OLCF 25 YEARS

OLCF ANNUAL REPORT
2016–2017

Oak Ridge National Laboratory

U.S. Department of Energy | Office of Science

OLCF 25 YEARS

Oak Ridge Leadership Computing Facility
Annual Report 2016–2017

Outreach manager – Katie Bethea

Writers – Eric Gedenk, Jonathan Hines, Katie Jones, and Rachel Harken

Designer – Jason Smith

Editor – Wendy Hames

Photography – Jason Richards and Carlos Jones

Stock images – iStockphoto™

Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory
P.O. Box 2008, Oak Ridge, TN 37831-6161

Phone: 865-241-6536
Email: help@olcf.ornl.gov
Website: https://www.olcf.ornl.gov

Facebook: https://www.facebook.com/oakridgeleadershipcomputingfacility

Twitter: @OLCFGOV

The research detailed in this publication made use of the Oak Ridge Leadership Computing Facility, a US Department of Energy Office of Science User Facility located at DOE's Oak Ridge National Laboratory. The Office of Science is the single largest supporter of basic research in the physical sciences in the United States and is working to address some of the most pressing challenges of our time.

For more information, please visit **science.energy.gov**.

Contents

In a Record 25th Year, We Celebrate the Past and Look to the Future

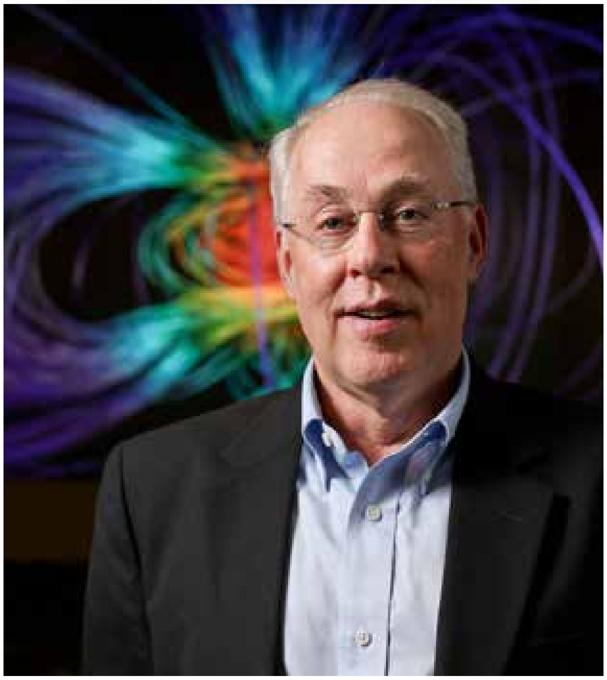

Jim Hack, Director NCCS

Excitement is building at the Oak Ridge Leadership Computing Facility as we prepare for the 2018 delivery of Summit, the facility's next leadership supercomputer, which will be up to 10 times more powerful than Titan, the OLCF's 27-petaflop flagship system.

Although there is a lot of buzz around Summit, including the very real sounds of construction on the new computer room, I am just as pleased to report on the strong scientific productivity enabled by Titan over the past year. Projects using Titan and other OLCF resources in 2016 resulted in the largest number of research publications since the system went into full production in 2013. Furthermore, record-high attendance at the annual user meeting attests to the involved engagement of our users.

For the OLCF, the past year was leadership computing at its best.

Now in 2017, we have the unique opportunity to look back—before Titan, before the 3-petaflop Cray XK6 Jaguar, even before Eagle, a teraflop IBM system installed at the turn of the new millennium—to the founding of the Center for Computational Sciences at the US Department of Energy's Oak Ridge National Laboratory. The establishment of CCS—the OLCF's predecessor—in May 1992 marks this year as our 25th anniversary. For a quarter-century, the OLCF has delivered some of the world's earliest breakthroughs in gigaflop, teraflop, and petaflop performance, progress that has transformed computational science.

In this issue, you will learn a little more about the history and growth of the OLCF ("OLCF Celebrates 25 Years," page 24). You will also read about the outstanding research taking place on Titan today and about how the OLCF is helping lead high-performance computing into the future through big data analytics and a partnership with DOE's Exascale Computing Project ("Big Data," page 18).

In 2016, 405 publications resulted from OLCF resources like Titan, and 59 appeared in high-impact journals, including Nature, Science, and Proceedings of the National Academy of Sciences—a 31 percent increase in high-impact publications from 2015. We are pleased at the exceptional productivity of our user community and the significant role of HPC at high levels of scientific discourse.

Research on the 50-year mystery of the origin of the "sigma" particle, which used OLCF resources, earned a team from DOE's Jefferson Lab the cover of Physical Review Letters, and a University College London team earned a Journal of the American Chemical Society cover for new insight gained through supercomputing simulations of large molecular adsorption at the atomic level.

Published in Nature, the first study to simulate atomic-level magnetism in regions of a magnetic nanoparticle was made possible through unprecedented experimental data and the award-winning Locally Self-Consistent Multiple Scattering quantum mechanical structure code developed at ORNL and optimized for Titan by ORNL computational scientists (Eisenbach, page 12).

Researchers from Imperial College London scaled their accelerated, industry-standard computational fluid dynamics code to more than 18,000 GPUs at 13.7 petaflops, earning a place as a finalist for the 2016 Association of Computer Machinery's Gordon Bell Prize (Vincent, page 6).

A team from the University of Michigan used Titan to study melting in 2-D systems—information that could improve predictions of phase transitions in nanoparticle applications from pharmaceuticals to thin films used in solar panels and batteries (Glotzer, page 10). The GPU-accelerated simulations provided extensive support for a well-known theoretical model of 2-D melting.

In this issue, you will also read about teams that are simulating a detailed 3-D picture of Earth's interior (Tromp, page 8), comparing hundreds of electrolyte materials to predict the best class of materials for lithium-ion batteries (Miller, page 14), and using supercomputing to study the type of photosynthesis that takes place in desert plants, which could show scientists how to make food crops more drought resistant (Jacobson, page 16).

In all, more than 1,100 users across 318 projects accessed OLCF resources. OLCF staff are dedicated to improving the user experience and increasing the breadth of our capabilities, and have responded to several challenges and opportunities this year.

For the second year in a row, our staff ensured Titan was available for user projects 97 percent of the time, exceeding our target availability by 7 percent.

Although Titan, now 4 years old, is aging and will eventually reach its operational lifespan, I am proud of our staff for finding proactive strategies for maintaining system reliability, providing 35 million additional core hours over last year.

We are striving to provide users with versatile services for data analysis integrated with our HPC systems. The OLCF will work more closely with ORNL's Compute and Data Environment for Science (CADES) to accelerate progress in developing machine learning, graph analytics, and other big data workflows and to connect data-centric resources with traditional HPC capabilities. Teaming with CADES will also provide ORNL researchers with access to HPC, cloud computing, and other systems so they can leverage modeling, simulation, and data analysis across the range of research domains at ORNL.

To explore new architectures and prepare for Summit, the OLCF has acquired several test bed systems. Examples include the DGX-1 NVIDIA system, the world's first purpose-built system for deep learning that has already attracted several health science projects that mine for connections in medical information. Another important acquisition is Summitdev, a 54-node system with IBM Power8 and NVIDIA Pascal GPUs that is one generation away from Summit's architecture. The OLCF early science Center for Accelerated Application Readiness teams are already hard at work using the Summitdev environment to optimize 13 diverse science applications so they are available for use on Summit on day one.

CAAR teams are also ensuring portability across diverse architectures through Percival, a Cray XC40 system with Xeon Phi processors, and computer science teams are exploring an ARM 64-bit programming architecture.

Our User Assistance and Outreach staff strive to help users focus on their science research by helping them navigate OLCF systems and accelerate codes for Titan's GPUs, and they once again have received high marks in user satisfaction. UAO staff hosted monthly user conference calls, conducted four hackathons, and created a new video training channel to help users maximize their time on OLCF resources. And the annual user meeting in May 2016 was the OLCF's largest user meeting ever.

In addition to facilitating capability-limited scientific investigations of our world, we continue to host hundreds of facility tours, informing groups ranging from middle school students to government officials of the science that Titan makes possible.

The impacts made here, year after year, through scientific discoveries, technological advances, and the work of purpose-driven staff, continue to raise the value of HPC for solving some of the world's toughest scientific and energy-related problems.

Peter Vincent and his research team used Titan to create a comprehensive simulation of a key jet engine component, called the low-pressure turbine. Shown here is a simulation of flow over three jet engine low-pressure turbine blades. Credit: Imperial College

Streamlining Accelerated Computing for Industry

Project: Development and Benchmarking of an In Situ Visualization Pipeline for Next Generation CFD Tools

PI: Peter Vincent

Institution: Imperial College

Total Usage: 39,000,000 hours

Building and testing physical prototypes of complex machines can be costly and time-consuming and can provide only limited feedback. For these reasons, companies in industries as diverse as aerospace, car manufacturing, and wind power have been turning to supercomputers to investigate complex design problems related to fluid flow, or how air and fluids interact with a machine.

But as supercomputers have increased in size and scale, many computational fluid dynamics (CFD) applications used by industry have struggled to keep pace with advances in high-performance computing (HPC), limiting their accuracy and ability to fully supplant physical testing.

To modernize CFD, a group of Imperial College researchers led by Peter Vincent developed open-source software called PyFR, a Python-based application that combines highly accurate numerical methods with a flexible, portable, and scalable code implementation that makes efficient use of accelerators like Titan's GPUs. Industry adoption of the code could allow companies to better exploit petascale computing to understand long-standing fluid flow problems, in particular unsteady turbulence—the seemingly random and chaotic motion of air, water, and other fluids.

To demonstrate PyFR's HPC prowess, the team ran a high-resolution, GPU-accelerated simulation of flow over a jet turbine linear cascade, scaling the simulation up to 18,000 GPUs. The team's highest-performing run contained 195 billion degrees of freedom—or independent variables—and operated at a sustained speed of 13.7 petaflops, or 13.7 quadrillion calculations a second. The work garnered Vincent's team a nomination for the 2016 Association of Computer Machinery's Gordon Bell Prize, one of the most prestigious prizes in supercomputing.

"Titan is the only supercomputer big enough to do the scale of simulation that we wanted to try."

Turbulent systems—picture a wave breaking on the beach or the curl of smoke rising from a campfire—feature a range of scales that emerge, mix, and move through space and time. PyFR balances these factors by employing a highly parallel numerical scheme known as flux reconstruction that ties together several high-accuracy methods.

Unlike current CFD codes, which typically use mathematical averages to approximate the unsteady features of turbulent systems, flux reconstruction allows researchers to do more fine-grained calculations of turbulent physics. This means researchers can run PyFR on large-scale accelerated architectures to accurately resolve fluid features that were previously out of reach.

For companies bent on creating the next lightweight, fuel-efficient jet engine, the resulting insights from high-resolution simulation can prove invaluable. Modern jet engine turbines are designed to use as few blades as possible. However, this arrangement can lead to unsteady airflow patterns that reduce engine efficiency. To study this phenomenon, Vincent's team used Titan to create a comprehensive simulation of a key jet engine component, called the low-pressure turbine. The team is currently collaborating with German aircraft engine manufacturer MTU Aero Engines to improve the technology.

"Titan is the only supercomputer big enough to do the scale of simulation that we wanted to try," Vincent said. "Our objective is to provide a capability that can act as a virtual wind tunnel for these low-pressure turbine cascades, where you can resolve all the physics and get an accurate answer without any tweaking to fit experimental results."

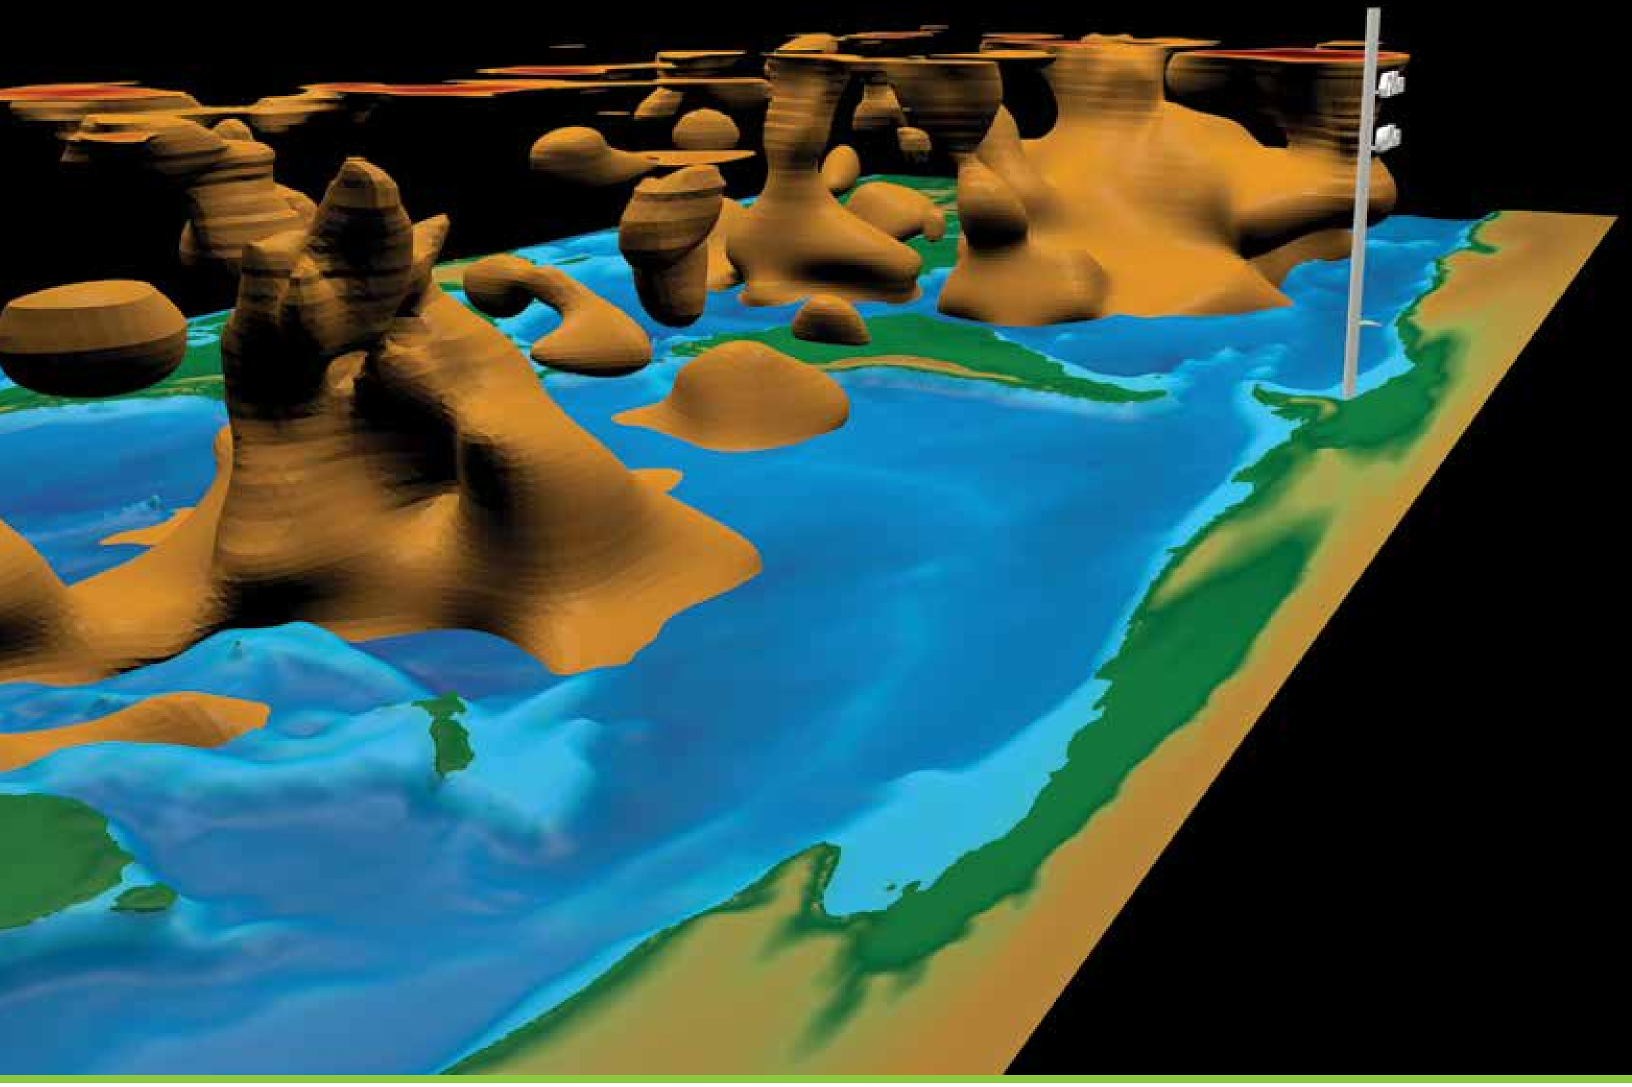

Shown here is a side view of a 3-D global map of Earth's interior created using seismic data generated by earthquakes and a mathematical method called adjoint tomography. As seismic waves travel, seismograms can detect variations in their speed. These changes provide clues about the composition, density, and temperature of the inner Earth. Credit: David Pugmire/ORNL

A Seismic Mapping Milestone

Project: Global Adjoint Tomography

PI: Jeroen Tromp

Institution: Princeton University

Total Usage: 73,000,000 hours

When an earthquake strikes, the release of energy creates seismic waves that often wreak havoc on life at the surface. Those same waves, however, present an opportunity for scientists to peer into the subsurface by measuring vibrations passing through the Earth.

Using advanced modeling and simulation, seismic data generated by earthquakes, and the Titan supercomputer, a team led by Jeroen Tromp of Princeton University is creating a detailed 3-D picture of Earth's interior. Currently, the team is focused on imaging the entire globe from the surface to the core–mantle boundary, a depth of 1,800 miles.

These high-fidelity simulations add context to ongoing debates related to Earth's geologic history and dynamics, bringing prominent features like tectonic plates, magma plumes, and hotspots into view. In September 2016, the team published a paper in Geophysical Journal International on its first-generation global model. Created using data from 253 earthquakes captured by seismograms scattered around the world, the team's model is notable for its global scope and high scalability.

"For the first time, we showed people the value and feasibility of running these kinds of tools for global seismic imaging," said Ebru Bozdag, a coprincipal investigator of the project and a professor at the Colorado School of Mines.

As seismic waves travel, seismograms can detect variations in their speed. These changes provide clues about the composition, density, and temperature of the medium the wave is passing through. For example, waves move slower when passing through hot magma, such as mantle plumes and hotspots, than they do when passing through colder subduction zones, locations where one tectonic plate slides beneath another.

"For the first time, we showed people the value and feasibility of running these kinds of tools for global seismic imaging."

Each seismogram represents a narrow slice of the planet's interior. By stitching many seismograms together, researchers can produce a 3-D global image, capturing everything from magma plumes feeding the Ring of Fire, to Yellowstone's hotspots, to subducted plates under New Zealand.

This process, called seismic tomography, works in a manner similar to imaging techniques employed in medicine, where 2-D x-ray images taken from many perspectives are combined to create 3-D images of areas inside the body.

Running its GPU version of the SPECFEM3D_GLOBE code, Tromp's team used Titan to apply a seismic tomography method called adjoint tomography, an iterative full-waveform inversion technique. This technique leverages more information than competing methods, using forward waves that travel from the quake's origin to the seismic receiver and adjoint waves, which are mathematically derived waves that travel from the receiver to the quake.

The team then compared these "synthetic seismograms" with observed seismic data supplied by the Incorporated Research Institutions for Seismology (IRIS), calculating the difference and feeding that information back into the model for further optimization. Each repetition of this process improves global models.

For its initial global model, Tromp's team selected earthquake events that registered between 5.8 and 7 on the Richter scale—a standard for measuring earthquake intensity. Tromp's team is working with Oak Ridge Leadership Computing Facility (OLCF) staff to extend that range to include more than 6,000 earthquakes in the IRIS database—about 20 times the amount of data used in the original model.

"For our first-generation model, we completed 15 iterations, which is actually a small number for these kinds of problems," Bozdag said. "Despite the small number of iterations, our enhanced global model shows the power of our approach."

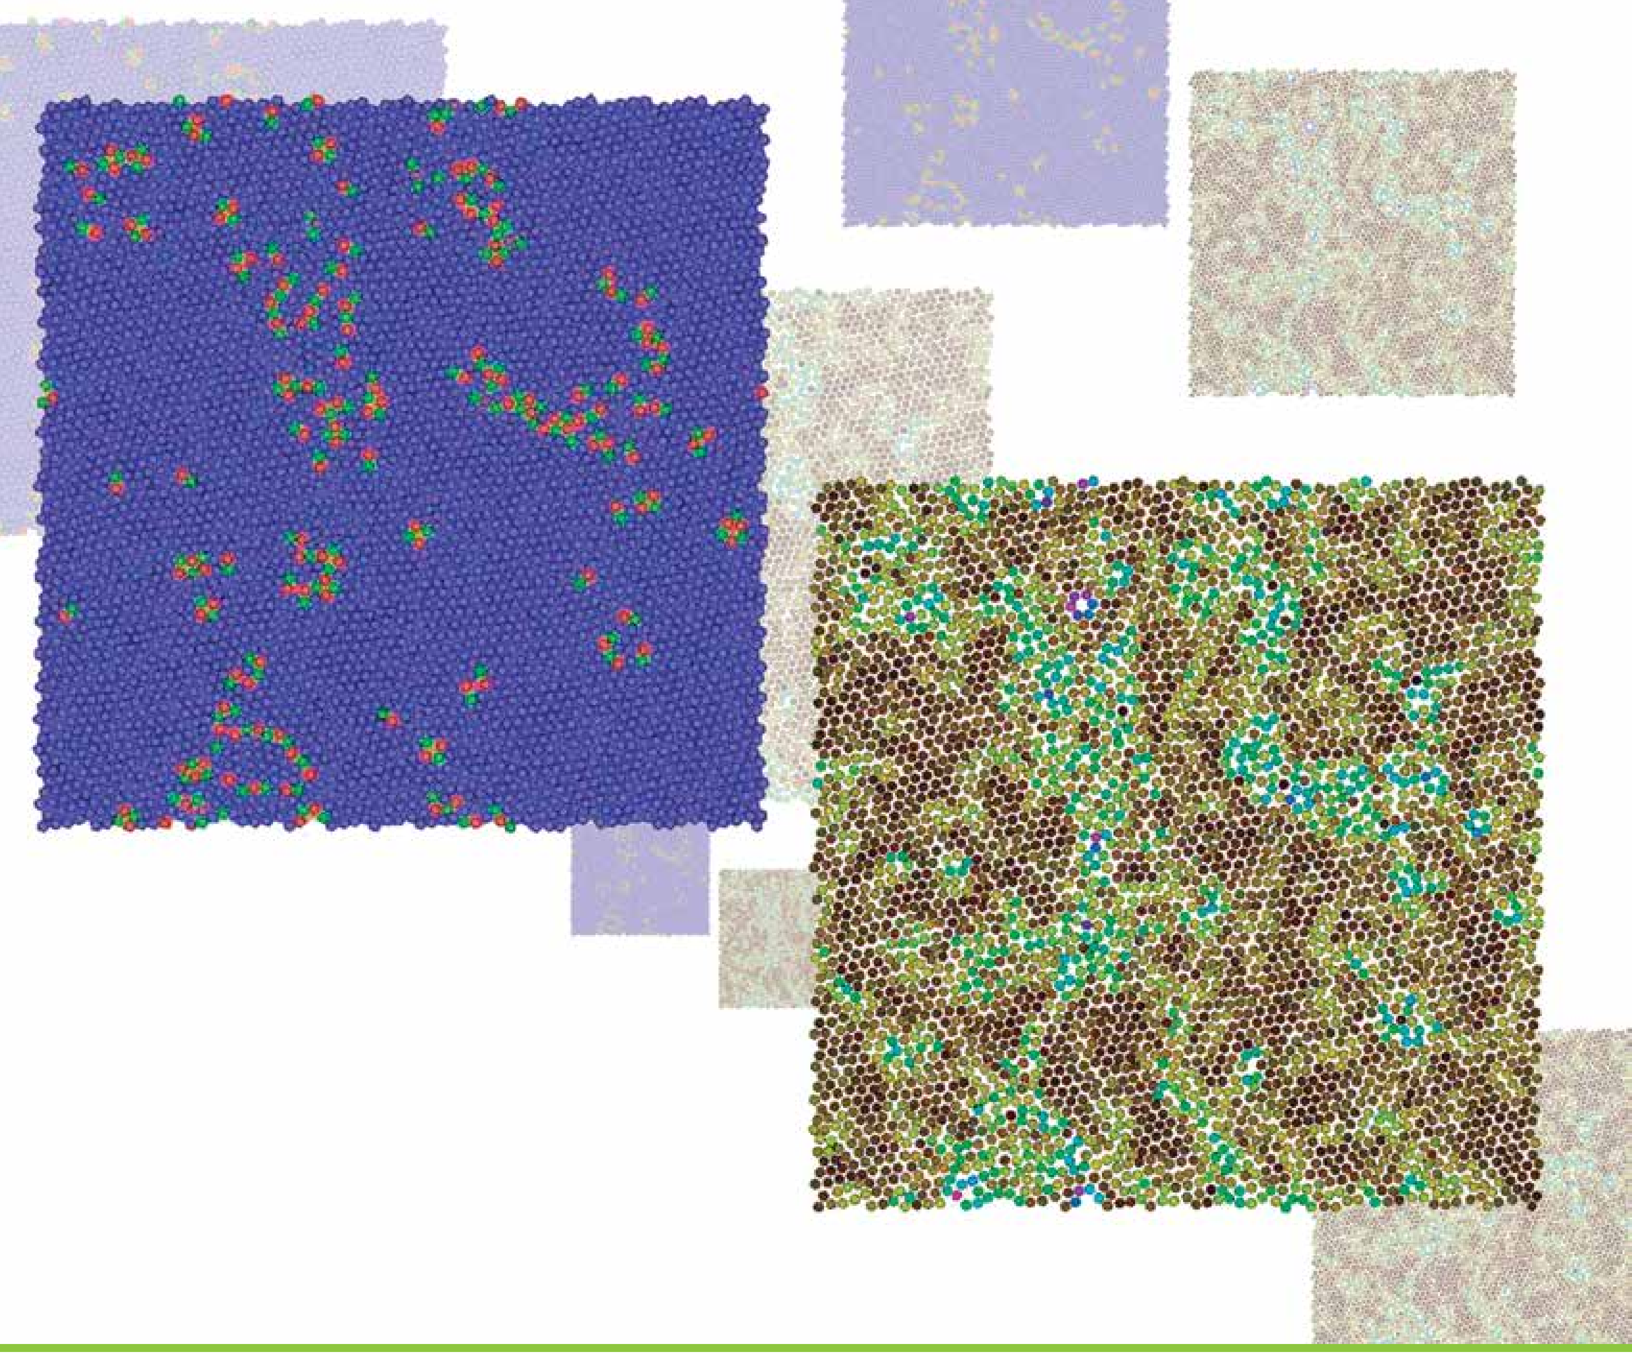

Using Titan, Sharon Glotzer's team took a close look at melting in 2-D systems, exploring how particle shape affects the physics of a 2-D solid-to-fluid melting transition. Shown are visualizations of a two-dimensional hard particle system of hexagons under external pressure. Credit: J. A. Anderson and J. Antonaglia/University of Michigan

The Shape of Melting in Two Dimensions

Project: Nucleation and Growth of Colloidal Crystals Using Highly Scalable Monte Carlo

PI: Sharon Glotzer

Institution: University of Michigan

Total Usage: 66,000,000 hours

An ice cube in room-temperature water. Wax near the flame of a candle. An ice cream cone on a hot summer day. Melting is a familiar phenomenon encountered in everyday life. Like any phase transition, such as a solid changing to a liquid or a liquid changing to a gas, it involves a transformation from one state of matter into another.

Inducing and controlling phase transitions is one way humans have become adept at manipulating their environment and creating new technologies. Although melting is a common occurrence, much remains to be discovered about this transformation at a fundamental level.

At the University of Michigan, Professor Sharon Glotzer leads a computational research team dedicated to studying matter's tendency to self-organize, a line of inquiry with applications in advanced materials and nanosystems. Using Titan, Glotzer's team took a close look at melting in 2-D systems, exploring how particle shape affects the physics of a 2-D solid-to-fluid melting transition.

The team's work revealed that the shape and symmetry of particles can dramatically affect the melting process, a fundamental insight that could help guide researchers in search of nanoparticles with desirable properties for energy applications such as solar panel materials. Additionally, this work paves the way for Glotzer's team to tackle pressing phase transition problems in three dimensions, such as how fluid particles crystallize into complex colloids—mixtures in which particles are suspended throughout another substance.

To sufficiently tackle the 2-D melting problem, Glotzer's team used its GPU-enabled HOOMD-blue code to simulate 11 different shape systems, ranging from triangles to 14-sided polygons, of up to 1 million particles. Each system was simulated at 21 different densities with the lowest densities representing a fluid state and the highest densities a solid state.

The researchers identified three distinct melting scenarios dependent on the shape of the systems' polygons. The most significant scenario emerged from hexagon systems, which Glotzer's team found to perfectly follow the phase transition described by a well-known theoretical model of 2-D melting called the KTHNY theory, first posited in the 1970s by researchers John Kosterlitz, David Thouless, Burt Halperin, David Nelson, and Peter Young.

In this scenario, the particles exhibit an intermediate phase between a solid and a liquid known as the hexatic phase. This state is characterized by systems of atoms exhibiting orientational order, a general—yet partially fragmented—configuration. As the theory predicts, the particles shift from solid to hexatic and hexatic to fluid in a perfect continuous phase transition, meaning the particle system changes at a constant rate in response to a changing external pressure.

Because the hexatic phase generally contains sixfold orientational order, "it seems natural that the hexagon, with its six sides, and the honeycomb-like hexagonal arrangement would be a perfect match for this theory," said University of Michigan research scientist Joshua Anderson. "This is something that hasn't been described until now."

"This is something that hasn't been described until now."

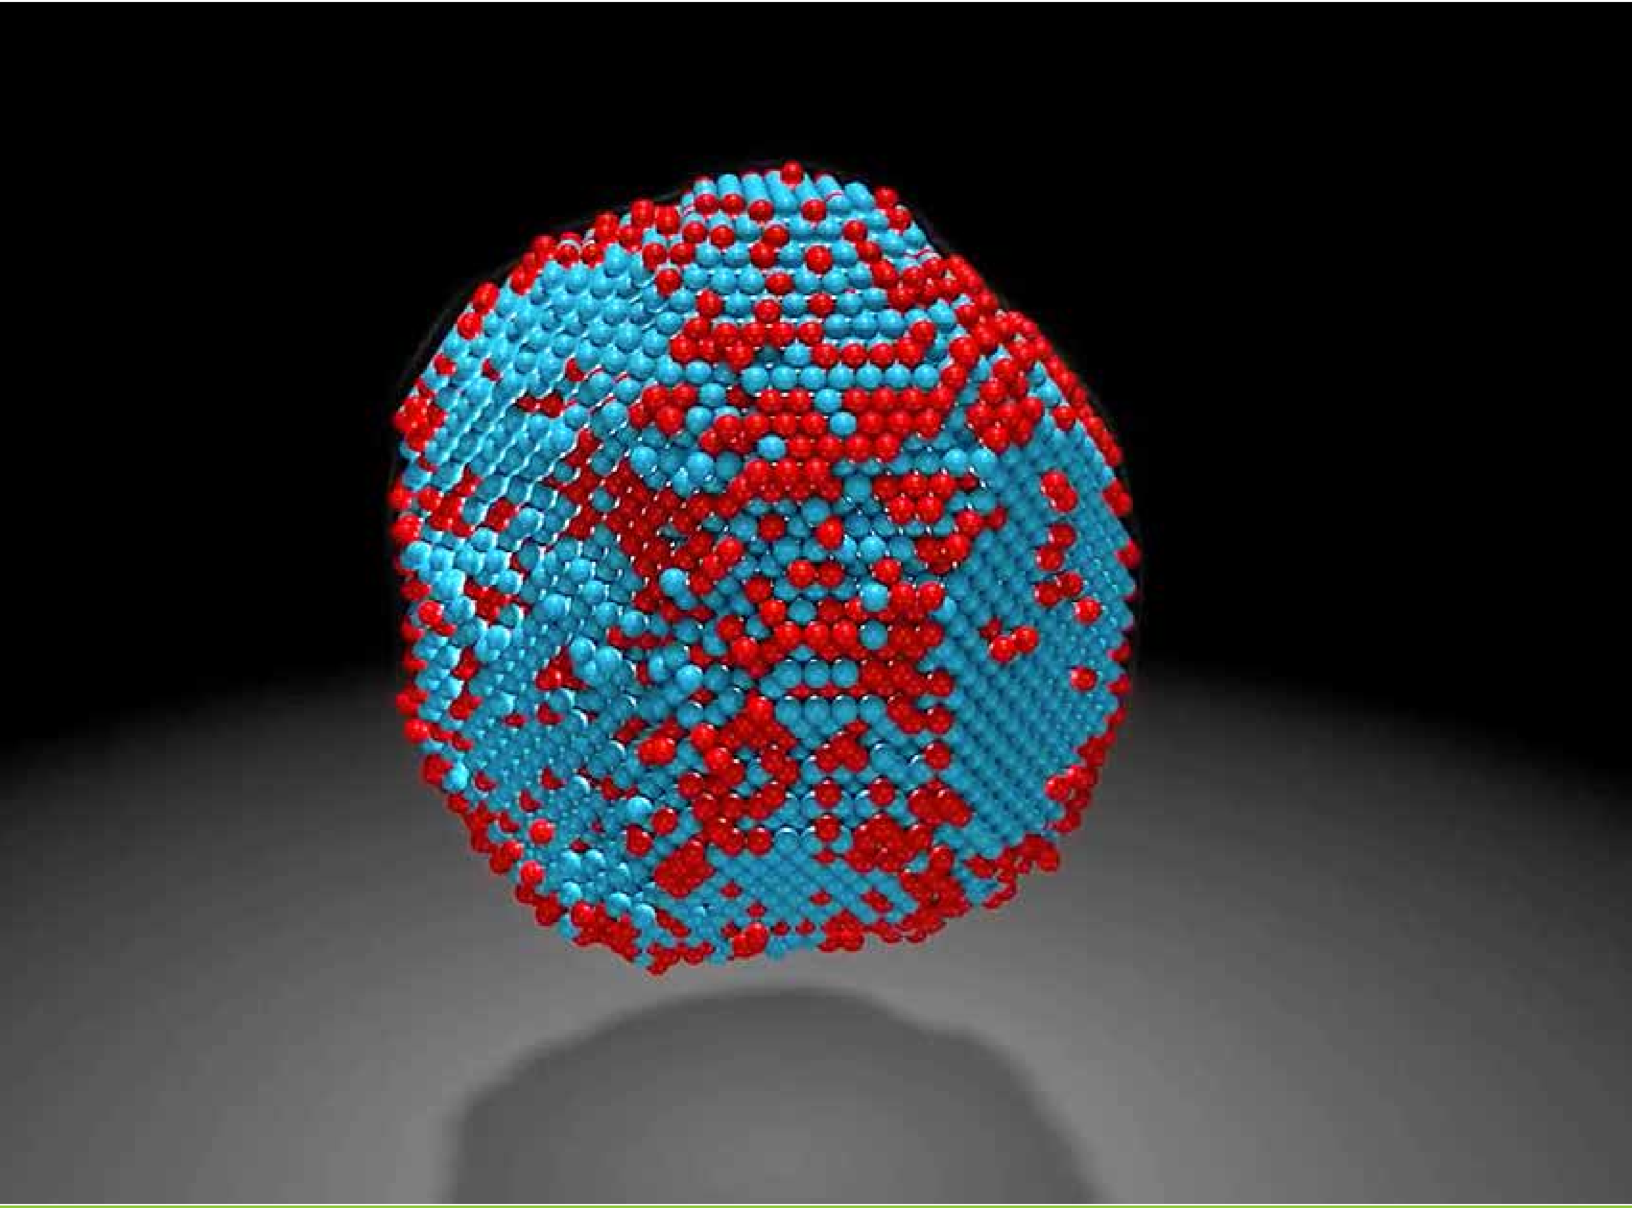

Researchers were able to use Titan to simulate magnetism atom by atom in a real nanoparticle. Shown here is a visualization of chemically-ordered phases in an iron-platinum (FePt) nanoparticle. Credit: C. Ophus/Lawrence Berkeley National Laboratory, Nature

A Supercomputing First for Predicting Magnetism in Real Nanoparticles

Project: First Principles–Based Statistical Physics of Alloys and Functional Materials

PI: Markus Eisenbach

Institution: Oak Ridge National Laboratory

Total Usage: 61,500,000 hours

With the potential to increase storage capacity and density, magnetic nanoparticles are promising materials for next-generation storage devices. Understanding how magnetism works at the atomic level will be critical for optimizing the performance of nanodevices, yet observing magnetism at such small scales is extremely difficult.

That's why researchers working with iron–platinum (FePt) nanoparticles at the University of California, Los Angeles (UCLA), and the US Department of Energy's (DOE's) Lawrence Berkeley National Laboratory (Berkeley Lab) worked with computational scientists at DOE's Oak Ridge National Laboratory (ORNL) to do something new—simulate magnetism atom by atom in a real nanoparticle.

The results were published in Nature in February 2017.

"These types of calculations have been done for ideal particles with ideal crystal structures but not for real particles," said Markus Eisenbach, OLCF computational scientist, who develops quantum mechanical electronic structure simulations that predict magnetic properties in materials.

Electronic structure codes solve for magnetic properties based on 3-D atomic and chemical structures. However, these structures are typically derived from many 2-D electron microscopy or x-ray crystallography images averaged together, resulting in a representative, but not true, 3-D structure.

"In this case, researchers were able to measure the precise 3-D structure for a real particle," Eisenbach said.

The West Coast teams used an electron microscope at Berkeley Lab and a 3-D reconstruction algorithm developed by UCLA to trace the positions of about 6,500 iron and 16,500 platinum atoms in a single FePt nanoparticle. Their sophisticated experimental data revealed 3-D chemical disorder and other defects at the atomic level.

using the Linear Scaling Multiple Scattering (LSMS) code—a first-principles density functional theory code developed at ORNL that received Association for Computing Machinery Gordon Bell Prizes in 1998 and 2009.

The team also ran VASP, a simulation package that is better

"These types of calculations have been done for ideal particles…but not for real particles."

suited for smaller atom counts, to simulate regions of about 32 atoms.

"VASP results were consistent with LSMS results, so we have a high confidence in the simulations," Eisenbach said.

The unprecedented simulations revealed that magnetic anisotropy energy suddenly transitions at grain boundaries, an important result for focusing future studies.

Although first-principles calculations are currently too intensive to solve small-scale magnetism for regions larger than a few thousand atoms, researchers hope that advances in computing and simulation will make a full-particle simulation possible in the future.

"There's a hope going forward that one would be able to use these techniques to look at nanoparticle growth and understand how to optimize growth for performance," Kent said.

Using the new data, Eisenbach and Paul Kent, a computational materials scientist at ORNL, precisely modeled the nanoparticle's atomic structure and simulated its magnetic properties on the OLCF's Titan supercomputer.

FePt nanoparticles are grouped into different regions or "grains" of iron and platinum atoms. Scientists want to understand how magnetism changes from grain to grain and how those transitions might influence the performance of magnetic devices. In particular, researchers are interested in magnetic anisotropy, or what direction magnetism favors from atom to atom.

"If the anisotropy is too weak, a bit written to the nanoparticle might flip at room temperature," Kent said.

To solve for magnetic anisotropy, Eisenbach and Kent used two computational codes to compare and validate results.

The ORNL team simulated a supercell of about 1,300 atoms from strongly magnetic regions of the 23,000-atom nanoparticle

Using Titan, researchers are creating models to screen electrolyte materials that could make lighter, longer lasting batteries. Credit: iStockphoto

Researchers Flip Script for Lithium-Ion Electrolytes to Simulate Better Batteries

Project: Next-Generation Nanostructured Polymer Electrolytes by Molecular Design

PI: Thomas Miller

Institution: California Institute of Technology

Total Usage: 49,000,000 hours

Ever since Italian physicist Alessandro Volta invented the first battery out of a stack of copper and zinc disks separated by moistened cardboard, scientists have been searching for better battery materials.

Lithium-ion batteries—which are lighter, longer lasting, and more functional than standard batteries under a wider range of temperatures—power everything from cell phones to aircraft carriers to electric cars. Their ubiquitous use makes their stability, efficiency, and safety important for businesses and consumers alike.

To improve lithium-ion batteries further, researchers must find novel, nonflammable materials for the electrolyte, the crucial battery component that shuttles lithium ions during charging and discharging. Promising electrolyte candidates must be not only stable but also conductive to lithium ions. This allows batteries to maintain efficiency during charge cycles.

A team led by the California Institute of Technology's (Caltech's) Thomas Miller used Titan to identify potential electrolyte materials and predict which ones could enhance the performance of lithium-ion batteries. The researchers ran hundreds of electrolyte simulations—each consisting of thousands of atoms. Their results led them to a class of polymer materials that may help experimentalists looking to invent the next battery technology.

Typically lithium-ion batteries feature liquid electrolytes, but new research is focusing on solid polymeric electrolytes, which are known to be more stable, less flammable, and less volatile. An ideal electrolyte is one that readily dissolves and then conducts lithium ions, but the problem with commonly used solid electrolytes is that they conduct positively charged lithium ions poorly and negatively charged anions too rapidly. Brett Savoie, a postdoctoral fellow at Caltech, said the team started its search by looking for solid polymers with the opposite qualities.

Using Titan, the team first created a coarse-grained simulation protocol called the chemically specific dynamic bond percolation model to screen electrolyte materials based on short molecular dynamics trajectories. Running its simulations on LAMMPS, a classic molecular dynamics code, Miller's team analyzed several dozen polymer–salt combinations under different salt concentrations. About 400 simulations at a time were run in parallel, each consisting of around 3,000 atoms periodically replicated in 3-D space.

The team ultimately found that a class of polymers called Lewis-acidic polymers not only conducted anions more slowly than previous solid electrolytes did but also conducted the positive lithium ions more quickly. Because the positive regions of Lewis-acidic polymers are contained in a small amount of space and their negative regions are spread out over a large amount of space, they give lithium ions more opportunities to dissolve, Miller said.

> "The predictions indicate that these polymers might exhibit a substantial increase in conductivity."

"It was known that Lewis-acidic molecules slowed down anions," Savoie said. "What was surprising here was that by using a purely Lewis-acidic system, we also sped up the lithium."

The simulations showed that these polymers may be capable of producing an eightfold increase in desired lithium conduction and a marked decrease in the unwanted anion conduction. This would be—given the historically slow pace of discovering new polymer materials—a very large jump.

"These new polymers are exciting because they seem to overcome some of the main problems with other polymer materials," Miller said. "The predictions indicate that these polymers might exhibit a substantial increase in conductivity. It would be a tremendous improvement from the current lithium-ion conductivity."

Miller and his team have applied for a patent based on this work.

Here an ORNL researcher studies desert plants. These plants use crassulacean acid metabolism—a photosynthesis method that holds promise for helping plants of all kinds conserve water. Credit: Jason Richards/ORNL

A Real CAM-Do Attitude

Project: Scaling Up of Parallelized Ortholog Detection Algorithms for Comparative Genomics of Bacterial Genomes

PI: Dan Jacobson

Institution: Oak Ridge National Laboratory

Total Usage: 100,000 hours

Photosynthesis, the method plants use to convert energy from the sun into food, is a ubiquitous process many people learn about in elementary school. Almost all plants use photosynthesis to gather energy and stay alive.

Not all photosynthetic processes are the same, though. In recent years, researchers have grown increasingly interested in desert plants' preferred method of photosynthesis—crassulacean acid metabolism (CAM).

These plants caught researchers' attention because of their seemingly opposite photosynthetic schedule, and understanding this process may be the genetic key to helping plants of all kinds conserve water. With a more fundamental understanding of CAM, scientists aim to help the plants upon which society relies for food and fuel become more drought resistant, thereby expanding the area where crops can grow and thrive.

To that end, ORNL computational biologist Dan Jacobson works with a large group of experimentalists and computational scientists to more fully understand the CAM process. This cross-omics team (combining expertise in metabolomics, proteomics, and genomics) uses computing resources at the OLCF to catalog how plants' CAM processes vary and ultimately uncover how CAM processes may be genetically engineered into feed stock, food crops, and crops for bioenergy applications.

When most people think of photosynthesis, they are actually thinking of a specific form, called C3 photosynthesis. This process follows the Calvin Cycle, in which plants capture light energy during the day and convert it into energy-bearing adenosine triphosphate (ATP).

ATP helps plants split water atoms into their hydrogen and oxygen constituent particles. Meanwhile, a C3 photosynthetic plant opens up small pores—called stomata—to absorb carbon dioxide from the atmosphere. Then at night, the newly freed hydrogen particles combine with carbon dioxide absorbed during the day to create the carbohydrates plants use to live and grow.

> "I think the ability to use supercomputing infrastructure enabled things that wouldn't have been possible otherwise."

CAM photosynthesis works the same way, but stomata open for respiration at night and stay tightly closed during the day, allowing plants to conserve more water. This helps plants like cactus and Agave survive in climates where water is scarce. Researchers hope that by understanding how CAM works, they can apply this water-saving method to other plants. To do that, though, researchers need to understand how molecules interact during CAM photosynthesis and how metabolites and proteins change over time.

For this project, researchers from ORNL, the University of Tennessee, Newcastle University in the United Kingdom, and the University of Nevada, Reno, gathered photosynthesis data from Agave (a CAM plant) and compared it with the Arabidopsis genus of plants (C3 plants). To conduct a study between Agave and a C3 plant, the team selected the Arabidopsis genus plant thale cress, one of the first plants to have its genome sequenced and a good candidate for plant studies.

The team then studied what gene expressions control stomata opening and closing in both CAM and C3 plants and how proteins regulated this process. Collecting this data in both a common CAM and a C3 species allowed the team to distinguish traits that are ubiquitous to CAM plants from species-specific traits.

Jacobson indicated that without access to high-performance computing, the team would not have been able to find these meaningful connections in a timely manner. "This is the first study looking at a cross-omics, time-course experiment to try and explore CAM at this molecular detail," he said. "I think the ability to use supercomputing infrastructure enabled things that wouldn't have been possible otherwise. We were able to have a pretty big impact on the analysis of this work because of those resources."

Big Data Emphasis and New Partnerships Highlight the Path to Summit

In 2014, the OLCF announced its next-generation supercomputer, Summit, set to enter full production in 2019. Summit will offer up to 10 times the computational power of Titan. However, these advancements also present new challenges. Since Summit's announcement, OLCF staff members have been hard at work preparing users and the facility for these impending changes, through both the scientific discovery process (algorithmic or computational and data-oriented developments) and infrastructure upgrades.

Each of Titan's 18,688 nodes has one GPU and one CPU. Summit's architecture will consist of six GPUs and two CPUs per node, meaning that to achieve maximum performance, application developers will need to find additional parallelism to help applications get the most out of Summit's complex architecture.

New computing architectures and associated new technologies always present challenges, and as a long-time leader in HPC, the OLCF has always been at the frontline of developing technologies to address these challenges. For OLCF staff, disruptive technologies are opportunities to innovate and help the entire HPC community move forward.

To that end, the OLCF formed several new partnerships and collaborations in 2016 that will help proactively address new HPC challenges and offer researchers opportunities to prepare their applications to make the best use possible of next-generation supercomputers. With these new partnerships, the OLCF is positioning itself as the leading

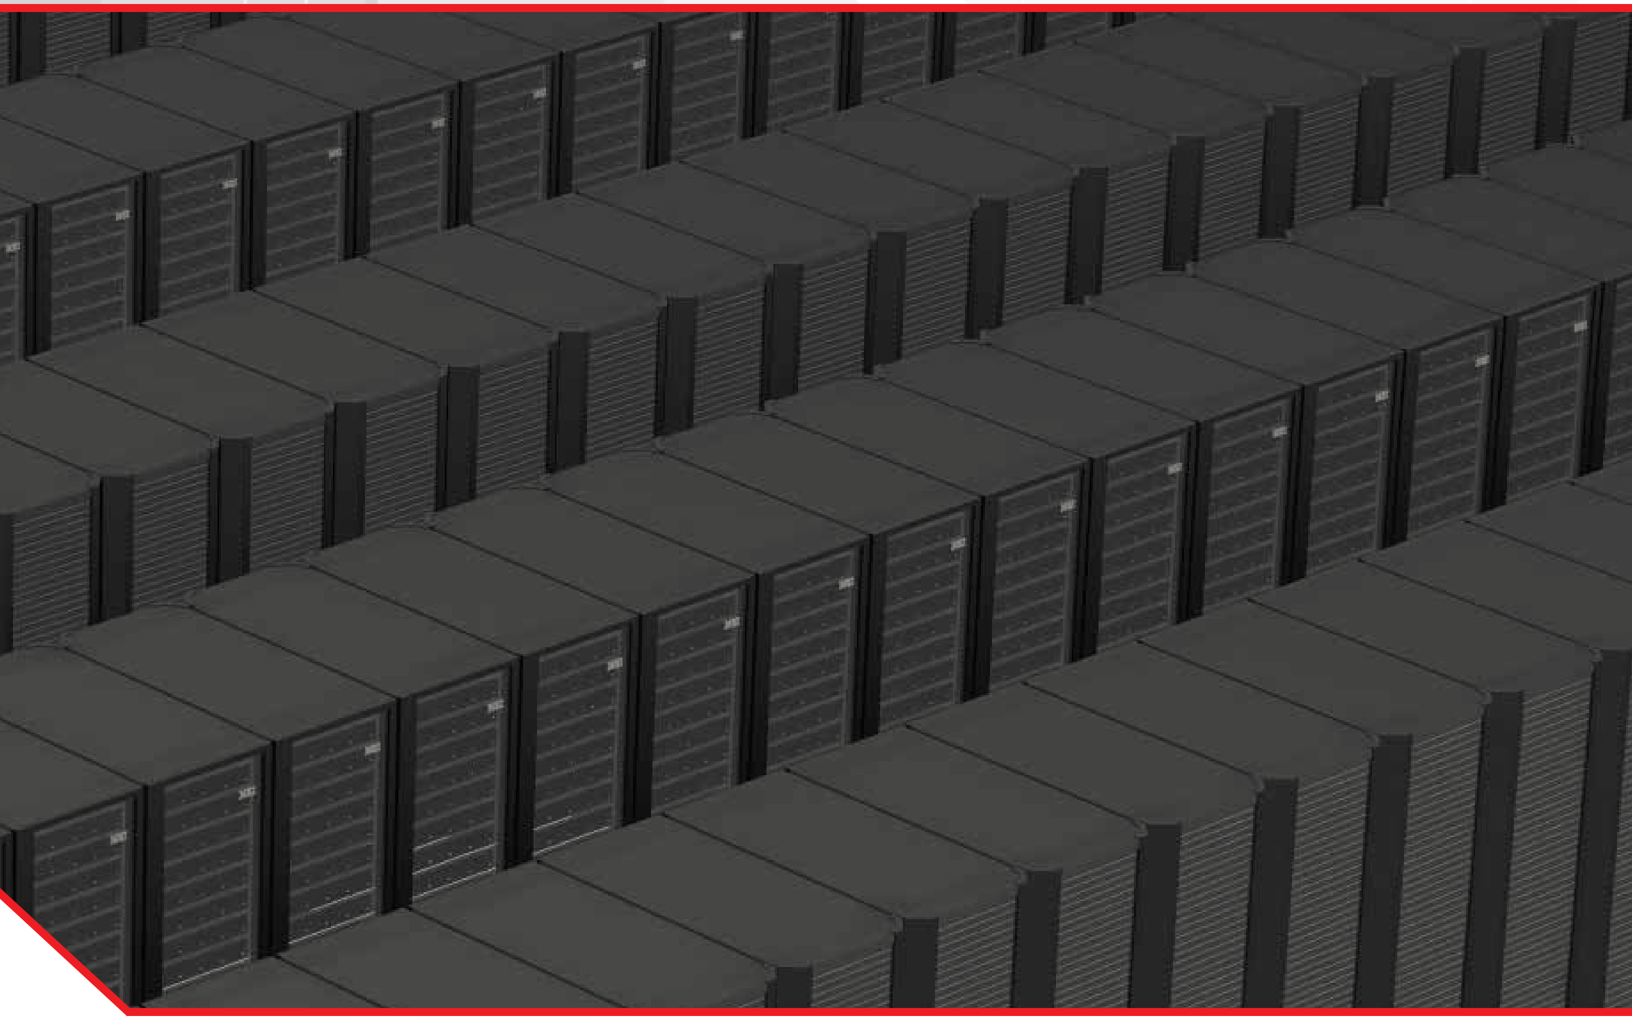

An artist's rendering of the Summit system. Credit: Andy Sproles/ORNL

problem-solving and scientific discovery environment in HPC.

One of the major initiatives to address these challenges came from DOE. In the second half of 2016, DOE announced funding for the Exascale Computing Project (ECP). In funding ECP, the United States recognized the importance of maintaining HPC competitiveness and allocated resources to enable delivery of an exascale computer—a supercomputer offering a thousandfold increase over current-generation petascale machines, or 50 times faster than the OLCF's 27-petaflop Titan—by 2021.

Efforts supporting ECP are dispersed across the national laboratory system, and ORNL experts are playing leading roles in the ECP project. Doug Kothe, deputy associate laboratory director in ORNL's Computing and Computational Sciences Directorate, was named director of ECP's application development initiative. The first round of application development awards fully funded 15 science projects, with several others being funded at the "seed" level. The project awards totaled $34.8 million, and ORNL researchers Paul Kent, John Turner, and Thomas Evans are each leading projects. In addition, ORNL is home to the ECP Project Management Office, led by Kathlyn Boudwin.

ECP also awarded $34 million to 35 software development projects, including those led by ORNL researchers David Bernholdt, Scott Klasky, Wayne Joubert, and Jeffrey Vetter. In addition to serving in leadership roles, several OLCF and ORNL staff members are

Center for Accelerated Application Readiness researchers, Thomas Papenbrock, Gaute Hagen and Gustav Jansen, discuss improving codes that must run on increasingly more powerful supercomputers. Credit: Jason Richards/ORNL supporting most of the application and software development awards in some capacity.

Both the application and software development awards help researchers proactively address challenges for next-generation HPC and beyond.

"A number of the new software technology projects have been informed by our experiences with OLCF users. These projects build on various innovative efforts which the OLCF has been pursuing," Bernholdt said. "The Exascale Computing Project will allow us to take these efforts to the next level, with a larger team and the ability to more easily place them in the broader context of preparation for exascale systems and applications. We're very excited about the ECP and expect it to provide significant future benefits to OLCF users."

These awards and collaborations are at the center of the United States' efforts to proactively address the challenges associated with next-generation supercomputers and ensure that US researchers have access to the best high-performance computing infrastructure in the world. In advance of its last two leadership-class systems, Titan and Summit, the OLCF started the Center for Accelerated Application Readiness (CAAR) to help prepare users for architectural changes that could affect performance. Teams were selected based on how their research supports DOE's science and energy missions, and they worked closely with OLCF and vendor experts to optimize their codes for Titan's new architecture.

The Titan-era CAAR initiative resulted in a handful of applications running at scale on Titan from the beginning. These codes were designed to study everything from radiation transport (the Denovo code, relevant for studying nuclear energy and other nuclear technology applications) to molecular dynamics (the LAMMPS code, used for simulating large-scale atomic systems).

"A number of the new software technology projects have been informed by our experiences with OLCF users. These projects build on various innovative efforts which the OLCF has been pursuing."

After announcing Summit, the OLCF began a second phase of the CAAR initiative for prospective Summit users and selected 13 teams to participate. Summit-focused CAAR teams are paired with an OLCF computational liaison and offered technical support through the IBM/NVIDIA Center of Excellence, housed at ORNL.

In addition to OLCF expertise, CAAR teams can use early-access platforms that more accurately represent the computing environment users will experience on Summit. In the fall of 2016, the OLCF began installation of the Summit Early Access development platform, known as Summitdev.

Summitdev's current architecture is one generation removed from the processors that will be installed in Summit, but having IBM's POWER architecture on Summitdev allows users to fine-tune their codes to run on IBM processors before Summit arrives. Summitdev contains 36 nodes, with each node consisting of two IBM POWER8 CPUs and four NVIDIA Tesla GPUs.

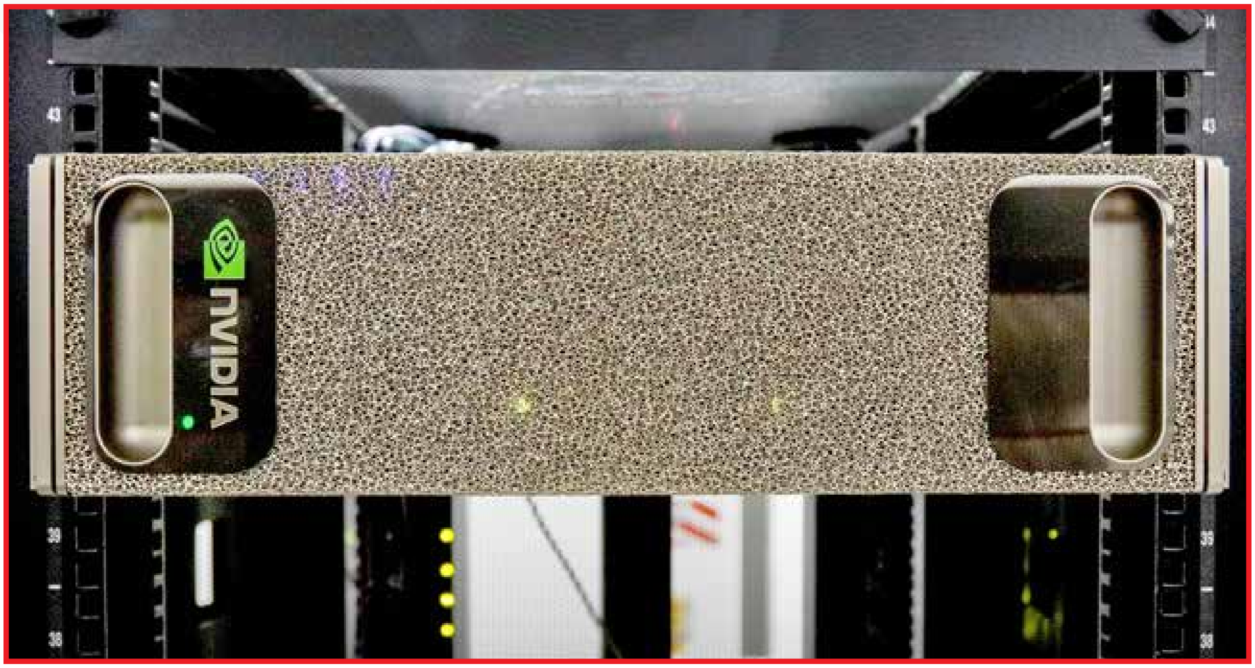

The OLCF recently deployed a DGX-1 artificial intelligence supercomputer by NVIDIA alongside ORNL's Compute and Data Environment for Science (CADES). The DGX-1 can be accessed by users to explore scientific workflows. Credit: Jason Richards/ORNL

In addition to supporting CAAR teams, Summitdev was expanded to 54 nodes to accommodate additional application and software developers.

Summitdev gives users a head start on refactoring their codes for Summit, but to do so, users need to be informed of how to take full advantage of the resource. To that end, CAAR teams were invited in January 2017 to participate in a Summitdev training workshop, where OLCF staff members and vendor partners jointly delivered training on the new platform.

"We want to provide our CAAR teams with the expertise and resources they need to take full advantage of Summit when it comes online," Tjerk Straatsma, Scientific Computing group leader, said. "Having access to Summitdev provides application developers the opportunity to port and optimize their codes on an architecture that closely resembles the Summit architecture. This will help get these applications ready for use by our user programs as soon as Summit becomes available and starts delivering science outcomes."

Being a part of CAAR not only allows users to learn how to debug and efficiently scale their codes but also enables them to become more familiar with the whole computing ecosystem at the OLCF, including the file and storage systems. This end-to-end support will allow users to focus on scientific productivity when Summit comes online rather than waste valuable time with debugging, dealing with I/O problems, or getting slowed down in the data transfer process.

In addition to the central mission of getting users ready for next-generation machines, the OLCF is also a leader in scalable data science and analytics.

For several years, ORNL's Compute and Data Environment for Science (CADES) has offered compute and data services that help users maximize their scientific output. Resources of CADES—ORNL's de facto big data center—allow ORNL user facilities such as the Spallation Neutron Source (SNS) and the Center for Nanophase Materials Science, ORNL staff members, and research teams to process, manage, and analyze large amounts of data. Staff members in CADES help tailor scalable infrastructure for a research team's needs so that users can deploy scientific workflows to accomplish their scientific discovery objectives.

CADES provides users access to a variety of open and moderate research resources and offers several heterogeneous compute resources. A Scalable HPC multi-tera-op computing cluster offers program-driven projects burst-capable high-performance computing. To provide leadership computing access to staff at ORNL, CADES also hosts a hybrid CPU–GPU Cray XK7 platform set up in a manner similar to Titan. CADES operates a several-thousand-core expandable cloud computing infrastructure to offer much-needed flexibility and data-analytics capabilities to all research staff at the laboratory. These heterogeneous resources are backed by a high-performance storage environment including a parallel file system, volume management for the OpenStack-driven cloud environment, and an object-storage system.

To support powerful data analytics from generated data and empirically observed data from ORNL's facilities, CADES also enables analytics on a set of platforms including SGI's shared-memory clusters, NVIDIA's Deep Learning appliances, and Cray's Urika-GX. These platforms coupled with the flexible cloud offer a dramatically scalable and configurable analytics environment.

An important service CADES offers is enabling scalable computing and data services for scientific workflows that span data creation facilities (such as SNS). The heterogeneous and workflow-capable infrastructure offers researchers multiple paths to scientific discovery, much like the compute, data analysis, and storage capabilities at the OLCF.

To align CADES' mission with ORNL's core mission of continuing to lead in high-performance, scalable computing, ORNL leadership created a new partnership between its flagship compute and data projects. In 2016, ORNL leadership moved the CADES team into the National Center for Computational Sciences (NCCS)—a leadership computing complex located at ORNL that houses the OLCF.

"CADES enables our diverse scientific staff to use the appropriate resources for their scalable computing and data-analytics needs," said Arjun Shankar, who directs CADES. "Ultimately, leadership computing and large-scale data analysis will accelerate scientific discovery for DOE core missions."

Summit will likely be the last pre-exascale machine at the OLCF. Executing CAAR and connecting with CADES will be vital to setting the stage for scientific success on Summit from the first day it is online. Meeting the challenge of delivering an exascale machine requires years of planning and preparation. To this end, OLCF leadership has already begun working with sister computing facilities—the Argonne Leadership Computing Facility, located at Argonne National Laboratory, and Lawrence Livermore National Laboratory—to begin the procurement process for Summit's successor.

Despite being several years away from full production, Summit has already inspired and encouraged the OLCF to adopt a forward-looking approach to solving next-generation computing challenges. The year 2016 brought new partnerships to the OLCF, and as the Summit era draws closer, OLCF staff members see these collaborative efforts not only as a way to get users ready to take full advantage of Summit but also as a springboard to continue OLCF leadership in high-performance computing toward the exascale horizon.


OLCF
25
YEARS

OLCF Celebrates 25 Years of HPC Leadership

In 2017, the OLCF is celebrating 25 years of leadership in high-performance computing. Since its founding as the ORNL Center for Computational Sciences (CCS) in May 1992, the center has consistently delivered supercomputers of unprecedented capability to the scientific community on behalf of DOE. Scientists, in turn, have used these versatile systems to solve critical problems in areas as diverse as energy, biology, and nanotechnology, among other domains.

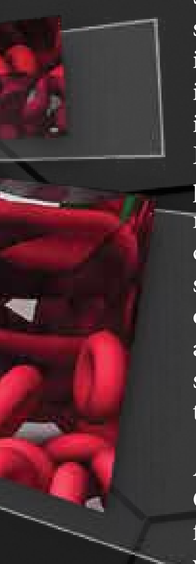

Since its founding with the launch of an Intel Paragon system, the OLCF has contributed to a rapid evolution in scientific computing that has produced a millionfold increase in computing power. This rise has included the installation of the first teraflop system (IBM Power3 Eagle) for open science, the science community's first petaflop system (Cray XT5 Jaguar), and two top-ranked machines on the TOP500 list, including the OLCF's current leadership-class machine Titan. Using OLCF systems, scientists have expanded the scale and scope of their research, solved complex problems in less time, and filled critical gaps in scientific knowledge. Today, simulation is considered on par with experiment and theory as an essential standard of modern science.

At the start of the 1990s, however, predicting the OLCF's role in elevating HPC would have seemed far-fetched. At the time, most of DOE's supercomputers were found in laboratories dedicated to national security and defense. Though ORNL's scientific computing group had been testing experimental architectures since the 1980s, the only HPC systems in Oak Ridge were a Cray X-MP with a single CPU and a Kendall Square computer with a novel shared-memory architecture.

An opportunity to leap to the front of the pack in scientific computing presented itself in 1991, when Congress passed the High Performance Computing and Communication Act, legislation that called for government funding of the technology and infrastructure that would come to be known as the Internet. Additionally, the act directed DOE's Office of Science to issue a call for proposals for the creation of new HPC research centers to serve scientists from national laboratories, universities, and private industry.

At ORNL, Laboratory Director Al Trivelpiece charged Ed Oliver, director of the Office of Laboratory Computing, with assembling a team that could put forth a compelling case for an HPC center to be located at Oak Ridge. The resulting document, titled "Partnerships in Computational Science" and submitted in collaboration with three national laboratories and seven universities, called for the construction of a large parallel computer to address three grand challenges: groundwater remediation, materials modeling, and quantum structure. The proposal also requested funds to develop new computational techniques, software, and tools to support this work and educational programs to create interest in supercomputing among students in kindergarten through college.

When DOE announced on May 24, 1992, that ORNL had won the bid, the newly formed CCS immediately went to work with partner Intel to stand up a relatively new type of machine: a massively parallel supercomputer. This type of system solved large problems by dividing the work among many processors, allowing calculations to be carried out simultaneously. The approach represented a shift from older architectures, which relied on less flexible, increasingly expensive vector processors.

In 1992, CCS unveiled the Intel Paragon, a 66-core system capable of 5 gigaflops (5 billion calculations per second). A series of ever-more-powerful parallel computers at ORNL culminated in the Intel Paragon XP/S 150. In 1995, the Paragon XP/S 150 was capable of 150 gigaflops, making it the fastest supercomputer in the world upon installation.

"At the time, big parallel supercomputers might have 100 processors. We brought in the Intel Paragon, which

Buddy Bland with the Intel Paragon XP/S 150 in 1995. Credit: ORNL

had over 3,000 processors," said Buddy Bland, who came to ORNL in 1984 to run the X-MP and was hired as the first CCS employee. Today, Bland serves as the OLCF project director, managing the deployment of a pre-exascale system called Summit, which will be capable of more than 100 petaflops (100 quadrillion calculations per second).

Working with scientific users, CCS staff helped showcase the benefits of parallel computing for research by getting scientific applications to run faster. World-class science results followed, with projects providing insights into problems such as how solids melt, how combustion occurs in engines, and how air flows over a plane's wing.

The Paragon was replaced in 1999 by an IBM Power3 system called Eagle, the first teraflop machine in DOE's Office of Science. Other machines followed, including an IBM Power4 (Cheetah), a Compaq AlphaServer (Falcon), and a Cray X1 (Phoenix). The arrival of each supercomputer presented new opportunities and challenges for CCS staff, who also contributed technical papers to the field to benefit the HPC community.

In the early 2000s, ORNL cemented supercomputing's place at the laboratory by setting aside space in a new office building for a 40,000-square-foot datacenter. The investment paid off in 2004 when Congress passed the DOE High End Computing Revitalization Act, a response to bolster the nation's position in HPC in the wake of Japan's Earth Simulator, a machine as powerful as 14 of the fastest supercomputers in the United States combined.

After a competitive proposal process, DOE awarded ORNL the first leadership computing facility, later named the Oak Ridge Leadership Computing Facility, to provide computing resources 100 times larger than those commonly available at the time. In addition to working with industry, universities, and other labs, the center was also charged with supporting other federal agencies, such as the National Science Foundation, the National Nuclear Security Administration, the National Institutes of Health, and the National Oceanographic and Atmospheric Administration.

The establishment of the OLCF led to a new Cray XT3 system called Jaguar, which, over the next few years, went through a series of upgrades—single-core processors to dual-core processors, then quad-core processors to six-core processors. In 2008, Jaguar, now a Cray XT5, became the first HPC system to run a scientific application at a sustained petaflop. The achievement, carried out using a quantum cluster code called DCA++, won an ORNL-led team the prestigious Gordon Bell Prize the same year. In 2009, Jaguar reached No. 1 on the TOP500 list after running the Linpack benchmark application at 1.75 petaflops.

"While I may have been the first employee of the Center for Computational Sciences, we couldn't have done all of this without the great people that we have."

At its apex in 2012, Jaguar boasted 300,000 cores and a theoretical peak of 3.3 petaflops, allowing researchers to investigate phenomena like high-temperature superconductivity in incredible detail. The machine was widely recognized as much for its scientific productivity as for its technical prowess.

"This was one of the first machines to demonstrate petascale performance on real applications and was a major asset to the laboratory," Bland said.

The OLCF's next move, however, transformed Jaguar into something else entirely. The combination of 16-core processors and GPU accelerators marked the birth of Titan, a hybrid machine that took parallelism to the next level.

Conceived by the OLCF and Cray as the answer to two seemingly incompatible objectives—a significant increase in computational power for a minimal increase in energy consumption—Titan debuted at No. 1 in November 2012 on the TOP500 list. Specifically, with a theoretical peak of 27 petaflops, Titan realized 10 times the performance of Jaguar with only a 20 percent increase in electricity usage. Using Titan, researchers have contributed to the development of better biofuels, more efficient utility-scale gas turbines, and hazard maps for earthquake-prone regions. This new research has led to thousands of published scientific journal articles made possible by OLCF systems and, more importantly, has contributed to a safer, more energy-secure future.

Underlying the OLCF's legacy of exponential growth in computing power and scientific productivity is a dedicated staff working with HPC vendors, scientific users, and others to serve the greater goals of science and the nation.

"While I may have been the first employee of the Center for Computational Sciences, we couldn't have done all of this without the great people that we have," Bland said. "There have been many people over the years who have worked in the Center for Computational Sciences and the Leadership Computing Facility to deliver on these new, first-of-their-kind machines, working very closely with the vendors to develop the systems, to get them to work, to fix the problems, and then with the scientists, to actually deliver the scientific results."

Groups within the OLCF

Advanced Data and Workflow

Group Leader: Mallikarjun (Arjun) Shankar

The Advanced Data and Workflow Group offers scientific, technical, operational, and thought leadership for building scalable data services for OLCF user needs. The group designs and develops creative data-science workflows, analytics, and visualization services to enable interactive data-driven discoveries that require scale and performance. These services empower users on leadership computing resources hosted by the OLCF and big data resources hosted by CADES. This team of data scientists, software engineers, and scientific visualization experts guides users through the data-to-knowledge discovery process by understanding user needs, designing scalable algorithms and workflows, and implementing and supporting tools to accelerate scientific discovery.

Scientific Computing

Group Leader: Tjerk Straatsma

The Scientific Computing Group (SciComp) works in partnership with users to help them obtain optimal results from the OLCF's computational resources and systems. The SciComp group is composed of computational scientists representing a wide spectrum of domain sciences, including astrophysics, biophysics, chemistry, climate science, combustion, computational fluid dynamics, computer science, mathematics, and numerical analysis. Each research team using OLCF systems for a project in the INCITE user program is assigned a SciComp liaison who is familiar with the field of research. SciComp liaisons actively participate in the research, help design and optimize code for the users' applications, streamline the computational workflow, solve computer issues that arise, and serve as the point of contact for expert help from data workflow and visualization specialists to assist in capturing data from the computational campaign in images and helping users analyze it.

High Performance Computing Operations

Group Leader: Kevin Thach

The High Performance Computing Operations (HPCO) Group keeps the OLCF leadership supercomputing and data systems running. Members of the group monitor all systems 24 hours a day, 7 days a week, 365 days a year, and are responsible for administration, configuration management, and cybersecurity. The staff work with infrastructure systems, with Titan, and with other OLCF supercomputers. The HPCO group tests the systems when they are installed and upgraded, uses diagnostic tools for continuous system monitoring, anticipates problems before they arise, identifies components that are near failure, and ensures cybersecurity of the systems, defending them from continuous attacks.

User Assistance and Outreach

Group Leader: Ashley Barker

The User Assistance and Outreach Group provides support to the OLCF users, serves as connections between the outside world and the OLCF, and acquaints the public with the work conducted at the OLCF. The group creates accounts for new users; provides technical support to research teams; generates documentation on OLCF systems access, policies, and procedures; and provides training opportunities to current and future HPC users. The group also creates science research highlights, writes articles for the public, produces videos and podcasts for scientists and the public, and connects the OLCF with universities across the nation.

Technology Integration

Group Leader: Sudharshan Vazhkudai

The Technology Integration (TechInt) Group is charged with delivering new technologies into the OLCF and other projects at ORNL by identifying gaps in the system software stack and working with the research and development community to develop, harden, and deploy solutions. The group's technology scope includes archival storage, parallel file systems, non-volatile memory, system architecture, high-performance networking, and data management. TechInt staff members work behind the scenes to develop the infrastructure that supports NCCS systems, keep the infrastructure ahead of the technology curve, research and evaluate emerging technologies, and provide systems programming to integrate technologies and tools seamlessly into the infrastructure as they are adopted. In addition, TechInt staff work on future system acquisitions and their preparations to support the OLCF mission.

Computer Science Research

Group Leader: David Bernholdt

ORNL's Computer Science Research Group supports the programming environment and tools for current and emerging OLCF computing systems through a combination of approaches. The group works to translate relevant research, carried out locally and in the broader community, into appropriate standards and high-quality implementations. The work often involves close interactions with vendors (of both compilers and performance and correctness tools) and the open source community. The group is driven and validated by in-depth engagement with OLCF users.

OLCF User Group and Executive Board

OLCF User Conference Calls

The OLCF User Group (OUG)—consisting of all facility users—meets once a month by conference call and webinar to discuss OLCF news, resources, policies, and timely HPC tutorials and techniques. The OLCF user conference calls cover comprehensive topics, including best practices, new tools, and how-to guides. The conference calls are either knowledge-based talks, during which staff members give presentations or tutorials that explain tools or policies, or user experience talks, which allow users to share their experiences with the facility's systems. During the calls, users give feedback, ask questions, and learn from staff and each other about the center's techniques and tools.

In 2016, more than 385 people attended the 11 user conference calls. Knowledge-based talks covered topics such as managing Python, using the Globus toolkit for grid computing, and visualizing data with VisIt. During one of the 2016 calls, Mike Zingale, associate professor in the Department of Physics and Astronomy at Stony Brook University, presented a user experience talk titled "Intro to Visualization with YT on Titan and Rhea."

OLCF training channel:
https://vimeo.com/channels/olcftraining

OLCF training events page:
https://www.olcf.ornl.gov/support/training-events/

OLCF help desk:
help@olcf.ornl.gov

OLCF's 2016 User Meeting

The OLCF hosts an on-site user meeting for the OUG each year to share computational science and engineering achievements and to encourage user and staff interaction. The 2016 User Meeting was the largest user meeting to date, with 133 people in attendance. The meeting included user, facility, and how-to talks focused on finer points of HPC at the OLCF and on the future of HPC.

The meeting coincided with the election of new OUG Executive Board members. The 11-member board provides the OLCF with in-depth feedback and guidance on topics such as training, facility resources, and policies. The board also organizes working groups to solicit user feedback on specific topics, such as the annual user meeting.

2016 OUG members

Katrin Heitmann, Chair
Argonne National Laboratory

Balint Joo, Vice Chair
Thomas Jefferson National Accelerator Facility

Mike Zingale
Stony Brook University

Brian Wirth
University of Tennessee

John Turner
Oak Ridge National Laboratory

Stephane Ethier
Princeton Plasma Physics Laboratory

Hai Ah Nam
Los Alamos National Laboratory

Thomas Maier
Oak Ridge National Laboratory

Mark Taylor
Sandia National Laboratory

Joe Oefelein
Sandia National Laboratories

Yifeng Cui
San Diego Supercomputer Center

INCITE, ALCC, DD

Researchers who are interested in using Titan can apply for the program path that best fits their research needs.

For large-scale, computationally intensive research projects, DOE's Leadership Computing Facilities lead the INCITE program. INCITE awards are sizeable allocations (typically, tens to hundreds of millions of processor hours per project) to address grand challenges in science and engineering. In 2016, the OLCF's 31 INCITE projects were awarded a collective 2.5 billion hours on Titan. Awards in this program are made annually. For more information about INCITE, visit www.doeleadershipcomputing.org.

The Advanced Scientific Computing Research Leadership Computing Challenge (ALCC) program is open to scientists from the research community in national laboratories, academia, and industry. The ALCC program allocates computational resources at the OLCF and other computing facilities for special situations of interest to DOE programmatic needs with an emphasis on high-risk, high-payoff simulations in areas directly related to DOE's energy mission. The OLCF's ALCC projects were awarded 1.2 billion hours through this program in 2016–17. For more information, please visit www.science.energy.gov/ascr/facilities/alcc.

Director's Discretionary (DD) projects are dedicated to leadership computing preparation, INCITE and ALCC scaling, and application performance to maximize scientific application efficiency and productivity on leadership computing platforms. The DD program also seeks to attract new and nontraditional applications to use leadership-scale resources. The OLCF Resource Utilization Council and independent referees review and approve all DD requests. Applications are accepted year-round at www.olcf.ornl.gov/support/getting-started/olcf-director-discretion-project-application/.

Among the three allocation programs, the total number of projects active in 2016 was 318.

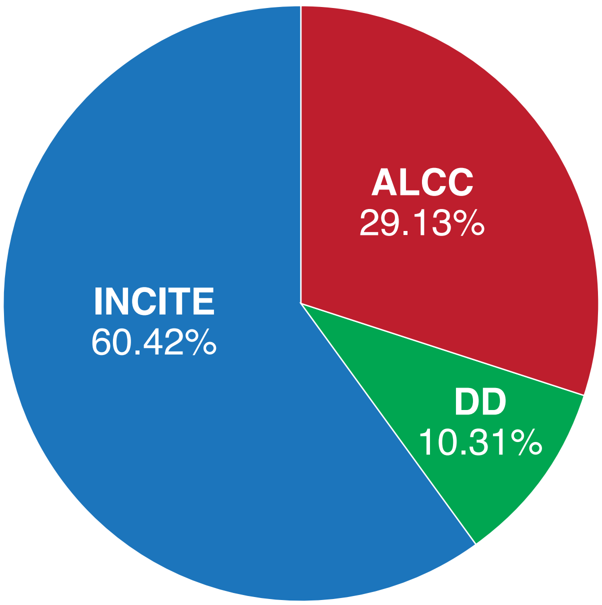

The 2016 utilization by allocation

Resource Overview

Titan Cray XK7

Compute

Titan—A hybrid-architecture Cray XK7 system capable of 27 quadrillion calculations per second. Titan is the OLCF's flagship system for leadership-class scientific computing.

Eos—A Cray XC30 cluster composed of 736 nodes and more than 47 terabytes of memory. Eos handles small-scale jobs that prepare users for running on Titan.


HPSS

Storage

HPSS—An archival file system consisting of disk and tape components for secure, long-term storage of scientific data. HPSS holds 61.77 PB of stored data.

Spider II—A center-wide parallel file system consisting of more than 30 PB of disk space. Spider II offers high-performance data transfer and simultaneous access to the OLCF's major platforms.


EVEREST

Visualization & Analysis

EVEREST—An analysis and visualization laboratory with 3-D capabilities for detailed visualization of simulation data.

Rhea—A 512-node data analysis cluster for scientific discovery dedicated to pre- and post-processing of simulation data generated on Titan.

CADES

Support Services

DTNs—Data transfer nodes (DTNs) field wide-area and local-area data transfers to and from the OLCF network. In most cases, DTNs improve transfer speed and decrease the load on a computing system's login and service nodes. Twenty DTNs are in production.

CADES—The Compute and Data Environment for Science is a fully integrated HPC ecosystem, offering compute and data services for ORNL researchers. Researchers can process, manage, and analyze large amounts of data using designated HPC resources, scalable storage, data analysis, and visualization tools.

Summitdev

Test Bed Systems

In 2016, the OLCF added four new test bed systems to the facility. These systems provide early access to a new architecture or offer exploratory environments for new HPC areas. Each will allow the OLCF to achieve new milestones in both research and technology as staff members dive into the Summit architecture and deep learning and ARM architecture–based systems. Projects that will use some of these systems include those involving CAAR, ECP, big data, and collaborations exploring the ARM architecture.

Summitdev—The Summit Early Access development platform used by CAAR and ECP to help prepare applications for Summit.

ARM1—An experimental cluster test bed that supports computer science research projects aimed at exploring the ARM architecture.

DGX-1—A small NVIDIA supercomputer that features a significant deep learning workflow in its software tools.

Percival—A Cray XC-40 based on the Knights Landing processor that is being used to explore application portability.

User Experience

The OLCF has established a support model based on continuous improvement, regular assessment, and a strong customer focus. The center can provide new services or upgrades based on customer use, demand, and feedback. OLCF leadership can also improve the center's support process based on the minimum satisfactory scores in its annual user survey, which measures the OLCF's support goals. By measuring performance based on a series of quantifiable metrics, the OLCF can ensure that users receive prompt and effective support from staff members.

In 2016, the OLCF upgraded its data transfer service to increase performance and speed while minimizing the risk of system outages. Upgrades to the DTNs included a major hardware refresh and the addition of 44 nodes. The center also increased the DTNs' Ethernet capabilities to 40-gigabit connectivity to enable faster transfers without extra steps for users. The upgrade was partially in response to increased utilization of DTNs for Globus-enabled GridFTP data transfers, which are used for moving large datasets to other facilities.

The OLCF also added a new software package—Allinea Forge—to enhance system performance for users. Forge combines the debugging capability of Allinea DDT with the MAP software performance tool. Allinea DDT allows users to control and examine data from hundreds of thousands of processing elements simultaneously, whereas MAP software provides in-depth analysis and bottleneck pinpointing and works as a profiler for both OpenMP and MPI.

The 2016 annual user survey demonstrated that users were satisfied with the OLCF and its support services. The mean rating for overall satisfaction was 4.6 out of 5 for the OLCF and 4.5 out of 5 for OLCF Support Services. Overall ratings for the OLCF were positive; 95% of users reported being "satisfied" or "very satisfied" with the center. Also, 92% of the center's tickets were resolved within 3 business days, exceeding the problem-resolution metric in 2016.

"The OLCF has some of the best computational resources in the world, managed by a very capable team of computer scientists and engineers who make freely available very useful supporting tools to take advantage of these resources," one user said.

The OLCF also delivered on its promise of providing complementary user support vehicles, including user assistance and outreach staff as well as scientific, data, and visualization liaisons.

Another user reported receiving "excellent support services from our INCITE liaison."

In 2016, the OLCF supported 1,186 users on 318 projects and addressed 2,404 reported user issues.

Education, Outreach, and Training

The OLCF hosts workshops, user conference calls, training events, and seminars to equip users with the tools they need for success with the OLCF's leadership-class machines. These activities not only alleviate difficulties in using large-scale parallel systems but also engage both the public and the user community. The training program takes a multifaceted approach, informing users about software tools and updates while also establishing software development best practices.

In 2016, the OLCF facilitated 11 user conference calls, 4 week-long hackathons, one mini-hackathon, and several other workshops and seminars.

The OLCF and partners expanded efforts in 2016 to train individuals and organizations from academia, industry, and government in GPU programming for current and next-generation HPC systems. In the fall, the OLCF conducted a GPU hackathon called ORNLHack, which included teams from St. Jude Children's Research Hospital, Duke University, and Georgia Institute of Technology. "Often companies either don't have a lot of access to experts or have to hire consultants," said Fernanda Foertter, user assistance specialist and hackathon coordinator for the OLCF. "Events like this can give them a push to the next level."

The OLCF also conducted its first mini-hackathon in 2016. The 3-day mini-hackathon was geared toward participants with little to no experience accelerating applications for GPUs.

Partnering with the Argonne Leadership Computing Facility, the National Energy Research Scientific Computing Center, and the Interoperable Design of Extreme-scale Application Software (IDEAS) project, the OLCF cohosted a series of webinars, "Best Practices for HPC Software Developers," in 2016. The series offered basic tools and processes that researchers could use to make software development easier. The series was presented through a teleconferencing service during which attendees could ask questions and experience interactive demonstrations.

As part of its program of collaborations, internships, and fellowships for young researchers, the OLCF brought in 29 students during summer 2016 with backgrounds ranging from computer architecture to mathematics and statistics to artificial intelligence.

The OLCF also strives to reach external populations. In the summer of 2016, the OLCF collaborated with ORNL's Women in Computing Group to host the "Introduce Your Daughter to Code" event, during which girls ages 10–16 ran code on Titan. The event provided daughters of staff in ORNL's Computing and Computational Sciences Directorate the opportunity to learn about code from programmers in several ORNL departments.

The OLCF also brought Tiny Titan—a mini parallel computer that helps educate students about the principles of parallel computing—to the Bay Area Maker Faire, regional schools, the USA Science and Engineering Festival, and the DOE National Science Bowl in 2016. Tiny Titan helps the OLCF reach student populations by showing how adding or subtracting processors can affect a simulation. Based in the overlook area of the OLCF where tours are conducted regularly, the tool continues to help make Titan and the power of supercomputers more understandable to numerous visitors each year.

'TitanWeek' Recognizes Contributions of Nation's Premier Supercomputer

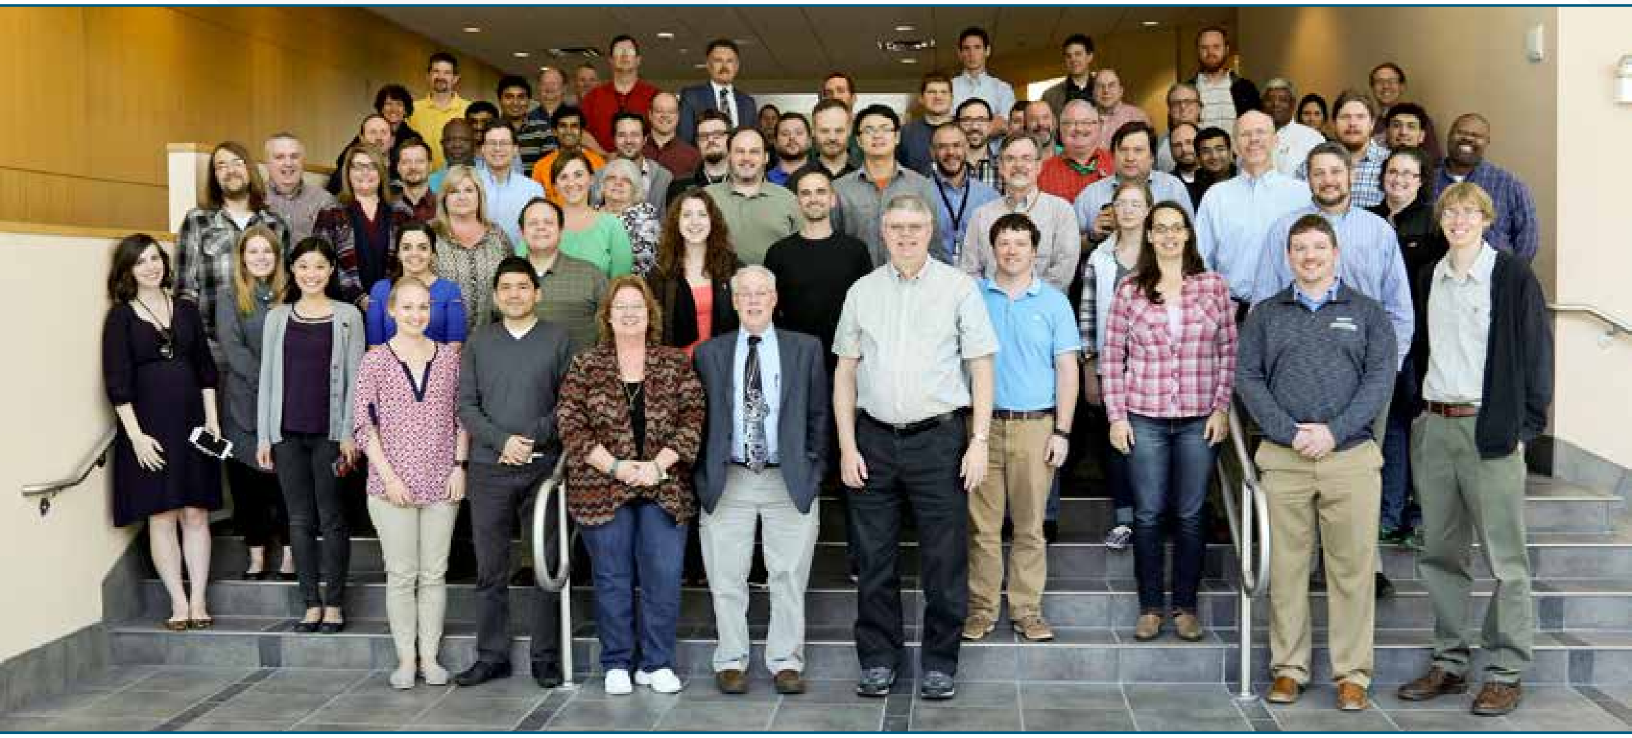

OLCF team members gathered to celebrate Titan Week, an event recognizing the 4th anniversary of Titan's operation. Credit: Jason Richards/ORNL

Since its 2012 debut, Titan has served as the nation's most powerful supercomputer for open science, enabling groundbreaking milestones in science and engineering. On Oct. 29, 2016, the Cray XK7 system reached a milestone of its own: 4 years of operation.

To recognize this occasion and the community of committed groups and individuals who have contributed to Titan's success, the OLCF celebrated "Titan Week" throughout the week of Oct. 24, 2016. The social media–driven event included interesting facts, science highlights, photos, and videos related to the operation, scientific productivity, and global impact of the machine.

Friends of Titan, including ORNL staff, vendor partners, national labs, and the public, participated by posting messages, pictures, and short videos on Twitter and Facebook wishing Titan a "happy birthday." Posts were accompanied with the hashtag #TitanWeek. More than 100 organizations and individuals contributed messages, including Ernest Moniz, who was Secretary of Energy at that time. To cap off the week, the OLCF staged a birthday party, complete with a Titan-themed cake.

In a Titan Week video tribute, Cray's chief strategy officer Barry Bolding commemorated the machine, which claimed No. 1 on the TOP500 list in November 2012. "When we began the project, this was one of the systems where we were going to bring together accelerator technology along with our scalable, productive supercomputers to build one of the most powerful systems in the world," he said. "Since that time, scientists at Oak Ridge National Laboratory and across the world have used this system to solve some of the most challenging problems we face."

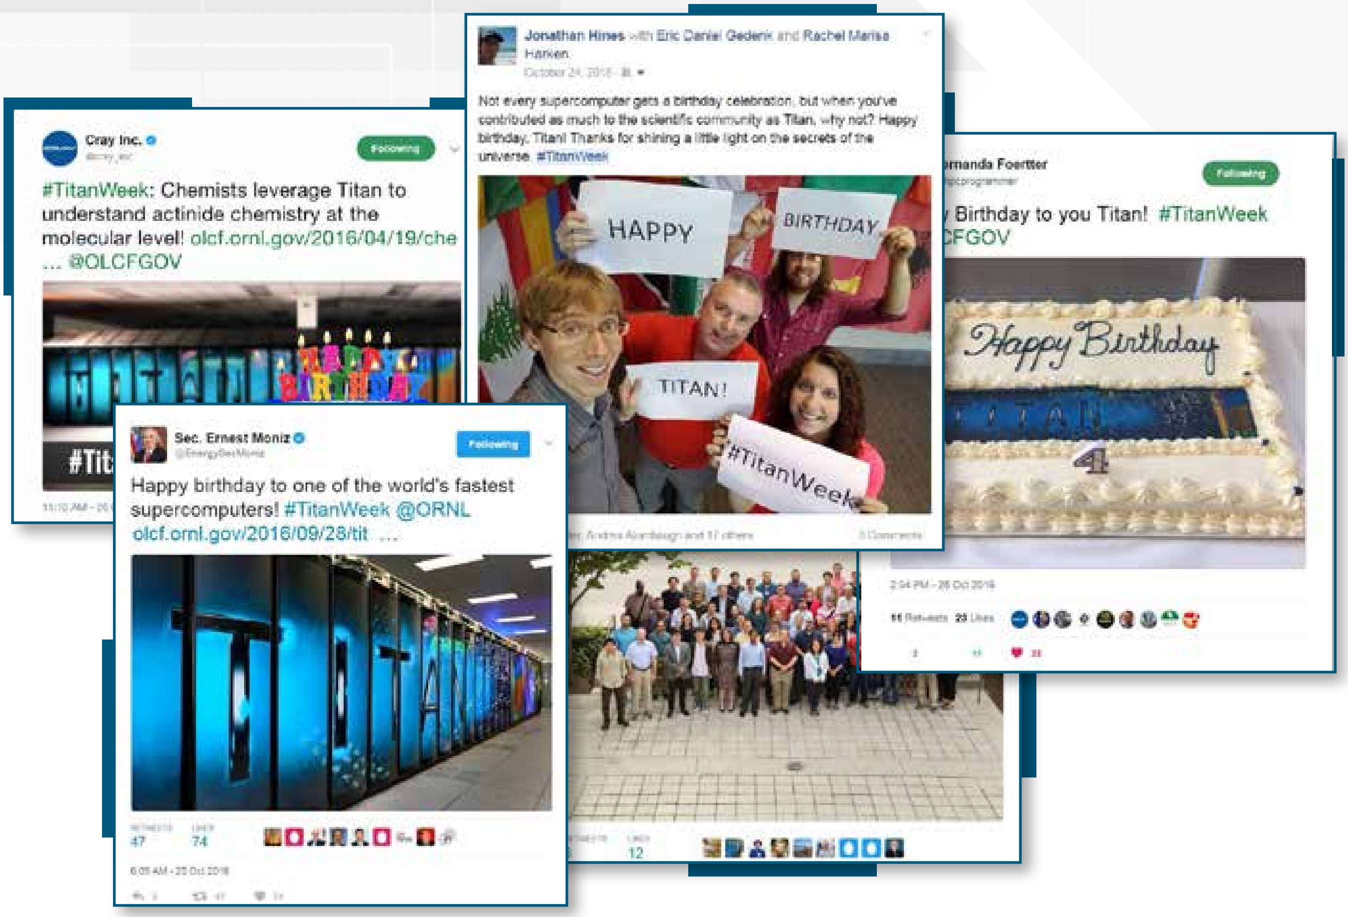

About Titan

Titan is a hybrid-architecture Cray XK7 system that combines AMD Opteron CPUs and NVIDIA Kepler GPUs to deliver world-class scientific computing capability. The system has demonstrated 17.59 petaflops, or 17.59 quadrillion calculations per second, and is theoretically capable of 27 petaflops.

Occupying about 5,300 square feet, Titan features 18,688 compute nodes, a system memory of 710 terabytes, and Cray's high-performance Gemini network. Its 299,008 CPU cores guide simulations, whereas the accompanying GPUs are capable of handling hundreds of calculations simultaneously. Titan is enabling researchers across the scientific arena to acquire unparalleled accuracy in their simulations and achieve research breakthroughs more rapidly than previously possible.

Titan delivered 4,323,608,405 core hours to OLCF users in 2016.

Selected Publications

For a complete list of all OLCF publications, please go to: https://www.olcf.ornl.gov/leadership-science/publications/

1. Eren, B., et al. (2016), Activation of Cu(111) Surface by Decomposition into Nanoclusters Driven by CO Adsorption, Science, Volume: 351, Issue: 6272. DOI: 10.1126/science.aad8868.

2. Perilla, J. R., et al. (2016), All-Atom Molecular Dynamics of Virus Capsids as Drug Targets, The Journal of Physical Chemistry Letters, Volume: 7, Issue: 10. DOI: 10.1021/acs.jpclett.6b00517.

3. LeVine, M. V., et al. (2016), Allosteric Mechanisms of Molecular Machines at the Membrane: Transport by Sodium-Coupled Symporters, Chemical Reviews, Volume: 116, Issue: 11. DOI: 10.1021/acs.chemrev.5b00627.

4. Zhao, Y., et al. (2016), Augmenting Drug–Carrier Compatibility Improves Tumour Nanotherapy Efficacy, Nature Communications, Volume: 7. DOI: 10.1038/ncomms11221.

5. Kalinin, S. V., et al. (2016), Big, Deep, and Smart Data in Scanning Probe Microscopy, ACS Nano, Volume: 10, Issue: 10. DOI: 10.1021/acsnano.6b04212.

6. Singharoy, A., et al. (2016), Binding Site Recognition and Docking Dynamics of a Single Electron Transport Protein: Cytochromec2, Journal of the American Chemical Society, Volume: 138, Issue: 37. DOI: 10.1021/jacs.6b01193.

7. Zuluaga, S., et al. (2016), Chemistry in Confined Spaces: Reactivity of the Zn-MOF-74 Channels, Journal of Materials Chemistry A, Volume: 4, Issue: 34. DOI: 10.1039/c6ta04388g.

8. Zuluaga, S., et al. (2016), Cluster Assisted Water Dissociation Mechanism in MOF-74 and Controlling It Using Helium, Journal of Materials Chemistry A, Volume: 4, Issue: 29. DOI: 10.1039/c6ta02501c.

9. Smith, M. D., et al. (2016), Cosolvent Pretreatment in Cellulosic Biofuel Production: Effect of Tetrahydrofuran-Water on Lignin Structure and Dynamics, Green Chemistry, Volume: 18, Issue: 5. DOI: 10.1039/c5gc01952d.

10. Yue, K., et al. (2016), Geometry Induced Sequence of Nanoscale Frank–Kasper and Quasicrystal Mesophases in Giant Surfactants, Proceedings of the National Academy of Sciences, Volume: 113, Issue: 50. DOI: 10.1073/pnas.1609422113.

11. Deb, N., et al. (2016), Harnessing Structure–Property Relationships for Poly(alkyl thiophene)-Fullerene Derivative Thin Films to Optimize Performance in Photovoltaic Devices, Advanced Functional Materials, Volume: 26, Issue: 12. DOI: 10.1002/adfm.201502653.

12. Mao, J., et al. (2016), Human-Induced Greening of the Northern Extratropical Land Surface, Nature Climate Change, Volume: 6, Issue: 10. DOI: 10.1038/nclimate3056.

13. Hu, A., et al. (2015), Impact of Solar Panels on Global Climate, Nature Climate Change, Volume: 6, Issue: 3. DOI: 10.1038/nclimate2843.

14. Meehl, G. A.; Hu, A.; & Teng, H. (2016), Initialized Decadal Prediction for Transition to Positive Phase of the Interdecadal Pacific Oscillation, Nature Communications, Volume: 7. DOI: 10.1038/ncomms11718.

15. Mostofian, B., et al. (2016), Local Phase Separation of Co-Solvents Enhances Pretreatment of Biomass for Bioenergy Applications, Journal of the American Chemical Society, Volume: 138, Issue: 34. DOI: 10.1021/jacs.6b03285.

16. Wieczorek, M., et al. (2016), MHC Class II Complexes Sample Intermediate States Along the Peptide Exchange Pathway, Nature Communications, Volume: 7. DOI: 10.1038/ncomms13224.

17. Raja, S. N., et al. (2016), Mechanisms of Local Stress Sensing in Multifunctional Polymer Films Using Fluorescent Tetrapod Nanocrystals, Nano Letters, Volume: 16, Issue: 8. DOI: 10.1021/acs.nanolett.6b01907.

18. Schunck, N. & Robledo, L. M. (2016), Microscopic Theory of Nuclear Fission: A Review, Reports on Progress in Physics, Volume: 79, Issue: 11. DOI: 10.1088/0034-4885/79/11/116301.

19. Curcio, D., et al. (2016), Molecular Lifting, Twisting, and Curling During Metal-Assisted Polycyclic Hydrocarbon Dehydrogenation, Journal of the American Chemical Society, Volume: 138, Issue: 10. DOI: 10.1021/jacs.5b12504.

20. Gupta, P., et al. (2016), Nascent Peptide Assists the Ribosome in Recognizing Chemically Distinct Small Molecules, Nature Chemical Biology, Volume: 12, Issue: 3. DOI: 10.1038/nchembio.1998.

21. Hagen, G., et al. (2015), Neutron and Weak-Charge Distributions of the 48Ca Nucleus, Nature Physics, Volume: 12, Issue: 2. DOI: 10.1038/nphys3529.

22. Karakalos, S., et al. (2016), Noncovalent Bonding Controls Selectivity in Heterogeneous Catalysis: Coupling Reactions on Gold, Journal of the American Chemical Society, Volume: 138, Issue: 46. DOI: 10.1021/jacs.6b09450.

23. Kattel, S., et al. (2016), Optimizing Binding Energies of Key Intermediates for CO2 Hydrogenation to Methanol over Oxide-Supported Copper, Journal of the American Chemical Society, Volume: 138, Issue: 38. DOI: 10.1021/jacs.6b05791.

24. Sener, M., et al. (2016), Overall Energy Conversion Efficiency of a Photosynthetic Vesicle, ELife, Volume: 5. DOI: 10.7554/elife.09541.

25. Maier, T. A., et al. (2016), Pairing in a Dry Fermi Sea, Nature Communications, Volume: 7. DOI: 10.1038/ncomms11875.

26. Carrillo, J.-M. Y., et al. (2016), Petascale Simulations of the Morphology and the Molecular Interface of Bulk Heterojunctions, ACS Nano, Volume: 10, Issue: 7. DOI: 10.1021/acsnano.6b03009.

27. Meldgin, C., et al. (2016), Probing the Bose Glass–Superfluid Transition Using Quantum Quenches of Disorder, Nature Physics, Volume: 12, Issue: 7. DOI: 10.1038/nphys3695.

28. Adelstein, N. & Wood, B. C. (2016), Role of Dynamically Frustrated Bond Disorder in a Li+ Superionic Solid Electrolyte, Chemistry of Materials, Volume: 28, Issue: 20. DOI: 10.1021/acs.chemmater.6b00790.

29. Shepherd, M. R.; Dudek, J. J.; & Mitchell, R. E. (2016), Searching for the Rules that Govern Hadron Construction, Nature, Volume: 534, Issue: 7608. DOI: 10.1038/nature18011.

30. Hiebel, F., et al. (2016), Self-Assembly of Acetate Adsorbates Drives Atomic Rearrangement on the Au(110) Surface, Nature Communications, Volume: 7. DOI: 10.1038/ncomms13139.

31. Schmitz, L., et al. (2016), Suppressed Ion-Scale Turbulence in a Hot High-Beta Plasma, Nature Communications, Volume: 7. DOI: 10.1038/ncomms13860.

32. Huang, W.-K., et al. (2016), Tailoring Kinetics on a Topological Insulator Surface by Defect-Induced Strain: Pb Mobility on Bi2Te3, Nano Letters, Volume: 16, Issue: 7. DOI: 10.1021/acs.nanolett.6b01604.

33. Hu, X., et al. (2015), The Dynamics of Single Protein Molecules is Non-Equilibrium and Self-Similar over Thirteen Decades in Time, Nature Physics, Volume: 12, Issue: 2. DOI: 10.1038/nphys3553.

34. Pedro, J. B., et al. (2015), The Spatial Extent and Dynamics of the Antarctic Cold Reversal, Nature Geoscience, Volume: 9, Issue: 1. DOI: 10.1038/ngeo2580.

35. Mahalik, J. P., et al. (2016), Theoretical Study of the Initial Stages of Self-Assembly of a Carboxysome's Facet, ACS Nano, Volume: 10, Issue: 6. DOI: 10.1021/acsnano.5b07805.

36. Han, Y., et al. (2016), Thermoreversible Morphology and Conductivity of a Conjugated Polymer Network Embedded in Block Copolymer Self-Assemblies, Small, Volume: 12, Issue: 35. DOI: 10.1002/smll.201601342.

37. Tan, K., et al. (2016), Trapping Gases in Metal–Organic Frameworks with a Selective Surface Molecular Barrier Layer, Nature Communications, Volume: 7. DOI: 10.1038/ncomms13871.

38. Muthukrishnan, S.; Albash, T.; & Lidar, D. A. (2016), Tunneling and Speedup in Quantum Optimization for Permutation-Symmetric Problems, Physical Review X, Volume: 6, Issue: 3. DOI: 10.1103/physrevx.6.031010.

39. Meehl, G. A.; Tebaldi, C.; & Adams-Smith, D. (2016), US Daily Temperature Records Past, Present, and Future, Proceedings of the National Academy of Sciences of the United States of America, Volume: 113, Issue: 49. DOI: 10.1073/pnas.1606117113.

40. Carlsson, B. D., et al. (2016), Uncertainty Analysis and Order-by-Order Optimization of Chiral Nuclear Interactions, Physical Review X, Volume: 6, Issue: 1. DOI: 10.1103/physrevx.6.011019.

41. Zuluaga, S., et al. (2016), Understanding and Controlling Water Stability of MOF-74, Journal of Materials Chemistry A, Volume: 4, Issue: 14. DOI: 10.1039/c5ta10416e.

42. Ruiz, R. F. G., et al. (2016), Unexpectedly Large Charge Radii of Neutron-Rich Calcium Isotopes, Nature Physics, Volume: 12, Issue: 6. DOI: 10.1038/nphys3645.

Acronyms

ALCC ASCR Leadership Computing Challenge

ATP adenosine triphosphate

CAAR Center for Accelerated Application Readiness

CADES Compute and Data Environment for Science

Caltech California Institute of Technology

CAM crassulacean acid metabolism

CCS Center for Computational Sciences

CFD computational fluid dynamics

CPU central processing unit

DD Director's Discretionary

DOE US Department of Energy

DTN data transfer node

ECP Exascale Computing Project

EVEREST Exploratory Visualization Environment for Research in Science and Technology

FePt iron–platinum

GPU graphics processing unit

HPC high-performance computing

HPCO High-Performance Computing Operations

HPSS High-Performance Storage System

I/O input/output

INCITE Innovative and Novel Computational Impact on Theory and Experiment

IRIS Incorporated Research Institutions for Seismology

LSMS Linear Scaling Multiple Scattering

NCCS National Center for Computational Sciences

OLCF Oak Ridge Leadership Computing Facility

ORNL Oak Ridge National Laboratory

OUG OLCF User Group

PI principal investigator

SciComp Scientific Computing Group

SNS Spallation Neutron Source

TechInt Technology Integration Group

UCLA University of California, Los Angeles

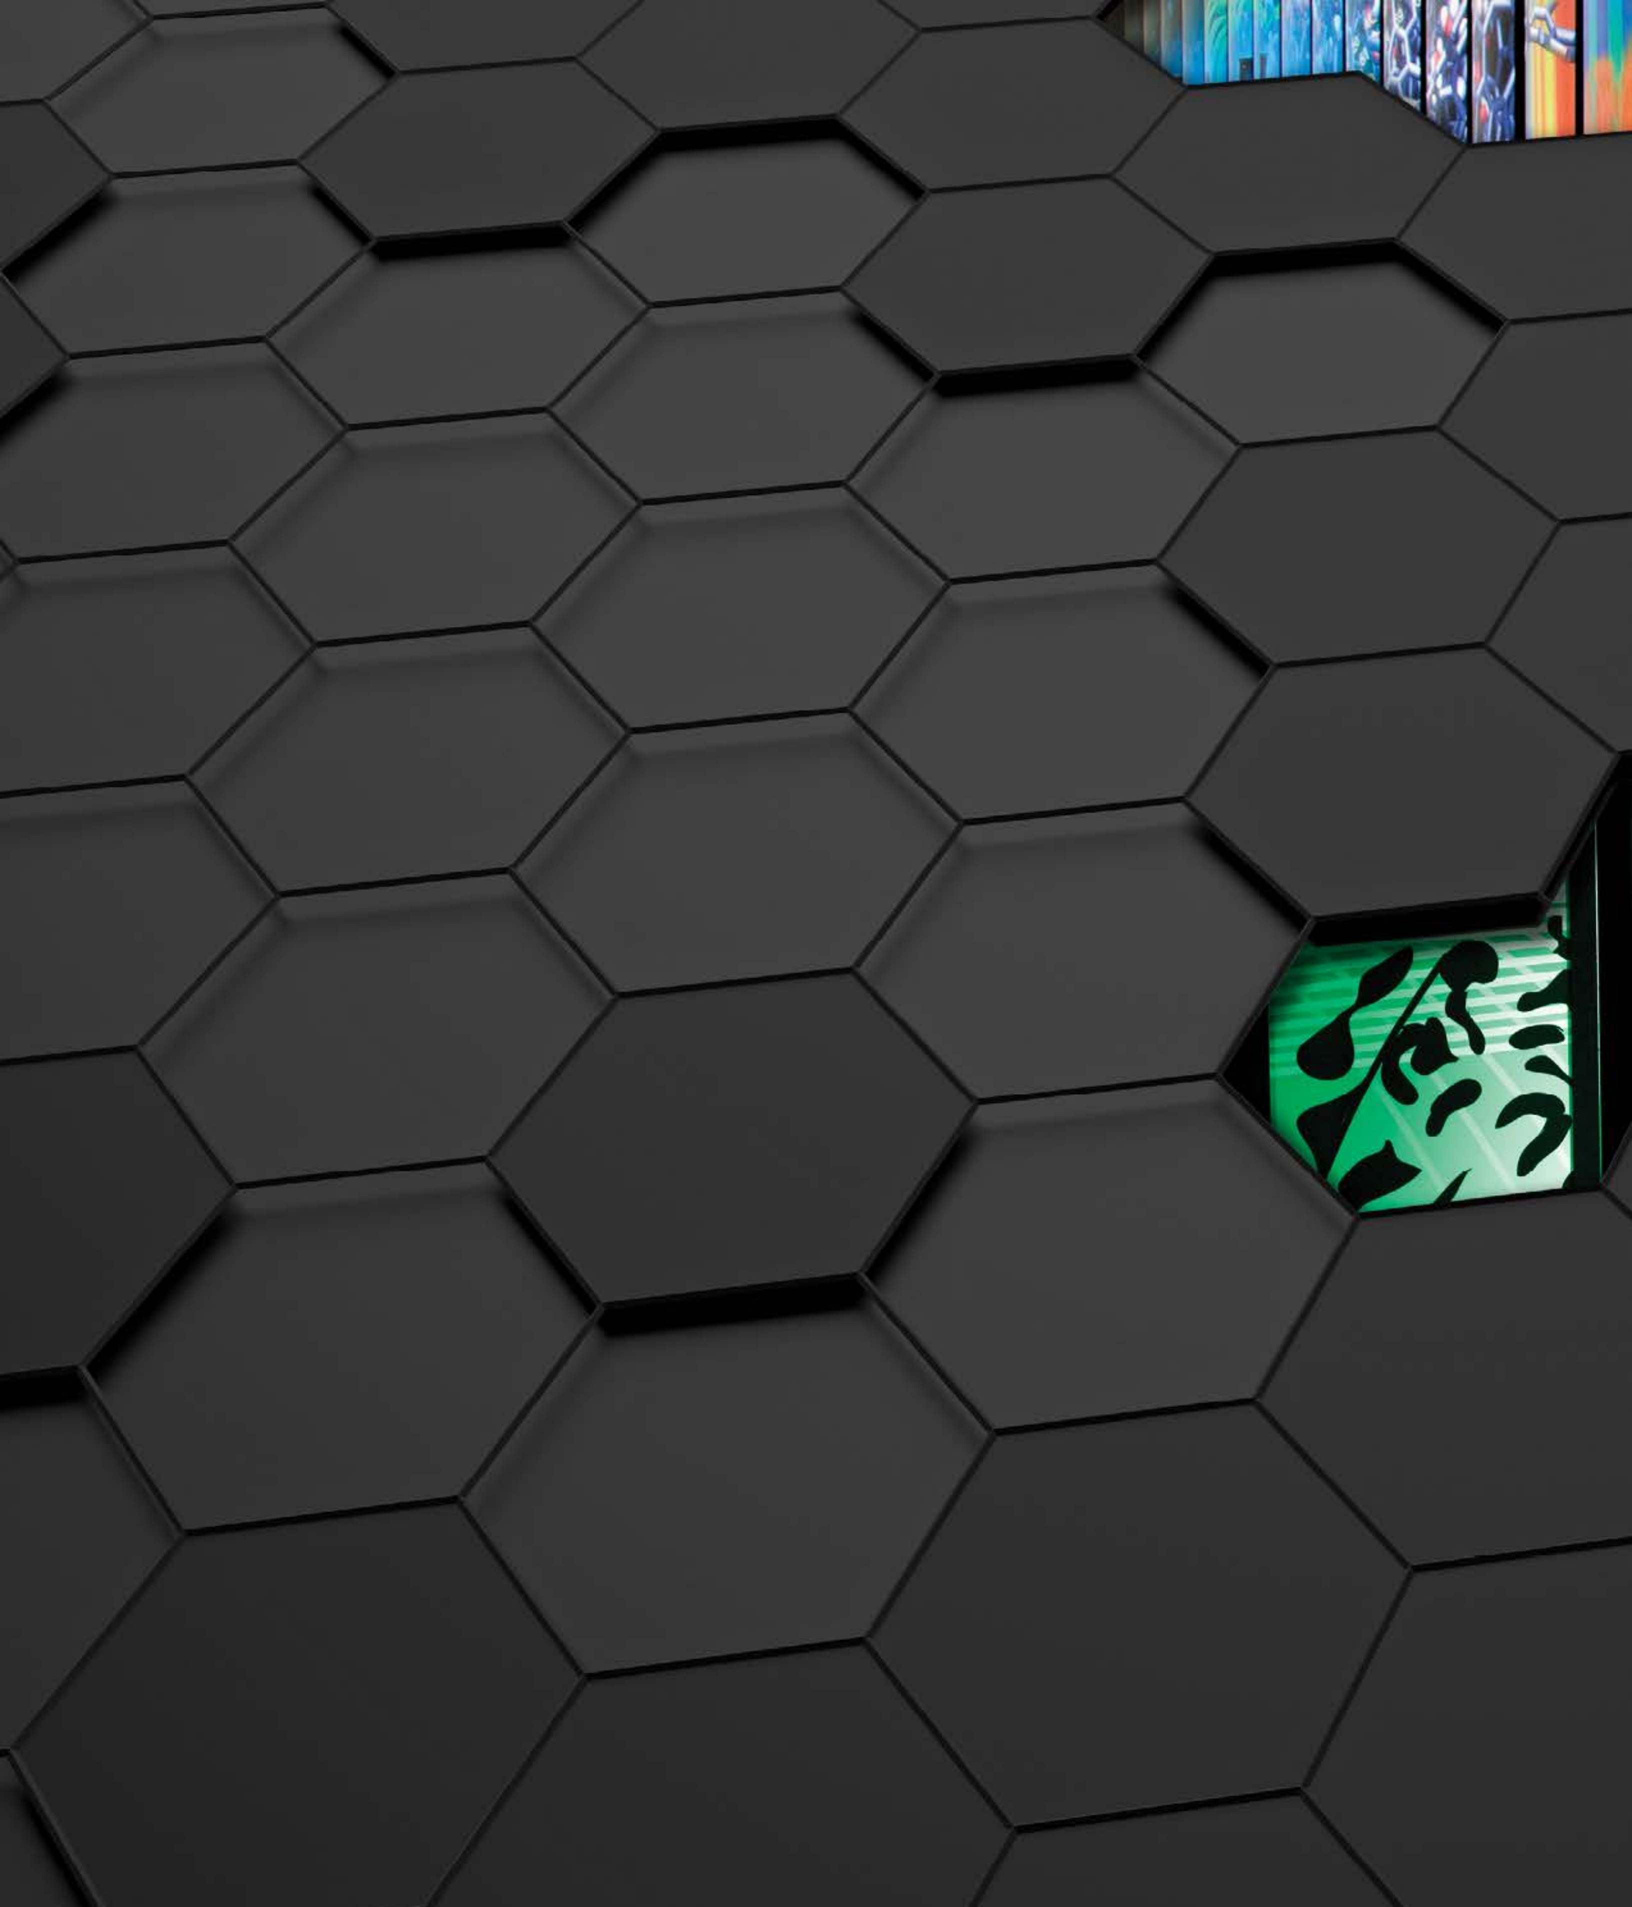

Managed by UT-Battelle for the US Department of Energy
ORNL 2017-G00112/jbs