OAK RIDGE LEADERSHIP COMPUTING FACILITY

2015 2016 ANNUAL REPORT

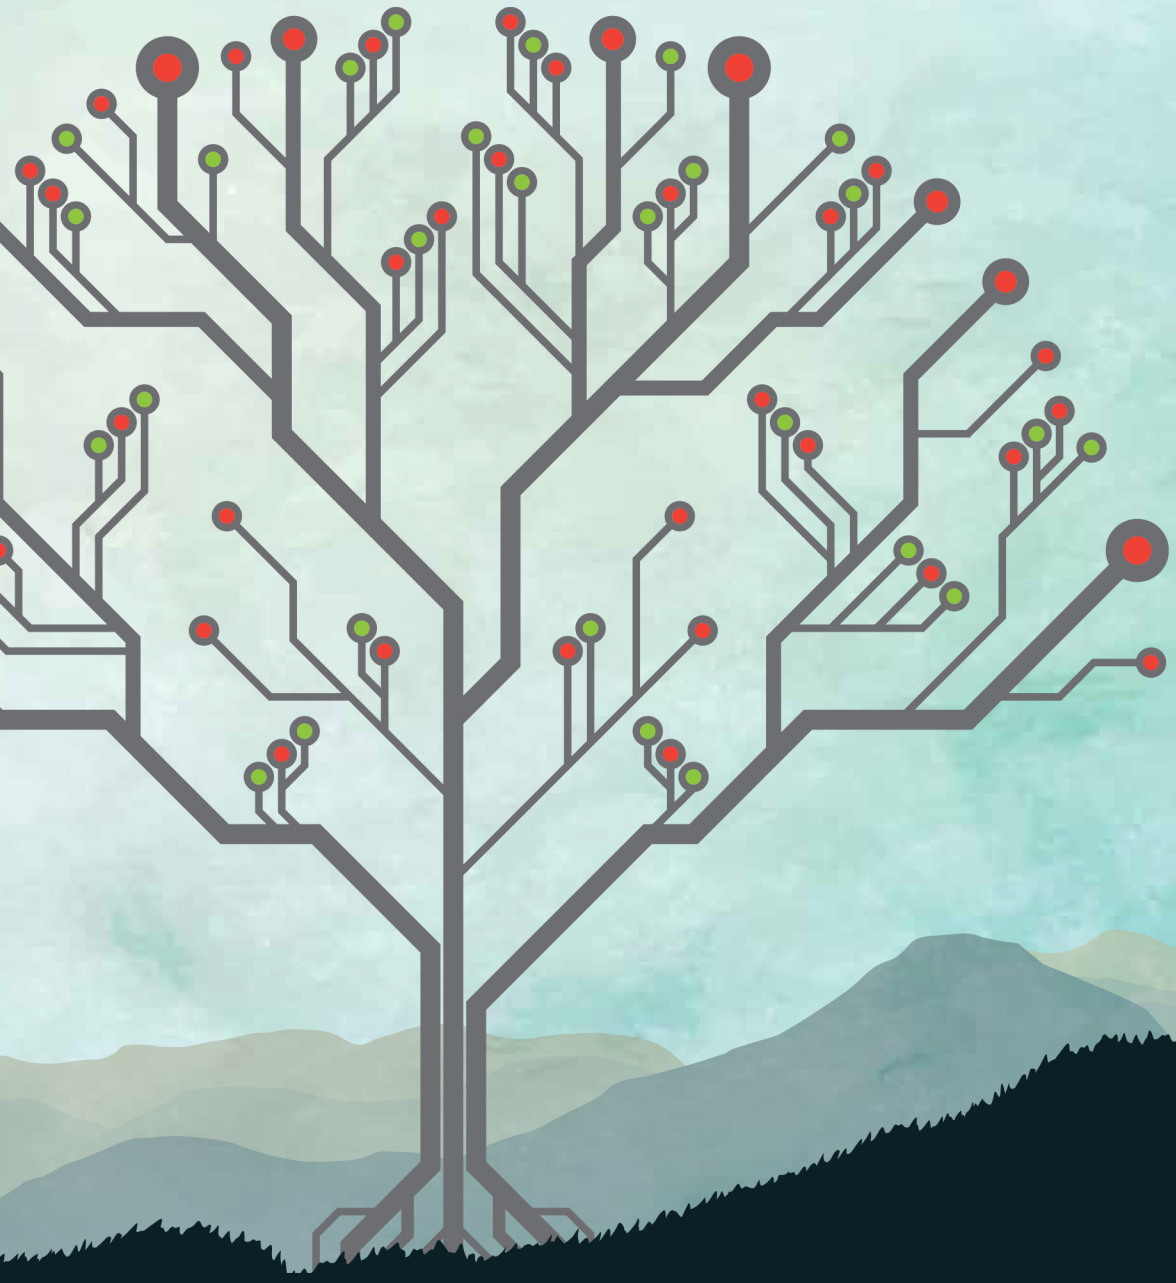

OAK RIDGE National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

Oak Ridge Leadership Computing Facility

Annual Report 2015–2016

Outreach project manager – Jeff Gary

Writers – Eric Gedenk, Jonathan Hines, Miki Nolin

Designer – Jason Smith

Editor – Wendy Hames

Photography – Jason Richards, Genevieve Martin, Carlos Jones, and Renee Manning

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

Here and Now (Director's Letter) 4

Science

Flowing Toward Red Blood Cell Breakthroughs 6
Researchers Mine Information from Next-Generation 8
Subsurface Flow Simulations 8
Fusion Researchers Use Titan to Burst Helium Bubbles 10
Sizing Up the Atomic Nucleus 12
Titan Takes on the Big One 14
Breaking Down Biofuel's Biggest Barrier 16

Featured Topics

Service, Customer Satisfaction Highlight Strong Year for OLCF 18
Stepping up to Summit 22
Early Science (CAAR) 24

People and Programs

OLCF Inaugurates Advanced Data and Workflow Group 32
Groups 34
OLCF Executive Board and User Group 36
INCITE, ALCC, and DD—the OLCF's Allocation Programs 37

Systems and Support

Center Overview Diagram 38
HPC User Support 40
Education, Outreach, and Training 42

High Impact Publications List 44
Acronyms 46

Here and Now

Jim Hack, Director NCCS

The arrow of science and technology may always point toward a distant horizon, but at the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, daily dedication to scientific discovery is illuminating solutions to global challenges in the here and now.

Much of that dedication in 2015 is captured in this report, including stories and infographics illustrating how our innovative staff members continue to push the boundaries of high-performance computing to better support our world-class community of scientific users. The latest fruits of these efforts address a wide array of basic and applied research questions—from the size of an atom's nucleus (Hagen, page 12), to the potential risks posed by earthquakes (Jordan, page 14), to effective methods for drug delivery (Karniadakis, page 6).

All of these problems—and many more—were addressed in Titan's second full year of production, in which the OLCF delivered 100 percent of its pledged allocations for our three major user programs—INCITE, ALCC, and Director's Discretionary. Furthermore, Titan's scheduled availability continued to hold high, at 99.4 percent, with an overall availability of 97 percent, and usage of Titan's available hours increased slightly to 90 percent.

When Titan debuted in late 2012, its GPU-accelerated architecture seemed novel—even risky. Perhaps the greatest impact of this machine on the HPC landscape is its ongoing validation of leadership-scale hybrid computing's value, a legacy that will continue and evolve with Summit, our next leadership-class machine.

On Titan, GPUs provide 90 percent of the computational power, a fact that users have exploited to find new and clever ways to produce more science in less time. In one such example through our INCITE program, a team from Virginia Tech leveraged Titan's NVIDIA GPUs to run the challenging calculations for their multiphase flow simulations of subsurface contaminates and other solids, liquids, or gases that mix below the Earth's surface (McClure, page 8). The approach freed Titan's CPUs to analyze the resulting data in real time and send relevant data sets to the OLCF's Rhea analysis cluster for visualization. The work has produced adaptable simulation tools that allow researchers to see geologic phenomena in motion, impacting research in carbon sequestration, oil recovery, and contaminant transport.

Other projects have used Titan to advance critical energy missions supported by the US Department of Energy. A Los Alamos National Laboratory-based team focused its ALCC project on scaling up its code to better understand how helium particles interact with the tungsten divertor wall of a nuclear fusion reactor (Sandoval, 10). The work is serving as the foundation for a larger DOE effort to comprehensively

simulate the extreme conditions of nuclear fusion and inform the development of the ITER experimental fusion reaction in France.

Another ALCC project led by a team at ORNL used modeling and simulation to complement DOE's development of next-generation biofuels (Smith, 16). By simulating complex, multicomponent biomass systems composed of thousands to millions of atoms, the ORNL team provided useful molecular-level context that researchers can apply to create renewable, cost-effective cellulosic ethanol.

This year, the OLCF undertook several initiatives to improve our world-leading resources and to continue to meet the evolving demands of the scientific computing community. We doubled the amount of memory per node on Rhea and added GPU nodes so users can analyze and visualize their data more efficiently. We also began a 2-year upgrade of our High-Performance Storage System to make it easier for users to store, move, and access their data when they need it most. The investment has started to pay dividends as data migration from disk cache to long-term tape storage has seen a four- to fivefold increase in speed.

To support users more, the OLCF formed a new team called the Advanced Data and Workflow Group. This group collaborates with users to help better understand the results of computations through the design, construction, and execution of leadership-scale data science solutions—including analytics, visualization, and other data services.

Participation in our user programs continues to grow. Building on the inaugural success of 2014, the OLCF partnered with the National Center for Supercomputing Applications and the Swiss National Supercomputing Centre to organize four GPU hackathons. The ORNL-hosted hackathon, held in October 2015, saw a 33 percent increase in attendance from its first year.

The OLCF is firmly committed to maximizing Titan's scientific productivity, but we're also gearing up to take a major step toward exascale computing. More than a year after announcing Summit as part of the CORAL procurement, preparations are well underway to ensure this next-generation supercomputer is ready for science upon arrival in 2018 (Summit Update, page 22).

Enhancements to the Computational Science Building at ORNL continue to progress, with crews working to add 10 megawatts of power and 20 megawatts of cooling capacity to the OLCF's infrastructure in anticipation of Summit. In April 2015, the OLCF accepted 13 application teams into the Center for Accelerated Application Readiness. In conjunction with the OLCF and vendor partners, these teams are readying their codes for Summit's novel large-node, hybrid architecture, which is expected to be more than five times faster than Titan is for applications.

A key emphasis of CAAR code preparation efforts is application portability across DOE architectures that are being deployed over the next few years. In September 2015, the OLCF took part in the DOE High Performance and Computing Operations Review, which brought together members across the DOE lab system to collaborate on this critical task. Supercomputing has never been more vital to America's global competiveness. That assertion was echoed in July 2015 when President Obama issued an executive order creating the National Strategic Computing Initiative. The NCSI is bringing together scientists and government agencies, including representatives from ORNL and our center, to create a formal plan for the first exascale ecosystem, a system about 30 times more powerful than today's supercomputers.

As we enter the dawn of the exascale era, it's satisfying to know that the OLCF is well-positioned to contribute to this historic effort. Moving forward, the OLCF will continue to connect users with the best computational and data resources in the world.

Flowing Toward Red Blood Cell Breakthroughs

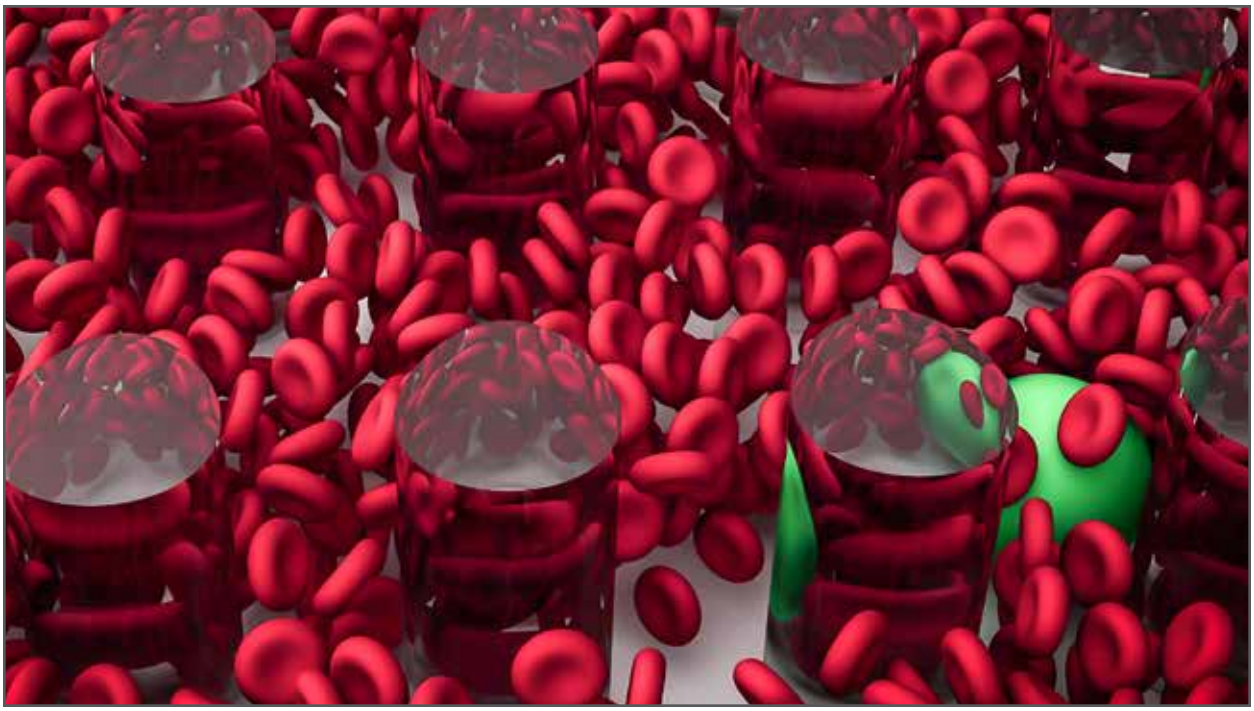

Red blood cells (red) and circulating tumor cells (green) traveling through a microfluidic cell sorting device as simulated by uDeviceX. Image credit: Yu-Hang Tang, Brown University

A team of researchers from Brown University, ETH Zurich, the Universita da Svizzera Italiana, and Consiglio Nazionale delle Ricerche is using America's largest, most powerful supercomputer to help understand and fight diseases affecting some of the body's smallest building blocks.

The team, led by Brown's George Karniadakis, is using Titan to simulate hundreds of millions of red blood cells in an attempt to develop better drug delivery methods and predictors to fight against tumor formation and sickle cell anemia.

Karniadakis' group approaches these simulations from a unique place—despite the research focus on medical issues, the team actually is based in Brown's Applied Mathematics department. During the first year of a 3-year Innovative and Novel Computational Impact on Theory and Experiment (INCITE) allocation at both the Oak Ridge and Argonne Leadership Computing Facilities, the team has been performing a suite of simulations related to different diseases and drug delivery methods to better predict, diagnose, and treat several mysterious hematological, or blood-based, diseases.

At the OLCF, the team primarily has focused its disease research on sickle cell anemia (SCA) and tumor cells, as well as on developing better drug delivery methods. The team uses dissipative particle dynamics in its simulations to study blood flow as a collection of individual particles rather than one fluid object. To model each individual particle's behavior accurately for any meaningful length of time, the team needed leadership-class supercomputing power.

In addition to its research on SCA, the team is also leveraging Titan to understand how diseases could be treated. Thus far, the team simulated blood and cancer cell separation using microfluidic devices. Microfluidic devices can manipulate extremely small amounts of fluids—typically microliters (1 millionth of a liter) or smaller.

Brown University post-doctoral researcher and project collaborator Yu-Hang Tang focused on how blood and cancerous tumor cells might be separated by microfluidic devices, and his simulations are up to 3 times larger—in terms of the number of simulated cells and computational elements—than the current state of the art within the field.

Such microfluidic devices would allow doctors to take a very small sample of blood and quickly identify whether someone had a malignant tumor. This "lab on a chip" could help doctors test for illness in the least invasive way possible.

Tang and his collaborators exploited Titan's GPU accelerators and developed uDeviceX, a GPU-driven particle solver—an important part of the team's code that helps plot individual particles in the simulation. Tang's new solver showed a forty-five-fold decrease in time to solution compared with competing state-of-the-art methods.

Moreover, Tang's extensive work with GPUs has led to the team's newest computational tool—the multiscale universal interface (MUI). The team's research interests require a variety of different codes, with certain computing architectures working better with certain codes and different architectures benefitting other codes. It also gives the team freedom to focus specific solvers for different scales simultaneously.

MUI allows the team to quickly connect its contrasting codes into one larger code, significantly cutting down the team's computational cost for running its simulations by focusing on different hardware configurations' strengths.

The team not only can target different solvers or other portions of a code toward certain parts of a supercomputer but also is able to offload different parts of a code on different supercomputers. These distinct pieces of code do not have to communicate during each time step but ultimately will share their results during the course of a simulation. OLCF staff member Wayne Joubert helped the team scale MUI to make efficient use of Titan's large node count.

Multiscale Simulations of Human Pathologies

George Karniadakis
Brown University and Universita da Svizzera Italiana
2015 INCITE project
25,000,000 hours

Researchers Mine Information from Next-Generation Subsurface Flow Simulations

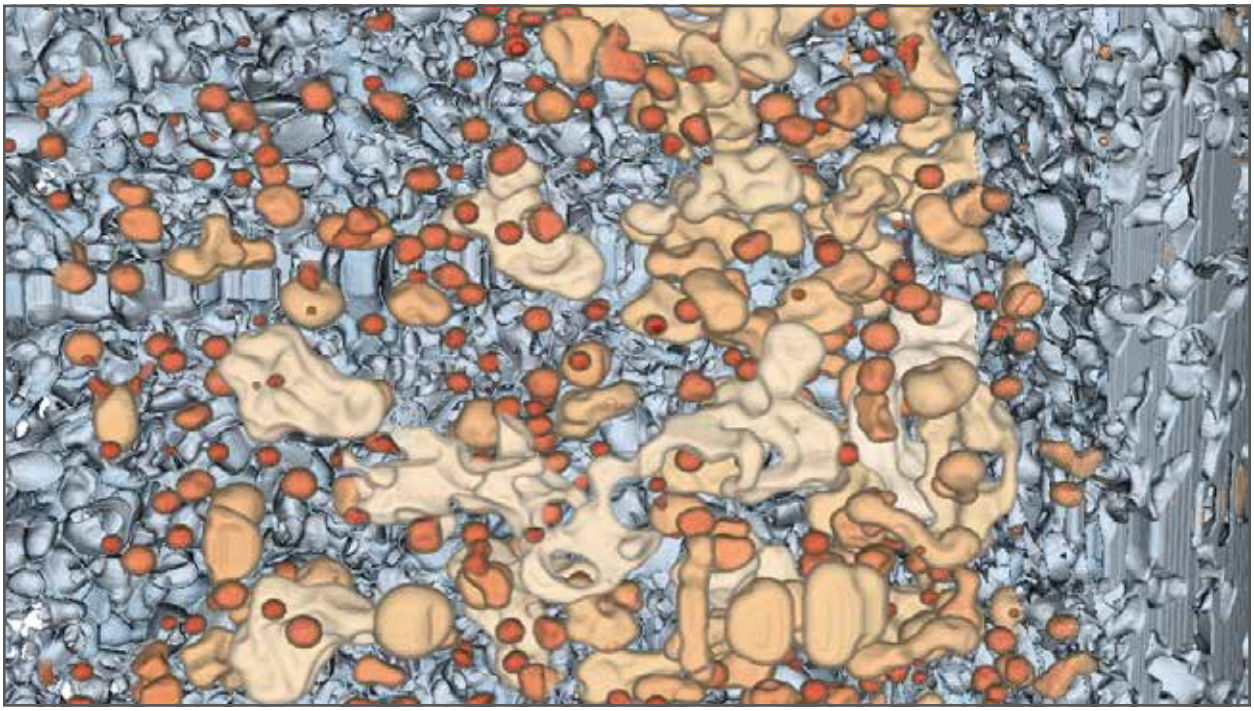

Simulations performed on Titan enable detailed tracking of two fluid phase systems during flow through porous media. In this case, the behavior of oil ganglia (yellow) can be simulated and understood as they are mobilized in experimentally imaged sandstone (the solid part of the sandstone is shown in blue). There is much societal benefit to understanding the conditions under which these oily phases can be removed from the subsurface rock or placed to reside permanently in rock structures. Image credit: James McClure

Ask a geoscientist how he or she feels about the phrase, "out of sight, out of mind," and you likely will meet some serious resistance. The scientist would likely explain that learning how different materials interact with one another below the Earth's surface is of the utmost importance.

Oil production, subsurface contamination, and carbon sequestration, among other research fields, share a common thread—they deal with multiphase flows, or situations where materials are flowing closely together in different phases (solids, liquids, or gases). They also pertain to when the flow is composed of materials that have a common phase with a different chemical makeup that prevents mixing (such as oil and water).

A research team led by Virginia Tech's James McClure is using resources at the OLCF to improve understanding of transport phenomena in multiphase systems. Through the INCITE program, the team has created detailed models for simulating two-phase flow systems and has produced a roadmap to simulate increasingly complex multiphase flows efficiently.

McClure noted that experimental advances such as microcomputed tomography (micro-CT) have allowed researchers to observe the structure of geologic materials with micron-scale resolution. Simulations are essential to connect the experimentally generated data sets to models that describe flow processes at length scales of kilometers, as they are needed to make predictions for a wide range of systems that affect society.

The team's Lattice-Boltzmann Methods for Porous Media with Internal Averaging code takes full advantage of Titan. Much of Titan's speed comes from NVIDIA Kepler GPUs, and the Lattice-Boltzmann calculations can take full advantage of those GPUs.

In fact, the ability to use GPUs has transformed the team's simulations. McClure said that the team does its calculations entirely on Titan's GPUs and uses the CPUs to do in situ (or real-time) data analyses.

OLCF computational scientist Mark Berrill helped optimize codes so that when running its largest simulations, the team can use Titan's CPUs to analyze the data, then send the relevant data sets to the Rhea data-analysis cluster to visualize the results from simulations as they are generated.

This novel approach has allowed the team to churn through hundreds of terabytes of data without any need for post-processing, providing continuous information about the simulation state.

Working with the OLCF, McClure has developed simulation tools that allow researchers to take the information from a 3-D micro-CT image and put it into motion.

With this approach, the team can study interfacial dynamics and other aspects of the system behavior quantitatively within a multiscale framework. The team can predict rocks' properties—such as relative permeability—from first principles to measure the resistance to flow that controls the movement of fluids at the reservoir scale. These results are vital for situations where the mobility of trapped, non-wetting phases—in this case, oil or liquid carbon dioxide—is of essential importance, such as carbon sequestration, oil recovery, and contaminant transport.

Advancing Models for Multiphase Flow and Transport in Porous Medium Systems

James McClure
Virginia Tech
2015 INCITE project
60,000,000 hours

Fusion Researchers Use Titan to Burst Helium Bubbles

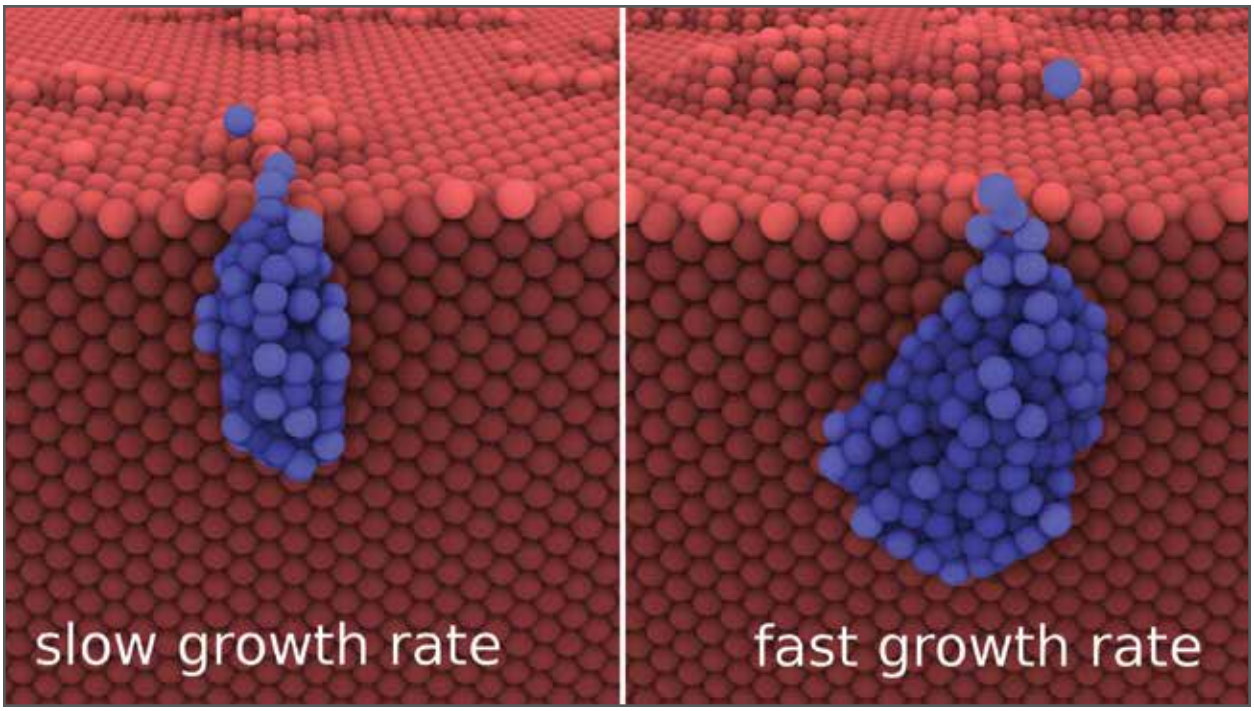

Snapshots of a helium bubble just before bursting when grown at slow (1 He/ 10 ns) versus fast (1 He/ps) rates. At slow rates, the bubble experiences the influence of the surface more readily and grows more asymmetrically toward the surface. At fast rates, the bubble grows more symmetrically and is larger when it bursts, creating more surface debris. The colors indicate helium atoms (blue) and tungsten atoms (red). Image credit: Luis Sandoval

Scientists look to the stars when it comes to developing clean, virtually limitless energy. Though humanity understands how stars power themselves—nuclei of hydrogen and its isotopes fuse together in extreme conditions, releasing bursts of energy in the process—they have been unable to replicate this massive fusion process on Earth in a practical way for power production.

As part of a Scientific Discovery through Advanced Computing (SciDAC) project, a partnership between DOE's Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC) and Fusion Energy Sciences programs, researchers are using Titan to get closer to producing sustainable fusion for electricity.

The project brings researchers from various organizations together to work on different aspects of the ITER experimental fusion reactor under construction in southeastern France. An international collaboration, ITER will be by far the largest fusion reactor ever built. Participating countries hope the reactor will serve as proof of concept for future fusion power plants.

The team's research focuses on the divertor, or the part of ITER that serves as a cache for heavier, more energetic particles from fusion byproducts and reactor material that can pollute the plasma, lower its temperature, and ultimately make it difficult to sustain the fusion reaction.

Researchers designing ITER plan to use tungsten—one of the toughest materials known—for the divertor. However, despite tungsten's relative strength, there is still concern about how the plasma could affect the divertor over time.

As helium particles bombard the tungsten wall, they begin to form clusters within the material. Once a helium atom is embedded in the wall, it attracts other helium particles. When enough helium is bunched together, it can "knock out" a tungsten atom from its normal position within the material, forming a nanoscale cavity, or hole, within the tungsten. This forms the nucleus of a helium bubble that can then grow very large, reducing the durability of the material. These bubbles also serve as traps for tritium, thus reducing the amount involved in the fusion reaction and introducing a radiological hazard.

A Los Alamos National Laboratory (LANL)-based team, led by Brian Wirth of the University of Tennessee, made up of Luis Sandoval, Danny Perez, Blas Uberuaga, and Arthur Voter is working to understand more fully how tungsten behaves in such harsh conditions. The group hopes that by better understanding the interactions between helium bubbles and tungsten, they can predict the evolution of the material and maybe even mitigate concerns over tungsten in the reactor.

By using ParRep, Sandoval and his group were able to simulate many microseconds of time rather than pico- or nanoseconds (10^{-12} and 10^{-9}, respectively). This speedup allows researchers to more fully understand how helium particles interact with the divertor wall. In particular, the LANL team found that the evolution of bubbles, when simulated over these realistic time scales, is qualitatively different from what is observed when traditional molecular dynamics is used. The team's findings were published in the March 11, 2015, issue of Physical Review Letters.

Ultimately, the LANL work will serve as the foundation for the larger SciDAC project. The larger SciDAC group is developing a suite of computer codes that will together create a more comprehensive fusion materials simulation.

Understanding Helium Plasma-Mediated Tungsten Surface Response to Better Predict Fusion Plasma Facing Component Performance in ITER

Brian Wirth
University of Tennessee, Knoxville
2015-2016 ALCC project
30,000,000 hours

Sizing Up the Atomic Nucleus

Conceptual art connects the atomic underpinnings of the neutron-rich calcium-48 nucleus with the Crab Nebula, which has a neutron star at its heart. Zeros and ones depict the computational power needed to explore objects that differ in size by 18 orders of magnitude. Image credit: Oak Ridge National Laboratory, U.S. Dept. of Energy; conceptual art by LeJean Hardin and Andy Sproles

On the list of essential elements for humans, calcium sits near the top. The element is the fifth most abundant in the human body, supporting bone structure, muscle movement, and blood function.

The significance of calcium, however, goes well beyond its role in supporting human life. By studying a rare isotope of calcium, calcium-48, using a first-principles computational approach, an international team of researchers led by ORNL's Gaute Haugen found answers to fundamental questions related to the size of the atomic nucleus and how the universe is put together.

Specifically, Hagen and his collaborators used Titan to gain insight into the size of calcium-48's atomic nucleus, knowledge that deepens scientists' understanding of objects as small as individual atoms and as large as neutron stars—a connection that bridges 18 orders of magnitude.

The team's calculations produced the first accurate predictive picture of the relatively large nucleus by tracking the complex nuclear interactions between protons and neutrons in pairs and groups of three, known as the two-body force and three-body force, respectively. The ab initio approach revealed that

the difference between the radii of neutron and proton distributions—called the "neutron skin"—is considerably smaller than previously thought.

"This is the first really reliable calculation of such a massive nucleus from first principles," Hagen said. "We reproduced basic observables for the first time by building this nucleus from scratch. We've answered a basic question—what is the size of the atomic nucleus?"

Additionally, the team made predictions for the physical characteristics—electric dipole polarizability and the weak form factor—that are current targets of precision measurements. The research helps close the gap between theory and experiment for light- to medium-mass nuclei and provides a fine-grained point of comparison to density functional theory, another method for calculating nuclei properties.

Calcium-48 contains a neutron-rich nucleus, requiring Titan's computing power to accurately account for 48 strongly interacting particles.

Whereas the distribution of electric charge inside the atomic nucleus is well known from experiments involving electron scattering, the distribution of the neutrons, which have no electric charge, is difficult to measure. In the nucleus of calcium-48, which has eight more neutrons than protons, the neutron distribution extends beyond the charge distribution and sets the actual size of the nucleus.

To validate the findings, Hagen's team also calculated radii, binding energies, and dipole polarizabilities for the isotopes helium-4, oxygen-16, and calcium-40.

"Many things had to come together—accurate nuclear forces, sophisticated computational algorithms, and a powerful tool such as Titan at ORNL—to achieve these results," Hagen said.

Nuclear Structure and Nuclear Reactions

Gaute Hagen
ORNL/ University of Tennessee
2015 INCITE project
25,000,000 hours

Titan Takes on the Big One

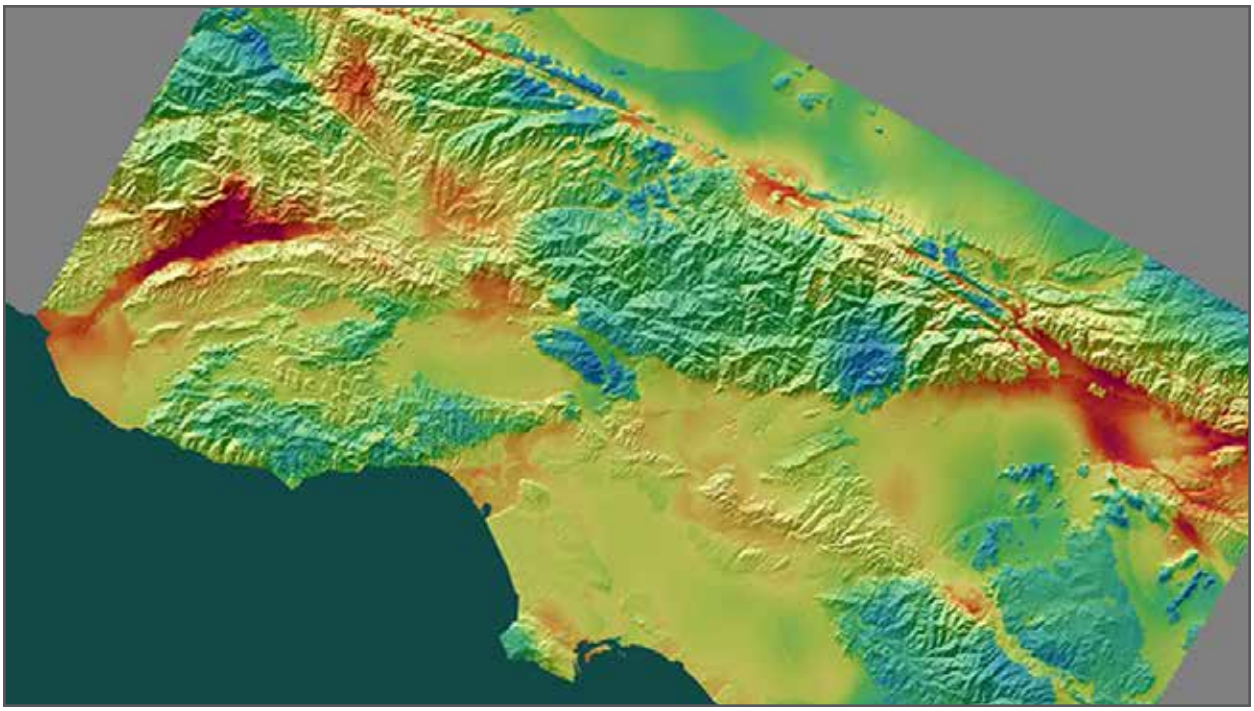

The CyberShake seismic hazard map shows the magnitude, or level of shaking, for the Los Angeles region, defined by the amount of change of a surface or structure in a 2-second period, with a 2% probability of increasing within the next 50 years. Image credit: Scott Callaghan, Kevin Milner, and Thomas H. Jordan (Southern California Earthquake Center)

Each year about 10,000 minor earthquakes emanate from the southern segment of the San Andreas Fault, which passes within 35 miles of Los Angeles. History tells us these tremors are just a preamble until the next "big one," an earthquake of a 7.5 magnitude or greater.

The last major earthquake along the fault hit San Francisco in 1906, registering a 7.8 magnitude. The quake took 700 lives and caused $400 million worth of damage, a figure equivalent to more than $10 billion today.

To help prepare emergency officials and structural engineers for the next large quake, Thomas Jordan of the Southern California Earthquake Center is leading an effort to create a highly accurate tool for assessing the risks earthquakes pose to urban areas and critical infrastructure, such as dams, nuclear power plants, and energy grids. The tool, called CyberShake, takes a physics-based, computational approach to the problem, integrating the many facets of an earthquake event into a single model. The result is a map of Southern California that can predict the intensity of ground shaking at specific sites during a quake.

Under its most recent allocation on Titan, Jordan's team produced the most accurate hazard map for the Los Angeles region to date, incorporating more

realistic physics and doubling simulation capability for seismic wave frequency—from 0.5 hertz to 1 hertz—one of the most important variables in predicting earthquake-caused building damage.

The team made the enhancements after identifying a threshold around 1 hertz at which their simulations diverged from well-recorded earthquake data. Building structures respond differently to different frequencies. Large structures like skyscrapers, bridges, and highway overpasses are sensitive to low-frequency shaking (1 hertz or less), whereas smaller structures like homes are more likely to be damaged by high-frequency shaking, which ranges from 2 to 10 hertz and above.

"One of the simplifications we use in low-frequency simulations is a flat simulation region," said Philip Maechling, a team member and computer scientist at the University of Southern California. "California isn't flat, of course. To improve these simulations and their results, we had to add new complexities, like topography."

In addition to topography—the roughness of the earth's surface—the team's simulations now include additional geometrical and attenuation (gradual dampening of the shaking due to loss of energy) effects, including near-fault plasticity, frequency-dependent attenuation, small-scale near-surface heterogeneities, near-surface nonlinearity, and fault roughness. Their enhanced anelastic wave propagation code, AWP-ODC, can now perform simulations using Earth models near 4 hertz.

Titan's GPU-accelerated nodes proved pivotal to improving the team's seismic hazard analysis capability, running the team's simulations 6.3 times faster than a CPU-only implementation and reducing the needed compute time by 20 percent. The team plans on pushing its code to even higher frequencies, above 10 hertz, to get a clearer picture of how large earthquakes may affect small structures in urban areas.

"This was impossible without high-performance computing," Jordan said. "We are at a point now where computers can do these calculations using physics and improve our ability to do the type of analysis necessary to create a safe environment for society."

High Frequency Physics-based Earthquake System Simulations

Thomas Jordan
Southern California Earthquake Center
2015 INCITE project
25,000,000 hours

Breaking Down Biofuel's Biggest Barrier

A model of pretreated lignocellulosic biomass—cellulose fibers in dark green, hemicellulose in light green, and lignin in brown. ORNL researchers are simulating biomass to help develop better plants and pretreatments for biofuel. Image credit: scistyle.com

When the Ford Motor Company's first automobile, the Model T, debuted in 1908, it ran on a corn-derived biofuel called ethanol, a substance Henry Ford dubbed "the fuel of the future."

Cheap fossil fuels undermined Ford's vision of wide-scale renewable transportation fuel, but recent advancements in biofuel research, including simulation and modeling conducted by ORNL's Jeremy Smith, have rekindled the dream of economically viable ethanol.

Today, biofuel research is focused on cellulosic ethanol—alcohol made from woody plants and waste biomass, such as corn stalks. Currently, breaking down these plant materials into simple sugars requires a lot of acid, water, and heat, adding cost and complexity to the process. Central to the problem is lignin, a critical component of the plant cell wall in nature and one of the primary roadblocks to cost-effective ethanol production.

DOE's BioEnergy Science Center (BESC) is dedicated to overcoming the lignin barrier, carrying out experiments to find ways to neutralize the molecule. In support of BESC's efforts, Smith's team uses computational biophysics to add molecular-level context that guides researchers' efforts.

With Titan, Smith's team has produced some of the largest and most complex biomolecular simulations to date. These multicomponent models provide

physics-based explanations for the microscopic interactions that take place during industrial biofuel production and experimental research.

To understand better why lignin is such a problem, Smith's team simulated a 23.7 million-atom biomass system consisting of lignin, cellulose, and enzymes. The simulation revealed that lignin not only blocks the enzymes from breaking down cellulose by occupying enzymes' preferred binding locations but also attracts and occupies the cellulose-binding domain of the enzymes.

While detailing the problem, the simulation also gave clues to help mitigate lignin's impact on biofuel production. Though lignin binds to cellulose in undesirable ways, the simulation demonstrated it does so selectively, preferring crystalline, or more ordered, cellulose fibers.

"That's the kind of rational insight we can provide using computer simulation," said team member and ORNL staff scientist Loukas Petridis. "There's nowhere else in the world where we could have run this simulation."

Other simulations carried out by Smith's team coordinated directly with BESC experiments and focused on developing new chemical pretreatments and improving biofuel yield.

A 250,000-atom model consisting of lignin and an experimental solvent mix of tetrahydrofuran (THF) and water gave researchers a detailed explanation of the cosolvent's effectiveness. The model showed how THF acts as a barrier between water and lignin, binding favorably to both and serving as a buffer that makes lignin easier to remove during biofuel processing.

In another project, Smith's team improved scientists' understanding of why a genetically modified lignin created by BESC researchers boosted biofuel yield. By comparing 100,000-atom simulations of a wild and genetically modified biomass system composed of lignin and hemicellulose, another major component of the plant cell wall, the team identified a strong explanation for the mutant lignin's effectiveness: hydrophobic, or water repelling, lignin binds less with hydrophilic, or water attracting, hemicellulose. This insight can be leveraged in future studies to make plants easier to deconstruct for biofuels.

"We're trying to reach the complexity that is found in nature and industrial conditions," Petridis said.

"It took a decade of work to determine all the steps of lignin biosynthesis and find ways to manipulate genes. In the future, we hope to circumvent some of the work by continuing to test our models against experiment and making good suggestions using supercomputers."

Molecular Simulation in Bioenergy

Jeremy Smith
ORNL/ University of Tennessee
2014-2015 ALCC project
25,000,000 hours

Service, Customer Satisfaction Highlight Strong Year for OLCF

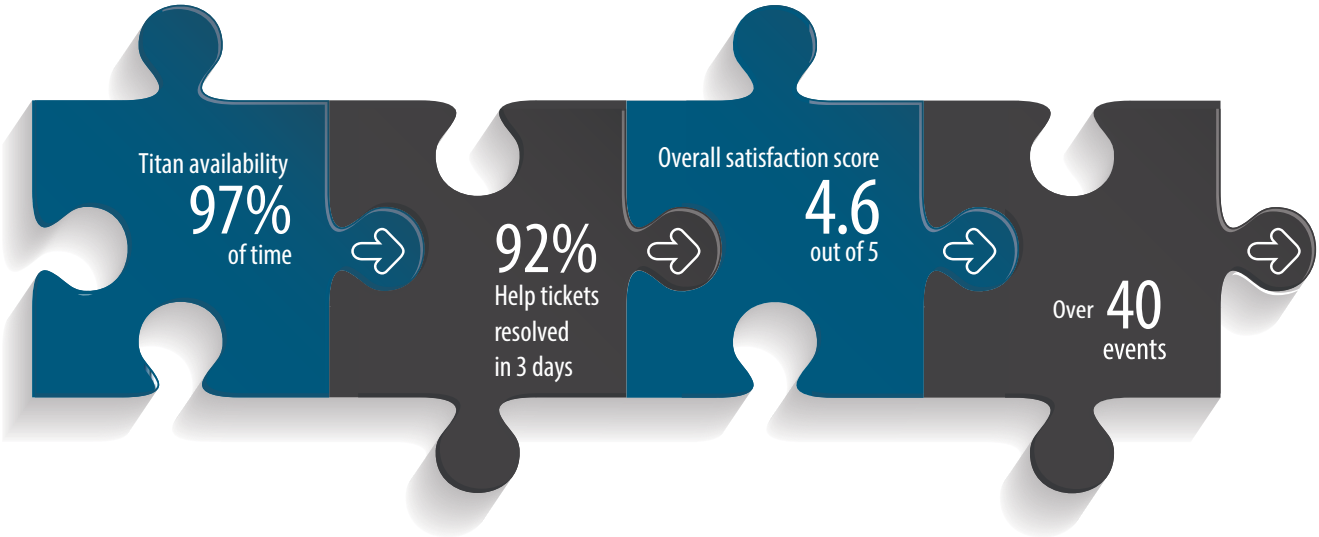

The OLCF offers users some of the world's most advanced high performance computing and data systems available for scientific discovery. However, these cutting-edge resources are not the only way we are providing valuable resources at the OLCF. The talented and dedicated staff of the OLCF—those who design, procure, provision, and provide support for these resources and the users of the systems—add a unique dimension.

Every year, the mission for OLCF staff is clear—partner with users to make the best possible use of computing resources.

In 2015, the OLCF took its annual charge to heart, vowing to maintain its very high customer satisfaction ratings from the previous year.

Staff in the various OLCF groups realized that improving customer satisfaction would require more than just a stable system—Titan was available 97% of the time in 2015, up from 95% in 2014. Improvement also would require staff members to train users, provide world-class customer support, and create tools to make using Titan and other OLCF resources as easy and efficient as possible.

To promote continual improvement, the facility conducts an annual survey to solicit feedback from users regarding their experience at the OLCF. The numbers spoke volumes—in overall satisfaction, the OLCF was rated 4.6 out of 5. Users' comments included the following: "The OLCF gives excellent support to industry," "The OLCF support pages have been improved in terms of finding helpful resources," and "The overall user experience is the best I've ever had."

In addition, the OLCF worked hard to resolve reported issues as soon as possible. In fact, staff exceeded all metric targets in addressing users' issues, resolving 92 percent of all help tickets within 3 days.

The OLCF conducted more than 40 training events, workshops, and seminars in 2015, inviting users and prospective users to learn from high-performance computing (HPC) experts on conference calls and webinars, come to the lab for workshops, and attend seminars and hackathons. Hackathons took place at ORNL as well as the University of Illinois at Urbana-Champaign and the Swiss National Supercomputing Centre.

The OLCF also began working with a new set of applications through its Center for Advanced Application Readiness (CAAR), an initiative to work with users to prepare their codes for the OLCF's next-generation supercomputer, Summit. Nearly 80 users took part in the 2015 CAAR workshops. In addition to preparing users for Summit, CAAR workshops also prepare users' codes to be as portable as possible to multiple computing architectures. Portability is one of the major challenges as next-generation supercomputers come online, and the OLCF is a leader in helping its users create maximum flexibility with their code bases.

Users took note of these training opportunities, and many project principal investigators considered them a catalyst for their respective team members' career advancements. ORNL researcher Chongle Pan said, "The OLCF staff trained our postdocs to be proficient in scaling up algorithms to supercomputers and executing large computing jobs. These research experiences launched many of their careers as tenure-track professors."

Despite overall positive grades from the user surveys, OLCF staff wanted to address several areas where certain users experienced difficulties. One major challenge for the OLCF is the volume of users—dozens of users run millions of computations per day, meaning that scheduling jobs to run on Titan can become complex.

OLCF user support specialists Adam Simpson and Matt Belhorn developed and released Wraprun; its name is a play on Aprun, the Cray job launcher Titan users have typically used.

The OLCF developed Wraprun to assist Titan users who wish to run ensemble jobs—ones that consist of multiple independent jobs that are run simultaneously—but had difficulty using Aprun to do these runs. By bundling the jobs, Wraprun allows users to launch more jobs in less time, increasing efficiency and productivity.

Fragmentation is another issue OLCF staff addressed in 2015 by fine-tuning how jobs are scheduled on Titan. New tests indicate that by differentiating between large, long-lived jobs and small, short-lived jobs, Titan is able to do more work in less time.

At any given moment, Titan runs between 30 and 60 jobs of various sizes—from 10,000-plus nodes to 1 or 2 nodes, such as debugging tasks. To regulate the flow of incoming jobs, Titan relies on the Moab workload management system. Traditionally, Moab keeps a list of all incoming jobs scheduled to run on Titan, irrespective of size. When space opens up on the machine, Moab scans from the top of the list to the bottom for a job to fill the vacancy.

"If you think of Titan as a big 3-D grid, you can have various portions of your job spread throughout the machine," said Chris Zimmer, an HPC systems engineer in the Technology Integration Group at the OLCF. "This is exacerbated by the fact that we have very small, short-lived jobs popping in and out of that list."

To reduce fragmentation, OLCF staff programmed Moab to schedule workloads differently. Instead of scanning a list from top to bottom, Moab now moves small, short-lived jobs (less than 2 hours) to the bottom of the list, increasing the probability that a large, intensive workload will receive a more favorable allocation.

As computing power has increased, users are creating larger simulations and generating even more data. To address bottlenecks coming in data transfer, a cross-departmental team of User Assistance Specialists Suzanne Parete-Koon and Belhorn as well as HPC Operations team member Jason Kincl have been working to improve Globus. Parete-Koon has led the OLCF data transfer working group, which works with users and center staff to understand data transfer needs. Globus, among other functions, facilitates data transfer both inside and outside a user facility.

"OLCF users have repeatedly expressed the need for a smoother Globus authentication process," Parete-Koon said. "The data transfer working group kick-started the effort for an OLCF certificate authority [CA] to help simplify the authentication for Globus transfers."

Parete-Koon and Kincl met with a Globus representative to set up a CA for the OLCF to host, and Kincl took the technical lead for establishing the CA. In the past, OLCF users had to obtain an Open Science Grid (OSG) certificate to authenticate to Globus. "It was a convoluted way of doing business," said Belhorn, adding that it could take up to a week for the OSG certificate to be approved, which would then slow down the transfer of data. Now users have a far simpler process for accessing the system remotely.

Another aspect of using supercomputers revolves around sending data to the various supercomputing resources of the OLCF; users can take data from Titan and send it between storage systems, analysis clusters, and transfer nodes using the Spider file system.

Spider contains roughly 20,000 disk drives divided among 72 controllers to handle supercomputer-sized data sets, making the task of monitoring Titan's system a challenge. For years, DataDirect Networks (DDN) and the OLCF have built the storage systems. However, because of the vast number of disk drives, the OLCF needed to develop a more efficient performance monitoring tool. OLCF staffer Ross Miller designed a tool to gather data from all 72 controllers and aggregate the information, which streamlines the process of monitoring the system.

In addition to alerting HPC Operations staff to potential drive failures, DDNTool has helped staff diagnose performance issues that users might experience with Titan. DDNTool feeds data into Splunk, a commercial software package for data mining, which then charts researchers' use of object storage targets, or OSTs, on Spider. OSTs are the smallest unit that researchers can use to save data at the OLCF. To get the maximum performance during file transfers, users need to structure the I/O in their codes to use as many of the 1,008 OSTs as possible. In many cases, users unwittingly use only a fraction of the OSTs, limiting their performance.

As supercomputers have continued to get faster, users are creating increasingly large amounts of data. To help manage that data more efficiently, the OLCF updated its High-Performance Storage System (HPSS) this year. Jason Hill, OLCF staffer, explained that the upgrades were user-centric. "Our users let us know that moving data into and out of HPSS was too slow, thus impacting their productivity," Hill said. "It's a resounding theme for us, and we took that to heart when we targeted building out infrastructure both in networking and with our disk cache."

To help users move data more efficiently between OLCF resources, staff doubled the number of data transfer nodes from 4 to 8. In addition, the OLCF team invested in 40-gigabit Ethernet connectivity for all the HPSS movers together across two separate switches. The HPSS movers are responsible for reading and writing the data in the disk cache as well as migrating data to the tape media. The Ethernet switches provide a high bandwidth path connecting the disk and tape movers together for moving the data quickly to the archival media. The two switches are now connected by twelve 100-gigabit-per-second links that are aggregated for performance.

The investment has already paid off, and researchers are taking note. "The upgrades at the OLCF make HPSS really useful for production," said Katrin Heitmann, OLCF user and computational scientist at Argonne National Laboratory. "There are so many things that we can do with the raw data that it's important to hold on to it for as long as possible."

OpenACC Hackathon October 2015

In addition to speeding up communication between the different storage tiers, the storage team also wanted to help users avoid retrieving data from tapes for as long as possible. The OLCF purchased 18 petabytes of disk cache to help expand the amount of data that could be recalled and accessed quickly.

Storing data on tape can offer more security than storing data on discs or array disc drives because of tape's increased tolerance of varied environmental conditions. In addition, tapes do not need power unless data is being moved on or off a tape. Until recently, though, tapes still had some risks associated with particularly large data sets.

For data that could fit on a single tape, OLCF staff used to "stripe" the information across two tapes. If one of the tapes was damaged or lost, though, there was no recourse for recovering that tape's data. Staff rarely striped data across more than two tapes because of that risk.

Enter the Redundant Array of Independent Tapes (RAIT) feature. HPSS's RAIT feature allows incoming data to be striped across four tapes, but it also has an extra tape—called a parity tape—that can allow the data to be reconstructed in the event that a tape is damaged or lost.

"RAIT uses the same algorithm that modern disk arrays employ," Hill said. "In our case, you can take a byte of data and split it into 8 bits. Put those 8 bits into groups of 2, then calculate a checksum of all those data points, and put it into a separate parity segment." This 4 + 1 approach allows staff to provide users with a method to recover from a tape failure—an option that was never before available for data saved on tape.

One of the other major requests from OLCF users dealt with the OLCF's Rhea cluster, a machine used for pre- and post-processing of simulation data generated on Titan. Users requested more RAM and GPUs for each of its 512 nodes. The OLCF responded by doubling the amount of RAM per node and adding GPU nodes to further accelerate data processing.

As the OLCF prepares to install its next-generation supercomputer, Summit, staff will continue to find new ways to offer users the most comprehensive training and support network for using one of the nation's most stable, reliable, and well-connected supercomputing centers.

Stepping up to Summit

More than a year after Summit's announcement, preparations are well underway to ensure the OLCF's next supercomputer is ready for science upon arrival in 2018.

Once completed, the large-node, GPU-accelerated system will provide researchers with at least a fivefold increase in computational power over Titan, along with improved memory, enhanced interconnection, and greater node performance. This leap in computing capability is expected to require 10 megawatts or less of additional power, only about 10 percent more than Titan requires.

Following the best path to a bigger, better, faster machine within a limited energy footprint requires a multipronged strategy that successfully integrates cutting-edge hardware, software, and infrastructure in novel ways. In 2015, the OLCF made noteworthy progress on all three fronts, as OLCF staff, vendor partners, application teams, and construction crews took important steps to help deliver the unprecedented system.

On the hardware front, OLCF staff and vendor partners became more familiar with Summit's heterogeneous architecture—to be composed of IBM Power9 CPUs and NVIDIA Volta GPUs—and storage platforms.

Two simple, yet realistic, test systems called Pike and Crest paved the way. Matched with IBM's parallel file system technology, the 14-node Pike provided a glimpse into Summit's high-speed data storage performance and capability. GPU-equipped Crest served as a platform for building and running early versions of software and identifying bugs.

"These two systems will give us a head start on Summit's next-generation compute and storage systems so we will be better prepared to support users," said OLCF staff member Dustin Leverman.

In 2016, test groups will get an even more realistic picture of Summit with the incorporation of NVIDIA's high-bandwidth interconnect, called NVLink, a critical component of the machine that will greatly improve data movement between CPUs and GPUs.

In April, the OLCF announced the 13 partnership projects selected to develop the first modeling and simulation software for Summit. Through the OLCF's CAAR initiative, teams from the United States and Europe began to work with staff from the OLCF Scientific Computing Group and the IBM/NVIDIA Center of Excellence at ORNL to redesign, port, and optimize their software to Summit's architecture.

The 13 CAAR projects represent a diverse range of scientific disciplines—from astrophysics to combustion engineering to seismology—and employ

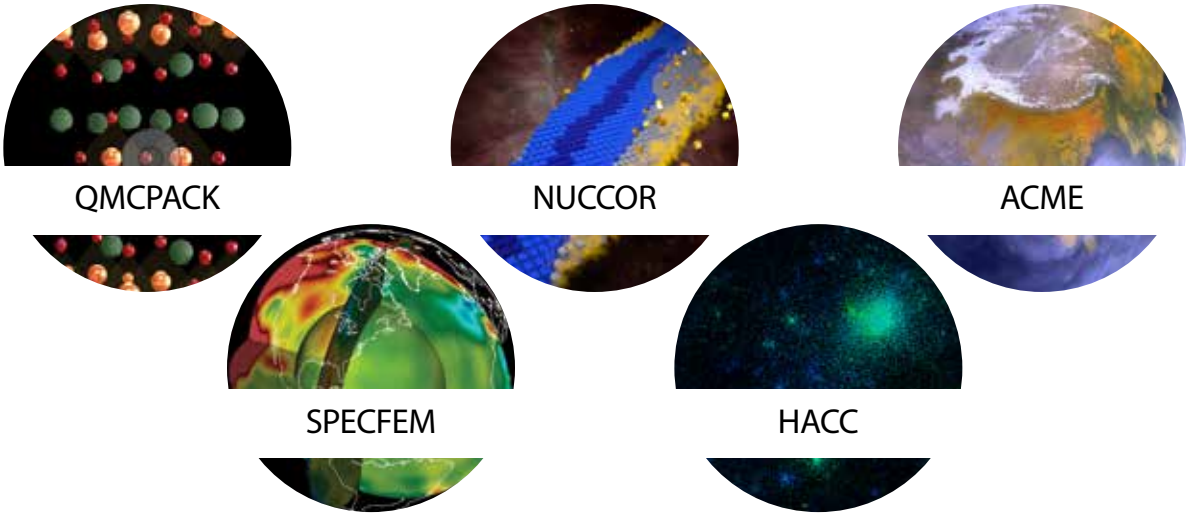

The 13 CAAR codes were selected based on how they aligned with DOE's scientific goals, showed potential to make good use of Summit's architecture, and represented diverse scientific domains

varied computational algorithms and programming approaches. The 3-year effort will culminate with a scientific grand challenge demonstration for each project upon the completion of Summit's acceptance tests, expected in 2018.

"We know that improved software is just as important as improved hardware to supercomputer performance and our facility's mission," said Tjerk Straatsma, ORNL's Scientific Computing Group leader. "CAAR brings together the people who know the science, the people who know the code, and the people who know the machine, so that cutting-edge science can be conducted on Summit."

For more information about the CAAR projects, see page 24.

While the OLCF continued to explore Summit's hardware and software, the facilities team began work on upgrading the building to house Summit.

To support future computing systems, ORNL's Computational Sciences Building is being equipped to supply 10 megawatts of additional power and 20 megawatts of additional cooling.

A new steel platform constructed outside the building will add the necessary infrastructure to install up to seven transformers. On the west end of the building, a subsurface duct bank and conduit bridge were installed to connect new transformers to the existing ORNL electrical distribution system. Additionally, a cooling tower platform raised south of the building will provide needed infrastructure to cool Summit.

In a break from the past, the Summit facility will feature a warm-water cooling system that operates at temperatures nearly 30 degrees Fahrenheit above the current cold-water cooling system used to cool Titan. The change, enabled by more temperature-tolerant hardware and new plate-and-frame heat exchange technology, will save nearly $1 million per year in Summit's operating cost.

"Because Summit's cooling water supply temperature is only about 70 degrees, it will be a whole lot easier to maintain optimal conditions in our environment," said Jim Rogers, Director of Computing and Facilities at ORNL's National Center for Computational Sciences.

Early Science (CAAR)

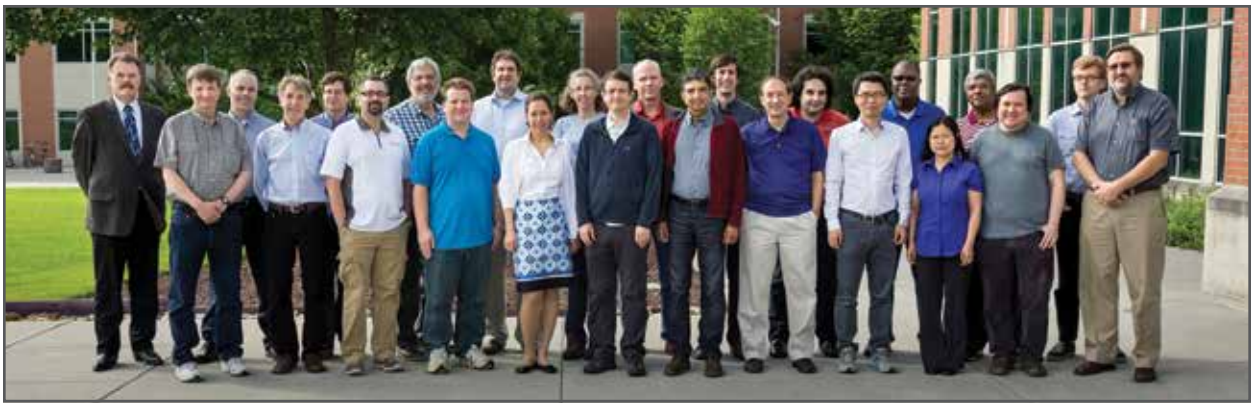

CAAR is a partnership of the OLCF's Scientific Computing Group, application development teams, and vendor partners, with the goal of readying a set of applications for the next-generation supercomputer Summit. CAAR focuses on redesigning, porting, and optimizing application codes for Summit's hybrid CPU–GPU architecture but also emphasizes the importance of performance portability across different architectures. The application software is expected to be of interest to a broad user base in its scientific community and to be open source or freely available through a license to the computational science community. Through CAAR, teams gain access to early software development systems, leadership computing resources, and technical support from the IBM/NVIDIA Center of Excellence at ORNL. The current CAAR program will culminate with each team's scientific grand-challenge demonstration on Summit. After review of 29 proposals, the OLCF selected 13 applications for the CAAR program.

Science Domain: Climate

Title: Climate Research: Advancing Earth System Models

PI: David Bader, Lawrence Livermore National Laboratory

A collaboration among DOE laboratories and Colorado State University aims to adapt its ACME code to next-generation supercomputers, with a goal of providing new insights related to the effects of climate change. Current work focuses on coupling components for atmosphere, land, ocean, sea and land ice, and river transport. New approaches leveraging Summit's GPU accelerators will allow researchers to better capture hydrological processes, notably critical land–ice interactions along the coasts of Greenland and Antarctica and complex convective cloud systems. The high-resolution models coupled with better physical representations of sub-grid scale processes will help scientists predict long-term changes in precipitation patterns and future sea-level changes.

DIRAC

Science Domain: Relativistic Quantum Chemistry

Title: CAAR Oak Ridge Proposal for getting the Relativistic Quantum Chemistry Program Package DIRAC ready for SUMMIT

PI: Lucas Visscher, Amsterdam Center for Multiscale Modeling, VU University Amsterdam

Lucas Visscher and his colleagues are developing predictive modeling tools for future studies of new materials containing heavy elements and their reactive properties. Such simulations are important in finding safer methods for handling nuclear materials during the electricity generation phase, as well as improving storage or possibly recycling these materials after use in power plants. Summit's GPU accelerators will enable DIRAC—the team's relativistic quantum chemical application—to increase its current capability from modeling small systems, with tens of atoms at most, to modeling systems with hundreds of atoms. This performance increase will enable a first-principles approach to investigating the electronic properties and molecular reactivity of uranium, platinum, gold, and other important heavy metals. With that capability, DIRAC can become a software tool sharable with the scientific community to test new methods, with the ultimate aim of limiting or even preventing the risks of handling radioactive material.

FLASH

Science Domain: Astrophysics

Title: Using FLASH for Astrophysics at an Unprecedented Scale

PI: Bronson Messer, Oak Ridge National Laboratory

By using the extremely scalable FLASH code, a team led by Bronson Messer plans to extend core-collapse supernova simulations to time and distance scales that can be compared with observations from, for example, space-based observatories like NuSTAR or XMM-Newton. By enhancing FLASH to make full use of next-generation computing power, the team plans an order-of-magnitude increase in the number of nuclear species tracked and evolved, increasing the physical fidelity of simulations significantly. Increasing the number of evolved nuclear species from 13 to 150 will allow accurate predictions of the amount of radioactive Titanium-44—an excellent indicator for incomplete nuclear burning—produced in core-collapse supernovae. The stretch goal of another order of magnitude, up to 2,000 species, will allow in situ simulation of r-process nucleosynthesis (rapid neutron capture). The r-process is thought to produce roughly half of all the neutron-rich nuclei heavier than iron in the universe. These improvements also will be useful in simulations of thermonuclear supernovae, novae, and x-ray bursts, among other phenomena.

Science Domain: Plasma Physics

Title: Particle Turbulence Simulations for Sustainable Fusion Reactions in ITER

PI: Zhihong Lin, University of California, Irvine

In an effort to create a sustainable fusion reaction at the ITER experimental fusion reactor in France, a team led by Zhihong Lin is using the extremely scalable GTC code to simulate billions of plasma particles inside the reactor. The team's simulations are capable of using the full computational capabilities of current-generation supercomputers. Lin's team should be able to gain more insight on particles' behaviors during a fusion reaction by further developing the code's ability to plot the microscopic detail of particle motion more accurately while also simulating whole-reactor conditions and the associated particle interactions at longer time scales. These simulations will give greater insight—such as further understanding plasma confinement properties and how plasma turbulence affects a fusion reaction—into the conditions for a sustainable fusion reaction.

Science Domain: Cosmology

Title: Cosmological Simulations for Large-Scale Sky Surveys

PI: Salman Habib, Argonne National Laboratory

Today's land- and satellite-based telescopes are able to capture huge amounts of data on galaxies throughout the observable universe, and researchers are using this data to answer questions about the nature of dark energy and dark matter and the formation of the universe. However, to understand how galaxies form, researchers need to compare observational data with cosmological simulations by modeling the large-scale structure and distribution of matter during the evolution of the universe over 13 billion years. Salman Habib's team is using its Hardware/Hybrid Accelerated Cosmology Code on DOE leadership computers to conduct unprecedented cosmological structure simulations at resolutions needed for modern-day galactic surveys. The results of these simulations will be used for a diverse set of scientific investigations, like strong and weak lensing maps and synthetic sky surveys. Over the next decade more realistic synthetic sky catalogs will be produced by adding the physics of gas dynamics and star formation along with dark matter—a level of realism that cannot be achieved on Titan.

Science Domain: Quantum Chemistry

Title: Large-Scale Coupled-Cluster Calculations of Supramolecular Wires

PI: Poul Jorgensen, Aarhus University

A team led by Poul Jorgensen is aiming to further develop its novel quantum chemistry code LS-DALTON, which provides researchers with the most accurate description of electronic structure in large molecular systems. On Titan, Jorgensen's team demonstrated that the code's time-to-solution closely correlates with system size, meaning Summit's greater per-node memory and increased CPU–GPU integration should result in a significant performance boost. The improved performance has the potential to benefit all areas of molecular science and engineering by enabling increases in both the maximum molecular system size that can be simulated and the overall accuracy achievable. Targeted applications include enzyme-catalyzed chemical reactions, carbon nanotubes and graphene, and the preferred crystal form of organic molecules.

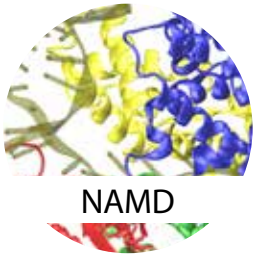

Science Domain: Biophysics

Title: Molecular Machinery of the Brain

PI: Klaus Schulten, University of Illinois at Urbana-Champaign

A team of researchers from the University of Illinois at Urbana-Champaign is searching for answers to challenging questions about the brain. Despite tremendous advancements in cellular imaging, the mechanism by which molecular processes give rise to cellular signaling is not clear. Running the biomolecular modeling application NAMD on Summit will enable the team to simulate 200-million-atom systems at 200 nanoseconds per day, or at eight times the speed of Titan. The team members will exploit Summit's GPUs and large per-node memory for interactive remote visualization and analysis using their program VMD. As a result, they will be able to simulate the synaptic transmission processes in the brain more comprehensively than ever before. The connection thereby established between cellular- and molecular-level signaling promises to reveal the molecular origins of neuronal disorders such as epilepsy, Alzheimer's disease, Parkinson's disease, and schizophrenia.

Science Domain: Nuclear Physics

Title: Nuclear Structure and Nuclear Reactions

PI: Gaute Hagen, Oak Ridge National Laboratory

Gaute Hagen leads a collaboration to advance a powerful tool to study a wide range of nuclei and their behaviors. NUCCOR, a nuclear physics application for deciphering the structure and reactions of atomic nuclei, was developed to help scientists expand understanding of fundamental interactions in the universe—interactions too difficult to observe in a laboratory. By optimizing the code on Summit, researchers could establish long-sought benchmarks based on first-principles computations that help guide and interpret experiments and gain insight into low-energy nuclear physics problems such as neutrinoless double-beta decay, a not-yet-observed form of atomic nuclei decay. This improved framework will inform research in nuclear energy, including the development of nuclear fusion and fission reactors, and energy generation in stars.

Science Domain: Computational Chemistry

Title: Developing Coupled Cluster Methodologies for GPUs

PI: Karol Kowalski, Pacific Northwest National Laboratory

Vital in myriad interactions with proteins, glycans (macromolecules with carbohydrate functionality) encompass and protect the cell, playing an important role in the recognition of pathogens. Plant glycans such as cellulose, starch, and chitin offer the promise for creating materials with new physical and chemical properties. Researchers from Pacific Northwest National Laboratory and IBM Research—Almaden aim to scale the NWChem application to use GPU accelerators. The accelerators will provide benchmark energies to allow for accurate parameterization of force fields for glycans as well as develop and disseminate an open-source database of accurate glycan conformational energies. New implementations of high-accuracy methods capable of taking advantage of Summit computational resources will significantly shift the system-size limit tractable by very accurate yet expensive methods accounting for the inter-electron correlation effects. The team's results promise to allow for the computational investigation of new carbohydrate-based materials from bulk polymers to nanoparticles with useful functionality, particularly in the areas of medicine, energy generation, and advanced materials.

Science Domain: Materials Sciences

Title: Materials Science Research for High-Temperature Superconductors

PI: Paul Kent, Oak Ridge National Laboratory

A team led by ORNL's Paul Kent plans to use its QMCPACK code to further research into transition metal behavior, ultimately hoping to gain insight and predictive capability for using materials as high-temperature superconductors and other energy-materials applications for which today's state-of-the-art methods are overly reliant on empirical information. The team has simulated strontium copper oxide on the OLCF's Titan supercomputer, but Summit's larger per-node memory and greater computational power will allow the team to simulate a range of heavier, more complex materials, moving researchers closer to understanding high-temperature superconductor properties at the subatomic level.

Science Domain: Engineering/Combustion

Title: Fluid Dynamics Research to Accelerate Combustion Science

PI: Joseph Oefelein, Sandia National Laboratories

Joseph Oefelein of Sandia National Laboratories heads a team focused on using the computational fluid dynamics code RAPTOR to explore the chemical and physical processes of combustion. Optimizing RAPTOR for Summit's hybrid architecture will enable a new generation of high-fidelity simulations that match engine operating conditions and geometries identically. Such a scale will allow direct comparisons to companion experiments, providing insight into transient combustion processes such as thermal stratification, heat transfer, and turbulent mixing. Ultimately, this work will help inform new design concepts of internal combustion engines and gas turbines, which have the potential to provide significant increases in fuel efficiency with minimal emissions.

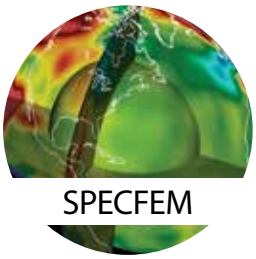

Science Domain: Seismology

Title: Mapping the Earth's Interior Using Big Data

PI: Jeroen Tromp, Princeton University

Jeroen Tromp of Princeton University and his team propose to create a high-fidelity picture of the Earth's interior using earthquake data gathered from around the world. Leveraging the code SPECFEM3D_GLOBE and a technique called adjoint tomography, Tromp's team has been able to image the Earth's entire mantle on Titan. Summit will allow the team to dig even deeper—all the way down to the Earth's inner core—and capture more detail by providing the extreme computational power necessary to optimize the code for seismic waves at periods approaching 1 second, the highest frequency signal that can propagate across the entire planet.

Science Domain: Plasma Physics

Title: Multiphysics Magnetic Fusion Reactor Simulator, from Hot Core to Cold Wall

PI: C. S. Chang, Princeton Plasma Physics Laboratory, Princeton University

A team led by C. S. Chang of Princeton Plasma Physics Laboratory is committed to delivering a more powerful tool for exploring the feasibility of sustainable fusion in a tokamak fusion reactor, the donut-shaped device being built by ITER. By using Summit, Chang's team expects its highly scalable XGC code, a first-principles code that models the reactor and its magnetically confined plasma, could be simulated 10 times faster than current supercomputers allow. Such a speedup would give researchers an opportunity to model more complicated plasma edge phenomena, such as plasma turbulence and particle interactions with the reactor wall, at finer scales, leading to insights that could help ITER plan operations more effectively.

OLCF Inaugurates Advanced Data and Workflow Group

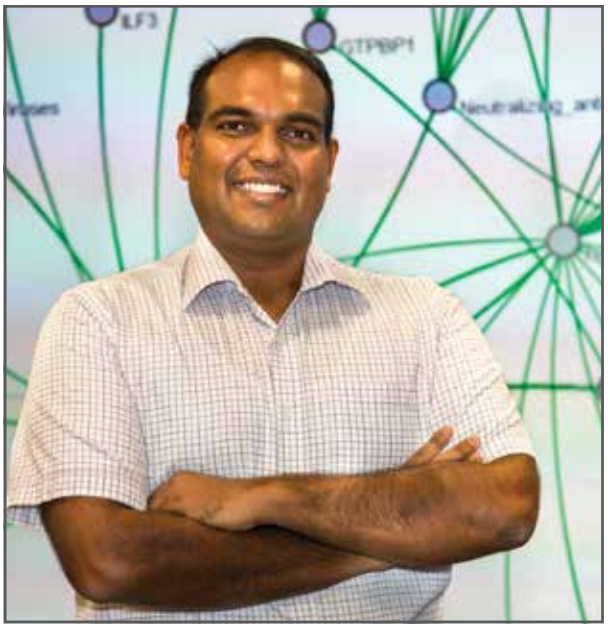

Advanced Data and Workflow group leader, Rangan Sukumar

As technology has improved researchers' abilities to create, simulate, collect, and store data, big data has become the name of the game. Based on the needs of our users, the National Center for Computational Sciences (NCCS) formed the new Advanced Data and Workflow Group to address data-related challenges. Sreenivas Rangan Sukumar began as group leader in October.

"I am very enthusiastic about this new group," Sukumar said. "Even in this decade of big data frenzy, there aren't many data science jobs that can make a difference across science domains spanning astronomy to zoology. It is exciting to be part of a fantastic team of data scientists, developers, and visualization experts that share a similar passion for science and are driven by the opportunity to apply their skills on problems for the greater social good while tackling tough computer science challenges at scale."

The group designs, builds, and offers creative data-science workflow solutions to enable interactive data-driven discovery. These solutions are designed for DOE-focused science that requires scale and performance on leadership compute architectures hosted by the OLCF and big data architectures hosted by the Compute and Data Environment for Science (CADES) at ORNL.

The group works primarily with OLCF users who participate in the INCITE, ALCC, and Director's Discretionary (DD) programs through liaisons who offer scientific, technical, and operational assistance for building data-based discoveries for their project-specific needs. The team also assists other facilities at ORNL, including the Center for Nanophase Materials Sciences, the Manufacturing Demonstration Facility, and the Spallation Neutron Source, to work with OLCF computing resources and deploy efficient workflows. So far, the group has piloted on-demand data, analytics, and visualization support services to accelerate the science missions at these DOE user facilities. The group aims to offer workflow productivity and automation tools, database and data visualization services, and data analysis/machine learning tools. This includes operating, maintaining, and offering the Exploratory Visualization Environment for Research in Science and Technology (EVEREST) Power Wall as a scientific visualization laboratory.

The group's software engineers and data scientists have expertise leveraging data analysis hardware and software tools; recommending suitable techniques for data analysis at extreme scales; designing analytic workflows that consider different aspects of parallelism toward optimizing performance; developing software services for seamless automation, data management, data transfer, and visualization to assist in discovery; and near real-time CPU–GPU-based rendering of large datasets using EVEREST resources for in situ experiment and simulation visualization. Ultimately, this group assists

users to draw insights from the massive amounts of data they generate, leading to scientific discoveries.

Sukumar's vision for the new group includes short-mid-, and long-term goals. In 2015, several group accomplishments met the short-term goals, including is the Accelerated Climate Modeling for Energy (ACME) project. OLCF staff members have been integral to the ACME project since its inception. In 2015, several data liaisons enabled sharing hundreds of terabytes of data through the ORNL Earth System Grid Federation node to the international community. The group also assisted in the largest-ever Vortex visualization, which allowed researchers to study complex turbulence characteristics that required a high-resolution blowup of specific highlighted regions. With off-the-shelf visualization tools failing, the group's data liaisons created a workflow using Titan and EVEREST resources—specifically Focus and Lens—to help users visualize their simulation results. This visualization was featured as the cover of an AIP Publishing physics journal. Furthermore, the data and visualization members also hosted a week-long hackathon involving participants from 13 international organizations in support of the BRAIN initiative.

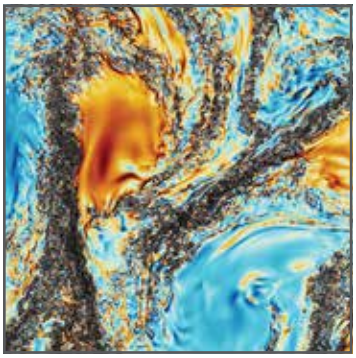

Image featured on cover of AIP's Physics of Fluids

In 2016, the group is continuing operational activities for the OLCF for data-intensive computing and visualization services. This includes supporting workflow engines such as Pegasus, Swift, and PanDA, as well as creating new on-demand analytic services such as Spark On-Demand. Group members also will develop and implement road maps for emerging analytics (e.g., deep learning) and workflow services to support future INCITE needs on Summit and beyond. The group is participating in training the user community on high-performance data analytics by creating new documentation and video tutorials for on-demand analytic services and remote visualization capabilities on EVEREST.

In the midterm, the group will develop expertise to understand scalability of algorithms on a variety of current and emerging architectures as well as develop integrated data strategies at the OLCF and CADES. In the long term, Sukumar plans to offer a collection of on-demand computing tools and applications that will scale seamlessly on different computer architectures for various DOE scientific challenges as well as other facilities of the future. These mid- and long-term goals include supporting on-demand analytics with software libraries—such as cuBLAS, MLLib, Torch, Caffe, and Theano—as well as offering an architecture-aware algorithmic suite of software for scalable data analysis and machine learning on HPC resources. The group also plans to design, develop, and deploy on-demand workflow infrastructures, including scripting tools for analyzing, filtering, mapping, and rendering custom visualization workflows.

"Science is becoming more and more data driven," Sukumar said. "And this group's charter is to leverage leadership-class computing resources to enable fundamental scientific discoveries that are going to change the world in the future. While doing the best with what we already have, we are also going to be thinking ahead on what will science workflows look like on the exascale machine, what problems can be solved then that we cannot solve today, and how can we help our science users solve such problems? To be the leaders in scientific computing, we have to be leaders in data. I'm looking forward to that."

Groups

Advanced Data and Workflow

Group Leader: Sreenivas R. (Rangan) Sukumar

The Advanced Data and Workflow Group offers scientific, technical, operational, and thought leadership for building data-driven computing environments for OLCF user needs. The group designs and develops creative data-science workflows (compute, analytic, and visualization) to enable interactive data-driven discoveries that require scale and performance 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 novel algorithms and workflows, and implementing and supporting convenient tools.

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, mathematics, numerical analysis, and computer science. 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 any 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.

User Assistance and Outreach

Group Leader: Ashley Barker

The User Assistance and Outreach (UAO) Group provides support to the OLCF users, serves as liaisons 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 R&D community to develop, harden, and deploy solutions. The group's technology scope includes archival storage, parallel file systems, non-volatile memory, data management, high-performance networking, and architecture. TechInt staff members work behind the scenes to develop the infrastructure that supports the NCCS systems, keep it 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.

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 works with infrastructure systems, with Titan, and with other OLCF supercomputers. The HPCO group tests the systems when they are installed and upgraded and 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.

Computer Science Research

Group Leader: David Bernholdt

The programming tools team of ORNL's Computer Science Research Group researches, tracks, and purchases a wide range of software tools that help science researchers access and improve the performance of their applications on current and emerging OLCF computing systems. The group also manages contacts with vendors for the purchase of new modeling tools, languages, middleware, and performance characterization tools. The group focuses primarily on issues that arise for research applications when running on very large-scale systems, such as the OLCF's Titan supercomputer.

OLCF Executive Board and User Group

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 hosted 11 user conference calls in 2015, with an average attendance of 35 per call. Once a year, the OLCF hosts a face-to-face user meeting to share selected computational science and engineering achievements emerging from the OLCF's user programs, enable direct interactions among users, advance the OLCF's relationship with its user community, and highlight users' future computational requirements. Ninety members participated in the 2015 OLCF User Meeting June 23–25 at ORNL.

The first and second days of the meeting focused on science produced on Titan and the computational challenges users overcame in the process. Presentations on the second day included an overview by NVIDIA's Jeff Larkin about GPU computing and the OpenACC programming standard. The third day's activities looked ahead to next-generation supercomputing applications and the OLCF's next petascale machine, Summit.

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 supply user feedback on specific topics, such as the annual user meeting.

The OUG Executive Board requested that the OLCF find opportunities for users to share user-driven content about lessons learned using OLCF resources.

2015 OUG members

- **Mike Zingale**
  State University of New York–Stony Brook, Chair
- **Katrin Heitmann**
  Argonne National Laboratory, Vice-Chair
- **Hai Ah Nam**
  Los Alamos National Laboratory
- **Thomas Maier**
  ORNL
- **Mark Taylor**
  Sandia National Laboratory
- **Stephane Ethier**
  Princeton Plasma Physics Laboratory
- **Balint Joo**
  Jefferson Lab (ex officio)
- **David Dixon**
  University of Alabama
- **Joe Oefelein**
  Sandia National Laboratories
- **Brian Wirth**
  University of Tennessee

The User Assistance and Outreach group provided this opportunity during two of the monthly OUG webinars for users to give short talks on their work to the group. The talks focused on methods to use OpenACC in user applications; OUG Board Chair Mike Zingale led the first one.

INCITE, ALCC, and DD—the OLCF's Allocation Programs

At the OLCF, 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 2015, 30 INCITE projects were awarded a collective 2.25 billion hours on Titan. Awards in this program are made on an annual basis. For more information about INCITE, visit www.doeleadershipcomputing.org.

The 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. For more information, please visit www.science.energy.gov/ascr/facilities/alcc. A total of 22 programs were awarded over 1 billion hours through this program in 2015–16.

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 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/. Awards in this program are made on an annual basis; 226 DD projects were active in calendar year 2015.

2015 Usage on Titan by Allocation Program

Resource Overview

TITAN
Cray XK7 Supercomputer

HPSS
High Performance Storage System

EVEREST
Exploratory Visualization Environment for Research in Science and Technology

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.

STORAGE

HPSS

An archival file system consisting of disk and tape components for secure, long-term storage of scientific data.

SPIDER II

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

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. Rhea is dedicated to pre- and post-processing of simulation data generated on Titan.

SUPPORT SERVICES

DTN's

Data Transfer Nodes 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 computing system's login and service nodes.

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.

For more detailed information visit: www.olcf.ornl.gov/computing-resources

HPC User Support

OLCF UAO specialist Matt Belhorn tested, gave user training on and wrote documentation for Globus. Here, he discusses some of these points in the user training session.

The OLCF has established a user support model for effectively supporting users that is based on continuous improvement, regular assessment, and a strong customer focus. In 2015, the OLCF supported 1,176 users on 316 projects.

OLCF leadership uses its annual user survey to measure how well the OLCF is meeting its customers' needs. The mean rating for overall satisfaction with the OLCF in 2015 was 4.6 out of 5. Overall rating for the OLCF were positive; 95% of users reported being "satisfied" or "very satisfied." The survey also asks users to rate their overall satisfaction with OLCF Support Services; the mean rating in 2015 was 4.5 out of 5.

The center measures its performance using a series of quantifiable metrics. The metric targets are structured to ensure that users are provided prompt and effective support and that the user support organization responds quickly and effectively to improve its support process for any item that does not meet a minimum satisfactory score. The center exceeded the problem-resolution metric in 2015 with 92% of tickets being resolved within 3 business days. In response to an open-ended question about the best qualities of the OLCF, user assistance was listed as the top choice by 50% of the survey respondents.

The OLCF provides complementary OLCF user support vehicles that include user assistance and outreach staff as well as scientific, data, and visualization liaisons.

In 2015, UAO handled and resolved 2,395 tickets

2395
RESOLVED TICKETS

Percent of problems addressed in 3 business days

Target 80% | Actual 92%

Average of problem resolution ratings

Target 3.5/5.0 | Actual 4.6/5.0

USER SURVEY RESPONSE RATE

35%

USER SURVEY RESULTS

Question	Mean rating
Overall satisfaction with the OLCF	4.6/5.0
Overall satisfaction with support services	4.5/5.0
Overall satisfaction with compute resources	4.6/5.0
Overall satisfaction with website	4.4/5.0
Overall satisfaction with user guides	4.5/5.0

Education, Outreach, and Training

Workshops, user conference calls, training events, and seminars are integral components of user assistance and outreach. Whereas training can alleviate difficulties in using large-scale systems, training events also can serve to engage both the public and the user community. In addition to instructing users on how to use OLCF resources, the training program focuses on software development best practices.

The OLCF facilitated 40 user events in 2015, including calls, workshops, seminar series, and conferences. Overall, 270 training hours were completed during these events.

The OLCF hosted two kick-off workshops for members of the new CAAR partnership projects during the spring of 2015. The events' 79 participants included core developers of the new applications, OLCF staff, and members of the IBM/NVIDIA Center of Excellence at ORNL. The workshops were designed to help the project teams begin preparations to ensure their scientific applications are ready to make effective use of Summit when it enters production.

The OLCF partnered with other institutions in 2015 to offer four GPU hackathons, where programmers from around the world gathered to gain application GPU portability expertise straight from the experts. The OLCF hosted the final two hackathons locally. The 2015 events included a diverse range of institutions and science disciplines. Following the formula for success, each of the nine scientific programming teams was matched with mentors from vendors IBM, Cray, The Portland Group, and NVIDIA. The lessons learned from this year's hackathons will aid programmers for years to come as they adapt continuously to newer, more diverse architectures.

The OLCF also maintains a broad program of collaborations, internships, and fellowships for young researchers. In 2015, the OLCF supported 29 faculty, student interns, and postdoctoral researchers.

OLCF staff members mentor and develop these students to become a part of the nation's next generation of scientists, technicians, and engineers through a variety of challenging projects. For example, Marcela Crosariol, who interned for the UAO group, developed tutorials to assist users in accelerating their program for both Titan and Summit, and Kevin Song developed methods for ensuring reliability on supercomputers as they approach exascale.

Additionally, the UAO group again offered science-writing internships to two students in 2015. They were tasked to research and write about OLCF technical innovations and the scientific accomplishments of OLCF users.

This is a part of the team's regular workflow resulting in feature packages, web content, and OLCF publications.

151+

OUTREACH PRODUCTS

Total outreach highlights **72**

- **23** Science
- **19** Technology
- **30** People

298
Total media pick-ups

241
Total OLCF group tours

WEBSITE TRAFFIC STATISTICS

109,569
Total visitors

- **61.6%** New
- **38.4%** Returning

173,943
Overall hits

332,321
Page views

TOP 5 TRAINING EVENTS

Event	Participants
BigNeuron Data Analytics and Visualization Design Meeting	**126**
2015 OLCF Users Meeting	**106**
2015 GPU Hackathon	**72**
Portability Meeting	**70**
2016 INCITE Proposal Writing Webinar	**60**

High Impact Publications List

The OLCF has identified 340 unique publications and 29 high impact publications for 2015. These selected high-impact publications are listed below. For a complete list of all OLCF publications, please go to: https://www.olcf.ornl.gov/leadership-science/publications/.

1. Rosenberg, J.; Li, C.; et al. (2015), A Laboratory Study of Asymmetric Magnetic Reconnection in Strongly Driven Plasmas, Nature Communications, Volume: 6 DOI: 10.1038/ncomms7190

2. Jesse, S.; He, Q.; et al. (2015), Atomic-Level Sculpting of Crystalline Oxides: Toward Bulk Nanofabrication with Single Atomic Plane Precision, Small, Volume: 11 Issue: 44 DOI: 10.1002/smll.201502048

3. Kalinin, S.; Sumpter, B.; et al. (2015), Big-Deep-Smart Data in Imaging for Guiding Materials Design, Nature Materials, Volume: 14 Issue: 10 DOI: 10.1038/NMAT4395

4. Hagen, G. (2015), Charge, Neutron, and Weak Size of the Atomic Nucleus, Nature Physics, DOI: 10.1038/nphys3529

5. Mitchener, M.; Hermanson, D.; et al. (2015), Competition and Allostery Govern Substrate Selectivity of Cyclooxygenase-2, Proceedings of the National Academy of Sciences of the United States of America, Volume: 112 Issue: 40 DOI: 10.1073/pnas.1507307112

6. Frieler, K.; Clark, P.; et al. (2015), Consistent Evidence of Increasing Antarctic Accumulation with Warming, Nature Climate Change, Volume: 5 Issue: 4 DOI: 10.1038/nclimate2574

7. Pronk, S.; Lindahl, E.; et al. (2015), Coupled Diffusion in Lipid Bilayers upon Close Approach, Journal of the American Chemical Society, Volume: 137 Issue: 2 DOI: 10.1021/ja508803d

8. Zhang, P.; Cohen, R.; et al. (2015), Effects of Electron Correlations on Transport Properties of Iron at Earth's Core Conditions, Nature, Volume: 517 Issue: 7536 DOI: 10.1038/nature14090

9. Gui, Z.; Wang, L.; et al. (2015), Electronic Properties of Electrical Vortices in Ferroelectric Nanocomposites from Large-Scale Ab Initio Computations, Nano Letters, Volume: 15 Issue: 5 DOI: 10.1021/acs.nanolett.5b00307

10. Mosier, A.; Li, Z.; et al. (2015), Elevated Temperature Alters Proteomic Responses of Individual Organisms within a Biofilm Community, ISME Journal, Volume: 9 Issue: 1 DOI: 10.1038/ismej.2014.113

11. Norman, M.; Reynolds, D.; et al. (2015), Fully Coupled Simulation of Cosmic Reionization. I. Numerical Methods and Tests, Astrophysical Journal Supplement Series, Volume: 216 Issue: 1 DOI: 10.1088/0067-0049/216/1/16

12. Gong, Y.; de Jong, W.; et al. (2015), Gas Phase Uranyl Activation: Formation of a Uranium Nitrosyl Complex from Uranyl Azide, Journal of the American Chemical Society, Volume: 137 Issue: 18 DOI: 10.1021/jacs.5b02420

13. Hu, A.; Levis, S.et al. (2015), Impact of Solar Panels on Global Climate, Nature Climate Change, DOI: 10.1038/nclimate2843

14. Yoon, J.; Wang, S.; et al. (2015), Increasing Water Cycle Extremes in California and in Relation to ENSO Cycle under Global Warming, Nature Communications, Volume: 6 DOI: 10.1038/ncomms9657

15. Ma, J.; Wang, Z.; et al. (2015), Interplay Between Plasmon and Single-Particle Excitations in a Metal Nanocluster, Nature Communications, Volume: 6 DOI: 10.1038/ncomms10107

16. Roth, N.; Kasen, D. (2015), Monte Carlo Radiation-Hydrodynamics with Implicit Methods, Astrophysical Journal Supplement Series, Volume: 217 Issue: 1 DOI: 10.1088/0067-0049/217/1/9

17. Nickles, J.; Cheng, X.; et al. (2015), Mechanical Properties of Nanoscopic Lipid Domains, Journal of the American Chemical Society, Volume: 137 Issue: 50 DOI: 10.1021/jacs.5b08894

18. Zhang, Y.; Zherebetskyy, D.; et al. (2015), Molecular Oxygen Induced in-Gap States in PbS Quantum Dots, ACS Nano, Volume: 9 Issue: 10 DOI: 10.1021/acsnano.5b04677

19. Ma, J.; Wang, L. (2015), Nanoscale Charge Localization Induced by Random Orientations of Organic Molecules in Hybrid Perovskite CH3NH3PbI3, Nano Letters, Volume: 15 Issue: 1 DOI: 10.1021/nl503494y

20. Li, C.; Hong, J.; et al. (2015), Orbitally Driven Giant Phonon Anharmonicity in SnSe, Nature Physics, Volume: 11 Issue: 12 DOI: 10.1038/NPHYS3492

21. Drummond, N.; Monserrat, B.; et al. (2015), Quantum Monte Carlo study of the Phase Diagram of Solid Molecular Hydrogen at Extreme Pressures, Nature Communications, Volume: 6 DOI: 10.1038/ncomms8794

22. Bhimanapati, G.; Lin, Z.; et al. (2015), Recent Advances in Two-Dimensional Materials beyond Graphene, ACS Nano, Volume: 9 Issue: 12 DOI: 10.1021/acsnano.5b05556

23. Shakun, J.; Clark, P.; et al. (2015), Regional and Global Forcing of Glacier Retreat During the Last Deglaciation, Nature Communications, Volume: 6 DOI: 10.1038/ncomms9059

24. LeBlanc, J.; Antipov, A.; et al. (2015), Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms, Physical Review X, Volume: 5 Issue: 4 DOI: 10.1103/PhysRevX.5.041041

25. Broennimann, S.; Fischer, A.; et al. (2015), Southward Shift of the Northern Tropical Belt from 1945 to 1980, Nature Geoscience, Volume: 8 Issue: 12 DOI: 10.1038/NGEO2568

26. Yang, S.; Ramirez-Cuesta, A.; et al. (2015), Supramolecular Binding and Separation of Hydrocarbons within a Functionalized Porous Metal-Organic Framework, Nature Chemistry, Volume: 7 Issue: 2 DOI: 10.1038/nchem.2114

27. Heitmann, K.; Frontiere, N.; et al. (2015), The Q Continuum Simulation: Harnessing the Power of GPU Accelerated Supercomputers, Astrophysical Journal Supplement Series, Volume: 219 Issue: 2 DOI: 10.1088/0067-0049/219/2/34

28. Nicholl, R.; Conley, H.; et al. (2015), The Effect of Intrinsic Crumpling on the Mechanics of Free-Standing Graphene, Nature Communications, Volume: 6 DOI: 10.1038/ncomms9789

29. Akyuz, N; Georgieva, E.; et al. (2015), Transport Domain Unlocking Sets the Uptake Rate of an Aspartate Transporter, Nature, Volume: 518 Issue: 7537 DOI: 10.1038/nature14158

Acronyms

ACME	Accelerated Climate Modeling for Energy
ALCC	ASCR Leadership Computing Challenge
ASCR	Advanced Scientific Computing Research
BESC	US Department of Energy's BioEnergy Science Center
CA	certificate authority
CAAR	Center for Accelerated Application Readiness
CADES	Compute and Data Environment for Science
CORAL	Collaboration of Oak Ridge, Argonne, and Lawrence Livermore National Laboratories
CPU	central processing unit
DD	Director's Discretionary
DDN	DataDirect Networks
DOE	US Department of Energy
EVEREST	Exploratory Visualization Environment for Research in Science and Technology
GE	General Electric
GPU	graphics processing unit
HACC	Hardware/Hybrid Accelerated Cosmology Code
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

LANL	Los Alamos National Laboratory
Micro-CT	microcomputed tomography
MUI	multiscale universal interface
NCCS	National Center for Computational Sciences
NCSI	National Strategic Computing Initiative
OLCF	Oak Ridge Leadership Computing Facility
ORNL	Oak Ridge National Laboratory
OSG	Open Science Grid
OST	object storage target
OUG	OLCF User Group
PI	principal investigator
RAIT	Redundant Arrays of Independent Tape
SCA	sickle cell anemia
SciComp	Scientific Computing Group
SciDAC	Scientific Discovery through Advanced Computing
TechInt	Technology Integration Group
THF	tetrahydrofuran
UAO	User Assistance and Outreach Group