-industrial partnership programs have let hundreds of companies, from Fortune 500 giants such as Procter & Gamble to innovative start-ups, solve otherwise intractable problems and build the base of advanced HPC expertise. In this way, ASCR has helped American industry become more internationally competitive.

A FUTURE OF EXASCALE—AND BEYOND

Rising from roots in the earliest origins of computing, today ASCR-led computational science—the third pillar of science—is a foundation for the largest such initiative ever, the Exascale Computing Project (ECP), a joint effort of the Office of Science and DOE's National Nuclear Security Administration.

ASCR, in close coordination with the ECP but supported separately, is on track to assemble the nation's first exascale systems at the Argonne Leadership Computing Facility and at the Oak Ridge Leadership Computing Facility in 2021.


Researchers are ready to use these breakthrough machines because of ASCR's accumulated decades of investment in creating the mathematical underpinnings, the computer science capabilities and the software tools for computational science. The ECP is accelerating the development of simulation codes for exploring scientific frontiers across a wide range of disciplines. And when it arrives, exascale computing—and all of the scientific and technical advances that emerge from it—will owe its success to ASCR's sustained and visionary investment.



The Cosmic Web, the calculated velocity field of matter and galaxy clusters in the universe, from a model run at Oak Ridge Leadership Computing Facility. The color corresponds to the cosmic matter density field. Credit: Miguel Aragon, Johns Hopkins University/Universidad Nacional Autónoma de México.

TO TAKE ON TURBULENCE, DOE APPLIED MATHEMATICIANS

DOE's most advance computers depend on constant algorithm improvements to solve fluid-flow equations—and depend on great mathematicians. Alexandre Chorin's career offers a case study.




Many people know turbulence mainly as a bump here and there, sometimes more, during a plane ride. To others, though, it's a complicated phenomenon that's challenged the best minds for centuries. Beyond some choppy air and a mental puzzle, turbulence matters in many practical ways, from predicting the weather to designing aircraft. If mathematicians and physicists could subdue this chaotic fluid flow, many realms of technology would be improved.

As early as 1509, Italian polymath Leonardo da Vinci drew pictures depicting the turbulence generated as fluid passed an object. Nobel laureate Werner Heisenberg, a German physicist and pioneer in quantum mechanics, took a crack at the mathematics behind the phenomenon but made little headway. Another Nobel winner, American theoretical physicist Richard Feynman, once call turbulence "the most important unsolved problem of classical physics."

For more than 40 years, turbulence has been on the mind of Alexandre Chorin, University Professor emeritus of mathematics at the University of California, Berkeley, and a longtime co-appointee at Lawrence Berkeley National Laboratory. "People say that

Nitric oxide emissions (the colors) in a lean hydrogen-air mixture simulation and the turbulence at the flame (gray) as a low-swirl injector's flow breaks down. Credit: Lawrence Berkeley National Laboratory.



Detail from a model of a nuclear reactor rod bundle's turbulent velocity magnitude. Warmer colors represent a higher velocity, colder colors a lower velocity. Researchers are particularly interested in multiple scales of turbulence. Credit: Argonne National Laboratory.

turbulence is an unsolved problem, but the equations of motion have been known for two centuries,” Chorin explains. “The obstacle is complexity.”

Chorin describes his long association with Berkeley Lab as optimal because it “gives you space and easy access to computers if you have a DOE grant.” Plus, financial support and joint appointments “give you the opportunity to work on a project continuously, of having students work without having to teach all that much, and it provides a whole slew of interesting collaborations.”

Although Chorin no longer teaches, his lab produced more than 50 Ph.Ds. “I really enjoyed having students, and I spent a lot of

time with them and talking to them,” he says. “Mostly, I tried to work with them by gauging what they were capable of doing, suggesting some problem and encouraged them to go in some new direction toward that question. I think that has been a successful strategy.”

SOLVING PROBLEMS

Chorin faced problems—not always mathematical ones—from an early age. He was born to a Jewish couple in Poland the year before Germany invaded in 1939. The family soon fled to Palestine. Amid that chaos, the boy stumbled into equations through a friend of the family, perhaps a distant relative. Chorin isn't sure.

‘I’ve been involved in a whole collection of projects on how to reduce the complexity of the fluid mechanical calculations.’

Either way, the mentor had a Ph.D. in mathematics but was a lawyer by trade. “He started asking questions of a mathematical type before I went to grade school,” Chorin recalls, “and I got this notion about being a mathematician before it meant anything concrete to me. I was telling people that I was going to be a mathematician when I was 8 or 9.”

After he turned 12, Chorin moved with his family to Switzerland, where academics got a little bumpy for him. Although he left high school in his junior year and took an equivalency exam instead, he went on to earn a degree in engineering physics at the École Polytechnique Fédérale de Lausanne. From there, Chorin moved to New York University’s Courant Institute to work on his Ph.D. with famed mathematician Peter Lax.

There, Chorin began working with the equations of motion that he mentioned—the Navier-Stokes equations, which describe how fluids (like air and water) behave. As he says today, “The essence of the turbulence problem is: How do you control the complexity of turbulent Navier-Stokes?”

REDUCING COMPLEXITY


In 1971, Chorin moved from the Courant Institute to Berkeley. He had developed a Navier-Stokes solver in New York and hoped it would help him get at turbulence. It didn't. The equations

MATH SPOTLIGHT: AIR FLOW

SIMULATING SHOCK WAVES

At supersonic speed, air flow over a jet plane turns chaotic, building shock waves—pressure surges that can cause deadly damage to the aircraft. To build structures that survive these bursts or, better yet, produce fewer of them, scientists simulate air flow across a range of possible lengths and times—micrometers to meters and microseconds to seconds, respectively.

In 2018, scientists used Mira, an IBM Blue Gene/Q supercomputer at the Argonne Leadership Computing Facility, to examine shock waves at a range of scales. The results provided suggestions for adjusting wings on the fly—really on the fly—to counteract some shock-triggering flow changes. But there's much more to learn.



A Rayleigh-Taylor instability caused by two mixing fluids, visualized in a pseudo-color plot.
Credit: Lawrence Livermore National Laboratory.

remained too convoluted. “So I’ve been involved in a whole collection of projects on how to reduce the complexity of the fluid mechanical calculations,” he says.

Even early in his career, grants funded Chorin’s efforts. At New York University, they came from the energy and defense agencies to support his Navier-Stokes research. In the early 1970s at Berkeley, Chorin received another DOE grant, which together with support from the National Science Foundation has sustained his work ever since.

Computing has been one of Chorin’s key tools almost from the beginning. “I started programming back in 1960,” he says. “In fact, I first programmed in machine language, and you fed the computer by tape, which you had to push into a hole.”

Some of his early computing included visualizing data. “At the time, if you wanted to make pictures, you had to write your software.” In 1964, he wrote packages for software portrayals of vortex motion as it passed an object. “So I visualized the centers of the vortices as dots—very primitive. It seemed to be fine, because journals accepted it without complaining.”

Despite easy access to powerful computers—even supercomputers—Chorin uses a desktop PC for his trial-and-error programming. He turns to bigger machines for production runs to address a particular question. “The calculation is interesting in general, but the really large version of that is not what I do. It’s what other people do.”

MATH SPOTLIGHT: CO-DESIGN

PUTTING MANY EXPERTS AT THE TABLE

As computational science grows ever more complex, especially in the move toward exascale, ASCR advances many programs through co-design. This approach assembles a broad team of experts, including hardware developers, programmers, applied mathematicians like Chorin, and others, including subject area specialists such as physicists and chemists. The scientists on a team bring their own perspectives and expertise, so together they can solve problems that none of them could tackle alone.

Co-design can be used for a range of projects, up to the biggest ones. For example, ASCR is using co-design in the runup to the Frontier supercomputer, which Cray Inc. is building for Oak Ridge National Laboratory (See “Laying the Foundation for Computational Science”, page 6.). This supercomputer will rely on various technologies, including some programming tools, that were developed via co-design.

By using a connected team of disparate experts to build parts and platforms, ASCR and others can enhance the odds that the resulting technology is applicable to the desired questions. Moreover, the process develops algorithms, software tools and hardware architecture in parallel to make them work together. Otherwise, even the most beastly supercomputer might not be tamed.

From the 1960s to today, Chorin and his colleagues watched every aspect of computational science evolve, sometimes via gigantic leaps that a team approach called co-design enabled.

That said, the programs he writes are complex, requiring thousands of lines of code. Tracking what they all do depends on taking notes along the way and adding careful comments in the programs themselves. “I’m not making much of an effort to optimize the program.” Other scientists focus on making computations more efficient. (See “More from Less,” this page.)

From the 1960s to today, Chorin and his colleagues watched every aspect of computational science evolve, sometimes via gigantic leaps that a team approach called co-design enabled. (See “Putting Many Experts at the Table,” previous page.) Both algorithms and hardware have improved over his career, Chorin says, “but my impression is that the algorithms have improved more than computers. Most people have a hard time believing that.” For example, he notes that the approach he used in the mid-1960s to solve the Navier-Stokes equations “is unbelievably slow, by modern standards.” Even with today’s far-faster computers, he says that programs would still run much slower—some big projects taking seemingly forever—if not for the speed better algorithms deliver.

STILL TRACKING TURBULENCE

While living through and contributing to some of those computational advances, Chorin has kept his thoughts on turbulence. “I’ve been working on it off and on for a very long time, and I haven’t solved the problem, but I’ve found applications of some of the tools—in data assimilation, model reduction, things of that kind.”

When asked why he’s gone back and forth between turbulence and other topics, Chorin says, “You don’t want to constantly work on something that doesn’t work. You look for things you can do with it, which are feasible with the machinery.” For instance, he developed one of the earliest algorithms for the Boltzmann equation, which can be used to analyze transport in fluids and other things. He’s also worked on less turbulence-related topics, like shock-wave simulations. (See “Simulating Shock Waves,” page 13.)

And maybe—just maybe—he’s near a step forward in turbulence, he says. “I’m working on something that I’ve been pursuing for many years, which may come to fruition fairly soon.”

MATH SPOTLIGHT: MULTIGRIDS

MORE FROM LESS

In their quest to simplify the computations behind many simulations, multigrids are one method scientists turn to.

Imagine mathematically simulating a land mass by first throwing a net over it. That makes a grid, and physics equations are used to collect data at each crossing point. Multigrid adapts the net’s coarseness in response to how sharply and quickly the elevation is changing.

The coarser grids reduce the amount of computation, making it easier for a computer to work on the problem or simulation. But going willy-nilly to a coarser grid with no plan won’t work right. Instead, programmers use a mathematical solver to keep the error as low as possible while making the simplifications.

There are all sorts of multigrid methods, including algebraic and geometric ones. DOE’s Advanced Scientific Computing Research (ASCR) program funded a project in which scientists from Columbia University and Sandia National Laboratories used a multigrid method to simulate an ice shelf breaking up—something happening more frequently as our world warms. The researchers also used an ASCR-backed program, Scientific Discovery through Advanced Computing (SciDAC), to develop the multigrid methods they used. It’s another example of the value of making tools and techniques as widely available as possible.



SUPERCOMPUTING

TO COMPETE

The rise of DOE leadership computing facilities is helping drive American commercial innovation and global competitiveness.

A team of researchers from U.S. insurer FM Global watched as a devastating fire raced through a new, state-of-the-art mega-distribution center. In under two minutes, intense flames reached seven stories, incinerating inventory worth tens of millions of dollars. Yet for all the destruction, this conflagration was completely under control—and helpful. It was a precisely detailed simulation on what was at the time one of the world's most powerful supercomputers, Titan, at DOE's Oak Ridge Leadership Computing Facility (OLCF).

Results from this first-of-its-kind virtual fire modeling have let FM Global, one of the world's largest property and casualty insurers, better protect clients and save tens of millions of dollars in payouts.

It's also one example among hundreds of how businesses from start-ups to Fortune 500 corporations are partnering with ASCR for access to DOE high-performance computing (HPC) capabilities and to competitive advantages in aerospace, life sciences, finance, manufacturing and more.

The ASCR program leads the nation and the world in supercomputing, high-end computational science and advanced

networking for research. By providing access to and support for these leadership-class HPC resources, ASCR enables major commercial research breakthroughs leading to transformative products, services and technologies—ones that are helping U.S. companies maintain a secure foothold in today's competitive global market.

PARTNERSHIPS FOR GLOBAL COMPETITIVENESS

Today's dynamic partnerships between DOE supercomputing centers and American companies spring from a deliberate decision—and vision—to extend the department's world leadership in HPC beyond the nation's pressing security and energy issues to include economic competitiveness.

In 2004, Congress' DOE High-End Computing Revitalization Act declared that, besides university and federal scientists, ASCR resources could be made available to researchers in U.S. industry on a competitive, merit-reviewed basis.

The act recognized that HPC and related advanced computing techniques drive leading-edge innovation, including in virtual prototyping, high-resolution modeling and data-driven business

and technology, yet few U.S. companies can afford a leadership-scale supercomputer costing hundreds of millions of dollars.

The act also reflected HPC's new place in global competitiveness. As a 2017 NSF-funded report, "Worldwide Best Practices in Partnerships between HPC Centers and Industrial Users," concluded: "In the race to make HPC more pervasive in industry in order to boost economic competitiveness, no nation or global region can afford to be complacent because no nation or region has a large, sustainable lead over all the others."

In response, ASCR created three mechanisms to give government, academic and industry researchers access to DOE's best open-science HPC facilities: INCITE (Innovative and Novel Computational Impact on Theory and Experiment), the ASCR Leadership Computing Challenge (ALCC) and the Director's Discretionary Program. The three pathways enable researchers to compete for time allocations at OLCF, Lawrence Berkeley National Laboratory's National Energy Research Scientific Computing Center (NERSC) and the Argonne Leadership Computing Facility (ALCF). Each center, designated Office of Science user facilities, has an Industrial Partnership Program and point person to guide companies through these access opportunities. Applications undergo a rigorous peer-review process to identify those with the best prospects to advance key areas in science and engineering.

HPC: MAKING THE IMPOSSIBLE DOABLE

Firms that win access to these HPC user facilities are focused on strategic, high-risk and high-return industrial applications. While solving these problems, companies large and small also advance their ability to use leadership-class HPC and its derived technologies.

For example, when General Electric (GE) and Pratt and Whitney researchers were separately seeking ways to create jet engines that are quieter and more energy efficient, they both took their problems to the ALCF for virtual prototyping. Physical tests of complex systems such as jet engine combustion are difficult, expensive and time-consuming. High-resolution, three-dimensional virtual trials enable companies to rapidly predict design performance and thus adjust schemes on the fly before committing to physical prototypes.

A jet engine combustor combines air flowing faster than a hurricane with swirling fuel to generate the extraordinary release of heat that ultimately powers the aircraft. Using ALCF's IBM Blue Gene-class supercomputers, GE researchers modeled turbulent flow features—measuring quantities such as velocity, temperature and pressure—to predict the noise different exhaust nozzle designs would produce. Similarly, Pratt and Whitney researchers reduced the solution times for their 3-D combustor simulations and provided key insights into the design of their next-generation engines. Quieter, more fuel-efficient engines are aiding these

ASCR's laboratory-industrial partnership programs enable firms to solve otherwise intractable problems and to build a base of advanced HPC expertise.

ASCR's laboratory-industrial partnership programs enable firms to solve otherwise intractable problems and to build a base of advanced HPC expertise. The ASCR collaborations provide DOE scientific and technical experts in a range of computational fields who help industry researchers learn about, troubleshoot and maximize their use of leadership-class computers. Meanwhile, through technology transfer and both open-source and fee-based licensing, ASCR-derived HPC software is similarly boosting U.S. companies' profitability and competitive edge.

companies, and their customers, in reaching greater heights in the extremely competitive airline and jet engine global market.

When companies face new market challenges, they often must scale up existing in-house computer models onto leadership-class HPC systems to gain the insights necessary to address them. When FM Global scientists wanted to simulate fires in their client's new mega-warehouses they expanded their in-house FireFOAM code on OLCF's Titan through a Director's Discretionary allocation.

ASCR supercomputers produce world-leading airflow simulations, like this one done at the Oak Ridge Leadership Computing Facility (OLCF), to improve efficiency in power-generation and transportation. A small South Carolina company, SmartTruck, used OLCF to show that when airflow is forced to veer off a semi's trailer toward the wake behind, drag is greatly reduced. Credit: SmartTruck.

THE DIFFERENCE LEADERSHIP HPC MAKES

For many American companies, the most complex, mission-critical challenges they face require true leadership-class computers—machines no company in the world possesses.

Small and medium-sized companies simply can't afford to invest hundreds of millions of dollars. And at America's largest companies, including those with extensive in-house HPC resources, "(even) for their advanced R&D, the kind that targets breakthrough innovations capable of transforming entire industries, [they] cannot justify buying these expensive (leadership-class) computers," a 2017 NSF-funded study explains.

As a result, ASCR's Industrial Partnership Programs at Oak Ridge, Argonne and Lawrence Berkeley National laboratories provide qualifying U.S. companies access to computational resources that aren't merely a little faster and more powerful than in-house ones; they are in a stratospheric category at the frontier of computing, innovation and competitiveness. For example, Summit, an IBM-built system at the Oak Ridge Leadership Computing Facility (OLCF), an Office of Science user facility, is the nation's fastest supercomputer.

Two even more powerful systems are planned for deployment in 2021: Frontier, at OLCF, and Aurora, at ALCF. Both are projected to be capable of a quintillion (one billion billion) calculations per second—exascale—giving U.S. companies who successfully compete for access a new leadership-class computing edge.

FireFOAM is the company's flagship software for simulating in fine time and spatial resolution the complex physics that occurs during an industrial fire. An OLCF combustion computational scientist provided expertise on how to efficiently simulate a seven-tier, 35-foot high storage area—something that would have been impossible on FM Global's in-house HPC system. Company scientists discovered that stacking storage boxes on wooden pallets slows horizontal flame spread, enabling them to offer better fire protection for clients (a third of all Fortune 1000 companies) and saving both insured and insurer tens of millions of dollars.

Similarly, environmental risk start-up KatRisk used Titan to scale-up its software and create the first high-resolution 10-by-10-meter flood risk maps for the United States and 90-by-90-meter maps worldwide. The model helped establish KatRisk as one of the nation's leading catastrophic risk modeling companies, giving insurance providers and U.S. public agencies critical information for analyzing flood danger at the scale of city blocks.

SCALING NEW FRONTIERS, FROM FUSION TO HAIRCARE

In many cases, ASCR HPC time allocations and expertise enable companies to push the boundaries of current technologies by first virtualizing and then optimizing designs toward someday achieving what's now physically impossible.

Such is the case in the push for fusion energy. California-based TAE Technologies' ambitious goal is to develop the world's first commercial fusion-powered generator for carbon-free electricity. A key part of TAE's plan is creating a first-of-its-kind Whole Device Model (WDM), or virtual model, of field-reversed configuration plasmas. To get there, TAE used time on ALCF's Theta supercomputer to build and test the model's two most computationally intensive components.

Fusion energy's central challenge is containing the plasma, the power source's key component, at sun-like temperatures for long enough to sustain reactions. On Theta, TAE computational scientists ran massively parallel particle-in-cell (PIC) high-fidelity micro-turbulence simulations vital to understanding, and thus predicting, how heat loss scales with plasma temperature. The results are letting TAE optimize the design of its next-generation experimental device, bringing the company a step closer to its overall goal.

The breakthrough benefits of HPC modeling and simulation on the world's fastest computers extend to seemingly simple everyday goods like shampoo and laundry detergent. Scientists

A single failed well can cost up to \$100 million and six months of lost time.

at consumer products giant Procter & Gamble, in collaboration with Temple University researchers, used OLCF's Jaguar supercomputer (Titan's predecessor) to address a key problem in product shelf life and performance. Many of P&G's cleaners are made of lipids, molecules that combine to form hollow, sphere-shaped vesicles. These are used to carry perfumes, dyes and other active ingredients. However, over time, vesicles can fuse, causing phase separation and product performance problems. Using 69 million core-hours on Jaguar over two years, company scientists made the first large-scale detailed and realistic simulations of vesicle fusions. The landmark simulation has enabled P&G scientists to further tie the microscopic properties of vesicles to the macroscopic properties of products and thus improve the performance and shelf-life of fabric softeners, lotions and more, helping the company wash its hands of poor performers and identify those that gleam with potential.

LEADERSHIP-CLASS SOFTWARE FOR INDUSTRY

In supporting companies' global competitiveness with leadership-class HPC, sometimes it's the advanced software, not the raw computing speed, that counts. ASCR is the world leader in developing such codes and sharing them to provide solutions to a panoply of commercial challenges.


For example, until recently the seismic imaging used to find fossil fuel resources under the sea couldn't distinguish between a reservoir of oil, gas, water or brine. It was a costly problem: A single failed well can cost up to \$100 million and six months of lost time. Now, the Electromagnetic Geological Mapper (EMGeo) software that researchers at Lawrence Berkeley National Laboratory (LBNL) developed lets the world's largest oil companies, including ExxonMobil, ConocoPhillips and Chevron, pinpoint hydrocarbons with unparalleled accuracy. To develop EMGeo, researchers used NERSC supercomputers to interpret seismic data in combination with huge data sets of electromagnetic measurements collected at the same time. This produces a subsurface view with unprecedented scale and resolution. EMGeo, which can be run on smaller in-house computers, has been licensed to more than 10 companies, reducing the environmental impact of unnecessary drilling while boosting bottom lines.

MAKING PAPER WITH LESS ENERGY, MORE PROFIT

How can U.S. companies more efficiently squeeze out water when manufacturing paper, saving tens of millions of dollars in energy spent drying it? The answer: use an ASCR supercomputer. A consortium of U.S. paper manufacturers, the Agenda 2020 Technology Alliance, is collaborating on the problem with computational scientists at Lawrence Berkeley National Laboratory's National Energy Research Scientific Computing Center (NERSC) and with Lawrence Livermore National Laboratory. The partnership is part of the DOE-funded HPC4Mfg program, in which DOE researchers work with industry partners to solve computationally challenging manufacturing problems in clean energy, energy efficiency and resource conservation.

The paper-making industry ranks third among the nation's largest energy users, behind only petroleum refining and chemical production, and drying wet paper is a main factor. Using computational fluid dynamics expertise and 60,000 cores on a NERSC supercomputer, laboratory computational scientists developed a full-scale microflow model. It simulates the complex pore structures in felts used to absorb water from paper before air-drying—information that's guiding the design of new manufacturing equipment.

HPC-derived insight is helping the Agenda 2020 Technology Alliance reach its goal of 20 percent energy reduction—which would save U.S. paper manufacturers as much as \$250 million a year.



PETAFLUPS FOR THE PEOPLE

ASCR opens its user facilities and expertise to researchers working on some of the world's most difficult scientific computing challenges.

EXTREME-WEATHER NUMBER-CRUNCHING

Certain problems lend themselves to solution by computers. Take hurricanes, for instance: They're too big, too dangerous and perhaps too expensive to understand fully without a supercomputer.

Using decades of global climate data in a grid comprised of 25-kilometer squares, researchers in Berkeley Lab's Computational Research Division captured the formation of hurricanes and typhoons and the extreme waves that they generate. Those same models, when run at resolutions of about 100 kilometers, missed the tropical cyclones and resulting waves, up to 30 meters high.

Their findings, published in *Geophysical Research Letters*, demonstrated the importance of running climate models at higher resolution. Better predictions of how often extreme waves will hit are important for coastal cities, the military and industries that rely upon shipping and offshore oil platforms. The study depended upon NERSC's data-crunching power.

In related research, NERSC is enabling extreme-scale deep learning, a relatively new form of artificial intelligence based on neural networks, to advance many data-rich scientific fields. This includes the ability to model climate change with unprecedented speed and accuracy. At the scale of just one petaflop, NERSC's Cori supercomputer can perform a quadrillion floating-point operations per second. But Cori can go much faster when using hardware for lower-precision arithmetic developed for machine learning applications, having attained a deep-learning milestone of 15 petaflops in 2017.

Deep learning can help detect extreme weather patterns in climate simulations to understand how those patterns will change under warming conditions. Climate scientists need simulations that span many centuries, and supercomputers trained to find patterns with deep learning can sift through those tsunamis of data to gain insight into real-life tsunamis.

Thousands of researchers have used facilities of the Advanced Scientific Computing Research (ASCR) program and its Department of Energy (DOE) computing predecessors over the past four decades. Their studies of hurricanes, earthquakes, green-energy technologies and many other basic and applied science problems have, in turn, benefited millions of people. They owe it mainly to the capacity provided by the National Energy Research Scientific Computing Center (NERSC), the Oak Ridge Leadership Computing Facility (OLCF) and the Argonne Leadership Computing Facility (ALCF).

These ASCR installations have helped train the advanced scientific workforce of the future. Postdoctoral scientists, graduate students and early-career researchers have worked there, learning to configure the world's most sophisticated supercomputers for their own various and wide-ranging projects. Cutting-edge supercomputing, once the purview of a small group of experts, has trickled down to the benefit of thousands of investigators in the broader scientific community.

Today, NERSC, at Lawrence Berkeley National Laboratory; ALCF, at Argonne National Laboratory; and OLCF, at Oak Ridge National Laboratory, serve as DOE's primary scientific computing centers. They provide researchers in industry, academia and national laboratories world-class computing capabilities that enable advances that can't be done anywhere else. In addition, ESnet—also based at Berkeley Lab—delivers the networking resources needed to share vast amounts of data between scientists around the world. (See "Big-Data Networking," page 24.)

NERSC was the first ASCR computing facility but started in 1974 with another name at another lab under the auspices of a sister organization—with a borrowed computer. Founded as the Controlled Thermonuclear Research Computer Center at Lawrence Livermore National Laboratory, the installation used a 10-year-old Control Data Corporation CDC 6000 to simulate plasma behavior in a fusion reactor in support of DOE's newly created Magnetic Fusion Energy program. As the first national computer center dedicated to unclassified research, it became the model for others that followed.

NERSC's name changed in 1976 to the National Magnetic Fusion Energy Computer Center. Its role expanded in 1983, when it was renamed again and began providing general computer services to all DOE programs funded by what is now DOE's Office of Science. NERSC acquired its current name in 1996 when it moved to Berkeley Lab.

The name change reflected a new philosophy. The idea was to improve the productivity of scientific computing, also known as high-performance computing (HPC) and supercomputing, while making it accessible to the broader Office of Science community.

Now NERSC makes its home in state-of-the-art Shyh Wang Hall with the rest of Berkeley Lab's Computing Sciences organization, which includes ESnet and the Computational Research Division. The center supports scientific computing as it applies to a wide range of problems in combustion, climate modeling, fusion energy, materials science, physics, chemistry, computational biology and other disciplines. Each year it provides high-performance scientific computing and data resources to 7,000 researchers working on more than 700 different projects at national laboratories and universities. In 2018, its users accounted for more than 2,500 scholarly publications.

In 2004, the Office of Science created the Leadership Computing Facility (LCF) program within ASCR to, as the name suggests, establish U.S. supremacy in HPC. The facilities would eventually deploy world-leading systems and give far-flung researchers access through peer-reviewed proposals. The Office of Science selected an ORNL proposal and ANL as an LCF partner. Together in 2017 they served nearly 2,300 users. The resulting output of each facility reaches about 500 scholarly publications annually.

The OLCF's predecessor was ORNL's Center for Computational Sciences, founded in 1992. From the beginning, the Oak Ridge facilities have contributed to a rapid evolution in scientific computing that has produced a million-fold increase in computing power. They deployed a large, vector-based Cray X1 in 2004 and a 25-teraflops Cray XT3 in 2005.

Scientists have used OLCF systems to expand the scale and scope of their research, solve complex problems in less time, and fill critical gaps in scientific knowledge. Today, simulation has achieved parity with experiment and theory as an essential element of modern science.

ALCF enables large-scale computing projects aimed at solving some of the world's most complex problems.

OLCF has hosted the fastest supercomputer in the world multiple times since 2009, beginning with Jaguar, a Cray XT5 system. In June 2018, Summit was the latest OLCF system to take the title as the world's speediest scientific supercomputer. It has a peak performance of 200,000 trillion calculations per second for modeling and simulation. The OLCF's next system and Summit's successor, Frontier, is anticipated to debut in 2021 with a performance of greater than 1.5 exaflops, certain to be one of the world's fastest computers both for modeling and simulation as well as machine learning.

The ALCF began operations in 2006 with an IBM Blue Gene system. Through substantial awards of supercomputing time and user-support services, the ALCF enables large-scale computing projects aimed at solving some of the world's most complex science and engineering problems. Research conducted there has spanned such diverse scientific areas as studying exploding stars, designing more efficient jet engines and exploring the molecular basis of Parkinson's disease. As a key player in U.S. efforts to deliver exascale computing capabilities, the ALCF also has helped advance scientific computing through a convergence of simulation, data science and machine-learning methods.

One of the nation's first exascale supercomputers, Aurora, is scheduled to arrive at the ALCF in 2021. It will be capable of performing a quintillion calculations per second—about five times faster than today's most powerful machines.

PETAFLOPS SPOTLIGHT: OLCF

SHAKING UP EARTHQUAKE PREPAREDNESS

A multi-institutional team led by the Southern California Earthquake Center (SCEC) has used Oak Ridge supercomputers for much of the past decade to simulate earthquakes at higher wave frequencies and to assess hazards to cities and critical infrastructure.

Beginning with Jaguar, a Cray XT5 system, and continuing on Titan, a Cray XK7, the team has conducted large-scale simulations to predict the intensity of ground shaking at specific sites in earthquake-prone Southern California. The results from individual simulations have informed and helped to improve SCEC's CyberShake platform, the first physics-based probabilistic seismic hazard model, which requires large modeling ensembles. City planners, structural engineers and emergency preparedness officials, in turn, consider CyberShake results when updating infrastructure, building codes and emergency response plans.

SCEC's projects can help mitigate the effects of large earthquakes (including those on the well-known San Andreas fault), saving lives and curtailing damage costs by providing more accurate seismic hazard information for California. Furthermore, the modeling tools the team creates can be applied to other at-risk regions.

As more people move to cities in seismically active areas, the potential costs of devastating earthquakes will continue to increase. The techniques the SCEC team are pioneering improve scientists' understanding of earthquakes, helping civic leaders prepare for rare, disruptive events, and giving engineers new tools to predict strong ground motions and their effects on the built environment.

Impact studies following SCEC's simulations helped prompt Los Angeles officials to develop a new citywide earthquake resilience and preparedness plan, Resilience by Design. And, although seismic prediction is still in its infancy, the simulation data the SCEC team produces is being used to train machine-learning algorithms as part of California's Earthquake Early Warning system, designed to increase the time citizens have to respond to the threat of an imminent quake.

PETAFLOPS SPOTLIGHT: ALCF

GE TURBULENCE SIMULATIONS CUT THROUGH THE NOISE

GE Global Research has devised a potential solution to quiet the screech of jet engines and thrum of wind-power turbines: reduce noise at the source. The team tapped the Argonne Leadership Computing Facility (ALCF) to simulate the turbulent flow of air as it passes through jet-engine exhaust nozzles and over wind turbine blades. The project exploited the capabilities of the ALCF's 10-petaflop Mira supercomputer.

The GE team used large-eddy simulations to improve predictions about the noise and aerodynamics of flow around airfoils under various conditions. Besides its supercomputing facility, ALCF contributed its considerable data expertise to the project and developed visualizations.

Together, the GE and ALCF teams have applied these new data toward improving old noise prediction models. The LES approach has already influenced new designs that may help revitalize a wavering alternative-energy market. As turbine size increases and cost-reduction programs encourage lighter, more flexible structures, the ability to make accurate predictions for advanced blade design has become increasingly important. But addressing key issues such as noise source and flow field variations comes at a computationally expensive cost.

The GE Global Research team continues to make major advances, especially in maximizing wind turbine efficiency and energy output, harnessing both its own supercomputing resources and those at the ALCF. The team ran computational fluid dynamic simulations that led to more aerodynamic, productive, and efficient designs. Additional simulations also helped determine where turbines should be sited to capture the most wind and boost productivity.

BIG-DATA NETWORKING

DOE's path-breaking ESnet moves science at the speed of light.

In September 2018, Department of Energy (DOE) scientists set a big-data speed record. Researchers at DOE's SLAC National Accelerator Laboratory took just 29 hours to transfer a massive 1 petabyte of data—the equivalent of the upload and download of about a quarter-million high-definition movies—across a 5,000-mile fiber optic loop. This beat their previous speed record by an impressive five hours.

The test was in preparation for the huge data sets that SLAC's Linac Coherent Light Source II (LCLS II) will generate when it launches, probably in 2020. Although the current LCLS produces up to 120 atomic-level images per second, LCLS II will produce up to 10,000 times that many. It's a data treasure, containing

potential discoveries in biology, chemistry and materials science. But its worth will be realized only if it's effectively transmitted to DOE supercomputing facilities and researchers across the nation for analysis, modeling and simulation.

The SLAC speed record was one more demonstration that after almost 35 years of accelerating discovery, ESnet (for Energy Sciences Network), a DOE Office of Science user facility, is ready to drive the next generation of the department's science mission. ESnet is the nation's and the world's most powerful network for research. It's custom designed, built, optimized and maintained to help investigators meet their goals from experiment to discovery in large-scale, data-intensive science. ESnet is the

equivalent of the DOE's data circulatory system, connecting the department's entire science complex with collaborating academic researchers and hooking the nation into more than 150 research and education (R&E) networks around the world.

Managed by Lawrence Berkeley National Laboratory (LBNL), ESnet supports tens of thousands of scientists in domains from fusion energy to climate and cancer genomics, enabling them to remotely control experiments, transfer data, and collaborate unrestrained by geography. For a researcher anywhere in the nation using ESnet to pursue DOE science mission goals, it appears as if the world's most powerful linac and supercomputer are in the room next door.

DEDICATED TO SCIENCE

In the early 1980s, leaders at DOE's Office of Energy Research (today's Office of Science) recognized that the advent of large-scale research facilities, high-performance computing (HPC) and distributed, collaborative science teams required a new kind of specialized communication: a dedicated scientific network. The internet is optimized for individual and enterprise operations, from email and movie streaming to e-commerce, with profit as the bottom line. As such, it's neither designed for, nor capable of, supporting experiments and HPC simulations that produce petabytes of data that must be transferred, shared and accessed remotely.

Founded in 1986, ESnet grew from the merger of two specialized DOE networks already serving scientists in fusion and high-energy physics research. In 1987, dual satellite links connected contractor General Atomics and Los Alamos, Argonne and Oak Ridge national laboratories at 56K modem speed. ESnet5, the network's fifth and current iteration, is a multi-hundred gigabits

Since 1990, ESnet's traffic has increased by 10 times every four years, a rate roughly double that of the commercial internet. Currently, every month, ESnet moves about 100 petabytes of data—the equivalent of 100 billion books—at the speed of light, with an excellent reliable network uptime.

Through its R&E partnerships, ESnet has enabled the Office of Science's Advanced Scientific Computing Research (ASCR) program supercomputers and hundreds of U.S. researchers to play pivotal roles in major international collaborations. Notably, this includes the ATLAS and CMS experiments at the Large Hadron Collider (LHC) at CERN in Switzerland. As a result, ASCR scientists and HPC resources helped discover the Higgs boson in 2012. Two years later, ESnet deployed four new high-speed transatlantic links with a combined capacity of 340 Gbps (since upgraded to 400 Gbps), giving researchers at national laboratories and partnering universities ultra-fast access to scientific data from the LHC and other European research sites.

PUSHING THE BOUNDARIES OF NETWORKING


In supporting the DOE science mission, ESnet's dozens of scientists, engineers and technicians are in a continual cycle of network research, development and innovation. A critical part of their mission is to determine scientists' networking needs, including their instruments and workflow, to design and deploy infrastructure before demand causes performance issues. As a result, ESnet has emerged as a world leader in redefining smart solutions for science-network optimization.

At a hardware level, ESnet has always maintained a technological edge, not only to facilitate existing science but also to enable researchers to imagine the next frontier. To achieve this, ESnet has worked with dozens of leading U.S. networking, chip, fiber

Since 1990, ESnet's traffic has increased by 10 times every four years, a rate roughly double that of the commercial internet.

per second (Gbps) optical network connecting more than 50 DOE, government and academic facilities and integrated into other R&E networks. ESnet5 moves data about 15,000 times faster than the average home network.

optic and communications companies as an early adopter and tester of fresh-from-the-lab technologies. This includes being one of the first networks to use Asynchronous Transfer Mode (ATM) technology and software-defined networking (SDN). ESnet was



ESnet connects DOE leadership-class computers with academic and national laboratory collaborators. Credit: ESnet.

the nation's first 10-Gbps Ethernet and among the first few to deploy 100 Gbps Ethernet between continents. In 2015, ESnet and the National Energy Research Scientific Computing Center (NERSC) at LBNL built a 400 Gbps super-channel, the first 400 Gbps production link ever deployed by a national R&E network.

DOE's networking expertise has helped make today's internet the robust global system it has come to be. For example, in 1988, after the web's first major traffic jam slowed e-mail message transmission from seconds to hours, LBNL network scientist Van Jacobson developed the TCP (transmission control protocol) algorithms that remain the foundation for avoiding internet congestion. ESnet continues to embody this spirit of engaging network researchers for unique solutions and incorporating their innovative ideas.

Just as important as adopting hardware and protocol advances, ESnet scientists and engineers have pioneered R&E network

architectural and software innovations, many of which have been adopted by other R&E networks. For example, when SLAC scientists achieved their big-data speed record they did so thanks to the traffic-directing and control provided by ESnet's award-winning On-Demand Secure Circuits and Reservation System (OSCARS) software. Developed by ESnet engineers, OSCARS is a pioneering software service that lets users make a big-data network reservation. OSCARS creates dedicated, pre-arranged bandwidth channels for scientists who need to move massive, time-critical data sets around the world.

What makes OSCARS so useful is that it can automatically create end-to-end circuits, crossing multiple network domains. As such, it avoids data traffic jams while optimizing transfer speeds. Before OSCARS, this was a time-consuming process; in 2010, ESnet engineers needed 10 hours of phone calls and about 100 emails over three months to do what one person can now do in five minutes using OSCARS.

OSCARS set the stage for ESnet engineers' next-generation versions—the automation of this bandwidth control through a technique known as Software Defined Networking (SDN). In its latest version, ESnet scientists are extending SDN to include optimization of both the network backbone and systems at DOE exascale computing facilities via the End-to-end Networked Science at the Exascale (SENSE) ASCR-supported research project.

Another major networking innovation was ESnet's Science DMZ, now a best practice at major research institutions worldwide. A key challenge in big data transfer at some universities and laboratories is that passing through the institution's general enterprise network slows, disrupts or even blocks scientific applications handling that information. First proposed by an ESnet network engineer in 2010, the Science DMZ physically separates big-data science applications, such as data for simulations and visualizations and remote experiment control, from an institution's general enterprise network.

In a Science DMZ system, a science-specific bypass network is built at or near a campus or laboratory's local network perimeter. The Science DMZ optimizes big-data science applications, including security features, while allowing optimization of the institution's enterprise system for its own purposes.

Since 2011, the National Science Foundation has provided more than \$120 million to support Science DMZ implementations at more than 100 U.S. universities, and other federal agencies, including the National Institutes of Health, have adopted the approach. It's also used in countries from Australia to Brazil.

NETWORKING SPOTLIGHT: PARTNERSHIPS

SCIENCE ENGAGEMENT TEAMS

ESnet's success relies on its highly effective science engagement teams, which teach and support myriad users to employ the system efficiently. This is critical, since 80 percent of ESnet traffic either originates or ends outside of a DOE facility, where there are often first-time users, such as the thousands of DOE-supported graduate and post-doctoral students at U.S. universities.

To support these users, in 2007 ESnet launched fasterdata.es.net, a website providing users tips and tools to ensure the trouble-free transfer of large amounts of data over scientific networks. The site is a step-by-step guide on how to configure computers to optimize file transfer speed, select the proper software and monitor the data movement. In addition, the Faster Data Knowledge Base provides proven methods for troubleshooting and solving performance issues.

In 2011, ESnet launched My ESnet, a portal for real-time, customizable information on network utilization and performance. The interactive site lets users have a real-time, detailed view of traffic traversing the various sections of the network as it connects 40 research sites around the United States.

ESnet's dozens of scientists, engineers and technicians are in a continual cycle of network research, development and innovation.

And to assist scientific collaborations, particularly international ones, ESnet collaboratively developed the open-source perfSONAR, the first software tool that lets engineers test and measure network performance, as well as archive data to pinpoint and solve service problems that may span multiple networks and international boundaries. Users can diagnose where packets are dropping or identify data logjams in real time and immediately solve problems.

Almost whenever users have deployed perfSONAR, it has revealed previously undetected, significant bandwidth-limiting problems, many of which are then relatively easily resolved. Developed through an international collaboration led by ESnet, Internet2, Indiana University, and the European R&E network GÉANT, perfSONAR is deployed in more than 2,200 locations worldwide.

THE FUTURE OF NETWORKING

In May 2019, an independent panel of networking and science experts met at LBNL for the final design review of ESnet6. This next-generation version is a major upgrade—designed to support exabyte-scale science. It anticipates a snowball effect of exponential data growth in the next several years, driven by new technologies centered on machine learning and artificial intelligence (AI), as well as sensors and upgrades to facilities such as SLAC's LCLS II, the High-Luminosity Large Hadron Collider and the Vera C. Rubin Observatory, which will come on line in Chile in 2022. Not only will ESnet6 have significantly more bandwidth than ESnet5, but this capacity also will be more efficiently allocated, thus providing greater consistency and resiliency.

Even as ESnet completes the roll-out of ESnet6, it's already looking ahead to its successor. Leaders are working with scientists to support research on emerging quantum networks and helped organize a recent workshop to discuss developing such systems, including quantum interconnects, quantum local-area networks, quantum metropolitan networks, and quantum

wide-area networks. Similarly, ESnet engineers are exploring cognitive, or smart, network applications that use machine learning and artificial intelligence (AI) to automatically grasp how users interact with one another, then employ this information to improve the network.

ESnet6 will be a remarkable next chapter for a network whose history began in 1986 with 56 Kbps modem links and today is imagining a future in which DOE researchers tackle the nation's biggest science challenges while connected at light speed via a quantum AI network.

NETWORKING SPOTLIGHT: 100 GBPS PROTOTYPE

PUTTING THE FUTURE TO THE TEST

In 2012, with \$62 million from the American Recovery and Reinvestment Act, ESnet built a 100-Gbps, long-haul prototype network and a wide-area network testbed to support development of next-generation, disruptive technologies. The testbed is an advanced networking laboratory, a place where researchers from DOE, academia and private industry can assess new hardware, software and network control tools—in areas from cybersecurity to data transfer—at scale without worrying about breaking the system. The testbed extends from Berkeley to Chicago and New York. To date, more than 70 groups have used it.



DISCOVERY FROM DATA

ASCR research produces detailed science simulations and the tools to sift through the massive data files they generate.

Building a Google for cancer treatment. Tackling trillion-particle computer simulations. Developing the first exascale supercomputers capable of performing a quintillion calculations per second. Combining the power of computing with the power of human perception. Moving data at ultrahigh speeds over long distances.

These and many other developments have been made possible by more than 40 years of sustained investment from the Office of Advanced Scientific Computing Research (ASCR), part of the Department of Energy (DOE) Office of Science. These wide-ranging efforts have advanced both separately and together for a common purpose: to discover, develop and deploy integrated computing, networking and analysis capabilities that enable knowledge-creation from data.

In one especially ambitious effort, ASCR researchers are collaborating with the National Cancer Institute (NCI) to apply supercomputers' number-crunching power to cancer research. The project combines high-performance computing (HPC) expertise to enable smarter choices about cancer prevention, discovery and treatment.

This unprecedented data-driven approach is pushing toward the exascale—nearly 10 times faster than today's most powerful supercomputers—and the promise of predictive oncology: custom therapies to treat an individual patient's cancer.

Translating information gathered from large-scale experiments to new cancer treatments will require pushing data-crunching to levels unavailable today. When the National Institutes of Health

and DOE began their collaboration, ASCR-supported researchers already were building computing infrastructure to handle big data on the scale of what the NCI generates. After months of talks, the agencies agreed that making progress means developing advanced algorithms that incorporate machine learning to automate functions that now require human intervention.

Deep learning is at the center of a data-driven cancer pilot project at Argonne National Laboratory (ANL). It builds predictive models for drug treatment based on biological-experiment and drug-response data.

Oak Ridge National Laboratory (ORNL) leads another pilot that assembles tools to better mine cancer surveillance data from millions of patient records. A third pilot centered at Lawrence Livermore National Laboratory investigates complex gene and protein interaction patterns in a molecular pathway implicated in many types of cancer but difficult to target with standard techniques. And Los Alamos National Laboratory leads a crosscutting effort to provide uncertainty quantification and optimal experimental-design support across all projects.

TRILLION-PARTICLE SIMULATIONS

Simulations based on trillions of particles generate highly challenging big-data problems. Yet such computations are exactly what researchers need to study fundamental science problems such as plasma physics or the history of the universe.

Science is rife with simulations that generate large data sets, including climate, groundwater movement and molecular dynamics.

They're unwieldy because with every step forward in time, the computer pumps out tens of terabytes (a trillion bytes) of data.

To manage and analyze these huge datasets, DOE scientists developed big-data clustering at trillion-particle scale (BD-CATS), a program to analyze the results of trillion-particle problems from start to finish on the nation's most powerful supercomputers.

Particles in a cosmology simulation, for example, can represent quantities of dark matter, a phenomenon that looms as one

of the biggest mysteries of modern science. As the model progresses through time, dense dark matter regions exert a gravitational pull on surrounding visible matter, forming it into clumps. The model reveals structures in the universe, including galaxy clusters that develop in the densest dark-matter regions. The more particles the model contains, the more accurate the results. To test the precision of the theory behind such simulations, scientists compare the number and size of clusters they show with real astronomical data.

Plasma physicists, meanwhile, simulate natural phenomena such as the interaction of charged particles streaming from the sun with the Earth's magnetic field. This leads to the visual delight of northern lights, but also can interfere with power systems. As in cosmology simulations, the more particles modeled, the better.

Running BD-CATS, the Edison supercomputer at the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory, for example, analyzed the results from a 1.4 trillion-particle plasma-physics simulation. The entire process took approximately 30 minutes on roughly 100,000 Edison processor cores.

Science is rife with simulations that generate large data sets, including climate, groundwater movement and molecular dynamics. These problems might require tracking only billions of particles today, but data analytics that work at the leading

edge—as in cosmology and plasma physics—will improve future simulations across many domains.

VISIONS OF EXASCALE

The Aurora supercomputer—scheduled to begin operations at ANL in 2021—will build on work done on previous-generation supercomputers, ranging from materials science simulations of molecular defect formation to fluid dynamics calculations for experimental engines.

Aurora will be based at the Argonne Leadership Computing Facility (ALCF), an Office of Science user facility. It will be one of the nation's first exascale computers, capable of performing a quintillion (a billion billion, or 10^{18}) calculations per second. That's about seven times faster than Summit, currently rated as the nation's most powerful supercomputer, at ORNL's Oak Ridge Leadership Computing Facility (OLCF), also an Office of Science user facility.


Supercomputing resources usually emphasize simulation capabilities. More recently, however, large-scale computing's focus has expanded to comprise three main pillars: simulation (simulation-based computational science), data (large-scale data-intensive computing), and learning (machine learning, deep learning, and other artificial intelligence capabilities). This trend will only accelerate in the exascale era.

Numerous efforts across the HPC community emphasize developing tools to help users of leadership-class resources cope with the information flood the machines generate. In one such collaboration, Argonne computer scientists work to incorporate into computations data analysis methods known as *in situ* (as the calculations happen) and *in transit* (as data move).

Researchers traditionally have visualized and analyzed data only after simulations end. They write out the information as a program runs, and then review it when they're done. Unfortunately, researchers have already reached the point where they can't save all their material.

Scientists now must examine their data while the information is still in memory to ensure that no insights are lost. The high-energy physics community has addressed this challenge for decades. The major research collaborations conducting experiments at the Large Hadron Collider at CERN, the European particle physics laboratory, generate such vast quantities of data that only a tiny fraction can be archived. But physicists have developed theories that help them determine which phenomena to seek in the LHC's subatomic particle collisions. This lets them write algorithms to instantly identify which data must be saved to long-term storage for further analysis.

Myriad scientific disciplines will tap Aurora for research using a variety of computational tools. They will need a generic interface—a scalable *in situ* analysis and data visualization tool called SENSEI—to build upon. SENSEI lets researchers prepare their simulation code just once—and in a fairly easy way—enabling them to analyze a given infrastructure on the computer system's back end. The SENSEI bridge will let researchers choose which data generator—what simulation code, for example—and which back-end visualization and analysis tool they want to run.



The universe at 7.68 billion years old, as depicted from a model run on the Mira supercomputer at the Argonne Leadership Computing Facility (ALCF). Credit: Argonne/Los Alamos/Lawrence Berkeley national laboratories.

SCIENCE IN MOTION

From the scout of a prehistoric hunting party scratching a herd's location in the dirt to a team of astrophysicists animating an exploding star in three dimensions, expressing abstract ideas through visualization seems innately human.

Visualization uses images to promote understanding and communication. Scientific visualization harnesses the power of computing and human vision as full partners in scientific reasoning. Simulations can result in highly detailed and explorable visualizations that provide scientists with unprecedented insight.

Computers can construct increasingly detailed duplicates of nature's processes in action, generating visual understanding that sometimes exceeds experimental observation, all in a virtual world where the discoveries are very real. But there is a caveat. Scientists must take extra care to remember that all simulations include some degree of error and uncertainty. (See "Embracing Uncertainty," page 33.)

Like any motion picture, a scientific visualization presents a final cut: selected and processed data from much larger troves. Visualization can reveal underlying meaning and structure in science and engineering data, typically generated in simulations that incorporate mathematical computation and geometric models.

Argonne developed one of the first approaches to advanced parallel I/O, called the Parallel Virtual File System (PVFS). PVFS could pull several DVDs worth of data into a computer in a second.

A single visualization now can display data in the terabytes and petabytes (a quadrillion bytes). Compare that to Google, which processes about 20 petabytes a day. A few seconds of animation may consume a million processor hours, but the payoff can be an exquisitely detailed vision of realms no experiment can probe. And increasingly powerful computers are leading to higher-resolution views of simulated reality at extremely fine grain—even at the atomic level.

A simulation showing an enzyme found in the white rot fungus digesting a million atoms of plant cellulose, for example, models a process that humans would like to recreate to make fuel from wood waste.

Scientists collaborating on the simulation are at the National Renewable Energy Laboratory in Golden, Colorado, and at Cornell and Pennsylvania State universities. NREL is a partner in the BioEnergy Science Center under DOE's Office of Biological and Environmental Research (BER). The simulation considers the angles and distances of all the chemical bonds among atoms, along with physical forces that govern reactions—a visualization *tour de force*.

HIGH-SPEED DATA MOVEMENT

Try to visualize bytes as pennies. A petabyte-sized stack, with one penny corresponding to one byte, would stretch from Jupiter to the sun and back.

Given such a tower of data—an amount of information appearing in modern scientific computing with increasing frequency—the trick becomes moving it and moving it quickly, whether within a single research facility, across the United States or around the world. This can be done, but only with the proper computer systems for sending data over high-speed networks and specialized input-output (I/O) and file systems for reading and writing data on storage devices.

To fill that need, DOE's ASCR-funded supercomputer centers have established Globus, developed at ANL, as their primary option for transferring data between far-flung locations. These centers include ALCF, NERSC and OLCF, and users stretch from coast to coast—from Brookhaven National Laboratory in New York to Fermilab in Illinois and the SLAC National Acceleratory Laboratory in California, with its Linac Coherent Light Source, a free-electron laser that takes snapshots of working atoms and molecules.

Researchers at thousands of other organizations around the world, including universities, government facilities and a growing number of commercial operations, use Globus as they work with data. The non-profit research data platform, operated by the University of Chicago, addresses many of the common challenges researchers face in moving, sharing and archiving large volumes of data. Among its many endeavors, Globus is part of the National Institutes of Health Data Commons, a multi-institutional consortium that makes biomedical research data and tools accessible to more scientists.

I/O, however, often creates time-consuming bottlenecks in large-scale, parallel-processing computers, every aspect of which must be balanced to attain top speed. To overcome these pinch points, I/O libraries and computer file systems must be able to read and write many data simultaneously. That is, they must exploit parallelism.

Argonne developed one of the first approaches to advanced parallel I/O, called the Parallel Virtual File System (PVFS). PVFS could pull several DVDs worth of data into a computer in a second.

Much like supercomputer systems in general, I/O includes many components, from bandwidth and programming efficiency to interfaces and real-world assessments. Consequently, computer scientists also have developed other approaches to improve the simplicity of controlling I/O and the speed of data transmission—even when the pile of data could go into outer space and back.

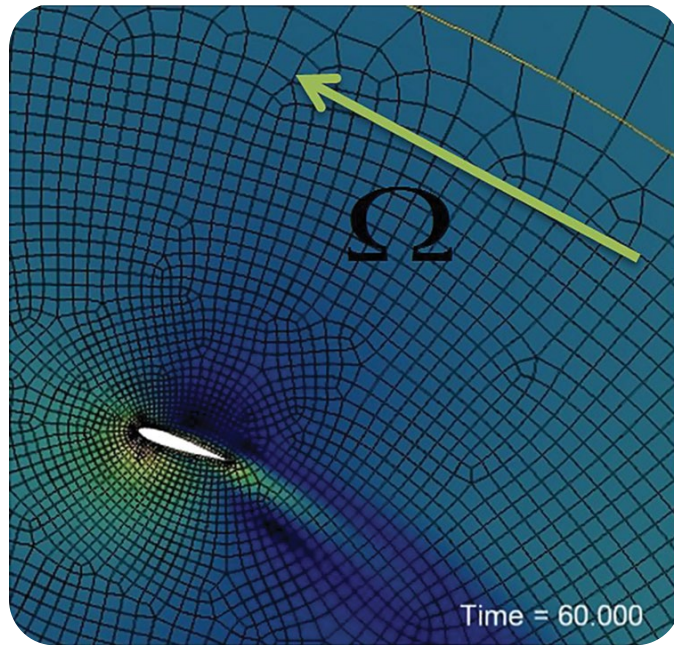


DOE has pioneered the fields of uncertainty quantification, improving simulations' reliability on problems with large economic, environmental and security implications.

The gold standard of scientific truth has always been experimentation and observation. Unfortunately, many systems and questions evade these tools, either because the experiments are unethical (genetic engineering of human embryos, for instance), potentially hazardous (climate engineering), difficult to instrument (materials inside nuclear reactors) or costly (laser fusion), or the observations are simply impossible (radiation transport in the moments following the Big Bang). Computer simulation of these and other phenomena with major scientific, economic and security implications has long been a goal of the Department of Energy's (DOE's) Office of Advanced Scientific Computing Research (ASCR).

For 40 years, DOE has invested substantially in the mathematics and computer science that support realistic simulations. Slowly but surely, ASCR-funded research advances have improved models' scale and detail and researchers' ability to tackle ever more complex systems; simulations barely possible in one decade became routine in the next.

But as simulations have become more and more powerful, their limitations have become harder to ignore. As statistician George Box once said, all models are wrong, but some are useful. The same is true about computer models. Each of their elements carries a degree of uncertainty—from imprecise, indefinite input parameters



Uncertainty analysis can be applied to a wide range of simulations, including this ASCR project that modeled wind turbines. Credit: Michael Eldred/Sandia National Laboratories.

but also from poorly understood physical properties and questionable verisimilitude of their recreation of reality. As models have matured over the past 20 years, the big question has been how scientists can calculate the degree of possible error in them.

The need for confidence in models addressing important national questions such as new energy sources, climate change and nuclear stockpile stewardship has inspired a generation of DOE computer scientists, physicists and statisticians to innovate methods for assessing the reliability of predictive simulations. ASCR has supported fundamental research in these various methods, which have coalesced into the new field of uncertainty quantification, or UQ.

UQ now plays a role in every important economic, environmental, and security issue, from designing microbes for biofuels and predicting climate change to developing fusion-power technology and buttressing the nation's power grid. Simple models and approximations aren't good enough to provide the

science-based information necessary to predict how complex systems will behave.

UQ has grown into a key capability in the modeler's toolkit, and now is routinely built into simulations. It's also become de rigueur in industry and other government agencies, and it has inspired numerous new journals, conferences, professional groups and changes in computer architecture and, potentially, processor design.

UNCERTAIN TIMES

The need to quantify uncertainty in predictions regarding the nuclear stockpile and other critical national security questions drove UQ's rise in the 1990s. Advances in three other areas—algorithms, data collection, and computer hardware—have converged to make UQ ubiquitous.

In the 1960s, algorithms underlying computer models of complex physical systems, including fluid mechanics, electromagnetics and acoustics, proliferated. But decades passed before models for these physical phenomena matured enough to become predictive—enabled by an ability to collect huge amounts of high-quality data and exponential advances in computing power.

Uncertainty can be described as a mathematical distribution of values within which a particular observed value could fall. To quantify uncertainty in a simulation's predictions, scientists first identify variability sources in parameters that affect the model's output. If it has just one uncertain parameter, it can be run as many times as that parameter varies, yielding a distribution of potential outputs. That's often hard enough, but now suppose there are two uncertain parameters. Doubling the number of parameters does not just double the number of runs needed; it squares that number. In other words, if assessing a single parameter's variability requires running a model 100 times, two uncertain parameters requires 10,000 runs—that's 100 squared.

Now, imagine a problem with a million uncertain parameters, like the Antarctic Ice Sheet. A model of that size requires sophisticated new mathematics.

In addition, a model's uncertainty can increase according to its degree of nonlinearity—meaning that small changes in an input can spawn large, unexpected changes in the output. Other factors that contribute to uncertainty include a model's

some slamming into the Gulf states, some marching up the coast before making landfall, and some curling harmlessly back out to sea.

Inverse UQ is another level up in complexity. Here, researchers go backward from a known output to calculate uncertainty in the parameters or improvements in the model form. This uncertainty can then be fed back into a forward model to

The question then becomes how to optimize a system given such uncertainty.

approximation of the real-world system's geometry, unknown detailed properties of a modeled object (grain structure, for example) and each variable's initial conditions. Models often rely on parameters that are estimated from sources ranging from empirical data to other calculations.

UQ comes in two broad forms: forward and inverse.

In forward UQ, uncertain inputs come into a model, which then spits out a distribution of uncertain outputs. Uncertainties can arise from a variety of sources, including the model form (independent of any parameters), external parameters (such as boundary or initial conditions), or model parameters (material properties, for example). Imagine a hurricane path model with a grid of uncertain parameters such as wind speed, temperature and humidity feeding into it. The output is not a single path, perhaps one that hugs the U.S. eastern seaboard. Rather, the forward model produces a whole family of possible paths, with

produce a distribution of predictions. Using inverse UQ, scientists can measure the exact path a hurricane took, then go back to estimate the distribution of possible local wind speeds that led up to it. These speeds can then be used as an input for future path predictions.

Perhaps UQ's ultimate use is in decision-making under uncertainty—a model with uncertain inputs inferred from measurements, with those uncertain inputs propagated through the model to produce uncertain outputs. The question then becomes how to optimize a system given such uncertainty. Optimization incorporates inverse UQ, which incorporates forward UQ—presenting a huge challenge in a stochastic system, with its random variables.

Take, for instance, groundwater contaminant remediation. If a subsurface flow and transport model accounts for uncertainty in soil properties and contaminant plume, it would be possible

to pose a question like “Where do I place remediation wells and how do I control them to minimize contamination risk to a city’s drinking water supply, within a budget?” An optimization procedure would automatically account for all those outcomes so long as the uncertainties were plugged in.

IN WALKED ASCR

In the late 1990s, ASCR managed several single-investigator UQ efforts, with joint support from the National Science Foundation (NSF), to get a pulse on what researchers were doing in the field. Those efforts turned out to be quite diverse. Some had to do with examining parametric uncertainties, while others explored Bayesian statistics (in which probabilities are based on prior beliefs and updated as new evidence comes in) or interval analysis (in which computation rounding errors or uncertainty in input parameters are dealt with by expressing outputs as a range of possible values). The early years were similar to the current effort in machine learning and artificial intelligence, with ASCR leading much of the push, particularly to develop the underlying mathematics.

ASCR started investing heavily in UQ in the early-to-mid-2000s, supporting several multi-investigator projects and launching a few notable center-level programs, including SciDAC (Scientific Discovery Through Advanced Computing) and MMICCs (Multifaceted Mathematical Integrated Capability Centers). SciDAC was launched in 2001 as a partnership

UQ’S FUTURE

Just as machine learning can be applied to almost any field, so can UQ. However, two fields stand out as early adopters: energy and transportation.

The oil industry increasingly uses UQ to seek resources and to manage reservoirs. Exploration presents a challenging inverse problem because decisions on where to drill rely on analyzing reflections of seismic waves off putative oil-bearing rocks. Companies want to know how much they can trust their models before they spend millions on drilling.

UQ also is deeply embedded in the aerospace and automotive industries. Every aircraft Boeing builds begins as a blueprint on a computer, and UQ is critical to the simulation-based design process. When Ford plans a car today, it crash-tests it on a computer using simulations employing UQ.

The biomedical industry has trailed other sectors in adopting UQ, but recently companies have explored the concept in earnest for, among other things, supporting decisions about surgery. Oncologists preparing to perform laser ablation therapy, for example, want to concentrate the beam on a tumor while minimizing damage to surrounding tissue. Their decisions about where to focus the device benefit from models explaining tumor growth and laser heat conduction through tissue as well as understanding the uncertainty inherent in the model’s predictions.

Just as machine learning can be applied to almost any field, so can UQ.

among DOE laboratories to facilitate the adoption of advanced computational tools and their deployment on scaled systems. MMICCs came on line in 2012, as the ASCR Applied Math Program created three large centers to cover all aspects of UQ, with a focus on basic mathematical research targeting applications in energy, security, and the environment. Work in the physics community converged with the optimization and statistics communities to develop the language of UQ and usher it into the computational science mainstream.

UQ is also inspiring changes in how computation is done. Researchers are already envisioning computer architectures for UQ that calculate with probabilities instead of precise numbers. This is prompting questions like whether the cost of computing accuracy to 16 digits is worth it if the output is a distribution.

Perhaps the greatest impact of DOE’s UQ investment is in how it’s changed the way computational scientists look at predictions, which many agencies, including DARPA, DOD and NSF, now view probabilistically. All now place uncertainty at the center of computational science.

RULES OF THE ROAD

Building a landscape for the high-performance computing community.

High-performance computing (HPC) presents unique computer science challenges. The field is constantly in flux with new computer architectures, applications and algorithms. How do you measure progress in the field? How do you ensure that today’s software will run on tomorrow’s machines?

To manage a constantly evolving field over the past four decades, researchers at the Department of Energy (DOE)

national laboratories have played key roles in developing rules of the road—the principles that help manage computation and communication and provide performance metrics for the world’s fastest supercomputers.

That challenge continues. Heterogeneous hardware systems and other emerging technology could require completely new algorithmic approaches. These research efforts will shape HPC in the coming decades.

BUILDING ROUTINES

Matching hardware and software has always been a key HPC challenge. For the earliest computers, vendors provided programming packages that were compatible with their machines. But when researchers and labs wanted to use different equipment or upgrade their systems, they often had to completely recode the software.

As early as the 1970s, computer scientists began developing strategies to streamline programming and make codes more portable. For example, Basic Linear Algebra Subprograms (BLAS)—supported by both DOE and the National Science Foundation—defined a standard for a core set of software that performs common mathematical calculations. Vendors implemented fast versions of these routines on their hardware, allowing programmers to achieve good performance across multiple platforms.

More than half of the TOP500 supercomputers have hybrid architectures.

Today BLAS remain part of vendor software and are incorporated within all HPC platforms, and the underlying idea has been broadened considerably. BLAS are a foundation for LINPACK, a software package for solving systems of linear equations on HPC systems that the DOE developed under Advanced Scientific Computing Research (ASCR) program funding. Over the years, LINPACK also has provided a clear, quantifiable measure of a supercomputer's speed in solving mathematical operations. As a result, it became the basis of the High-Performance LINPACK (HPL) benchmark, the primary metric for the TOP500, the industry's biannual rankings of the fastest supercomputers worldwide since 1993.

HPL solves linear equations and provides an assessment of a system's floating-point speed—the number of mathematical calculations it can perform per second. Today's fastest supercomputers work at speeds of hundreds of quadrillions of flops, and the LINPACK benchmark provides a comprehensive history of computer performance over the past 25 years.

FOSTERING COMMUNICATION

For decades Moore's Law—which says microprocessor transistor density should double every 18 to 24 months—has supported proportional boosts in computational capability. In addition, until the mid-2000s computer designers steadily increased the clock speeds which further enhanced supercomputer performance.

In the 1970s, an innovation called vector processing let computers crunch through streams of data rather than performing single calculations. By the 1990s, the ability to link many microprocessors to work simultaneously on large jobs led to massively parallel processing, the basis of today's HPC. But this dramatic change forced computer scientists to adapt software in a dramatic fashion. ASCR funded some of the early proof-of-concept work on parallel processing and played a pivotal role in navigating the transition from vector to parallel computing.

With the computational workload distributed over many processors, communication between processors became a key software portability challenge. Vendors provided their own message-passing software to facilitate communication, but as with earlier programming platforms these tools were specific to the manufacturers' hardware. To make applications more portable, ASCR supported key work to develop the message-passing interface (MPI) standard, a portable standard that allows millions of processors to exchange data and computation. Researchers have adapted MPI into many implementations, including MPI Chameleon (MPICH) from Argonne National Laboratory, which translates code to match a system's hardware needs.

EXTREME HETEROGENEITY

Today computer scientists wrestle with another wave of software portability challenges as hardware becomes extremely heterogeneous. With HPC approaching the exascale, researchers have examined a range of computer architectures to support growing simulations and data-driven applications.

Graphics processing units (GPUs), the chips first used to boost speed and 3-D graphics in gaming consoles, are a key part of this shift. Today, more than half of the TOP500 supercomputers have hybrid architectures that combine GPUs and CPUs, including the DOE supercomputers, Summit and Sierra, that sat atop the list in 2019. GPUs have hundreds of cores and can support increased simultaneous computation while consuming less power. Those advantages are useful for many traditional applications and for emerging applications such as artificial intelligence.

Until recently, NVIDIA GPUs dominated these architectures, but other vendors are developing competing platforms. For example, one of the first DOE exascale systems, Aurora at Argonne National Laboratory, will use Intel's graphics architecture. DOE labs also have pursued other hardware strategies. For example, Astra, a Sandia National Laboratories system entirely built with Cavium ThunderX2 ARM processors, debuted on the TOP500 in November 2018. But with a growing diversity of platforms, many

of them sporting heterogeneous computing elements, portable software is once again a huge challenge. The libraries and tools of the recent past are insufficient to let researchers develop software that is portable to all these diverse machines. GPUs can boost computation speed by a factor of 100 in some cases, but programmers often must rewrite codes to take advantage of that acceleration. Each architecture can require a new version. Reworking millions of lines of a climate code for one of these new systems, for example, is an arduous and time-consuming undertaking. Because the DOE supercomputers often have five-year commissions, this obstacle threatens to limit the sustainability of computational research.

ASCR is supporting work on several approaches to make programming more adaptable across heterogeneous hardware. In January 2018, ASCR sponsored a workshop on extreme heterogeneity in HPC to explore the key challenges and outline research strategies to address them.

RULES SPOTLIGHT: MEASURING PERFORMANCE

IN SEARCH OF BROADER BENCHMARKS

Since the early 1990s, the supercomputing industry has measured performance using High Performance LINPACK (HPL)—a program that tallies computer performance in floating-point operations per second (flops) while using a particular method to solve millions of equations.

HPL's design overwhelmingly favors computers that are proficient at numerical operations, but it is a poor tool for measuring other elements of computer performance that are critical to many real-world applications. The HPL metric doesn't heavily stress data movement and memory system performance, says Michael Heroux of Sandia National Laboratories' Center for Computing Research. But today's supercomputers must focus on those factors, moving information from the system's memory to the processor and back, if they're to address a different, broad set of science and engineering applications, including modeling and simulating such things as automobile crashes, aerodynamics and oil recovery. These applications all are based on a different method for solving systems of equations.



Credit: Sandia National Laboratories.

Heroux, working with Jack Dongarra (an original LINPACK author) and Piotr Luszczek from the University of Tennessee, developed a supercomputing benchmark called the High Performance Conjugate Gradients (HPCG). It executes an algorithm distinct from the one HPL uses but one more representative of a wide range of applications. It has become a proxy test code for much larger programs and a novel benchmark for HPC systems worldwide.

If computer scientists are comparing two systems that both can compute at 10 petaflops per second, for example, but one handled memory twice as fast as the other, an HPL test wouldn't necessarily distinguish between the two. HPCG would.

Previous page: A model of mechanical strain in nickel oxide made with QMCPACK, an open-source software package for electronic structure calculations. The code is optimized to run on the most advanced supercomputing architectures. QMCPACK's development is part of the Exascale Computing Project, a collaboration between the DOE Office of Science and the National Nuclear Security Administration. Credit: Argonne National Laboratory.

To tackle some of these HPC software problems, researchers make the code more abstract but include directives—high-level instructions—that tell the program how to compile and execute code for different architectures. Such solutions annotate the various loops and functions within the program, providing instructions to parallelize the code for that architecture, allocate data and manage execution. The two primary directive strategies are OpenACC and OpenMP. OpenMP is a broad tool for parallel programming while OpenACC helps to implement codes on accelerating hardware, such as GPUs. Partners from the DOE national laboratories, academia and industry are developing these tools collaboratively.



The MPI development team celebrated its 25th anniversary in 2017. Credit: NERSC/Lawrence Berkeley National Laboratory.

With funding from ASCR and from the National Nuclear Security Administration's Advanced Simulation and Computing program, computer scientists at the DOE national labs also are developing even higher-level approaches for adaptable programming. Taking advantage of C++ language standards, they can write code that is even more isolated from changes in supercomputer hardware. Kokkos, developed by Sandia researchers, includes template code to handle the memory space, where data resides on various processors, and execution space, which describes how computation is parallelized across processors. Similarly, RAJA, developed at Lawrence Livermore National Laboratory, uses C++ programming abstractions that create connections to hardware- and software-specific compilers and data-management routines.

EMERGING HARDWARE, LOOMING CHALLENGES

Today's heterogeneity challenges, though significant, are primarily linked to making software match hardware. As ASCR contemplates a future that includes quantum and neuromorphic computing, researchers will face new opportunities and programming challenges. Current coding structures could be difficult to map onto these emerging platforms and will force computer scientists to rethink how they write algorithms.

Quantum computing harnesses the unique properties of subatomic particles. In quantum computers, qubits—electrons or photons—aren't restricted to the 1s and 0s of binary bits that form the foundation of conventional computing. They can exist in both states simultaneously, a property known as superposition. Superimposed qubits can compute multiple scenarios simultaneously. Quantum particles also can be entangled with others, even those that

aren't nearby, which means that altering one simultaneously affects many others. In traditional computing, adding processors increases computational speed linearly—a computer with twice as many processors works twice as fast. In quantum computers, entanglement means that adding more processors can boost the computational speed exponentially for some applications.

Neuromorphic processors are modeled after the ways that brain cells process information. Instead of simply being on or off, such devices weigh incoming information and propagate it once it reaches a critical threshold. Such processors would consume far less power than today's CPUs and GPUs.

These new approaches could be incredibly useful for problems that today's machines can't tackle. Quantum processors could speed up the simulation of quantum systems, and neuromorphic processors seem ideal for artificial intelligence applications. But some problems such as simulating neutron transport or fluid dynamics might be better suited to traditional CPU or CPU/GPU architectures. Computational scientists might need to tailor the problem and, therefore, the algorithms and code, to match the best architecture to address it. These complex, highly-heterogeneous, future architectures will require new software insights. Although the correct approaches aren't yet known, history suggests that ASCR will play a central role in inventing them.

Over the past 40 years, DOE computer scientists have built performance benchmarks and tackled many software and hardware challenges. New computer architectures and processors offer new obstacles and opportunities for achieving faster, more realistic scientific results than ever, with ASCR guiding the path.

THE HUMAN ELEMENT

The rapid growth of scientific computing and workforce training to fuel it go hand in hand.

Computational science has changed dramatically over the past four decades. High-performance computing (HPC) uses parallel rather than serial processing. Faster chips and new architectures give computers power to work on ever more complex and data-intensive problems. But capitalizing on these developments requires an accomplished and nimble workforce that can devise new algorithms, codes and software for more realistic simulations and harness emerging strategies to analyze data through artificial intelligence.

However, demand for computational scientists greatly exceeds the supply. To address this shortage, the Department of Energy (DOE) Office of Science, the Office of Advanced Scientific Computing Research (ASCR) and the DOE national laboratories have devised several ways to boost interest and training in computational science at all levels.

CSGF SPOTLIGHT: JUDITH HILL

FROM ALGORITHMS TO LEADERSHIP



Judith Hill leads the Scientific Computing Group at the National Center for Computational Sciences at Oak Ridge National Laboratory (ORNL) and manages the Innovative and Novel Computational Impact on Theory

and Experiment (INCITE) program at the ORNL and Argonne National Laboratory leadership computing facilities. But her work on these programs is rooted in her DOE CSGF experience from 1999 to 2003, when she was a Ph.D. student at Carnegie Mellon University.

A frame from an animation of ocean currents carrying dissolved organic carbon (yellows are lower concentrations, purples higher). The work stemmed from Computational Science Graduate Fellowship recipient Riley X. Brady's 2018 practicum at Los Alamos National Laboratory. Credit: Riley X. Brady/University of Colorado Boulder; and Stephanie Zeller, Annie Bares/University of Texas at Austin.

(Continued on page 42)

(Continued from page 41)

Her practicum at Sandia National Laboratories led to Hill's first job at its Computation, Computers, Information and Mathematics Center from 2005 to 2008. The breadth of experience the fellowship provided has proved valuable in her ORNL roles as a computational scientist and coordinator for INCITE projects allocated on Summit, the lab's IBM AC922 system. "As a computational scientist in the facility," she says, "it's useful for me to have not only the domain expertise, applied math in my case, but also the algorithmic computer science expertise: all the things that the DOE CSGF program emphasizes in the program of study, in the practicum experience and in the experiences of the fellows."

OUTREACH, INTERNSHIPS AND RESEARCH

DOE supports the Oak Ridge Institute for Science and Education, a program that rose from peacetime efforts just after the Manhattan Project ended. It supports a range of science, technology, engineering and math (STEM) research opportunities at the DOE laboratories and other federal facilities for undergraduate, graduate, postdoctoral and faculty scientists, including special outreach to historically black colleges and universities and minority-serving education institutions.

In the 1990s ASCR supported the Computational Science Education Project, which provided outreach and online educational materials in computational science. At the same time, the office launched Adventures in Supercomputing, which for six years gave high school students in five states opportunities to work on DOE HPC systems. One of these programs, the New Mexico Supercomputing Challenge, continues with support from Los Alamos National Laboratory.

The DOE national laboratories back multiple educational and research opportunities for K-12 students, college faculty and everything in between. For example, Argonne National Laboratory (ANL) sends laboratory scientists into area schools. Its on-site initiatives include learning labs, camps and high school research programs to support young scientists. At Lawrence Berkeley National Laboratory (LBNL), community college students, undergraduates, graduate students and college faculty can participate in internships and research programs.

GRADUATE TRAINING IN COMPUTATIONAL SCIENCE

Graduate fellowships that provide Ph.D. training are a centerpiece of the broader national effort to expand the pool of talented scientists and engineers in many disciplines. Outside DOE, for instance, the National Science Foundation Graduate Research Fellowship and National Defense Science and Engineering Graduate Fellowship provide tuition and stipends for up to three years of study in many fields, including computational science.

In 1991, ASCR created the unique and prestigious Computational Science Graduate Fellowship (DOE CSGF). The program distinguishes itself from other graduate fellowships by requiring multidisciplinary training across computation, mathematics, and science and engineering. Fellows receive support for up to four years and complete at least one practicum at a DOE national laboratory. That requirement provides access to advanced research, an expanded network of colleagues and mentors, and the world's fastest supercomputers. The DOE CSGF encourages graduates to pursue national-lab careers and fosters collaborations with researchers at the labs and in academia and industry. In 2019, the program received more than 400 applications for just 26 available fellowships.

The DOE CSGF also provides a stipend for research equipment and conference travel. Fellows meet with alumni and present research at annual conferences, exchanging ideas, expanding networks, learning about career opportunities and developing professional skills.

James Corones, formerly of Iowa State University and DOE's Ames Laboratory, developed the DOE CSGF. The Krell Institute, the non-profit organization he founded, continues to administer it. Corones recognized that computational scientists need specialized training that makes them fluent in mathematics and computer science and in scientific disciplines such as physics, chemistry, biology and engineering where these tools are applied.

CSGF SPOTLIGHT: ASEGUN HENRY

PHYSICS, ENGINEERING AND COMPUTING



As a graduate student, mechanical engineer Asegun Henry initially worried that the computer science courses in his DOE CSGF plan of study might detract from his physics and engineering work at the Massachusetts Institute of Technology. But today he recognizes how that expertise has benefitted his studies of phonon transport. "Just about every paper I've published since graduate school leverages this computing knowledge," he says. His laboratory has a competitive edge because everyone in his group writes HPC-scalable codes from the beginning. "Other people in my field are sometimes limited to commercial or open-source codes, and some don't know how to write code for large parallel machines."

After completing his Ph.D. and DOE CSGF fellowship in 2009, Henry carried out postdoctoral research at Oak Ridge National Laboratory and Northwestern University and was a fellow at DOE's Advanced Research Projects Agency-Energy. He then started a tenure-track position at the Georgia Institute of Technology in 2012 before returning to MIT in 2018.

The fellowship taught him how to build partnerships with researchers in other fields and to appreciate their varying perspectives. "I now collaborate extensively with faculty in departments ranging from materials science and chemical engineering to electrical engineering and physics." These partnerships have yielded important advances in understanding energy and heat transfer, the world's highest-temperature pump, lower-cost ceramic metal heat exchangers for concentrated solar power, and new theoretical frameworks to more rigorously describe phonon thermal conductivity and interface conductance.

Henry has received a 2016 NSF Career Award, the 2018 Bergles-Rohsenow Award in Heat Transfer from the American Society of Mechanical Engineers and the 2018 World Technology Award in Energy.

CSGF SPOTLIGHT: SARAH RICHARDSON

INFRASTRUCTURE-BUILDER



In high school, Sarah Richardson already was learning molecular biology techniques with genetically engineered yeast and mice, experiences that led to her work as an entrepreneur today. As a DOE CSGF recipient

studying human genetics, she worked with Joel Bader at the Johns Hopkins University School of Medicine to develop algorithms for designing synthetic nucleotide sequences and engineer a synthetic yeast genome. The fellowship brought her cross-disciplinary training: an understanding of the problems in other sciences and the solutions that researchers have already developed in those areas.

Richardson served a Berkeley Lab practicum in 2008 and returned there in 2012 as a postdoctoral fellow, domesticating microbes so scientists can use the organisms' biology and chemistry to solve a range of problems in bioenergy and medicine. "Some of my work is literally building infrastructure in science," she says. That year she received one of five Women in Science \$60,000 postdoctoral fellowships from L'Oreal USA. In 2017, she founded a startup company, Microbyre, with support from Cyclotron Road, DOE's innovation program at Berkeley Lab.

Since its inception, the DOE CSGF has supported more than 500 fellows, developing their careers as computational scientists, applied mathematicians, scientists and engineers at the DOE national labs, in industry and in academia.

A 2017 study of the DOE CSGF's impact found 88 percent of 278 responding fellows and alumni said the program enhanced their HPC knowledge. More than 80 percent said it improved the overall quality of their research to a moderate or major extent. Nearly 60 percent worked on DOE supercomputers during their fellowships, and 46 percent have used those resources since completing the program.

The vast majority of alumni remain employed as computational scientists or engineers. Most start their careers as academic postdoctoral researchers or in positions at the national labs. Among 198 alumni surveyed, 57 percent had been employed in academia, 36 percent in industry and 36 percent at a DOE laboratory over their careers.

Nearly 90 percent of respondents were engaged in interdisciplinary research and had achieved their overall career goals. More than 80 percent had mentored others. More than half had received competitive research funding, and more than 40 percent had received a research award.

BEYOND THE PH.D.

Workforce development doesn't stop with graduate training. ASCR supports many postdoctoral researchers at universities and national laboratories through grants to principal investigators. In addition, the largest DOE laboratories have prestigious named postdocs that back promising early-career researchers. ANL founded the J.H. Wilkinson Fellowship in Scientific Computing, the first of these programs, in 1988. Others include Sandia's John von Neumann Fellowship in Computational Science, Oak Ridge's Alston S. Householder Fellowship in Applied Mathematics and Scientific Computing, and LBNL's Luis W. Alvarez Postdoctoral Fellowship in Computing Sciences. Many of the recipients—about half—opt for careers at the DOE labs; the rest pursue opportunities in academia and industry.

Such fellowships attract applications from many high-quality Ph.D. scientists. DOE labs often hire unsuccessful candidates into other positions, boosting the impact these programs have on recruitment and keeping the nation's tech pipeline filled with talented computational scientists.

CSGF SPOTLIGHT: ANUBHAV JAIN

THE SEARCH FOR NEW MATERIALS



The DOE CSGF was a critical component in launching Anubhav Jain's career in computational materials science. From 2006 to 2011, he conducted Ph.D. research in Gerbrand Ceder's MIT laboratory, working on high-throughput

density functional theory calculations that would eventually form the basis of the Materials Project. This open initiative, now led by Kristin Persson at LBNL and for which Jain acts as a thrust lead, uses supercomputers to model new materials that can be useful for a range of energy challenges, including better battery storage.

The DOE CSGF helped Jain leverage large supercomputers in this challenge and gave him a training advantage. "Computational materials scientists usually don't receive much formal education about large computers, how they work, and what limits they face," he says, inhibiting researchers' ability to apply HPC. "The DOE CSGF brought many of these issues to the forefront and helped me develop an appreciation for and understanding of computer science issues in large calculations."

Jain has continued that work and launched related projects at LBNL, first as a Luis Alvarez Postdoctoral Fellow and today as a staff scientist and chemist. In 2015, he received a \$2.5 million DOE Early Career Research Program award focused on using high-throughput computations to discover new materials.

ASCR@40



ASCR@40

Credit: Rakesh Yadav/Harvard University.



U.S. DEPARTMENT OF
ENERGY

Office of
Science