



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.

PETAFLUPS SPOTLIGHT: NERSC

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: 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.

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.