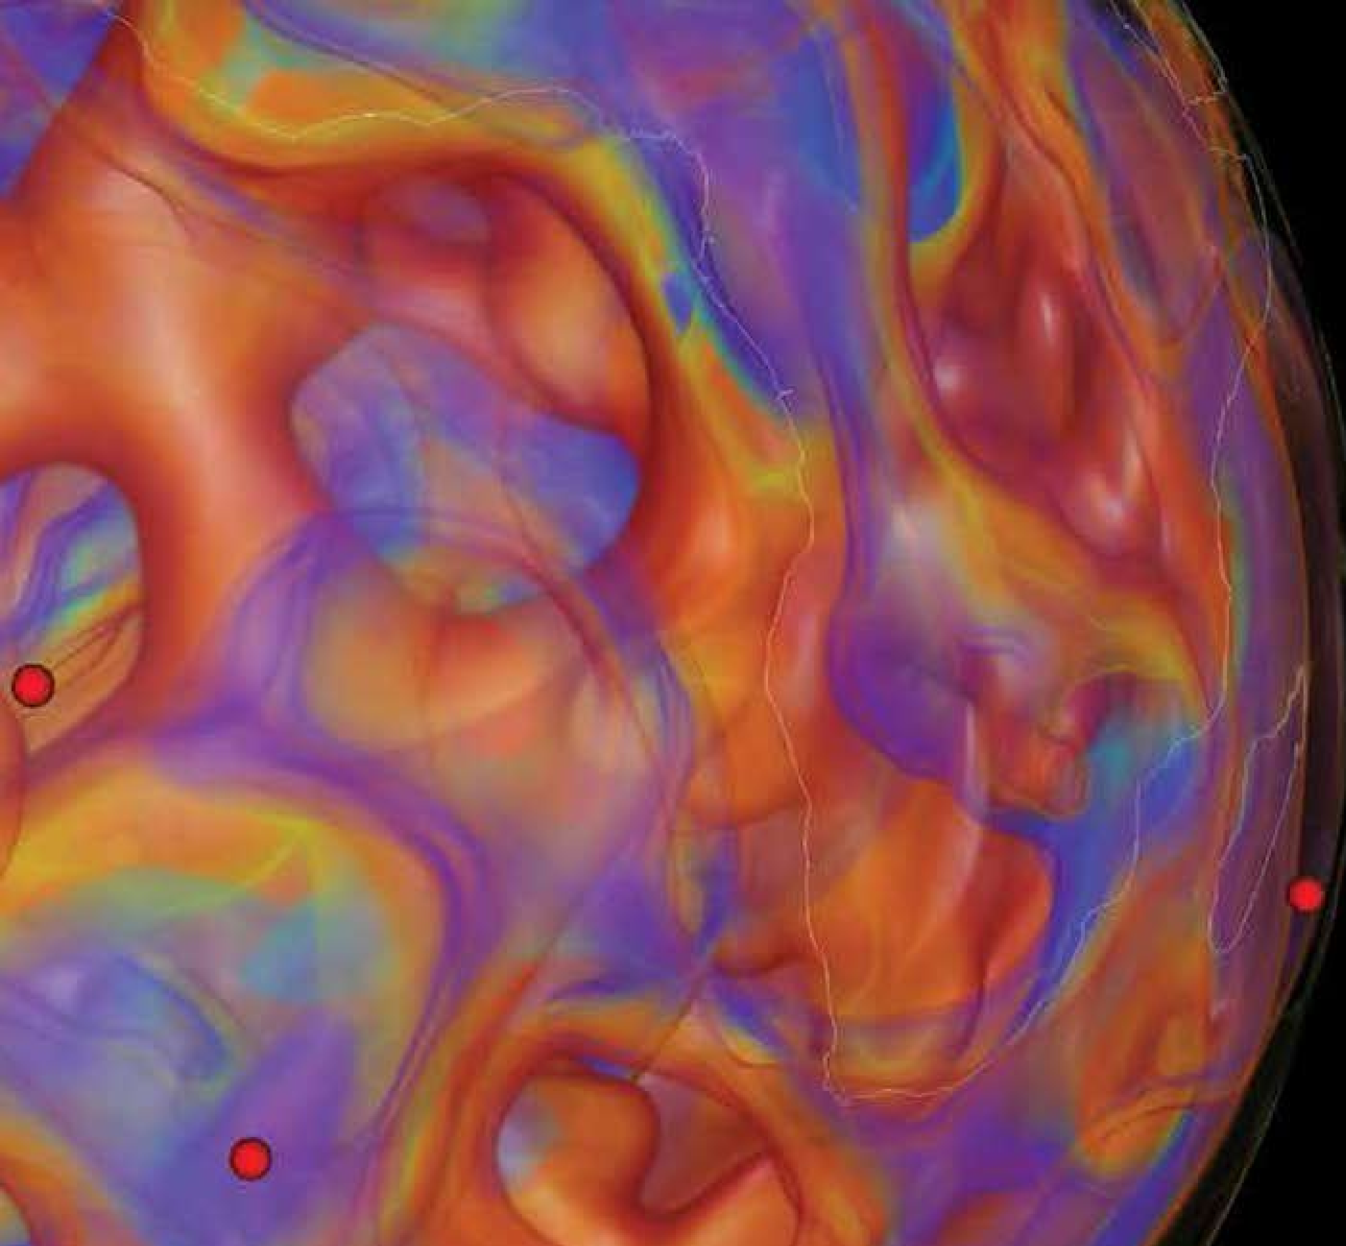

Oak Ridge Leadership Computing Facility

Annual Report 2013–2014

OAK RIDGE
National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

Contents

Oak Ridge Leadership Computing Facility

Annual Report 2013–2014

About the Cover:
A group of researchers led by Professor Jeroen Tromp of Princeton University has used adjoint tomography, or the interactions of forward and reverse wavefields, to reveal the Earth's inner workings. With Titan the team hopes to map the interior of the entire globe to a depth of 2,900 kilometers. Image courtesy of Dave Pugmire, ORNL.

The research and activities described in this report were performed using resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC0500OR22725.

Titan Pays Off

Titan pays off in scientific discoveries, engineering innovations, industrial competitiveness

James Hack
Director, National Center for Computational Sciences

Titan debuted at the end of 2012. Deploying this computing colossus was like climbing K2. Having reached the summit, it's natural to look back at challenges overcome in the delivery of America's fastest supercomputer, capable of 27 quadrillion calculations per second (petaflops). It's also exciting to enjoy the spectacular panorama of potential scientific discoveries, engineering innovations, and industrial competitiveness that accelerated computing makes possible.

The path to Titan was uncharted, bumpy, and risky. But demands to extract more parallelism within the same power and space envelope drove the Oak Ridge Leadership Computing Facility experts to architect a novel, energy-efficient, high- performance computing system for the scientific community. This new hybrid architecture would include both graphics processing units and conventional central processing units. Building and operating the hybrid system would challenge Oak Ridge Leadership Computing Facility (OLCF) staff, partnering vendors, and users alike. Nonetheless, a rolling upgrade of Jaguar's 200 cabinets was executed with as little disruption to users as possible, transforming that supercomputer into Titan. Once the upgrade was complete the GPU/CPU hybrid demonstrated 10 times the performance on standard benchmarks in the same footprint with only a small increase in power, making it far more energy-efficient than its CPU-only predecessor.

The lack of mature tools and training platforms for hybrid systems initially made the transition challenging for users, who nevertheless embraced the challenge of converting scientific application codes in anticipation of the reward of faster time to solution, or more complex representations of their scientific problem. Dedicated OLCF staff members worked during early stages with "power" users to prepare accelerated scientific codes for research communities. Through the OLCF's Center for Accelerated Application Readiness, they improved performance of codes simulating combustion, nuclear physics, molecular dynamics, advanced materials, and climate. Some early-science results emerging from this effort include demystifying the mechanisms by which liquid-crystal thin films rupture and characterizing the efficiency of organic photovoltaic materials for solar energy. These are difficult science problems. Enabling major computational applications provided a suite of exemplars that demonstrated to our broader usercommunity the benefits of investing in this new class of computer architecture. OLCF staff members have continued to develop tools, workshops, and guides for a growing population of users.

Able to calculate in an hour what would take five centuries or more on an average personal computer, Titan has taken the scientific community by storm, running models of supernovas, hurricanes, biofuels, clean combustion, and much more. It has allowed researchers to consider more realistic abstractions, and with new tools to implement these models into a predictive means of solving large, complex scientific and engineering problems.

Like other experimental devices, supercomputers produce data, and big problems generate big data. The speed of scientific progress increasingly depends on capturing, moving, sharing, integrating, analyzing, curating, and preserving data. And the volume, variety, and velocity of that data only grow.

Typically simulation experiments are configured to look for specific phenomenological behavior, where the experiment is conducted and the results are published. But there are times where there are other insightful features of a simulation that are not obvious to standard diagnostic techniques and therefore overlooked. Other examples are unexpected phenomena or mechanisms that have previously been unexplored theoretically or experimentally. . In 2014, for example, Nature published an observation of stellar shock-wave dynamics that verified the results of a simulation done at the OLCF a decade ago. This highlights the importance of curating and archiving the products of these state-of- the-art calculations for later examination in the context of ongoing scientific investigations and the development of more elaborate diagnostic methods.

As part of accommodating big data, in 2013 we deployed Spider II, a parallel file system of 32 petabytes and an input/output bandwidth of more than 1 terabyte per second—three times larger and four times faster than its predecessor. Twenty six million books, equivalent to 10 terabytes, would fit in less than one percent of Spider II. Being able to keep data in its entirety allows for more efficient workflows and analytics. Indeed, prioritized data services and sophisticated data policies can't come soon enough, as data has joined experiment, theory, and simulation as the fourth paradigm of the discovery process.

In 2013 a former user of Oak Ridge supercomputers, Martin Karplus, shared the Nobel Prize in chemistry for developing multiscale models for complex chemical systems. This is important recognition of how computational research has become an essential element of scientific investigation. Karplus' body of work is an example of setting the stage for what we now do routinely in our virtual exploration of scientific phenomena.

The Department of Energy (DOE) continues to provide investigators like Karplus with resources needed to advance their science. That's what facilities like the OLCF are all about. The OLCF provides access to these unique capabilities through three user programs: Innovative and Novel Computational Impact on Theory and Experiment (INCITE), the Office of Advanced Scientific Computing Research Leadership Computing Challenge (ALCC), and Director's Discretion (DD). In addition to our more traditional academic and government users, the OLCF's industrial partnership program helps companies like Ramgen, SmarTruck, and General Electric (GE) supercharge their competitiveness with access to Titan computational resources through the user programs.

Growing evidence shows scientists are making the most of Titan. Four of six finalists for the 2013 Gordon Bell Prize, awarded for the novel application of high-performance computing to science, engineering, and large-scale data analysis, used Titan. In 2013 the Adaptable I/O System, developed with ORNL leadership for efficient analysis of big data such as that generated on Titan, was named one of R&D Magazine's Top 100 Technologies. In 2013 the use of OLCF resources resulted in 346 peer-reviewed publications. And Titan user Masako Yamada of GE Global Research, who studies non-icing surfaces for cold-climate wind turbines, made HPCwire's People to Watch list in 2014. You can read about many of our other users, of whom we're exceptionally proud, in this report. Their accomplishments range from simulation of earthquakes to the discovery of a molecular switch with implications for drug design.

We should look at Titan as a preview of future exascale platforms, machines that are 40 times faster than Titan. But before that is the OLCF's next system, Summit, slated to arrive in the 2017–18 time frame. At 150 petaflops or more, it promises fresh challenges and rewards for the user community. As with Titan, we'll work with our vendors to tailor the architecture and programming environment to meet the needs of the scientific community. That architecture will embody many of the abstractions required to fully exploit Titan's potential: data locality, fine-grained parallelism, and power and cooling innovations. We'll also work with users to provide tools—compilers, algorithms, performance analyzers and the like— and training. We believe our experience with Titan will make this next transition smooth for users. In the meantime we are all looking forward to the scientific progress that will be enabled in the coming years by novel exploitation of the Titan architecture.

Choong-Seock Chang,
Princeton Plasma Physics Laboratory
2013 INCITE Project
Allocation: 100M core-hours

High-fidelity Simulation of
Tokamak Edge Plasma Transport

The Bleeding "Edge" of Fusion Research

Few problems have vexed physicists like fusion, the process by which stars fuel themselves and researchers hope to create the energy source of the future.

By heating hydrogen isotopes tritium and deuterium to more than five times the temperature of the Sun's surface, scientists create a reaction that could eventually produce electricity. However, confining the engine of a star to a manmade vessel and using it to produce energy is a tricky business.

Now, a team led by Princeton Plasma Physics Laboratory's C.S. Chang has found a solution in Titan, increasing the performance of its fusion XGC1 code fourfold.

"In nature, there are two types of physics," said Chang. The first is equilibrium, in which changes happen in a "closed" world toward a static state, making the calculations comparatively simple. Unfortunately, plasma physics falls in the second category, in which a system has inputs and outputs that constantly drive the system to a nonequilibrium state, which Chang refers to as an "open" world.

Most magnetic fusion research is centered on a tokamak, a donut-shaped vessel that magnetically confines the extremely hot and fragile plasma. But when it comes to studying the plasma edge confined in the tokamak, which is critical to understanding the plasma as a whole, Chang said "the effort hasn't been very successful because of its complexity."

Because the plasma is constantly contacting the vessel wall and losing mass and energy, which in turn introduces neutral particles back into the plasma, equilibrium physics generally don't apply and simulating the environment is difficult using conventional computational fluid dynamics.

Another reason the simulations are so complex is their multiscale nature. The distance scales involved range from millimeters (what's going on among the gyrating particles and turbulence eddies inside the plasma itself) to meters (looking at the entire vessel that contains the plasma). The time scales are also complex, evolving from microseconds in particle motions and turbulence fluctuations to milliseconds and seconds in its full evolution.

Chang's team is shedding light on a little-understood phenomenon known as "blobby" turbulence in which formations of strong plasma density fluctuations or clumps flow together and move around large amounts of edge plasma, greatly affecting edge and core performance in the DIII-D tokamak at General Atomics in San Diego, Ca. DIII-D-based simulations are considered a critical stepping-stone for the full-scale, first principles of the ITER plasma edge. ITER, a billion-dollar, multi-national effort, is a tokamak reactor to be built in France to test the science feasibility of fusion energy.

Blobby turbulence was discovered more than 10 years ago, and is one of the "most important things in understanding edge physics," said Chang, adding that people have tried to model it using fluids (i.e., equilibrium physics quantities). However, because the plasma inhabits an open world, it requires first-principles, ab-initio simulations.

Now, researchers have verified the existence and modeled the behavior of these blobs for the first time using a gyrokinetic code (or one that uses fundamental plasma kinetic equations with analytic treatment of the fast gyrating particle motions) and the DIII-D geometry.

The team has also developed a technique on Titan that allows them to simulate electron physics approximately 10 times faster than on Jaguar, Titan's predecessor, enabling the researchers to model physics such as electron-scale turbulence, which was out of reach as little as a year ago. Eventually, researchers plan on simulating the full volume plasma with electron-scale turbulence to understand how blobby turbulence affects the fusion core.—*Gregory Scott Jones*

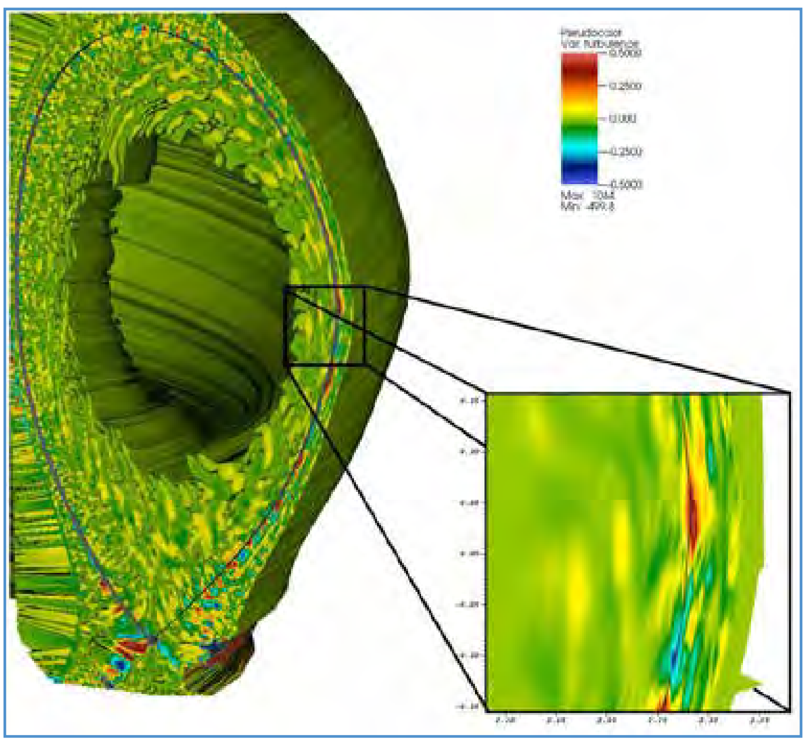

This visualization shows the turbulence front from the plasma edge being spread inward in multiscale interaction with the evolving background profile under the central heat source. Eventually, the whole volume becomes turbulent, with the spatial turbulence amplitude distribution being just enough to produce the outward heat transport to expel the centrally deposited heat to the edge. The edge turbulence source is continuously fed by the heat flux from the core. This is how the plasma profile, the heat source and the turbulence self-organize.Image Credit: Dave Pugmire, ORNL.

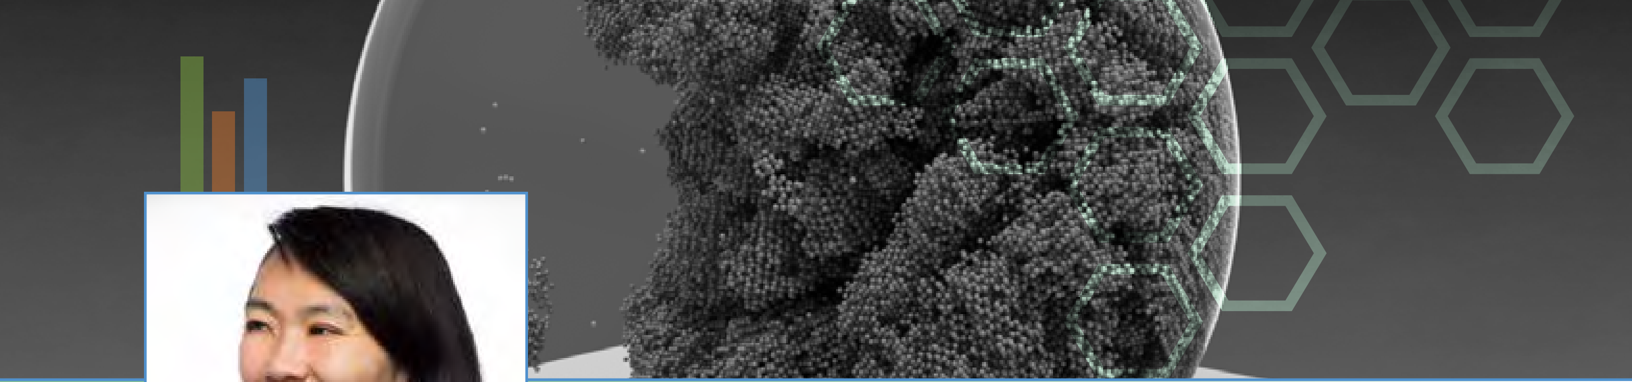

Masaka Yamada,
GE Global Research
2011-2014 ALCC Project
Allocation: 80M core-hours

Non-icing Surfaces
for Cold Climate Wind Turbines

Titan Propels GE Wind Turbine Research into New Territory

The amount of global electricity supplied by wind, the world's fastest growing energy source, is expected to increase from 2.5 to as much as 12 percent by 2020, but for some regions, going wherever the wind market blows, is not always practice.

"Many of the regions in the world that could benefit from wind energy are actually in cold climates," said Masako Yamada, a General Electric (GE) Global Research computational scientist. "If you have ice on a turbine, it reduces the efficiency of energy generation."

Rather than invest in turbines that might freeze at high latitudes or altitudes, these regions simply don't adopt wind power. But for a company like GE, wind turbines represent a lot of potential in a global market that could attract almost 100 billion dollars of investments by 2017.

To recruit cold regions to wind energy, GE researchers—including Yamada, principal investigator, Azar Alizadeh, co-principal investigator, and Brandon Moore, formerly of GE and now at Sandia National Laboratories—are using Titan to simulate hundreds of millions of water molecules freezing in slow motion to better understand the physical changes during ice formation on icephobic surfaces.

Since 2005, GE's anti-ice experimental research group has developed and tested a variety of engineered surfaces under freezing conditions.

"What we found experimentally is that surfaces that are great icephobic surfaces at one temperature may not be as effective at a different temperature," Yamada said.

Researchers realized something is happening on the atomic level that affects how ice forms on various surfaces at different temperatures. However, even in the best experimental setup researchers cannot observe individual water molecules freezing, only the general region where ice starts forming.

To pinpoint the time and place water molecules nucleate (or bud into ice), GE needed to enhance two things through simulation: the size of the system (individual molecules rather than droplets) and the increments of time (femtoseconds, or quadrillions of a second) in order to calculate their dynamics.

"We needed a million molecules per droplet to even be able to start guiding the experimentalists," Yamada said.

However, because molecular dynamics is a computational technique that takes a long time, most freezing simulations are still in the 1,000-molecule range. The team not only needed million-molecule droplets, they needed hundreds of them to simulate freezing across a surface.

The company looked to the DOE for help. Through an ALCC award, the team received 40 million core-hours, first on Jaguar then Titan. Moving to Titan, Yamada partnered with the OLCF's Mike Brown to make the most of Titan's GPUs by repurposing the LAMMPS molecular dynamics application, which was originally written for CPUs.

Through LAMMPS modifications, as well as changes to a model that simulates water molecules and Titan's powerful hybrid CPU-GPU architecture, Yamada's simulations accelerated at least 200 times. The simulations also increased from three different surfaces and three different temperatures to six surfaces and five temperatures, totaling just under 350 simulations— due to the acceleration, Yamada's team was able to carry out eight times as many simulations as originally projected.

What Yamada is able to take back to her experimental colleagues is a suite of conclusions about which surfaces are favorable under certain temperatures.

"I think I can go back and say that for temperatures between minus 10 and minus 20, this is the type of solution you want. For minus 30 to minus 40, try this type of solution," Yamada said.

By accumulating a lot of virtual ice on Titan, GE is reducing the number of time-consuming and costly physical experiments in their goal to open up cold climates to renewable power.—*Katie Elyce Jones*

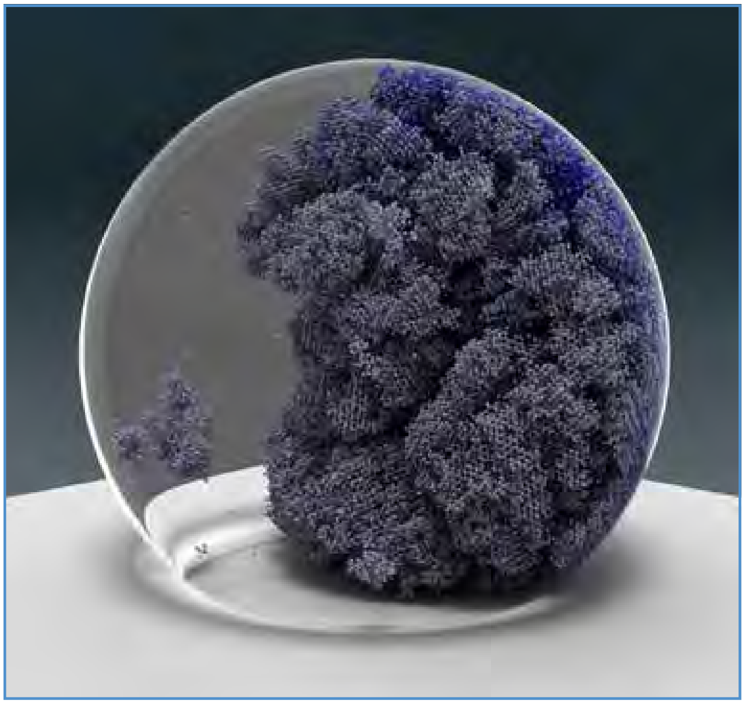

GE simulated hundreds of water droplets, each including one million molecules. Simulations accelerated at least 200 times over pre-GPU estimates permitting GE to study the nucleation of individual ice molecules. Image credit: Mike Matheson, ORNL.

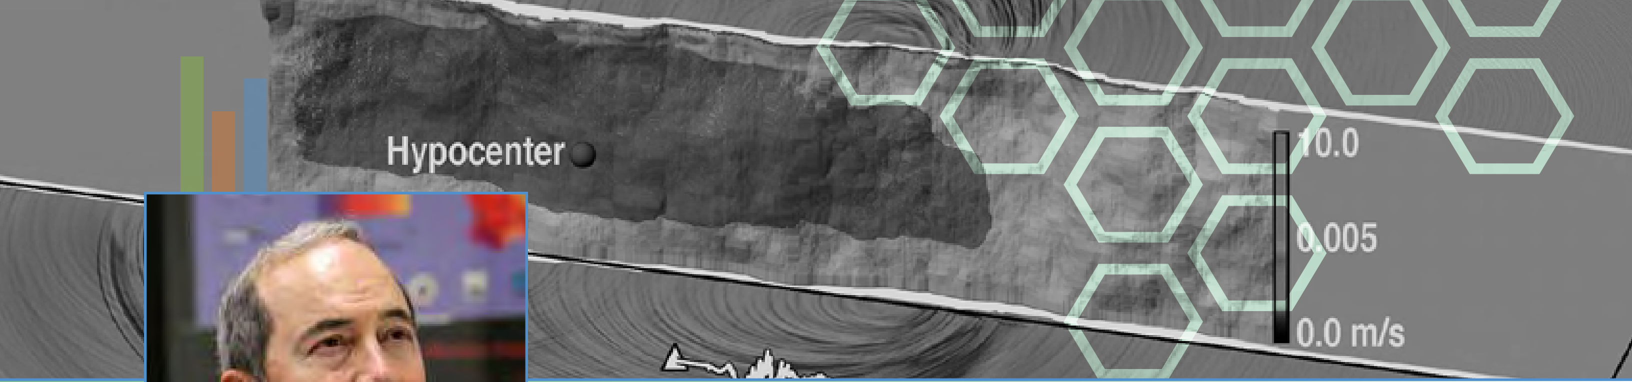

Thomas Jordan,
University of Southern California
2013 INCITE Project
Allocation: 45M core-hours

CyberShake3.0: Physics-based Probabilistic Seismic Hazard Analysis

Titan Simulates Earthquake Physics Necessary for Safer Building Design

When the last massive earthquake shook the San Andreas Fault in 1906, no one would hear about "plate tectonics" for 50 years, and the Richter scale was still a generation away. Needless to say, by today's standards, only primitive data survive to help engineers prepare southern California for an earthquake of similar magnitude.

"We could face an earthquake of 7.5-magnitude or bigger in the future," said Thomas Jordan, Southern California Earthquake Center (SCEC) director. "But the data accumulated from smaller earthquakes in southern California over the course of the last century is insufficient to predict the shaking associated with such large events."

To prepare for the next "big one," SCEC joint researchers—including computational scientist Yifeng Cui of the University of California, San Diego, and geophysicist Kim Olsen of San Diego State University—are simulating earthquakes at high frequencies on Titan. The team's physics-based model calculates wave propagations and ground motions radiating from the San Andreas Fault through a 3-D model approximating the Earth's crust.

Seismic wave frequency, measured in hertz (cycles per second), is important to engineers who are designing buildings, bridges, and other infrastructure to withstand earthquake damage. Low-frequency waves, which cycle less than once per second (1 hertz), are easier to model, and engineers have largely been able to build in preparation for the damage caused by this kind of shaking.

"It's mostly the big structures like highway overpasses and high-rises that are sensitive to low-frequency shaking," Olsen said. "But smaller structures like single-family homes are sensitive to higher frequencies, even up to 10 hertz."

High-frequency waves (in the 2–10 hertz range) are computationally more daunting because they move much faster through the ground, and there has been little information to give engineers on shaking up to 10 hertz.

In the case of the SCEC's Titan simulations, the perameters defining the ground were extremely detailed: representing a chunk of terrain one-fifth the size of California (including a depth of 41 kilometers) at a spatial resolution of 20 meters. The ground models also included detailed 3-D structural variations—both larger features such as sedimentary basins as well as small-scale variations on the order of tens of meters.

All in all, the simulation totaled 443 billion grid points. At every point, 28 variables, such as wave velocity and stress, were calculated.

Back in 2010, SCEC simulations of an 8-magnitude earthquake along the San Andreas Fault on Jaguar peaked at 2 hertz. At the time the Jaguar simulations were conducted, doubling wave frequency would have required a 16-fold increase in computational power.

But on Titan in 2013, the team ran simulations of a 7.2-magnitude earthquake up to their goal of 10 hertz, which can better inform performance-based building design. By modifying their code originally designed for CPUs—the Anelastic Wave Propagation by Olsen, Steven Day of San Diego State University, and Cui, known as the AWP-ODC—for GPUs, they significantly improved speed up. The simulations ran 5.2 times faster than they would have on a comparable CPU machine without GPU accelerators.

And considering the "big one" could shake California anytime in the next few decades to the next few years, accelerating our understanding of the potential damage is crucial to SCEC researchers.—*Katie Elyce Jones*

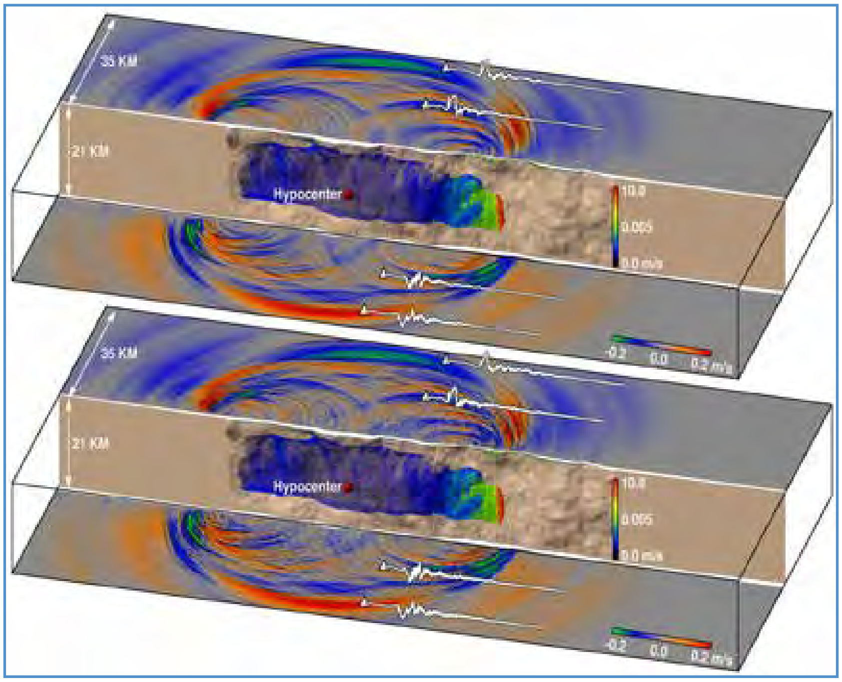

From an unprecedented study on GPUs of high-frequency earthquake shaking: snapshots of 10-Hz rupture propagation (slip rate) and surface wavefield (strike-parallel component) for a crustal model without (top) and with a statistical model of small-scale heterogeneities (bottom). Image courtesy Amit Chourasia, San Diego Supercomputer Center.

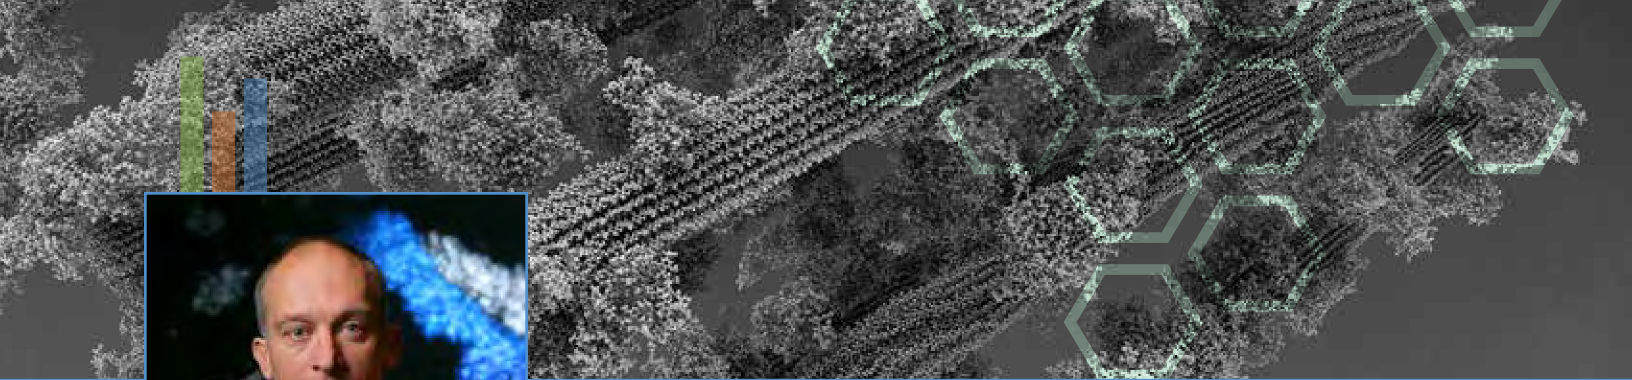

Boosting Bioenergy

Jeremy Smith,
Oak Ridge National Laboratory
and the University of Tennessee
2013 INCITE Project
Allocation: 78M core-hours

Cellulosic Ethanol: Physical Basis of Recalcitrance to Hydrolysis of Lignocellulosic Biomass

A team led by Oak Ridge National Laboratory's Jeremy Smith, the director of ORNL's Center for Molecular Biophysics and a Governor's Chair at the University of Tennessee, has uncovered information that could help others harvest energy from plant mass. The team's conclusion—that less ordered cellulose fibers bind less lignin—was published in the August 2013 edition of *Biomacromolecules*.

The team used simulations on the Oak Ridge Leadership Computing Facility's Jaguar supercomputer—a 2.3-petaflop machine that in 2012 morphed into Titan, America's fastest open-science supercomputer—along with neutron scattering to seek ways to make ethanol as cheap as gasoline.

"We are trying to figure out how to effectively break down plant materials like grass or wood chips cheaply enough to make biofuels economically viable," said Loukas Petridis, a researcher in the Biosciences Division. "We are investigating the two main features that make biomass recalcitrant, or resistant to breakdown—the presence of lignin and the tightly ordered structure of cellulose."

All plants contain a sticky molecule called lignin that intertwines with cellulose and hemicellulose in their cell walls. Lignin provides strength to the stalks of plants so that these organisms can stand. But during biofuel production—a process that converts plant mass into alcohol—lignin blocks enzymes from breaking down cellulose into the sugars necessary for fermentation.

Neutron imaging and supercomputer simulation allow scientists to resolve the structure of lignin aggregates down to 1 angstrom, about 1 million times smaller than what the naked eye can see.

Using neutron beams at ORNL's High Flux Isotope Reactor, researchers have discovered that cellulose fibers that are less organized, or noncrystalline, are easier for enzymes to break down. Simulations run on Jaguar helped explain this phenomenon.

"Jaguar has shown us that not only is noncrystalline cellulose more easily broken down, but it also associates less with lignin," Petridis said. "This was the first simulation that has looked at the interaction of lignin with specific types of cellulose."

Smith's team was awarded 23 million core-hours on Jaguar in 2012 through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program.

The team ran a classic molecular dynamics simulation with a code called GROMACS. The code monitored 3 million atoms and used 30,000 of Jaguar's cores, which means that each core was responsible for 100 atoms. The application provides information on the interaction of water with the cellulose, degree of lignin aggregation, shapes of lignin molecules, diffusion constants, etc.

The researchers received another 78 million INCITE hours in 2013 to run an adapted version of GROMACS that can take advantage of Titan's hybrid nodes, making the updated application 10 times bigger and much faster than the one run on Jaguar.

Current simulations track around 30 million atoms, which include crystalline and noncrystalline cellulose, lignin, enzymes, and water molecules. The models also account for the long-range interactions that lignin molecules experience with other surrounding molecules.

The team is now studying how lignin behaves in different types of biomass, which will help identify the plant characteristics best suited for biofuel production.

"The scientific insights we gain help improve the biofuel production process," Petridis said. "We give engineers hints about how the production process works, which we hope will allow them to design new pretreatment methods and engineer different types of biomass and enzymes that can harvest more energy from plant materials."
—Jennifer Brouner

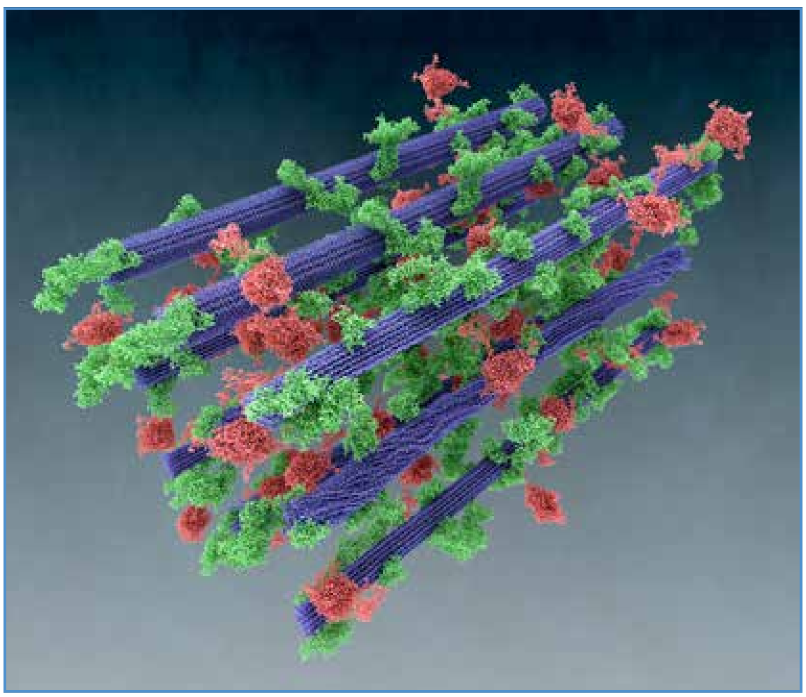

Interaction between cellulose fibers (blue), lignin polymers (green), and cellulase enzymes (pink). Visualization by Michael Matheson (ORNL)

Jeroen Tromp,
Princeton University
2013 INCITE Project
Allocation: 100M core-hours

Global Seismic Tomography based on Spectral-Element and Adjoint Methods

Investigating the Earth's Inner Workings

When it comes to the Earth's interior, researchers have only "scratched the surface." To better understand our planet and better prepare for tomorrow's calamities, researchers need a more detailed picture of our Earth's anatomy, specifically the mantle. This means imaging about halfway down to the Earth's center.

Jeroen Tromp of Princeton University is part of a team using Oak Ridge National Laboratory's (ORNL's) Titan supercomputer, to reveal the Earth's inner workings via adjoint tomography simulations, or inversion analysis of seismic waves. "We want to image the physical state of the Earth's interior," he said.

The research will provide a better idea of what might happen when earthquakes occur. This information is extremely important to city planners and structural engineers in metropolitan areas aligned with fault zones.

Much as medical tomography uses waves to image sections of the human body, seismic imaging uses waves generated by earthquakes to reveal the structure of the Earth. By mapping the speeds of waves generated after earthquakes, a process known as seismic tomography, Tromp's team can arrive at better estimates of temperature and rock composition for the Earth's interior.

Individual earthquakes are captured by several thousand seismographic instruments, each recording three-component data with time series typically 2 hours long. These data are then used to simulate seismic wave propagation in three dimensions, and the resulting "synthetic seismograms" are compared to observed data. The differences between the synthetic and observed seismograms help the team improve its models for the Earth's mantle, said Tromp.

With the help of Titan, an approach known as "full waveform inversion" can achieve full three-dimensional (3D) images of the earth mantle. "The number of calculations involved in doing these operations is so overwhelming it can only be done on systems like Titan," Tromp added.

The team made a preliminary trip to the Oak Ridge Leadership Computing Facility in December 2012 and visited with Judy Hill, a liaison in the Scientific Computing Group. Hill helped the group compile its SPECFEM3D_GLOBE spectral element wave propagation code and took the researchers through the basics of running on Titan. The team is now running simultaneous "ensemble" simulations that will further improve its model of the Earth beneath our feet.

In traditional seismology, every seismogram is one file, and each earthquake involves thousands of seismograms, with the team eventually studying thousands of earthquakes. Tromp calls the amount of I/O "phenomenal." Opening such a large number of files could hurt Titan's file system.

Enter ADIOS, a parallel I/O library used to describe data in code that may need to be written, read, or processed outside of the running simulation. The middleware, developed at ORNL with partners from academia, was chosen to create a new data format for this project to aggregate a large amount of information into large files and improve the I/O performance.

The team's SPECFEM3D_GLOBE simulation now uses ADIOS to write its output for checkpointing and visualization. Furthermore, the OLCF has developed a parallel data reader for SPECFEM3D_GLOBE simulation output that is being used inside VisIt, a tool for parallel, interactive visualization and analysis of scientific data.

Using the new ADIOS data reader, the OLCF has produced visualizations of these recent simulation runs, giving Tromp's team yet another major advantage in its quest to find out exactly what goes on in the planet's underworld.
—*Gregory Scott Jones*

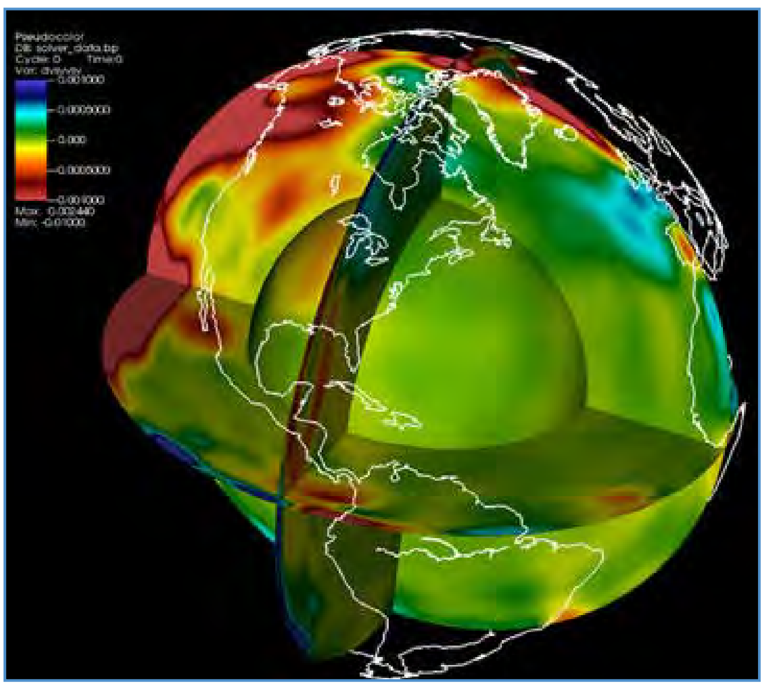

Researchers led by Jeroen Tromp are using Titan to map the speeds of waves generated after earthquakes. This process, known as seismic tomography, provides the team with unprecedented temperature and composition (type of rock) estimates for the Earth's interior. Image Credit: Dave Pugmire, ORNL

Michael Klein,
Temple University and
Russell Devane, Procter & Gamble
2013 INCITE Project
Allocation: 65M core-hours

Advanced Modeling of the Human Skin Barrier

ORNL Supercomputer Turns the Tide for Consumer Products Research

Consumer products giant Procter & Gamble (P&G) generated more than $85 billion in revenue in 2013 from the global sale of everyday products like lotions, shampoos, toothpaste, and soaps, whose success in competitive markets depends on outstanding product performance. To maintain brand dominance and generate new market leaders, P&G researchers need a deeper understanding of their product formulations. With colleagues at Temple University, they turned to ORNL's Jaguar supercomputer (since upgraded to become Titan) in an attempt to model and simulate microstructural arrangements of molecules that dictate the flow, thickness, performance, and stability of P&G products comprised of systems of fat-soluble molecules called lipids.

"For P&G the motivation is to understand how products that are based on lipid systems perform and age," explained P&G research scientist Russell DeVane.

Temple's Michael Klein, a National Academy of Sciences member who has collaborated with P&G for 15 years, led the simulation project. In addition to DeVane, Klein's research partners include Temple's Giacomo Fiorin and Axel Kohlmeyer and P&G's Peter Koenig and Kelly Anderson.

Lipid vesicles are relatively large aggregates that form when lipid bilayers "roll" into spheres. In manufactured products such as lotions and fabric softeners, vesicles encapsulate and suspend perfumes, dyes, and active ingredients. Over time vesicles can fuse with each other, causing phases to separate and otherwise threatening consistent product performance.

Changes in the viscosity, or thickness, of a product can destabilize it. For many lipid-based products, viscosity largely depends on the size and quality of the vesicles. Due to their greater surface curvature, small vesicles suffer more mechanical stress than do large vesicles and over time fuse with other vesicles to grow and relieve that stress. As these systems reorganize, viscosity can shift outside the desired range.

Moreover, problems can pop up when perfumes or dyes are added. The addition of a lemon-fresh scent, for example, may alter the molecular structure of a formulation and turn a gel into a liquid, or vice versa. It takes time and money to design, manufacture, and test a formulation, only to find a flaw that forces an ingredient's removal, replacement, or concentration readjustment.

If P&G researchers could use simulations to predict performance attributes before a product is manufactured, they would be able to deliver high-quality products faster than competitors and earn greater revenues in the marketplace.

“We needed extremely long simulations of extremely big systems,” DeVane said. “The Oak Ridge Leadership Computing Facility is one of the very few places where we could accomplish that.”

Through the INCITE program, the team used 69 million core hours on Jaguar over 2 years, setting the stage for simulations of large, complex systems of lipid assemblies. To date no one has ever been able to simulate vesicle fusion in realistic systems.

With Jaguar’s computational power, Klein and collaborators succeeded in artificially triggering the deformations of a single vesicle that happen before it fuses with another vesicle. At Temple, where researchers share an 800-core cluster, the same simulation would have taken 5 or more years.

The first-time accomplishment of simulating one vesicle contorting sets the stage for simulations of two vesicles fusing and eventually systems of vesicles large enough to show how the fusion process plays out in aging products.—*Dawn Levy*

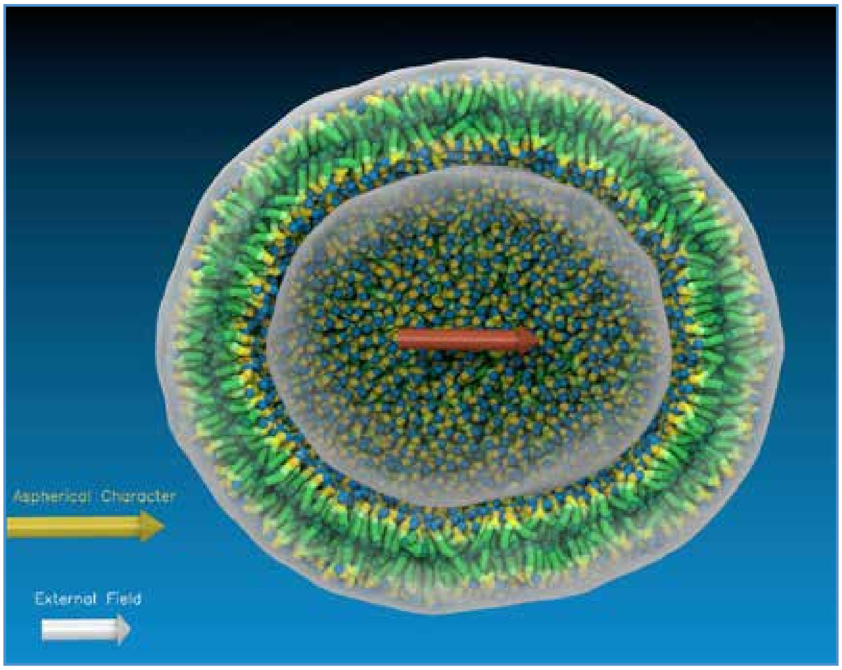

This simulation depicts a slice through a liposome subjected to an oscillating electric field—an artificial trigger of behavior that precedes vesicle fusion. This baseline provides a critical first step to simulating the fusion of two or more vesicles, a process that is both a bane (it can jeopardize lipid-containing shampoos and detergents) and a boon (it enables transport of biological agents to and through cells). Image credit: Axel Kohlmeyer

OLCF Early Science Projects

At the OLCF, select projects are identified for the early science program. The early science phase consists of a select group of research teams granted early access to Titan's GPUs in an effort to quickly produce science results after a system is delivered. Below are brief summaries of four of the 2013 early science projects.

Titan Sheds Light on Unknowns in Organic Photovoltaic Research

Researchers from the OLCF and ORNL's Center for Nanophase Materials Science used their Early Science project to pursue a promising technology for generating solar energy.

Organic photovoltaic cells are not as efficient as inorganic (typically silicon) cells, but they are less expensive and easier to produce. You can, in fact, roll them out in flexible sheets with the same technology used for printers and place them on a variety of surfaces. The ORNL team hopes eventually to double the performance of these cells, to 20 percent efficiency.

Using Titan, team members simulated the energy-producing interaction between two molecules blended together. The first, a long, semi-flexible chain of the P3HT polymer, absorbs sunlight and transforms into an excited state (known as an exiton). The second, a tiny, clustered molecule called PCBM, splits the exciton into an electron and a hole (which has the opposite charge as an electron).

Previous simulations were limited by the size of system they were able to tackle. As a result, the team's understanding of the phase segregation between P3HT and PCBM was compromised.

The ORNL team, led by Jan-Michael Carrillo, optimized the popular LAMMPS molecular dynamics code to take advantage of Titan's GPUs. As a result, they simulated millions of molecules and reached an equilibrium in which the system had been completely rendered. It was no longer dependent on run time. Team members believe the two molecules can be configured to improve efficiency.

Magnetic Simulations of Nickel Scale to 14.5 Petaflops on Titan

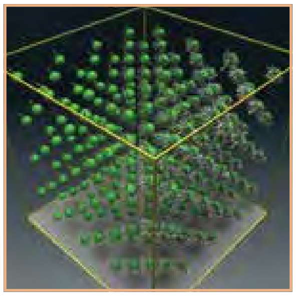

A team of researchers led by ORNL's Markus Eisenbach is using the power of Titan to calculate the temperature at which nickel loses its magnetism.

The problem is both challenging and important. Magnetism is indispensable in a wide variety of technologies, from electric motors to disc drives. The better understanding and reliably simulated magnetic properties of materials will produce numerous benefits.

Specifically, the team is calculating the material's Curie temperature. This is the temperature above which the magnetic directions of specific atoms (known as atomic moments) go from pointing in the same direction to pointing randomly. It uses an application known as WL-LSMS, which calculates the magnetic properties of materials from basic laws of physics rather than from models that must incorporate approximations.

The team did an impressive job calculating the Curie temperature of iron on Titan's predecessor, Jaguar, reaching 1.84 petaflops. With WL-LSMS, the team calculated a value of 980 degrees Kelvin. The experimentally determined value was 1,050 degrees Kelvin.

Nickel, however, presents a much more complicated problem. Not only must the application calculate the direction of each magnetic moment; it must also calculate the size of the moments. Such a challenge had to wait for Titan and advances in the algorithm (Replica-Exchange Wang-Landau) that make it even more suitable for parallel computers.

In fact, WL-LSMS scaled up to 14.5 petaflops on Titan by utilizing the hybrid nodes. With these advances the calculations for magnetism are ongoing and results are expected during 2014.

Titan Helps Explain Neutron Behavior in Reactor Startups

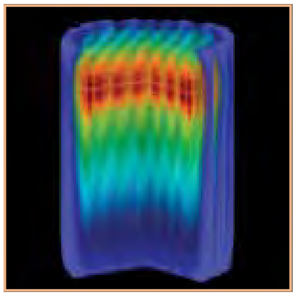

Researchers with the Consortium for Advanced Simulation of Light Water Reactors used their Early Science allocation to simulate the startup conditions of a next-generation nuclear power reactor.

The project applied three separate applications to the same problem: the behavior of neutrons after the reactor has achieved criticality but before it is powered up. The project is providing greater detail into neutron behavior inside the reactor as well as helping researchers better understand the tradeoff between an application's computational demands and its accuracy and reliability.

The three codes are

- Denovo, which models the reactor as a fine mesh of rectangles, or cells. Denovo scales to the entire Titan system.
- Shift, a Monte Carlo transport code that models the reactor geometries exactly, without dividing them into a mesh. Shift can use all of Titan's 299,008 CPU cores, but it has not yet been optimized to use the system's GPUs.
- SPN, a coarser neutron transport application that typically uses about 1,000 processor cores. The payoff for SPN's coarseness is that it can run on smaller clusters and does not require the supercomputing muscle of a system like Titan.

The applications are modeling the startup conditions of a Westinghouse AP1000, one of a new group of reactors known as generation III+. The design of generation III+ fuel rods is more complex than that of earlier models, which creates a challenge for computational simulation.

A team representing Westinghouse and CASL received an International Data Corporation HPC Innovation Excellence Award for applied simulation on Titan. The award recognizes achievements made by industry users of high-performance computing technologies and was presented to awardees on June 24 at the International Supercomputing Conference in Leipzig, Germany.

Understanding Why Liquid-Crystal Thin Films Disintegrate

Thin films made from liquid crystals—the mojo behind today's flat-panel electronic displays—may provide the might behind tomorrow's nanoscale coatings, optical and photovoltaic devices, biosensors, and more. First, however, we must better understand why they rupture.

Researchers from ORNL used Titan supercomputer to enhance that understanding, testing two theories behind the process known as "dewetting." The first points to heat-related movement of atoms in the film, while the other focuses on small undulations on the thin film's surface that grow in amplitude over time.

The Titan project, led by Trung Nguyen and W. Michael Brown of ORNL, verified that these processes both take place— but that one dominates over the other depending on the thickness of the film. Their contributions to the discussion were made possible by Titan, which allowed them to simulate the system with unprecedented scale; in fact, the team's simulation is the first to study liquid-crystal thin films at experimental length- and timescales. "No one has ever simulated such big films for such a long time," Nguyen said. "Another first is to relate all the dewetting process to the molecular-level driving force."

The work appeared as the cover story in the March 21, 2014, print edition of *Nanoscale*, a high-impact journal of the Royal Society of Chemistry, and was also published online.

OLCF Gordon Bell Finalists

The real measure of a scientific supercomputer, such as Titan, is its ability to model the physical world at unprecedented size, scale, and detail. Four of six finalists for the most recent Gordon Bell Prize, awarded during November's SC13 conference in Denver, ran on Titan.

Those four finalists are discussed below.

Superconductor Simulation Tops 15 Petaflops

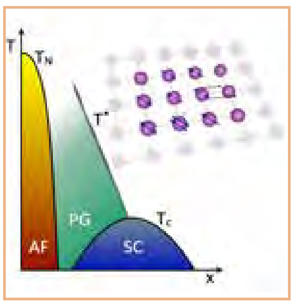

Researchers from ORNL and Switzerland's ETH Zurich were nominated for overcoming two major roadblocks to realistic superconductor modeling, while topping 15 petaflops on Titan.

By conducting electricity without resistance and generating an especially powerful magnetic field, superconducting materials show enormous promise in technologies as diverse as power transmission and MRI scanners. Their use is limited, however, because they must be kept very cold. The discovery—or creation—of superconductors that needn't be cooled would revolutionize power transmission and the energy economy.

The Swiss and American team approached the problem with DCA++, an application that describes the behavior of electrons in a solid using a quantum Monte Carlo technique, which involves repeated random sampling.

Unfortunately, every solution contains its own problems; two inherent in the DCA++ approach are known as the fermionic sign problem and the cluster shape dependency.

The sign problem points to changing signs (i.e., plus or minus) in the wave functions of electrons. It limits the size of an atom cluster that can be modeled and, therefore, the accuracy of the solution that can be calculated. The shape dependency acknowledges that the answer given for a cluster of atoms can depend on the shape of the cluster.

The team approached these problems with a new algorithm called DCA+, developed by Peter Staar at ETH. This algorithm gets to a solution nearly 2 billion times faster than its predecessor, allowing the team to model realistically large atom clusters and overcome the shape dependency altogether.

Peter Staar, ETH Zurich, doctoral student of Thomas Schulthess, ETH Zurich

Plasma Jet Simulations Tell Us about Faraway Matter

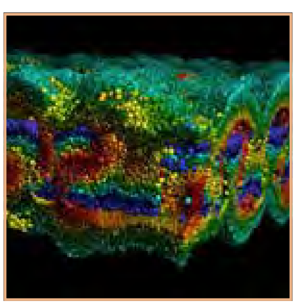

A team from Germany's HZDR–Dresden earned a nomination for simulating the passing plasma jets emitted by faraway objects such as stars and galactic nuclei (the swirling matter surrounding galactic black holes).

The team used Titan to model a property of turbulent plasma flow called the relativistic Kelvin-Helmholtz instability (KHI), which occurs where passing plasma jets collide. By focusing on this instability, the researchers were able to make out patterns of individual particle behavior.

Radiation produced by the KHI can be observed by telescope, so the team used observed radiative signatures to correlate plasma dynamics with radiation emitted during turbulence. This approach required the team to perform computationally demanding kinetic simulations, which follow individual particles, rather than a more traditional hydrodynamics approach.

The team used a particle-in-cell code to model two streams of unmagnetized hydrogen plasma totaling 75 billion particles per simulation, including protons and electrons.

"To our knowledge, this simulation on Titan is 46 times larger, and the spatial resolution is 4.2 times higher than the largest kinetic KHI simulation to date," said HZDR–Dresden

Computational Radiation Physics group leader Michael Bussmann.

Ultimately the KHI tells physicists about the properties of passing plasma jets through comparison: Is one plasma stream denser than the other? What are their velocities? In what directions are they traveling?

With the data from Titan, researchers can begin to apply the results to actual plasma jets. And understanding plasma jet dynamics could reveal information about their objects of origin, such as active galactic nuclei.

Michael Bussman, HZDR Dresden Computational Radiation Physics Group Leader

Biology Simulations Peer into Cells One GPU at a Time

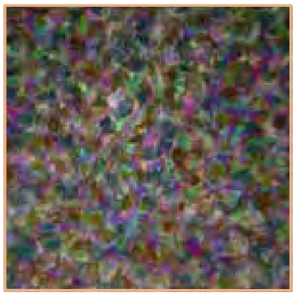

A simulation of the internal workings of cells reached a sustained 20 petaflops, earning a nomination for researchers from the Italy's Consiglio Nazionale delle Ricerche (CNR), or National Research Council.

"We are simulating the crowded protein solution that is representative of our cell compartments," said Simone Melchionna from CNR's Institute for Chemical and Physical Processes. "If we want to target a specific cell for treatment, we need to understand how that cell works."

Scientists have predicted the activity of isolated proteins, but nature isn't an isolated system. The CNR team is interested in studying the cell interior, which is packed with proteins—an environment that can't be studied in a laboratory using dilute protein solutions.

Proteins perform a vast number of functions in living organisms, from catalyzing reactions to replicating DNA to transporting molecules. In living cells proteins interact with other proteins and with surrounding fluids by changing their shapes, movements, and behaviors, sometimes dramatically. These dynamic interactions complicate studies that examine biological systems.

Knowing the general pattern for protein movement and interaction could help scientists understand specific diseases like Alzheimer's, which results in a clumping of proteins in the brain.

The research team developed a code called MUPHY, for MUlti PHYsics simulator, to study the two-way interactions that exist between proteins and fluids within a cell. The code used 18,000 of Titan's 18,688 nodes, reaching a peak performance of almost 27.5 petaflops by using a mixed-precision calculation.

Massimo Bernaschi, Chief Technology Officer at CNR's Institute for Applied Computing

Cosmologists Replicate the Evolution of the Universe

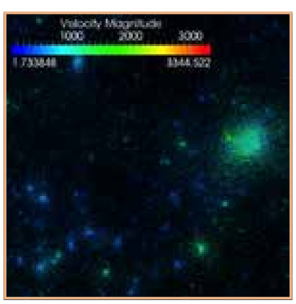

A team from Argonne National Laboratory (ANL) took Titan to a peak of 25 petaflops and sustained performance estimated at 10 petaflops as it modeled the multibillion-year development of the universe. Its efforts resulted in a sky catalog to which scientists can compare instrument observations.

The project is aimed at helping physicists corroborate observational data with their understanding of the fundamental processes that govern the structure of the universe.

In order to tackle this project, the team developed an application known as HACC, for Hardware/Hybrid Accelerated Cosmology Code. Simulations began with an extremely dense and very uniform universe, which expanded over thousands of time steps into a complex cosmic web that shows a detailed clustering of mass across large distances.

While HACC tracks trillions of mass points, these points don't correspond to specific physical objects. Instead, they are "tracers," which represent conglomerations of mass. The smallest clump distinguishable in the code is 100 billion solar masses (roughly equivalent to the Milky Way galaxy), and the code resolves distances from kiloparsecs to gigaparsecs—from galaxy-scale distances to the entire swath of the observable universe.

HACC is the first large-scale cosmology code that can run on a hybrid CPU/GPU supercomputer, as well as on multicore or many-core architectures.

Salman Habib, project leader and high-energy physicist and computational scientist at Argonne National Laboratory

Big Risks, Big Rewards

In April 2009 when the OLCF decided to upgrade its CPU-based Jaguar supercomputer to the hybrid Titan, the organization understood full well that the road ahead would be full of obstacles.

It would be difficult, no doubt, yet necessary if the OLCF was to continue to meet the needs of our users while making energy efficiency a top priority.

Titan, like any other disruptive technology, carried with it a risk-versus-reward ratio. The risks were well known and substantial: Would the GPUs work, and would programmers take the time to port their codes to an unknown technology? Could all of the parts be delivered on time and on budget? After all, a machine the scale of Titan has to be built from multiple manufacturers located literally around the world.

The reward, however, was equally large—a computer ten times faster than Jaguar that consumed only moderately more power. When completed Titan would likely be the fastest computer in the world and, more importantly, the most powerful tool for scientific simulation the world had ever seen. But first it had to be built.

In recognition of the numerous risks inherent in standing up a machine of Titan's stature, the OLCF set out very early to develop a number of contingency plans for each potential obstacle.

Despite all of the OLCF's foresight and risk mitigation strategies, Titan's journey to becoming the world's fastest computer would not be without its challenges.

For instance, the Center could never have predicted what would cause the boards to become unstable.

It was also impossible to predict that Tropical Storm Nock-ten would strike Thailand, a country that happens to manufacture a disproportionate number of the world's disk drives, setting in motion a series of events that would ultimately delay the upgrade of Titan's Spider file system.

Yet despite the bumps in the road, the OLCF's foresight largely paid off. Titan was delivered on schedule and is now up and running with a robust user base of more than 150 project teams tackling everything from biofuels to climate change to astrophysics, all at resolutions only dreamed of just years ago.

But behind every interesting destination lies a far more interesting journey. As the Chinese philosopher Lao Tzu said, "A journey of a thousand miles begins with a single step."

Accelerating the State-of-the-science

Of all the foreseeable risks inherent in standing up Titan, one in particular loomed large above the rest: Titan's novel GPU technology. The potential was obvious, but whether leading scientific applications could be modified to take advantage of that potential was still unproven.

Historically, researchers moving from single- to many-core computers had to break their calculations into smaller problems that could be parceled out separately to different processing cores, an approach referred to as parallel computing and one taken by users of previous OLCF systems, including Titan's predecessor Jaguar.

Titan, however, was developed with the recognition that achieving even greater parallelism by adding hundreds of thousands of computing cores has its limits, primarily for one reason—CPUs are fast but very power hungry. If high-performance computing (HPC) systems are to move beyond the petascale and into the exascale, they must become more powerful in terms of calculations while maintaining or only marginally increasing their power consumption.

Accelerators, or in this case GPUs, can push parallelism even further by allowing researchers to continue to divide their larger problems into even smaller parcels than is customary on traditional supercomputers like Jaguar.

The reason is simple: GPUs have many threads of execution. While each one may run slower than traditional threads, there are simply so many threads that a much higher performance can be achieved—and with a minimal increase in power consumption.

Because they handle hundreds of calculations simultaneously, GPUs can perform many more than can CPUs in a given time. By relying on its 299,008 CPU cores to guide simulations and allowing its new NVIDIA GPUs to do the primary calculations, Titan would, in theory, enable researchers to run scientific calculations with greater speed and increased fidelity while maintaining the reasonable demand for power.

ORNL originally had the option of building an equivalent Cray system that did not incorporate GPUs, using those 18,688 sockets instead to hold additional 16-core CPUs (however, a CPU-only machine of this size would not have met OLCF mission needs for 2012-2013). Essentially, whereas each Titan node contains one NVIDIA Kepler GPU plus an AMD 16-core Opteron CPU, the other option would have simply replaced the GPU with another Opteron CPU. The resulting system would have contained nearly 600,000 processing cores, but it would nevertheless have paled in comparison to Titan in terms of speed while demanding a much greater supply of electricity.

Despite the simplicity of such an alternative system, the performance would not meet the Center's needs. Something had to be done to more efficiently push the envelope in HPC system design, and the OLCF was well positioned to lead the way.

More than 5 years after the idea of Titan was born, the machine is exceeding expectations as it hosts more than 600 users across numerous scientific domains. The GPUs, it seems, are paying off.

A greener machine

One of the primary reasons for incorporating GPUs was the quickly accelerating power demands required by CPU-only processors.

By simply adding processors, machines can continue to grow faster, but energy efficiency suffers.

At some point the benefits of high-performance computers are trumped by the energy they consume. Titan was the OLCF's effort to show the world that innovation can solve scientific computing's energy problem.

"It's a totally different machine [than Jaguar] in the same cabinet," said Jim Rogers, the OLCF's Director of Operations. "It's a powerful machine, and it's a power-efficient machine."

Jaguar's run, in its original configuration as a Cray XT3 supercomputer, came to an end at the close of 2011, and it was upgraded before Titan's debut. From October 2009 to September 2011, Jaguar averaged using 3.63 million kilowatt hours per month. That translates to an average of 5.04 megawatts.

Meanwhile, Titan came online on May 31, 2013. In the period from June 2013 to February 2014, the Cray XK7 supercomputer averaged 3.52 million kilowatt hours per month (or an average energy consumption of 4.89 megawatts).

In short, Titan was using an average of 3 percent less electricity to run its computations at tenfold the performance. Where the new machine gathered its extra bang for the buck was through the configuration of the nodes, Rogers said.

"The NVIDIA Keplers only use power when needed," Rogers said. "They're so sophisticated that they can see work coming in, schedule it, and power up immediately. They're ready to go in an instant. When they're near idle, they almost power themselves down—going to a low of 20 watts. The beautiful thing is they can switch between those states seamlessly . . . The net effect is we are actually using less energy to get a lot more done."

Rogers said his team is tracking stats to show just how efficient the GPUs are at performing their work. "They decrease the wall clock time to half in some cases and at least 25 to 30 percent in other cases.

"We're operating the same number of cores with less power on the same footprint," he said. "It's not just the processors that are producing this energy efficiency, but they are, by far, the dominant force."

In fact, Titan's hybrid architecture helped it to earn the number three on the Green500 list—with a performance more than 100 times greater than the number one system on the list—in November of 2012, the same month it was recognized as the world's most powerful system.

Photo courtesy of Tennessee Valley Authority

Organized by Virginia Tech's Wu-chun Feng and Kirk Cameron, the list takes the world's 500 most powerful supercomputers—as ranked by the Top500 list—and reorders them according to how many calculations they can get per watt of electricity.

"With Titan, ORNL provides an exponential increase in scientific computing power toward predictive capabilities to address complex challenges without significantly increasing energy consumption," noted Jeff Nichols, ORNL associate laboratory director for computing and computational sciences.

Courting the user community

Despite the potential of the GPUs to accelerate science while reigning in energy consumption, there was still an enormous problem; people had to learn to use them.

"Would anybody be willing to put in the time to port their code to the GPUs?" asked Bronson Messer of the OLCF's Scientific Computing Group, alluding to the early concerns over the transition of applications from CPUs to Titan's hybrid architecture. Titan was literally the first system of its kind at such a large scale available to a wide user base, and its final reception was unknown during the OLCF's early conceptual stages. Working with users was always a top priority.

The OLCF needed to demonstrate to its user base that the porting work would be worth it no matter where the individual users ended up running their simulations. In essence, they needed to ensure that any tweaks to an application to take advantage of the GPUs would likewise help, and certainly not hurt, its performance on the CPUs.

"The code changes made to applications so that they can take advantage of Titan's GPUs generally result in faster run times on CPUs as well, so researchers tweaking their codes for Titan will also, intentionally or not, optimize them for other systems that they might run on," said Messer, who headed up the OLCF's Center for Application Acceleration Readiness, charged with bringing users' codes up to spec on Titan. "The changes made are good on any other machine, not just Titan."

This fortunate relationship was critical to getting researchers on board to perform the enhancements necessary to port their codes. With the knowledge that the code enhancements would remain on nearly every other system of similar size, it was a win-win for research teams.

"We had to communicate really closely with our users, and we needed users to work closely with us to tell us about their struggles," said OLCF Director of Science Jack Wells.

The result: numerous applications now perform better than before, enabling unprecedented breakthrough potential across a range of computing platforms, both hybrid and CPU-based.

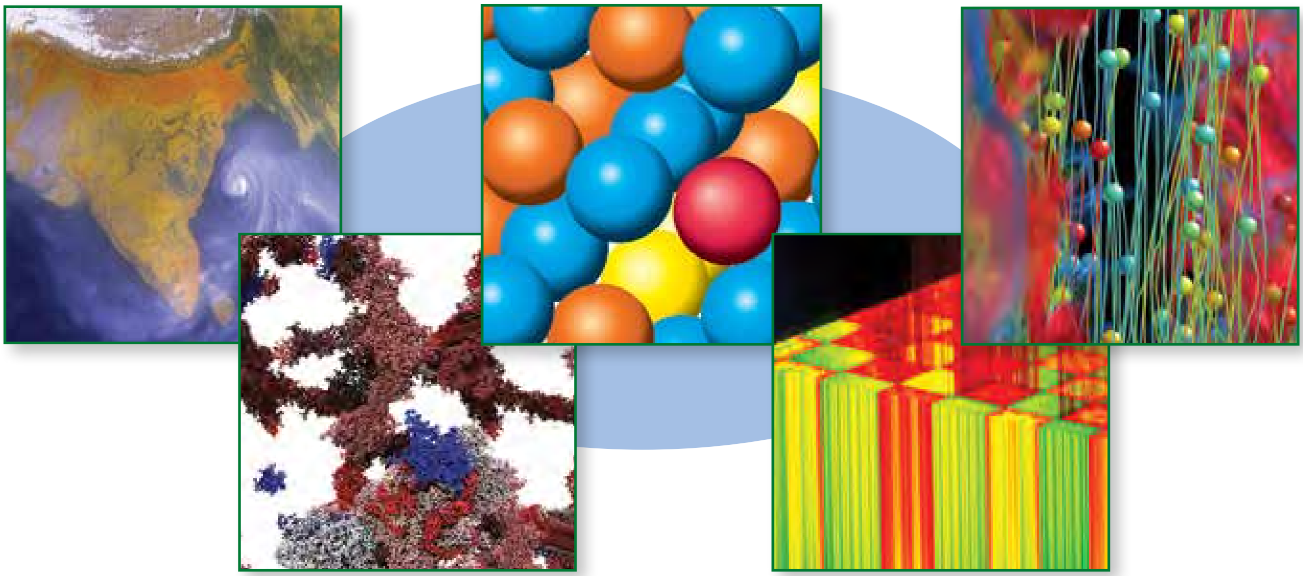

The importance of integration – CAAR

Because Titan's hardware is only as good as the research applications that exploit it, preparing users for Titan's unique architecture was a top priority.

Recognizing codes' hot spots, or places with maximum computational intensity, and manipulating them to respond to the GPUs were always top priority for the OLCF, but how to do so most effectively wasn't always as clear as the objective itself. Integrating Titan's hardware with users' applications was one of the greatest challenges facing the OLCF from the beginning.

Enter CAAR, a collaboration among application developers; Titan's manufacturer, Cray; GPU manufacturer NVIDIA; and the OLCF's scientific computing experts, to identify and port a select group of applications that showed GPU-ready potential.

Headed by the OLCF's Bronson Messer, CAAR worked for nearly 2 years to establish best practices for code-writers. CAAR was divided into small development teams, and each team worked exclusively with one of the OLCF's most advanced and representative applications.

CAAR included five applications chosen to quickly make use of Titan's GPUs and show that existing apps could be ported quickly and efficiently. These included the combustion code S3D; WL-LSMS, which studies magnetic systems; LAMMPS, a bioenergy and molecular dynamics application; Denovo, which investigates nuclear reactors; and CAM-SE, a code that explores climate change.

While each of the five codes realized dramatic improvements, one code in particular showed enormous early promise. LAMMPS, a molecular dynamics code that simulates the movement of atoms through time and is a powerful tool for research in biology, materials science, and nanotechnology, saw more than a sevenfold speedup on Titan for one problem's compute intensive interaction potential, as compared to its performance on the comparable CPU-only Titan system (before the GPUs were available).

WL-LSMS, another CAAR application, also proved to be an excellent match for Titan's hybrid architecture. It calculates the magnetic properties of a promising material, including its Curie temperature, from basic laws of physics rather than from models that must incorporate approximations.

WL-LSMS ran 3.8 times faster on the GPU-enabled Cray XK7 Titan than its XE6 CPU-only equivalent on a problem that consumed 18,600 of Titan's compute nodes, 88 short of the full machine. Equally impressive was the fact that even with this dramatic increase in performance, the GPU version of Titan consumed

7.3 times less energy than the CPU-only incarnation. This combination of accelerated performance and reduced energy consumption is precisely what the addition of the GPUs was intended to accomplish.

All of the CAAR applications achieved significant speedups using Titan's GPUs, paving the way for the rest of Titan's users, who are now beginning to ramp up their individual codes to scales never dreamed of just a few years ago, proving that not all changes are necessarily bad.

The accelerated performance of these leading applications was the direct result of years of collaboration and hard work among the OLCF and its user community. After all, codes are only as good as the people who write and use them.

And for that reason, bringing the scientific computing community up to speed on GPUs was of the utmost importance.

The importance of training

Because Titan was such a revolutionary technology, the OLCF knew early on that training its user base would be critical to its eventual success. Workshops, user conference calls, training events, and seminars have always been integral components of the OLCF user support model, but were more important than ever during the Titan transition.

This meant not only hosting its own workshops, but also working with other HPC centers to engage and educate the broader user community. In 2013 the OLCF training program focused on increasing overall engagement with users and the greater community by creating more opportunities for participation.

The OLCF itself hosted 30 events, with more than 750 participants. With Titan coming on line, user training focused primarily on programming and using accelerators. The OLCF collaborated with other HPC facilities with the goals of increasing efficiency by sharing resources, the quality of the training material by involving more experts, and user participation in training events.

Throughout the massive training effort at the OLCF, however, two workshops in particular stood out in their importance and engagement with the user community:

East/West Coast Titan Workshops: To minimize travel costs for its users, the OLCF conducted two similar training events in January and February 2013 in opposite geographic locations. The format was intended to help facilitate user access to training. It also served to increase collaboration with the OLCF accelerator vendor/partner NVIDIA, which graciously offered a classroom for the first event within its headquarters in San José, California. The second was held at a downtown Knoxville, Tennessee, hotel in response to user feedback that indicated difficulties in reaching the ORNL campus and additional transportation costs.

Both events had an almost identical curriculum, which included hands-on exercises on GPU programming, practicums on tools, and best practices developed by OLCF's staff and the CAAR team. As with all of our training events, both were broadcast live over the web and reached 176 people both remotely and at the two sites.

Processing and Analysis of Very Large Data Sets: In 2013 the OLCF broadened the training program to include data processing and analysis. It became apparent upon completion of the Application Requirements for Exascale report that users needed help with analyzing increased data sets being generated by Titan.

The workshop lasted three days and covered major aspects of data processing such as I/O, scalable data tools, and visualization. In addition to ORNL staff, the workshop also included speakers from Brookhaven National Laboratory (BNL), Lawrence Berkeley National Laboratory (LBNL), the National Energy Research Scientific Computing Center (NERSC), the University of Chicago, the Swiss National Supercomputing Centre, the University of Tennessee, and Kitware.

After all, the OLCF as a user-first facility must ensure that users can make the most of Titan and other OLCF systems must be a priority.

And that also means giving users the tools they need to make the most of their time on Titan.

The right tools for the job

While the engineering behind a system the size of Titan is impressive in its own right, its real value lies in its ability to enable scientific breakthroughs.

To further assist users in porting their codes and to get Titan breakthrough-ready as quickly as possible, the OLCF partnered with several outside vendors to ramp up the various tools necessary to get applications up to speed and ready to tackle the GPUs.

Because Titan's architecture had never been attempted at such a scale before, the community offered few mature tools that were out-of-the-box ready to help the application teams. The necessity for the proper tools meant the OLCF would have to form partnerships with vendors that would not only help to make Titan one of the world's leading computational research platforms but also help the larger HPC community.

Three tools in particular were identified as essential early on: compilers, debuggers, and performance monitoring tools. In fact, the OLCF has an entire group dedicated to forging and maintaining these partnerships and independently developing the tools necessary for Titan's entire user base.

"It's all motivated by the needs of our users," said Application Performance Tools Group Leader David Bernholdt. "At the end of the day, we are focused on ensuring that our users have the tools to program and debug effectively."

According to Bernholdt, all of these vendor partnerships are needed to accomplish three major goals for both the OLCF and its users: accelerate delivery of capabilities, enable developments that would be useful to the user community, and provide general enhancements aimed at the user base. In other words, vendors need users and users need vendors, and because Titan is only as good as its user base, these vendor relationships are critical to the machine's success.

And to ensure users are successful, the OLCF provides vendors with access to Titan to get to the root of any problems quickly, and staff members work closely and consult with external tools teams on a regular basis.

The result: a suite of tools that includes CAPS and GPU-manufacturer NVIDIA's work on OpenACC compilers, Allinea's DDT debugger, the Technische Universität Dresden's Vampir performance analysis tool, and compilers from Cray, PGI, and numerous other efforts, all of which enable users to get the most from their time on Titan. "We gave them [vendors] a whetstone to hone their tools against," said Messer.

Now code developers have access to an indispensable suite of tools at the ready to help researchers get up and running on Titan as quickly as possible, and to ensure that their time on Titan is as productive as possible.

A helping hand

While well-versed in porting and scaling applications for traditional CPU-based systems, the OLCF simply did not have the necessary tools to get users up and running on a unique technology such as Titan within the projected timeline.

Going it alone was not an option; the Center reached out to several third-party vendors to assist in the enormous task of getting Titan ready for users on day one. From debuggers to compilers to performance analysis tools, the OLCF searched far and wide for partners willing to go the extra mile in developing tools for the Center's user base.

"We weren't satisfied with single point of failure," said Wells. "We had backup plans. And in reality all of these tools get used. Some users use one, some use another, and some use more than one, but they are all being fully implemented . . . In the end it builds a richer ecosystem."

Take Allinea's DDT debugger, for instance. When an application containing hundreds of thousands of lines of code is running across 300,000 cores, spotting bugs is a tricky business. It's a case of sheer numbers and a problem that was anticipated. The Center knew it would have to create a revolutionary tool to allow applications to run smoothly after scaling to Titan.

"Sometimes there are problems that only occur at scale," said Bernholdt, "and we need tools like DDT to figure out what's going on."

For that reason OLCF staff began working with software developer Allinea on Titan's previous incarnation, known as Jaguar, and in preparation for Titan's launch. The product of this relationship is Allinea's distributed debugging tool, or Allinea DDT, created precisely for the world's leadership computing systems. In collaboration with the OLCF, Allinea was able to customize its large-scale debugger to Titan's hybrid architecture, enabling the supercomputer's first users to easily scale to large portions of the machine and assisting the OLCF during Titan's critical acceptance phase.

"Before we joined this project, tools weren't capable of getting anywhere near the size of the hardware," noted Allinea's COO David Lecomber. "The problem was that a debugging tool might do 5,000 or 10,000 parallel tasks if it was lucky, when the machines and applications wanted to write things that could do 200,000 plus. So the tools just got beaten up by the hardware."

This partnership has resulted in Allinea's DDT being among the most scalable debuggers in the HPC community. Furthermore, the OLCF helped to fund Allinea's effort to visualize data in arrays. This point-and-click approach allows users to simply click a mouse over any part of a visualization to pinpoint what areas of the code need work.

Without DDT, it would have been difficult, if not impossible, for applications to port to Titan in a reasonable timeframe. Catching problems at scale is critical for code performance on the Cray XK7.

Just ask ORNL's Markus Eisenbach. He works with an application known as WL-LSMS, which provides first-principles calculations of properties that are important for the understanding of materials such as steels, iron-nickel alloys, and advanced permanent magnets that will help drive future electric motors and generators. Titan is helping Eisenbach's team improve the calculations of a material's magnetic states as they vary by temperature.

Thanks to Allinea DDT, the transition from Jaguar to Titan was smoother than previously thought possible for Eisenbach's project. It didn't start out that way, however.

Whenever the team's application scaled to roughly 14,000 cores, the GPU-based version of WL-LSMS mysteriously crashed. The traditional CPU-based version, on the other hand, ran fine.

"It was puzzling," said Eisenbach, noting that the versions weren't noticeably different when it came to scaling. At such a scale, researchers like Eisenbach can't decipher all of the data in an application's core file, so it was difficult to pinpoint the error in the code.

DDT was instrumental in locating the errors and getting Eisenbach's code up to scale on Titan as quickly as possible.

Allinea representatives have worked with ORNL to extend Allinea DDT to a scale 40-plus times greater than previous high-end debugging tools. But debugging is only one piece of the porting puzzle. Compiling is, and always has been, a likewise critical component of porting applications. Typically, compilers convert an application's code to the binary language of computers; with the addition of Titan's GPUs, a difficult task immediately became more complex.

The OLCF recognized early on that compilers would play a huge role in getting applications up and running.

Based on the directive-based OpenACC standard and its predecessor, OpenHMPP, CAPS compilers enable developers to incrementally port their applications to various many-core systems such as NVIDIA and AMD GPUs and Intel Xeon Phi.

The OLCF and CAPS partnership "is pushing the language standard forward and piloting new capabilities that should be in the standard," said Bernholdt.

The same can be said for the development of various performance analysis tools, such as Vampir, developed to help researchers maximize their codes' efficiency on Titan.

Developed originally by the Technische Universität Dresden, the Vampir suite is a set of tools to investigate in detail the parallel performance of MPI, OpenMP, CUDA, and hybrid codes running at large scale. Although oriented toward tracing, Vampir allows profile-based output, filtering of information, and detailed analysis, handling large amounts of trace data using the parallel analysis of VampirServer.

Researchers usually don't begin their computations on Titan. They often start with scaled-down problems run on computing clusters at their home institutions or on partitions of supercomputers that provide access to hundreds of processors at most. They add detail to their models to achieve greater realism, and eventually they need to run the simulation on as many processors as possible to achieve the requisite resolution and complexity.

At the OLCF they may use a supercomputer with 100 times as many processors as they've used before, and when the performance of their code does not scale as expected, researchers need to know what went wrong and where.

Computer scientists have developed tools to help diagnose the millions of things that can go wrong when a program runs through its routines. In May 2012 a team from ORNL, Argonne National Laboratory, and the Technische Universität Dresden used Vampir to analyze a running application in detail as it executed on all 220,000 CPU processors of Titan's predecessor, Jaguar.

"Understanding code behavior at this new scale with Vampir is huge," said ORNL computer scientist Terry Jones, who works on system software for leadership-class supercomputers. Vampir traces events and translates the files into diagrams, charts, timelines, and statistics that help experts locate performance bottlenecks in codes. "For people who are trying to build up a fast leadership-class program, we've given them a very powerful new tool to trace events at full size because some things that happen at larger scales just don't happen at smaller scales," Jones said.

The research team first presented its findings in June 2012 at the High-Performance Parallel and Distributed Computing conference in Delft, Netherlands. The work was published online on June 20, 2013, in Cluster Computing. Before this research effort, Vampir could perform this kind of detailed analysis on a maximum of 86,400 cores. Large supercomputers use a version of MPI to communicate between processors. Vampir started as a tool for Visualization and Analysis of MPI Resources (hence the name) but today supports parallelization paradigms beyond MPI, such as OpenMP, Pthreads, and CUDA.

These vendor partnerships extend beyond the commercial community into the realm of open-source developers as well, said Bernholdt.

Other areas being explored include Open|Speedshop, an open source multi-platform Linux performance tool targeted to support performance analysis of applications running on numerous large-scale platforms; HPC Toolkit, an integrated suite of tools for measurement and analysis of program performance on computers ranging from multicore desktop systems to the nation's largest supercomputers; and Tau, a portable profiling and tracing toolkit for performance analysis of parallel programs written in Fortran, C, C++, UPC, Java, and Python, to name a few.

Bernholdt's group is also pursuing longer-term projects with the Open MPI community such as developing new implementations of collective communication primitives, working on fault tolerance, and understanding how applications can use these capabilities to enhance resilience.

"We provide access to machines for vendors so they can test products, and we have staff members who consult with external tools teams on a regular basis," said Bernholdt.

This relationship with vendors is a major reason behind Titan's success.

But like any disruptive technology, there is much more to Titan than just the machine. There is an entire computing ecosystem designed around Titan's revolutionary architecture, all of which had to be upgraded in order for users to get the most bang for their computing buck.

Support Systems

Titan requires a robust network of support systems and staff to reap the rewards offered by its immense computing power. If Titan was to be successful, portions of the support network would likewise need to be upgraded to ensure users got the most bang for their computing buck. Among the top priorities was Titan's aging file system, Spider.

Since 2008, Spider has been the primary production file system responsible for providing disk space for Titan's predecessor Jaguar, and now Titan itself, as well as other OLCF systems such as Eos, a 744-node Cray XC30; Rhea, a 196-node Linux cluster; and the data transfer nodes that provide connectivity to the High Performance Storage System (HPSS) and other institutions via the OLCF's connection to the 100 gigabit-per-second ESnet backbone network. Spider serves these and other systems by sharing data among OLCF systems with 20,000-plus compute nodes in aggregate.

Spider was a Center-wide shared resource that served all major OLCF platforms while providing the highest performance with simultaneous accessibility.

But as Jaguar transitioned to Titan, the I/O demands on Spider increased drastically, and an updated file system was the logical next step. As a result, the OLCF designed and implemented a new file system known as Spider II.

"We can now provide users with more than 1 terabyte per second file system performance at the top end," said Sarp Oral, the task lead for File and Storage Systems projects in the Technology Integration Group within the OLCF. "At that speed we expect Spider II to be safely in league with the top three parallel file systems in the world."

Assembling Spider II required an OLCF-wide effort. A diverse team that spanned several OLCF groups such as Technology Integration and High-Performance Computing Operations came together with ORNL procurement to work on hardware acquisition, deployment, installation, and software updates and integration.

Inside the four rows of cabinets is a clustered file system architecture, with I/O servers sitting in front of, and managing, high-end storage arrays, which in turn manage Spider II's 20,160 disks.

The cluster is connected to the OLCF systems via an FDR InfiniBand-based system area network providing up to 56 gigabits per second of connectivity to each node in the fabric. Every existing and future OLCF system will connect to this upgraded system area network, and data stored in Spider II will be instantly accessible to all.

With the new hardware and added space, the OLCF staff was able to increase the aggregate bandwidth from 240 gigabytes per second to over 1 terabyte per second, in addition to increasing the total storage capacity from 10 petabytes to 32 and the Dell I/O servers from 192 to 296. Despite these enormous increases in capacity, bandwidth, and server count, power and space utilization remained similar to that of the older Spider file system. Spider II's overall size is within 5 square feet of Spider, and the upgraded version has a power footprint that is only 20 kilowatts higher, a direct result of the Center's decision to deploy 208-volt power distribution networks and a cold-aisle containment solution to increase the efficiency of the data center.

"In essence," said Oral, "Spider II is at least four times more powerful in terms of I/O performance and provides three times more capacity than the previous version of Spider."

However, he noted, having top-of-the-line hardware is just one aspect of a powerful system; top-of-the-line software is the other.

"In addition to offering the newest and most capable hardware, we wanted Spider II to run the latest and greatest version of Lustre," said Oral. "We chose to deploy Lustre version 2.4 because it will improve our scalability and metadata performance and will also provide us with unique features which we have wanted and worked for quite some time to deploy for OLCF users."

Lustre 2.4 is the premier open-source software for high-performance computing and offers vast improvements and stability. The 2.4 release marks the first phase of the Distributed Namespace Environment (DNE); for the first time the ability to have multiple metadata servers is supported for a single file system.

"With DNE phase I, we will begin to deploy multiple metadata servers to start to relieve the pressure that a single metadata server has inflicted on Lustre users in the last decade," said Storage Team leader in the HPC Operations Group Jason Hill.

Spider II is the result of a Center-wide file system project that began in earnest in 2005. It was targeted to eliminate islands of data, allow gradual growth of I/O performance and capacity, make data available even if certain HPC systems (e.g. Jaguar, Titan) were down, and to improve reliability characteristics of production parallel file systems. The design and development took 3 years, and the first version, Spider, was deployed and available to OLCF systems and users in 2008.

This original deployment was viewed as risky, but with the success of over 5 years of production, deployment moving forward with Spider II was the obvious path forward, and after many months and a few more lessons, Spider II is now serving OLCF users.

"Spider II allows our parallel file system to keep pace with the newly-increased size and computational horsepower of Titan," said Messer. "The metadata improvement, in particular, enables our users to produce and analyze the kind of large, complex datasets we anticipate being produced on Titan. Spider II is both bigger and better." And that means bigger and better things for Titan's users.

But the OLCF didn't stop there. In order for users to truly harness the power of Titan, improving data management would likewise be key.

HPSS enhancements

With more than 600 users constantly storing and retrieving data from their simulations on Titan, managing that data is both critical and complex.

To address the enormous task that is data management, the OLCF employs HPSS, a series of tape and disc drive components, servers, and software that provides long-term storage for the massive amounts of data generated by users.

To ensure against data loss, the team updates the archive as often as possible with the latest software and storage technologies. This, however, can be a daunting effort.

"We need to employ the latest technology to keep up with the demands that users have for data integrity, space, and speed," said OLCF HPSS Developer Mitchell Griffith.

For instance, in 2006 the amount of data stored in HPSS surpassed 1 petabyte for the first time. Reaching this number took 8 and a half years. To reach the second petabyte, however, took under 2 years, and getting to the third took only 6 months.

Just as Jaguar's transition to Titan prompted the OLCF to upgrade the Spider file system, it likewise prompted the organization to rethink its data management and storage.

Titan represents much more than one big machine—it represents the entire OLCF support ecosystem that makes enabling breakthroughs possible. Data storage is just one of many of these components.

For starters, 2013 saw the OLCF streamline the day-to-day operations of the HPSS archive system by collocating six Oracle StorageTek SL8500 tape libraries and more than 40,000 media cartridges in a centralized location.

Each tape library can hold 10,000 individual media cartridges, with each cartridge capable of storing 1 to 8 terabytes of data. This sheer volume of information made the move from two locations to a single one extremely challenging because more than 30 petabytes are stored within the HPSS archive—roughly three times the size of the entire printed collection at the Library of Congress. Adding to this challenge, the team had to deal with the previous, very complex cabling plant and large array of fiber-channel and Ethernet switch gear.

After several months of preparation, though, the team was able to not only move the library itself, but also upgrade facilities and systems such as power, space, and cooling, as well as complete a new cabling plant and fiber/Ethernet network. The team also worked with ORNL fire engineers and vendor representatives to design a fire-suppression system to meet fire code requirements and further protect the archive media and data from damage.

The result is a tape library infrastructure better able to share the load of the HPSS archive's requests for tape resources. The libraries are also in a more controlled environment that regulates temperature, humidity, and air quality, leading to better resiliency of the tape media.

Lastly, the upgraded cabling plant and infrastructure moved the HPSS archive off older, more-expensive-to-maintain hardware, which will save tens of thousands of dollars each year through reduced maintenance expenses.

"We are always trying to implement the newest storage technologies so that our researchers know their data is safe," said HPC Operations' Kevin Thach. "But sometimes the technology is just not there yet or is too expensive at its current stage, so we have to come up with our own ways to make the HPSS archive more reliable."

With the new updated and centralized HPSS, the OLCF is better positioned in one more arena to ensure that its users can get the most from their time on Titan.

But data comes from all angles, and while one of the world's fastest file systems was a welcome addition, it was only a piece of the puzzle.

All systems go

The next order of business was ensuring that researchers had even more efficient means of analyzing the reams of data created by their simulations on Titan.

After all, computational scientists need all the help they can get. The amount of data generated by today's high-performance computing systems is enormous and growing proportionally with the size and capability of the most recent hardware. To quickly and easily make sense of that much information, visualization and data analysis tools are critical.

When the OLCF upgraded Jaguar to Titan, the Center knew that an upgrade to its visualization resource was necessary to complement Titan's computational horsepower. The result is a world-class visualization facility that allows researchers to view their data interactively and simply, without the help of the OLCF's visualization liaisons, who focus on very large datasets and high-end rendering.

When the OLCF set out to upgrade its visualization facility, it did so with the researcher in mind. In fact, the entire system is the result of conversations with the user community aimed at providing the best system possible for the vast spectrum of researchers and specialties that use OLCF resources such as Titan.

After extensive investigation into the users' needs—and literally meeting with representatives from all of the OLCF's scientific domains—visualization staff set about creating one of the world's premier scientific visualization facilities, a working laboratory that will allow researchers to analyze and decipher their data as efficiently as possible.

Despite the diverse and numerous requirements of the Center's user community, the OLCF quickly realized that several key qualities were critical to pleasing the user base as a whole. Five, in particular, stood out: interactivity for complex data, binocular depth perception, ease of use, ultrahigh resolution, and fast I/O performance.

To gauge the new laboratory's effectiveness, the OLCF opened the experimental system to a set of early users. The results were promising, to say the least. Researchers in biofuels, liquid crystals, and climate were just a few of the early visualization success stories thanks to the new lab.

The new system was officially "accepted," or put through a series of benchmark tests to confirm its functionality, performance, and stability, at the very end of fiscal year 2013. Because early users are already successfully taking advantage of the new capabilities, the potential for real scientific breakthroughs via visualization and data analysis seems greater than ever. "These resources are necessary for achieving the Center's scientific mission," said OLCF visualization specialist Jamison Daniel, adding that a new immersive tracking technology is now installed that allows users to follow, in 3D, their simulations in real time.

The new ARTTRACK 3 tracking system, developed to work in a fashion similar to that of the human visual system, also allows interactivity with large datasets using Titan and Rhea, the OLCF's new visualization and analysis cluster. Datasets generated on Titan are so large, they cannot be analyzed or visualized on a single workstation as they were in the past; the work must be distributed across several systems. The new visualization lab serves as the front end for observation and analysis as simulations unfold. The OLCF visualization team has worked with Kitware, Inc. to deploy this technology in both the new lab and the Center for the Advanced Simulation of Light Water Reactors, likewise located at ORNL.

Finally, the OLCF visualization team has worked closely with GPU-maker NVIDIA on driver support, deploying the highest-ever resolution on a single node and enabling Thornton's team to visualize the 88-variable parallel coordinate work. This capability allows a researcher to drive the entire resolution of each wall from a single shared-memory system, so there is no need to distribute the graphics across a cluster of machines. The OLCF was the first Center to deploy a shared-memory node that drives at this resolution.

With this revolutionary visualization capability, Titan has the perfect partner in world-leading scientific simulation. Data is only as good as our ability to understand it, and now that understanding can be visually obtained faster and more accurately than ever before.

But the news doesn't stop there. In addition to the new visualization lab, two new computing resources were recently installed at the OLCF to enhance the user experience and maximize project outputs. The new machines are helping scientists to consolidate the massive amount of data gathered from Titan, and thus allowing their research teams to find answers to some of the world's most puzzling problems.

Eos came online October 3, 2013 and is now prioritized as an additional INCITE resource. The Cray XC30 system has 744 nodes divided among four cabinets. Each node contains two eight-core Intel Sandy Bridge processors and 64 gigabytes of memory. The compute cluster also relies on the new Spider Lustre file system and has many of the same debuggers, profiling codes, and software packages that are used on Titan.

"Eos is the newest generation of Cray architecture," said Suzanne Parete-Koon of the OLCF User Assistance and Outreach group. "The machine initially will be used as an extra computing resource to help INCITE projects reach their goals."

In November 2013, OLCF users also gained access to Rhea, a replacement for the aging Lens visualization cluster, which was retired upon Rhea's successful deployment. The system—which is composed of 196 Dell PowerEdge C600 nodes—offers users a newer resource for post-processing analysis and visualization of information gleaned from Titan.

Rhea offers a full suite of post-processing software, including ParaView and iPython. "We are excited about bringing Rhea online so we can offer our users a dedicated analysis and post-processing resource that features the latest commodity hardware," said Robert D. French, also of the User Assistance and Outreach group. "We expect it to offer a faster time to solution for our users."

Looking ahead

The OLCF now boasts one of the world's most powerful scientific research tools along with one of the country's richest scientific computing ecosystems in terms of file systems, staff expertise, and energy efficiency.

Titan's deployment reinforces the notion that while the march toward the exascale will indeed be full of challenges, solutions are always within reach. And despite the delay, the OLCF still met its user metrics through a combination of working with DOE and the Center's user base, and achieved everything it set out to do within the established timeframe since deciding to upgrade Jaguar to the hybrid marvel that is now Titan.

In fact, the OLCF exceeded its commitments to the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, ALCC, and Director's Discretion programs for Titan core-hours, delivering more than 2.6 billion compute hours among all programs in 2013, including staff programs, workshops, and Director's Discretion projects. These programs produced numerous scientific success stories, explained in detail in pages 4-15.

This year promises to be Titan's most successful year yet, as applications continue to harness the power of its hybrid architecture to build on early breakthroughs. No one can be certain what the future holds, but from here it looks pretty bright.

In Memoriam

Whatever It Took: Remembering Ricky Kendall

Ricky Kendall, a self-avowed pig farmer from Indiana, achieved much in his life. A big man with a big personality, a big reputation, and a big love for his friends and family, the only thing he lacked was a big ego.

His whatever-it-takes attitude was a constant reminder to friends and coworkers of his perseverance and determination, and his passing was a tremendous loss to the supercomputing, Oak Ridge National Laboratory (ORNL), and Oak Ridge, Tennessee, communities of which he was a part.

Born in North Carolina, Kendall grew up in Indiana, where he graduated magna cum laude from Indiana State University before earning his PhD in computational chemistry from the University of Utah in 1988. In 1989, after participating in postdoctoral training at Argonne National Laboratory in Illinois, he became a staff scientist at Washington state's Pacific Northwest National Laboratory (PNNL).

After a decade in the Northwest, which saw him earn a number of awards, including a highly prestigious R&D 100 Award in 1999 for his work on the Molecular Science Software Suite, Kendall moved to Ames Laboratory at Iowa State University. There he taught practical parallel computing courses and performed a key role at the Ames Lab Scalable Computing Laboratory.

In 2004 he received an Outstanding New Professional Award from the College of Liberal Arts & Sciences at Iowa State University, and in 2005 Kendall left Ames lab for ORNL, where he became the founding leader of the Scientific Computing group (SciComp) at the National Center for Computational Sciences (NCCS).

The website insidehpc.com, a popular industry blog, designated him a "Rock Star of HPC" (high-performance computing) in 2010, and ORNL's management contractor, UT-Battelle LLC, named him Group Leader of the Year in 2011. He was appointed Chief Computational Scientist at the NCCS in 2011 and remained in that post until his death in early 2014.

Whatever it takes

Jeff Nichols, the associate laboratory director of the NCCS since 2009, knew Kendall back when he was a PhD candidate at the University of Utah. Every other year Kendall's advisor, Jack Simons, would take a group up to Kings Peak in the High Uintas Wilderness and spend 2 or 3 days hiking.

Kings Peak is the highest peak in Utah, with an elevation of 13,528 feet above sea level. "These were not day hikes," said Nichols. "These were significant events." According to Nichols, Kendall wasn't known as a fast hiker, but he did whatever it took to get there. "There were times we would be on this tall peak, and we would holler down to him about a mile back. We could see him gesturing back at us. That was just our way of encouraging him."

And Kendall kept walking. During one particular hike to Kings Peak, the group was about halfway there and decided to spend the night on a lake. Simons was worried that Kendall would not be able to handle the steep, slippery shale trek up to the summit the next day, but he wouldn't be denied. He made it up and back down just fine because, Nichols said, "he'd put his mind to it."

To Nichols, that singular event personified Kendall's whatever-it-takes mantra. If he committed to doing something, he did it. It might take him twice as long or twice as much work, but he completed the task he'd set for himself.

Whatever floats your canary!

Everyone in the SciComp group thought Kendall was a tough, but fair, group leader. They also knew that he was their biggest supporter. While demanding much from his staff, Kendall fought hard for promotions and awards to make sure that everyone was recognized for their contributions.

Perhaps Kendall's greatest asset, however, was giving his team the tools they needed. His first response to a team member asking for a new tool was always a resounding yes; he knew that to get the work done, his team needed access to an open toolbox.

His saying, "whatever floats your canary" was an offbeat way of acknowledging that fact. "He was incredibly results oriented," recalled Bronson Messer, a colleague of Kendall's at the OLCF. "He really didn't care how we got to the answer, as long as we got the answer. And perhaps most importantly for us, he treated us as autonomous professionals, capable of working independently to solve our own problems."

NCCS Director Jim Hack also valued Kendall's approach. "He was a critical thinker, a thoughtful person . . . creative in terms of problem solving," he said. "He did it all without a pretentious tone. He was one of the guys to lean on when you had hard problems. He knew computer science as well as anyone I know, despite coming from a computational chemistry background. He built the SciComp group from the ground up, and they became his family. He loved and took care of them as he did his own family. They felt like they belonged."

And Kendall treated everyone equally, from the computational scientists in his group, to ORNL management, to the leaders within the Department of Energy (DOE).

Barbara Helland, DOE facilities division director, had this to say about Kendall: "His efforts and whatever-it-takes attitude were some of the reasons that the OLCF is the success it is today. I counted on him to give me an honest opinion when I needed to test out an idea that I had. I will miss that."

Everyone who worked with Ricky Kendall agrees that he had many talents: He was a master of team building, of bringing people together from different academic backgrounds, who knew how to treat everyone fairly regardless of his or her level of education. He could translate the jargon of domain scientists, computer scientists, and mathematicians because he had hands-on experience at PNNL and Ames. He was an exceptional code developer who could write and debug code in multiple languages and had a deep understood of parallel computing. And he could just as easily have a conversation about parallel computing and algorithms as he could about computational chemistry.

But Kendall was much more than just an accomplished scientist. He was ultimately a people person, a first-rate collaborator who rejoiced at seeing colleagues fulfill their personal and career goals. His honesty and integrity were invaluable to his friends and coworkers. He was a passionate educator who wanted to give his pupils practical knowledge that they would use in the real world. He was a volunteer who enjoyed working with local youth bowling leagues. But most of all, he was a family man who loved his wife, children, and friends more than words could express. All will miss him.

Advancements

Introduction

The OLCF is among the world's leading computational research centers. Not only is it home to Titan; it also boasts one of the largest file storage systems in the world, premier visualization tools, and a knowledgeable staff.

The OLCF strives to develop new tools to make scientific discovery at the facility more efficient. Here are some of the facility's 2013 innovations.

Industrial partnering to use GPUs

The OLCF recognizes that many industrial HPC users may lack the experience with GPUs to efficiently run applications on Titan. To address this issue, the OLCF's outreach efforts include a tailored industrial partnership program, Accelerating Competitiveness through Computational Excellence (ACCEL), to help industrial partners expand their ability to harness the power of GPUs to accelerate problem solving.

In 2013, the OLCF invited select firms to submit proposals for this unique Titan access opportunity. The program provides both computing time on Titan and technical expertise, volunteered by Cray's Center of Excellence at the OLCF and by NVIDIA. Seven ACCEL proposals were selected in 2013: two from Procter & Gamble and one each from Boeing, the United Technologies Research Center (collaborating with Pratt & Whitney), Caterpillar, General Electric, and Simulia (an independent vendor of Abaqus software). This is a very exciting set of projects from firms that are making a commitment to GPUs to see how they perform.

New ways to respond to users' data needs

As leadership systems grow in size, so potentially does the amount of critical data generated during and after the simulations. These data are, in essence, the intellectual capital of the research communities. The OLCF continues to focus attention on ways to facilitate scientific accomplishments through efforts to improve data management: its movement, access, and analysis.

In 2013 the OLCF formalized the addition of data liaisons to its team. The data liaisons provide critical I/O and workflow assistance to enable users to complete the cycle of science discovery from simulation to data movement, storage, and analysis.

In response to user requests for continued access to data following successful completion of projects, a select number of data projects were created in 2013 in order to allow users to access data after the award of compute time was exhausted. The pilot includes common use cases: servicing these projects will allow the OLCF to develop a suite of quantifiable requirement metrics for data-centric projects, including measures such as storage capacity, bandwidth requirements, and latency requirements.

Through the data project beta-test and introduction of data liaisons, the OLCF seeks to understand the ways in which data creation, storage, curation, exploration, and technical support lead to scientific insight. Effective storage (including capacity, bandwidth, and security) and curation are essential to allow complete use of the data to produce outcomes. Additional requirements, both in infrastructure and policy, come to the forefront if those data are to be made available to a wider community of researchers. Importantly, the kind of requirement gathering undertaken by the assignment of data liaisons and the instantiation of data projects is very much in keeping with the successful methods the OLCF has used in the past to understand simulation requirements.

Development of the Spider II file system

In 2013, the OLCF successfully designed, acquired and deployed the Spider II parallel file system (PFS). Spider II is a centerwide Lustre PFS with a usable capacity of 32 petabytes and I/O bandwidth of more than 1 terabyte/second to OLCF users. Compared with Spider I (the previous centerwide Lustre PFS at the OLCF), Spider II is a significant leap forward, providing a 4× speedup in transfer rate and a 3× increase in capacity while operating under the same power envelope and a smaller footprint. The increased capacity and performance will enable OLCF users to perform scientific simulations much more quickly.

Technology transfer to vendors

During the installation and testing of the 32 petabyte block storage solution as part of the Spider II system, and during the introduction to production of the accompanying Lustre file system, the OLCF staff identified a series of critical performance and functionality issues. In each instance, OLCF staff was able to characterize the issue, provide a mechanism for reliably reproducing the issue, and then either fix it internally, or work directly with the vendor to identify the appropriate solution. These efforts led to greater functionality and performance on the block storage controller, a revised device driver for each of the Lustre file system's drives, with better sequential read and write performance, and a high-performing FDR client for the Cray XK7 LNET routers. As a result of these efforts, the Lustre file system has met all performance goals, and has been released to production for the Titan, Eos, and Rhea users.

The OLCF uses Adaptive Computing's TORQUE product as its workload manager. Several significant contributions to that code base were made in the areas of security, scalability, and reliability. Most recently, a security issue was discovered at OLCF that allowed a root exploit in TORQUE. OLCF provided a patch to Adaptive Computing to fix the issue, and all TORQUE sites were notified to fix the issue. Over the past year, OLCF also provided fixes to Adaptive Computing to address memory leaks and deadlocks found on Cray platforms, and a fix for phantom jobs that were stranded on the job queue servers that then prevented new jobs from starting.

Innovative data management

In order to facilitate free access to scientific data arising from taxpayer-funded research, the OLCF has initiated a feasibility study to assess the promise of digital object identifiers (DOIs) and their potential use within an OLCF context. A DOI offers a unique, permanent identifier to help track datasets that are produced by users. Metadata about the dataset is stored along with the DOI name; metadata may include information such as a location, a description of the dataset, and the processes used to produce the dataset. DOIs provide a key first step toward establishing the desired data-sharing ability for the scientific community.

OLCF's High-Performance Storage System (HPSS) installation holds more than 35 petabytes of scientific data, accumulated over a period of 20 years. This data's value is incalculable as intellectual capital for the scientific community. As computer architectures and software evolve, much of the data stored in HPSS cannot be reproduced. Ensuring that this legacy of scientific data is conserved effectively and remains reliably available, which depends on the reliability and robustness of the HPSS and its components, is a crucial interest for OLCF. OLCF staff began work on the HPSS Integrity Crawler in September 2013. Its purpose is to catch problems early, thereby avoiding data loss. Development toward production deployment is ongoing and is expected to be completed in 2014.

The Data Transfer Node (DTN) environment is a collection of 16 Linux nodes designed specifically to facilitate moving data between OLCF storage resources and as well as to and from other laboratories and institutions. The nodes are connected to Spider II and HPSS. They are configured and tuned for high-speed data transfers, and they provide users a dedicated resource to access and transfer data so that the workload does not disrupt OLCF compute resources. Interactive DTNs allow users to manually perform transfers. These nodes also function as Globus Online endpoints, which has become a very popular method of data movement between sites, due to its ease of use. This method of deploying DTNs for specific-use cases has allowed the OLCF to scale the number of nodes in a more targeted manner to meet user needs.

In order to more efficiently meet the I/O throughput requirements of scientific applications, OLCF staff has developed a tool, a monitoring infrastructure and a partition selection tool that is based on file system load. The custom tool, fs-select, logs file system activity from the RAID controllers on the out-of-band management network with negligible impact to real application workloads. Use of the partitions selected by fs-select has the potential to alleviate I/O bottlenecks.

The OLCF's Spider II PFS has more than 30 petabytes of storage capacity, and a data transfer performance of greater than 1 terabyte per second. Leveraging the high performance available from massive-scale file systems such as Spider II requires scalable software. The parallel tools effort, a collaboration among ORNL, LANL, LLNL, and DataDirect Networks, is aimed at employing the scalability available from current supercomputers and parallel file systems to increase performance for everyday file and data management tasks. Current tools under development include a Lustre-aware parallel data copy tool, a parallel tape archive (tar) tool, a parallel filesystem scanning tool, and a parallel filesystem search tool. OLCF staff has already produced initial versions of the parallel tar, parallel filesystem scanning, and parallel search utilities; these tools have been tested in the OLCF environment, and additional development of these tools continues among the collaborative partners.

Innovative application support

Titan's hybrid CPU-GPU environment brings a significant increase in capability to the OLCF, but it also requires users to adapt to the new hybrid programming environment in order to make effective use of the system. To facilitate adaptation to hybrid computing, the OLCF offers several different programming environments, which can be used in applications individually or in combination. CUDA is a C-based language for NVIDIA GPU accelerators. Well-written CUDA is generally expected to provide the best possible performance on a system such as Titan; however, code portability is sacrificed, as CUDA only works on systems with NVIDIA GPUs. OpenCL is based on a set of APIs rather than on specific architectural implementations. It is an open standard with multiple implementations targeting many-core CPUs, GPU accelerators from different vendors and other accelerators. Benchmarking shows that current OpenCL implementations lag significantly behind CUDA implementations in performance, which has led to limited interest so far among OLCF users. OpenACC is a higher-level approach, using directives (specially structured comments) and APIs to control the offloading of certain computations to the accelerator. It provides higher levels of both code portability and performance portability than CUDA or OpenCL at some cost in absolute performance. It is an open standard with multiple implementations targeting a variety of accelerator hardware types.

ORNL researchers, supported by the OLCF, have played a significant role in enabling application-level resilience with a fault-tolerant Message-Passing Interface (MPI). As part of the MPI Forum's Fault Tolerance Working Group, the ORNL researchers have been engaged in reviewing and evaluating HPC fault tolerance. In support of this effort, an ORNL team has developed the Scalable Tools Communication Infrastructure, a new low-level runtime infrastructure that provides the fundamental capabilities required to (1) allow an MPI job to "run through" a node failure without the entire job failing and (2) refine the application-facing interface that allows for the notification of node failures and recovery under control of the application.

The OLCF has developed aprun-usage, a new module that reviews a user's aprun flags while a job is running. The module provides feedback that could lead to improved performance by reducing walltime, run-time variability, or both. If jobs are properly spaced, they will run 1.4 to 1.7 times faster (i.e., walltime will be reduced by 30% to 40%). The aprun-usage module detects aprun calls that fail to properly place their processes and prints an advisory with a URL to an NCCS webpage with details on how to improve the performance. The module also suggests that the user turn on core specialization, when appropriate, to reduce jitter (i.e., runtime variability due to system noise). Lastly, the module also warns if the user specified more OpenMP threads than the user requests per process from aprun.

Innovative hardware monitoring and diagnosis

The OLCF developed a series of diagnostic tools that could identify underperforming or failing hardware. These tools were frequently used to screen the system after a reboot before allowing production jobs to start. Another diagnostic tool was developed that could isolate individual GPUs that were contributing to an incorrect result or residual. Because this fault was intermittent, the individual node that was causing the error was difficult to identify. The OLCF developed a method of isolating the offending hardware, using the concept of a virtual cabinet, placing a potentially faulty node in a virtual cabinet of known good nodes and executing a series of runs among the members of the virtual cabinet to identify the failing node.

The OLCF developed a benchmark and diagnostic application called TopoBW, a user-level MPI program that stresses the Cray Gemini network by saturating each link on the system and measuring the resulting performance. Using topology-aware algorithms, it can detect previously undiscovered performance problems. While Cray provides tools for assessing the health of the Gemini network, those diagnostics can only be run off line while the system is down for maintenance. TopoBW can be run while the machine is in production, and can even be run on a contiguous subset of the machine. TopoBW supersedes the capabilities of the offline tool as well, demonstrating an ability to identify failing connections that the offline diagnostics cannot.

The Liebert XDP protocol translation has been rolled into centerwide monitoring. The Liebert XDP units use a data communication protocol for building automation and control networks called BACnet. This is a global standard used in the refrigeration, air-conditioning, and heating industries, but with little or no adoption by available open source monitoring solutions. Because the OLCF primarily depends on the Simple Network Management Protocol (SNMP) for communicating with devices, a device developed by Chipkin Automation Systems provided the needed protocol translation between BACnet and SNMP. Once that device was configured and deployed, custom SNMP scripts and queries were developed, and now the cooling system data is presented to the OLCF's Nagios server to provide device status, alarm notification, threshold errors, and similar information. With access to this data, graphs were also generated to look for outliers and to trend the data.

OLCF Systems Overview

Cray XK7 Titan

Titan is now the nation's premier supercomputer for scientific discovery.

By combining GPUs with traditional CPUs, Titan is capable of achieving a peak performance of 27 petaflops, ten times faster than its predecessor while using only slightly more power. In Nov. 2012, Titan earned the number 3 spot on the Green500, a biannual ranking of the world's most efficient supercomputers.

Occupying less than 4,500 square feet, Titan comprises 200 cabinets that house 18,688 NVIDIA Kepler GPUs accelerators and 18,688 16-core AMD Opterons (299,008 cores total), with a total of 710 terabytes of memory.

Operating on the ESnet 100 gigabyte network, Titan can ship and receive massive amounts of data to and from users at other labs and institutions all over the world.

Because of its innovations, Titan has, in many ways, set a new bar for supercomputing standards. For the next several years Titan will be critical to solving some of the world's most pressing problems.

Data management

The growth in supercomputing power brings with it the need to both store more data and access it faster and more efficiently. The OLCF uses the HPC industry standard InfiniBrand high-performance network to quickly ship data to and from user locations. Titan utilizes several data storage systems, such as the High-Performance Storage System (HPSS), the Spider II file system, and more; it's InfiniBrand's job to keep Titan connected to each of those platforms. It won't be long until supercomputers transition from the petascale era into the exascale era—a move that guarantees the ever-continuing growth of OLCF data storage.

HPSS

The OLCF is home to one of the world's largest archival storage systems. HPSS uses high-speed data movers to write data onto disks; later, that data is transferred to tapes. Due to the growth of data collected from the center's resources, staff members are constantly adding more disk space and tape. For instance, in 2006, the HPSS offered approximately 1 petabyte of data storage. By mid-2013, the OLCF had over 30 petabytes stored in six Storage Tek SL8500 tape libraries containing 10,000 tape slots each. The libraries have 128 tape drives.

Everest

Many of the scientific breakthroughs are enabled by the data analytics and visualization capabilities provided to the OLCF users. In 2013, the OLCF coordinated the complete redesign, deployment, and management of the Exploratory Visualization Environment for Research in Science and Technology (EVEREST) visualization laboratory as a state-of-the-art scientific discovery facility. Within EVEREST, there are three computing systems and two separate state-of-the-art visualization display walls. The primary display wall spans 30.5 feet × 8.5 feet and consists of 18 1920 × 1080 stereoscopic Barco projection displays arranged in a 6 × 3 configuration. The secondary display wall contains16 1920 × 1080 planar displays arranged in a 4 × 4 configuration, providing a standard 16:9 aspect ratio. With the new upgrade, scientists will be better able to visualize their data to make discoveries. The larger stereoscopic display provides 37 million pixels; the stereoscopic design creates depth in an image for a 3D effect. While the previous systems could show movies, high-resolution images, and some live applications, the new system's 3D capabilities will allow for significantly greater detail in data visualization.

Sith

Sith is an Opteron-based InfiniBrand cluster running Linux. The system is provided as an end-to-end user resource for center users, meaning users can input data and expect a result without having to do any intermediary mechanics, or further data manipulation in order to get the output. It is used for workflow automation for jobs running from Titan and for advanced data analysis. The system contains 40 compute nodes. Each compute node contains four 2.3 GHz eight-core AMD Opteron processors with 64 gigabytes of memory, and a 512 gigabyte solid-state drive (SSD). In addition to mounting the Atlas file systems, Sith is configured with an 86 terabyte Lustre file system for scratch space.

New Systems

EOS – Cray XC30

In September 2013, the OLCF installed and accepted a four-cabinet Cray XC30, called "Eos." The Eos system was released to production on October 3, 2013. The system, with 744 Intel Xeon E5-2670 compute nodes and 47.6 terabybtes of memory, provides the OLCF user community with a substantive large-memory-per-node computing platform. The Eos nodes are connected by Cray's Aries interconnect in a network topology called "Dragonfly."

Spider II

In September, the OLCF released Spider II, its next-generation Lustre parallel file system, to production. Spider II is the architectural revision of the original Spider file system. The Atlas file systems are two instantiations of Spider II, with an aggregate usable capacity of more than 30 petabytes, and block-level performance of more than 1 terabyte/second. The Atlas file systems were mounted on all compute systems in the last quarter of 2013 and will be the default file systems for Titan beginning in the first quarter of 2014.

Rhea

In the fourth quarter of 2013, the OLCF completed the installation and testing of a 196-node linux cluster, called "Rhea". Rhea was released to production on January 8, 2014. This system, with 392 Intel E5-2650 processors and 12.5 terabytes of memory, provides a conduit for large-scale scientific discovery via data analysis and visualization of simulation data generated on Titan. Rhea's compute and I/O nodes are interconnected using fourteen-data-rate (FDR) InfiniBand; this same technology connects the cluster to the Atlas file systems.

Education, Outreach and Training

Last year, the OLCF staff was busy at the facility, and around the world, to provide super education, outreach and training opportunities for users and members of the team. Here are some of the highlights of center activities in 2013:

OLCF Staff Leadership in HPC Community

OLCF management talks Titan at International Supercomputing Conference

OLCF Leadership was among the 2,500 attendees at the International Supercomputing Conference held June 16-20, 2013 in Leipzig, Germany. System managers, researchers from academia and industry representatives were on hand. More than 50 countries were represented at the event.

Both Buddy Bland, project director for the OLCF, and Jack Wells, director of science for ORNL's National Center for Computational Sciences, gave multiple presentations during the conference.

OLCF wins big at SC13

The OLCF and industrial users Ford Motor Company and GE Global Research received 5 awards at SC13, the 25th meeting of the annual leading Supercomputing Conference.

Taking place from November 17–22 in Denver, SC13 brought the world's leading supercomputing centers together with high-performance computing (HPC) users, developers, and sponsors from academia, industry, and government laboratories around the world.

The conference allowed participants to share knowledge and plans through a wide variety of venues and events, including tutorials, workshops, panel discussions, invited talks, research poster sessions, technical paper presentations, and birds-of-a-feather sessions.

CUG 2013

"A New Vintage of Computing" was the theme for the annual Cray User Group (CUG) meeting, held May 6-9, 2013 in Napa Valley, Calif. A number of OLCF staff members were on hand to discuss a variety of HPC topics as well as the impact of climate on agriculture, developing new crops to fuel economic growth, and analyzing carbon mitigation techniques.

Staff members collaborate on winner of Gauss Award

A paper co-authored by three Oak Ridge Leadership Computing Facility (OLCF) staff members received international recognition. The team was given the Gauss Award at the International Supercomputing Conference (ISC) for the most outstanding paper in the field of scalable supercomputing.

The paper, "TUE, a New Energy-Efficiency Metric Applied at ORNL's Jaguar," was an initiative from the Energy Efficient High Performance Computing Working Group. ORNL's Stephen Poole, Chung-Hsing Hsu, and Don Maxwell contributed to the paper.

Titan webpage earns Gajus American In-house Design Award

OLCF web developer Brian Gajus was recently honored by the graphic design community.

Gajus took an American Inhouse Design Award for his work on the OLCF Titan page, which introduced America's most powerful supercomputer. Sponsored by the news magazine Graphic Design USA, the award recognizes the contributions made by designers working in institutional marketing and communications departments.

OLCF Users in the News

OLCF user earns early honor

OLCF user Gaute Hagen, a researcher in ORNL's Physics Division, recently received an Early Career Award from the Department of Energy's Office of Science.

Hagen's award-winning project, entitled "State of the Art Microscopic Computations of Weak Processes in Nuclei," aims to develop novel extensions of computational methods that address fundamental questions in physics such as "What is the nature of the neutrino?" and "What are the mechanisms for the production of elements heavier than iron?"

UT-Batelle Corporate Fellows

"Witek" Nazarewicz is an international leader in theoretical nuclear physics and is widely recognized as a principal driving force behind research on the physics of exotic nuclei. His team used Jaguar and density functional theory to determine that there are about 7,000 possible combinations of protons and neutrons allowed in bound nuclei with up to 120 protons (a hypothetical element called "unbinilium"), work that continues on Titan. The team's results are presented in the June 28, 2012 issue of the journal Nature. Witek also served as scientific director of the Holifield Radioactive Ion Beam Facility at ORNL from 1999 to 2011.

Bobby Sumpter leads both the Computational Chemistry and Materials Science group in the Computer Science and Mathematics Division and the Nanomaterials Theory Institute at the Center for Nanophase Materials Sciences at ORNL. He has been exceptionally productive at the confluence of theory and experiment, working across organizational boundaries to provide theory and modeling leadership across disciplines.

His recent work on Jaguar and Titan includes investigating a next-generation capacitor, a device that stores energy through charge separation at an electric double layer formed within porous materials.

As transportation systems are electrified, fast charging of vehicles with such devices will be necessary. Iterating between simulation and experiment to reduce intermediate models, within 2 years Sumpter's team and collaborators at Rice University arrived at a practical device for high-power energy delivery. Their supercapacitor stores—and quickly discharges—a thousand times more energy than a conventional capacitor.

GE, OLCF wins Editors Choice Award

The Editor's Choice Award for the best use of HPC in manufacturing went to GE and the OLCF for research to better understand ice formation at the atomic level. Using the OLCF's Titan supercomputer, researchers for the first time simulated hundreds of millions of water molecules freezing in slow motion. New insights into how ice forms will help GE develop wind turbines that are better able to withstand debilitating ice accumulation in cold climates.

Ford, OLCF wins Editor's Choice Award

Ford and the OLCF were recognized with an Editor's Choice Award for the best use of HPC in automotive research and a Readers' Choice Award for the best HPC collaboration between government and industry. Using the OLCF's Jaguar and Titan systems, Ford optimized for the first time the underhood airflow in automobiles to reduce cooling drag and increase fuel efficiency.

Ford and GE also received the International Data Corporation's HPC Innovation Excellence Award for this research done on the OLCF's supercomputing systems. This year's award marks the second win for GE which was recognized last year for modeling on the OLCF's Jaguar system the unsteady air flows in the blade rows of turbomachines.

OLCF Industry user named Person to Watch in High-Performance Computing for 2014

Titan user Masako Yamada of GE Global Research has been named one of HPCwire's People to Watch 2014. Through an Advanced Scientific Computing Research Leadership Computing Challenge award, she has been using the Oak Ridge Leadership Computing Facility's (OLCF's) 27-petaflop supercomputer to model million-molecule water droplets freezing onto various anti-ice surfaces.

Training & Workshops

The Future of Computing and Data Integration

Sponsored by Oak Ridge National Laboratory's (ORNL's) Computing and Computational Sciences Directorate, the annual Smoky Mountains Computational Sciences and Engineering Conference on September 3-5, 2013 brought together approximately 100 key high-performance computing (HPC) players from government, academia, and industry.

Users defeat 'Data Deluge' at OLCF's firstlarge data sets workshop

To assist users planning for data management, including modifying input/output (I/O) operations for efficiency, selecting short and long-term data storage, and building visualizations, OLCF HPC user assistance specialist Fernanda Foertter organized the August 6–8 workshop "Processing and Analysis of Very Large Data Sets," held in Knoxville, TN. 91 people attended the workshop.

East and West Coast workshops help prepare users for Titan

In attempt to minimize travel costs for the users, the OLCF conducted two similar training events in January and February of 2013 in opposite geographic locations. The format was intended to help facilitate access to training by users. It also served to increase collaboration with the OLCF accelerator vendor/partner NVIDIA, who graciously offered a classroom for the first event within its headquarters in San José, California. The second was held in downtown Knoxville, Tennessee, at a Hilton, in response to user feedback that indicated difficulties in reaching the ORNL campus and additional transportation costs. Both events had an almost identical curriculum, which included hands-on exercises on GPU programming, practicums on tools, and best practices learned by OLCF's staff and the CAAR team. As with all of our training events, both were broadcast live over the web and reached more than 176 people both remotely and on site.

Lattice QCD community visits ORNL to discuss challenges in the field

The workshop brought quantum chromodynamics (QCD) researchers together with OLCF computing experts. The first day included an introduction from OLCF Director of Science Jack Wells and a talk from Fermilab's Paul Mackenzie, who serves as chair of the USQCD Executive Committee. Mackenzie discussed QCD research at high-energy accelerators around the world, including Europe's Large Hadron Collider. QCD researchers make extensive use of Titan: this session gave them a chance to give input about necessary upgrades, and discuss QCD applications.

OLCF participates in GPU Technology Conference

From March 18 to March 21, 2013, staff members from OLCF were in San Jose, California, participating in the world's leading conference on graphics processing units, or GPUs: the GPU Technology Conference. OLCF representatives included Director of Science Jack Wells, Director of Operations Jim Rogers, and User Assistance Specialists Suzanne Parete-Koon and Fernanda Foertter. Foertter and Wells gave a talk describing the OLCF, Titan, and the early science phase ongoing on Titan. The early science phase consists of a select group of research teams granted early access to Titan's GPUs.

OpenACC headlines OLCF summer workshops

The OLCF offered three workshops this summer to fit diverse talents and interests.

The biggest draw was the "OpenACC Tutorials" workshop cosponsored by Cray and AMD. This three-day workshop, organized by the OLCF's Training Coordinator Fernanda Foertter, was hands-on and allowed users of Titan to port and optimize their applications. There were 70 participants in total, with 35 joining Foertter via live webcast.

OpenACC featured attendees from industry (Rolls Royce), universities (Georgia Tech, Princeton, and Stony Brook), national labs (Sandia and ORNL), and other government agencies (National Oceanic and Atmospheric Administration).

ADIOS code spring: A race for new technologies

Nearly two dozen software tool developers gathered recently at ORNL to boost the integration of the award-winning Adaptable I/O System for Big Data (ADIOS) into their technologies.

The ADIOS Code Sprint was held August 19–20, 2013 bringing 22 developers together from research organizations across the United States, including ORNL, Lawrence Berkeley National Laboratory, Sandia National Laboratories, the University of Oregon, Kitware, North Carolina State University, Rutgers University, Georgia Tech, and the University of California–Davis.

Education and Outreach

The Research Alliance in Math and Science Program

Last summer, a dozen or so minority college students were selected as interns with the Research Alliance in Math and Science (RAMS), a program hosted by the ORNL Computing and Computational Sciences Directorate. Internships lasted from June 2 to August 8. Each student was assigned to a STEM research mentor to work on a project of interest to the student, the student's professor(s), and the ORNL researcher. They received stipends while interning, and had several opportunities to present the results of their research to their professional communities. Their research covered topics as diverse as economics and climatology.

OLCF continues education leadership with summer students

Nearly 25 student interns from middle school to graduate school got the opportunity this summer to work with OLCF staff and boost their computing skills. One, William Walker Smith, was 13 years old—the youngest intern ORNL has ever had. Smith landed an internship working with OLCF researcher Hai Ah Nam by being one of 10 winners in a related essay contest. Benjamin Brock, a University of Tennessee (UT) Haslam Scholar and computer science major, was another exceptional intern who was full time at the OLCF from May until August.

NUFO event offers chance to spotlight OLCF

OLCF representatives were on-hand for the 2013 NFUO exhibit in Washington D.C. It was an interactive demo-type exhibit focused on how basic science has led to major changes in people's lives.

ORNL Intern Wins Best Abstract Award at Student Poster Session

Roisin Langan, an intern at ORNL, spent last summer developing climate models that are better able to predict the variability and extremes of precipitation. Her project titled "Stochastic Representation of Unresolved Processes in Climate Models" won the best abstract award at the Research Alliance in Math and Science (RAMS) banquet, a student poster session held at ORNL on August 8.

OLCF in the News

Slick science: Modeling surfaces to cast off ice

At the GE Energy Learning Center in upstate New York, workers learn to fix and maintain wind turbines. Physicist Masako Yamada works nearby, at the GE Global Research Center.

Whenever she passes the learning center, Yamada recalls a nacelle, a bus-sized hub she's seen on the training center floor. Generators housed in nacelles atop wind-turbine towers convert into electricity the rotation of blades longer than semitrailer trucks.

It's a reminder of the huge impact Yamada's studies of tiny processes could have. At the research center, she focuses on the molecular foundation of the freezing process, helping find surfaces to keep turbine blades ice-free and efficient.

Featured on the Department of Energy, Office of Science webpage, 10/17/13

http://ascr-discovery.science.doe.gov/synchronized/yamada1.shtml

Watch a supercomputer track trillions of particles in the latest universe simulation

This purple star map is our best guess at the shape of the universe at half its current age, according to a new simulation from a research team at Argonne National Lab. The project is called the Hardware / Hybrid Accelerated Cosmology Code (HACC), an attempt to use supercomputers to model the intricate physics of mass in an expanding universe. The project tracks roughly 1.1 trillion particles as they expand and cohere, nearly four times larger than the scale of previous simulations. The hope is that by understanding the patterns of matter in the universe, cosmologists can track down the dark matter that experts estimate composes roughly 25 percent of our universe.

To get there, the project needed help from some of the most powerful supercomputers on earth, drawing processor time from Titan at OLCF, Sequoia at Lawrence Livermore Lab, and Mira at Argonne Labs. Running on Titan, the simulation used over 25 petaflops of peak processing power, a speed that only one other computer in the world could manage. The simulation architecture also had to be unusually flexible to work across all three machines, and the code has been nominated for a high-performance computing award as a result. But researchers are quick to say that what's happened so far is still just a trial run, designed to test out the architecture's capabilities. "The hope is that starting in January, we'll run a mix of medium-sized simulations on Titan along with one or two really large ones," project leader Salman Habib, of Argonne National Laboratory, told The Verge.

The result of all that computing power is an intricate picture of how the structure of the universe formed, tracking the trillions of particles as they interact and form more complex structures. The simulation shows the universe at roughly 7.4 billion years after the big bang (or roughly 6.4 billion years ago), imagining the development of the so-called "cosmic web," while taking into account modern ideas of dark energy.

Featured on The Verge webpage, 11/11/13.

http://www.theverge.com/2013/11/11/5081024/new-supercomputer-visualization-shows-the-formation-of-the-universe

OLCF Groups

The OLCF is led by a management team experienced at guiding world-class systems. The management team has five principals: Director Jim Hack guides the overall vision of the center, while Project Director Buddy Bland supervises installation and upgrades of the OLCF supercomputers. Deputy Project Director Kathlyn Boudwin provides project management for a variety of computing and computing system installation and upgrade projects. Director of Science Jack Wells oversees the research teams that use the computing systems. Jim Rogers, as Director of Operations, manages day-to-day operations along with future systems and infrastructure planning. And Julia White, program manager for INCITE, manages one of the leading programs through which researchers gain access to OLCF resources. Team leaders are: Ashley Barker, David Bernholdt, T.P. Straatsma, Kevin Thach and Sudharshan Vazhkudai.

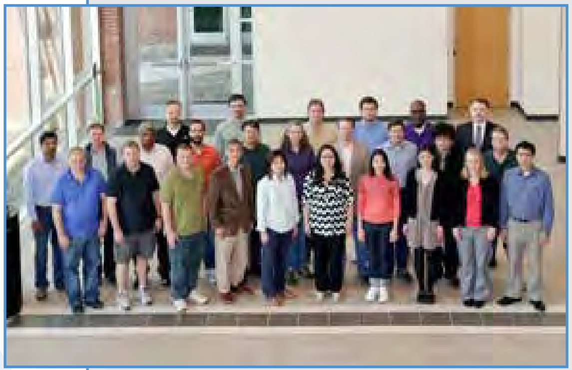

The **Scientific Computing Group (SCG)** works hand in hand with OLCF users to help them achieve optimal results from their time on the HPC systems. The SCG is comprised of research scientists, visualization specialists, and work-flow experts who are trained in chemistry, physics, astrophysics, mathematics, numerical analysis, or computer science. Each research team is assigned an SCG liaison who then actively assists users throughout the entire research process. Staff members help design and optimize code for applications, streamline the work flow, and solve any computer issues that may arise. Visualization specialists capture the resulting data in images and help the users analyze it.

The **Technology Integration (TechInt) group** is charged with delivering new technologies to 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, they keep OLCF ahead of the curve. Their scope includes archival storage, parallel file systems, data management, high-performance networking, and architecture. TechInt staff work behind the scenes to develop the infrastructure that supports the OLCF systems. They research and evaluate emerging technologies and tools in the aforementioned areas and provide the systems programming to seamlessly integrate them into existing networks, file systems, and archival storage infrastructures.

The **User Assistance and Outreach (UAO) Group** is one of the user community's primary points of contact with the OLCF and is responsible for delivering seamless access to OLCF resources, providing swift and effective technical support, and showcasing OLCF research in strategic communications activities. Specific activities include problem management and consulting; developing content for the OLCF website; coordinating training events and workshops; serving as advocates for the user community; supporting third-party applications; developing science highlights, annual reports, posters and other communication products; and maintaining the OLCF website and the Resource Allocation Tracking System.

The **High Performance Computing Operations (HPCO) Group** keeps the OLCF supercomputing systems running smoothly. Members of the group monitor all systems 24 hours a day, seven days a week, 365 days a year. They are tasked with overseeing the administration, configuration management, and cybersecurity of everything from infrastructure systems to the Titan supercomputer. The HPCO tests systems when they are installed and upgraded, and uses diagnostic tools to continually monitor them. This allows them to anticipate problems before they arise and identify components that are near failure. The HPCO's ceaseless monitoring and maintenance ensure that OLCF's important resources function smoothly.

The **Computer Science Research (CSR) Group** ensures that researchers' applications are performing at their peak. CSR researches, tracks, and purchases a wide range of software tools which enable scientists to perform their simulations in an effective manner. Members collect and analyze performance data of these tools, which they use to decide which pieces of software need to be enhanced or replaced. The group also works with vendors to facilitate the purchase of new modeling tools, languages, middleware, and performance-characterization tools. APT members seek to address the practical needs of scientists receiving allocations on HPC computing resources so that they are able to spend their time focusing on research, instead of ineffective tools and programs.

OLCF Computing Allocations for 2013

INCITE / www.doeleadershipcomputing.org

Biological Sciences

Advanced Modeling of the Human Skin Barrier
Michael Klein, Temple University
65,000,000 hours

Petascale Computing of Biomolecular Systems
Klaus Schulten, University of Illinois at Urbana-Champaign
110,000,000 hours

Simulations of Ribosome Biogenesis and Cellular Processes
Zan Luthey-Schulten, University of Illinois at Urbana-Champaign
51,410,000 hours

Cellulosic Ethanol: Simulation of Multicomponent Biomass Systems
Jeremy Smith, Oak Ridge National Laboratory
78,000,000 hours

Chemistry

High-Fidelity Simulations for Advanced Engine Combustion Research
Joseph Oefelein, Sandia National Laboratories
100,000,000 hours

Precision Many-Body Quantum Simulations of Functionalized Structures
Shiwei Zhang, College of William and Mary
27,000,000 hours

Non-Perturbative QED Study for Matter and Anti-Matter Collisions
Michael Pindzola, University of Auburn
30,000,000 hours

Computer Science

Performance Evaluation and Analysis Consortium (PEAC) End Station
Leonid Oliker, Lawrence Berkeley National Laboratory
45,000,000 hours

Collaborative Research into Exascale Systemware, Tools, and Applications (CRESTA)
Lorna Smith, University of Edinburgh
21,000,000 hours

Earth Science

Climate-Science Computational Development Team: The Climate End Station II
Warren Washingon, University Corporation for Atmospheric Research
60,000,000 hours

CyberShake 3.0: Physics-Based Probabilistic Seismic Hazard Analysis
Thomas Jordan, University of Southern California
45,000,000 hours

Global Seismic Tomography Based on Spectral-Element and Adjoint Methods
Jeroen Tromp, Princeton University
100,000,000 hours

Engineering

Multiscale Blood Flow Simulations
George Karniadakis, Brown University
51,000,000 hours

Explosive Hazard Predictions with the Uintah Framework
Martin Berzins, University of Utah
45,000,000 hours

Parameter Studies of Boussinesq Flows
Susan Kurien, Los Alamos National Laboratory
34,000,000 hours

The Solution of Three-Dimensional PWR Neutronics Benchmark Problems for CASL
Thomas Evans, Oak Ridge National Laboratory
21,000,000 hours

Materials

Computational Prediction and Discovery of Magnet Materials
Bruce Harmon, Ames Laboratory
45,000,000 hours

Scalable First Principles Calculations for Materials at Finite Temperature
Markus Eisenbach, Oak Ridge National Laboratory
105,000,000 hours

Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo
Dario Alfe, University College London
55,000,000 hours

Quantum Monte Carlo Simulations of Hydrogen and Water Ice
Richard Needs, University of Cambridge
75,000,000 hours

Predictive and Insightful Calculations of Energy Materials
Paul Kent, Oak Ridge National Laboratory
45,000,000 hours

Ab Initio Simulations of Carrier Transports in Organic and Inorganic Nanosystems
Lin-Wang Wang, Lawrence Berkeley National Laboratory
25,000,000 hours

Safety in Numbers: Discovery of New Solid Li-Ion Electrolytes
Boris Kozinsky, Bosch in U.S.A.
48,000,000 hours

Physics

Petascale Simulations of Type Ia Supernovae
Stan Woosley, University of California, Santa Cruz
55,000,000 hours

Three Dimensional Simulations of Core Collapse Supernovae
Anthony Mezzacappa, Oak Ridge National Laboratory
35,000,000 hours

Simulating Reionization of the Local Universe: Witnessing Our Own Cosmic Dawn
Paul Shapiro, University of Texas
40,000,000 hours

Lattice QCD
Paul Mackenzie, Fermilab
140,000,000 hours

Nuclear Structure and Nuclear Reactions
James Vary, Iowa State University
74,000,000 hours

Transformative Advances In Plasma-Based Acceleration
Warren Mori, University of California, Los Angeles
30,000,000 hours

Unraveling the Physics of Magentic Reconnection with 3D Kinetic Simulations
William Daughton, Los Alamos National Laboratory
55,000,000 hours

Magnetic Reconnection in High-Energy-Density-Laser-Produced Plasmas
Amitava Bhattacharjee, University of New Hampshire
35,000,000 hours

High-Fidelity Simulation of Tokamak Edge Plasma Transport
C.S. Chang, Princeton Physics Plasma Laboratory
100,000,000 hours

ALCC / science.energy.gov/ascr/facilities/alcc

Titan-Enabled Supersonic CCS Compressor and Engine Technology Development: Intelligently-Driven Optimization, Time-Varying Boundary Layer Flow Control, and GPU-Accelerated Convergence Augmentation
Allan Grosvenor, Ramgen
20,000,000 Hours

High-Fidelity Simulations of Transition and Turbulent Separation in Turbomachinery
Gorazd Medic, United Technology Research Center
10,000,000 Hours

Accelerating Design of Complex Fuel Injectors Through Petascale Computing
Madhusudan Pai, GE Global Research
46,000,000 Hours

Gyrokinetic Simulation of Energetic Particle Turbulence and Transport
Zhihong Lin, University of California, Irvine
50,000,000 Hours

First Principles Investigations of Adsorbate-Metal Interactions: Quantum Monte Carlo and Ab Initio Molecular Dynamics Simulations
Jeff Greeley, Purdue University
25,000,000 Hours

Accelerated Modeling of Non-icing Surfaces for Cold Climate Wind Turbines
Masako Yamada, GE Global Research Center
40,000,000 Hours

V/UQ Assessment of a Large Eddy Simulation Tool for Clean-Coal Technology
Jeremy Thornock, University of Utah
7,000,000 Hours

Projections of Ice Sheet Evolution
Phillip Jones, Los Alamos National Laboratory
2,600,000 Hours

The Role of Aerosols and Multi-scale Water Cycle Processes in Climate Change: Sensitivity to Aerosol Emissions and Modeling Frameworks
Philip Rasch, Pacific Northwest National Laboratory
36,000,000 Hours

High-Resolution Coupled Climate Simulations on Titan with GPU Acceleration
Mark Taylor, Sandia National Laboratory
64,000,000 Hours

Unraveling the Coupling of Radio Frequency Power to Fusion Plasmas
David Green, Oak Ridge National Laboratory
50,000,000 Hours

Understanding Helium Plasma Mediated Tungsten Surface Response that Controls Plasma Facing Component Performance and Lifetime
Brian Wirth, University of Tennessee
5,000,000 Hours

Structure and Dynamics of Nuclear Systems within Time-Dependent Density Functional Theory Approach
Aurel Bulgac, University of Washington
25,000,000 Hours

Transforming Modeling and Simulation for Nuclear Energy Applications
John Turner, Oak Ridge National Laboratory
20,000,000 Hours

High Impact Publications

Biology

D. Ortega, C. Yang, P. Ames, J. Baudry, J .Parkinson, I. Zhulin, "A phenylalanine rotameric switch for signal-state control in bacterial chemoreceptors," Nature Communications, **4**, 2881 (2013).

J. Ostmeyer, S. Chakrapani, A. Pan, E. Perozo, and B. Roux, "Recovery from slow inactivation in K+ channels is controlled by water molecules," Nature, **501** (7465), 121-124 (2013)

B. Lindner, L. Petridis, R. Schulz, and J. Smith, "Solvent-Driven Preferential Association of Lignin with Regions of Crystalline Cellulose in Molecular Dynamics Simulation," Biomacromolecules, **14** (10), 3390-3398 (2013)

Climate Change Science

F. He, J. Shakun, P. Clark, A. Carlson, Z. Liu, B. Otto-Bliesner, and J. Kutzbach, "Northern Hemisphere forcing of Southern Hemisphere climate during the last deglaciation," Nature, **494** (7435), 81-85, (2013)

G. Meehl, A. Hu, J. Arblaster, J. Fasullo, and K. Trenberth, "Externally Forced and Internally Generated Decadal Climate Variability Associated with the Interdecadal Pacific Oscillation," Journal of Climate, **26** (18), 7298-7310, (2013)

K. Evans, P. Lauritzen, S. Mishra, R. Neale, M. Taylor, and J. Tribbia, "AMIP Simulation with the CAM4 Spectral Element Dynamical Core," Journal of Climate, **26** (3), 689-709, (2013)

G. Meehl, W. Washington, J. Arblaster, A. Hu, H. Teng, and J. Kay, "Climate Change Projections in CESM1(CAM5) Compared to CCSM4," Journal of Climate, **26** (17), 6287-6308, (2013)

Materials

Q. Li, J. Owens, C. Han, B. Sumpter, W. Lu, J. Bernholc, Jerzy, V. Meunier, P. Maksymovych, M. Fuentes-Cabrera, and M. Pan, "Self-Organized and Cu-Coordinated Surface Linear Polymerization," Scientific Reports, **3**, 2102 (2013)

M. Wang, C. Zhang, X. Lu, G. Tan, H. Luo, Y. Song, M. Wang, X. Zhang, Xiaotian, E. Goremychkin, T. Perring, T. Maier, Z. Yin, K. Haule, G. Kotliar, and P. Dai, "Doping dependence of spin excitations and its correlations with high-temperature superconductivity in iron pnictides," Nature Communications, **4**, 2874 (2013)

E. Brown, B. Clark, J. DuBois, and D. Ceperley, "Path-Integral Monte Carlo Simulation of the Warm Dense Homogeneous Electron Gas," Physical Review Letters, **110** (14), 146405 (2013)

Q. Li, G. Cao, S. Okamoto, J. Yi, W. Lin, B. Sales, J. Yan, R. Arita, J. Kunes, A. Kozhevnikov, A. Eguiluz, M. Imada, Z. Gai, M. Pan, and D. Mandrus, "Atomically resolved spectroscopic study of Sr2IrO4: Experiment and theory," Scientific Reports, **3**, 3073 (2013)

J. Ren, N. Vukmirovic, and L. Wang, "Nonadiabatic molecular dynamics simulation for carrier transport in a pentathiophene butyric acid monolayer," Physical Review B, **87** (20), 81108 (2013)

Physics

B. Barrett, P. Navratil, and J. Vary, "Ab initio no core shell model," Progress in Particle and Nuclear Physics, **69**, 131-181 (2013)

S. W. Bruenn, A. Mezzacappa, W. R. Hix, E. J. Lentz, O. E. B. Messer, E. J. Lingerfelt, J. M. Blondin, E. Endeve, P. Marronetti, and K. N. Yakunin, "Axisymmetric Ab Initio Core-Collapse Supernova Simulations of 12-25 Mo Stars," Astrophysical Journal Letters, **767** (1), L6 (2013)

A. Ekstrom, G. Baardsen, C. Forssen, G. Hagen, M. Hjorth-Jensen, G. Jansen, R. Machleidt, W. Nazarewicz, T. Papenbrock, J. Sarich, and S. Wild, "Optimized Chiral Nucleon-Nucleon Interaction at Next-to-Next-to-Leading Order," Physical Review Letters, **110** (19), 192502 (2013)

G. Wlazlowski, P. Magierski, J. Drut, A. Bulgac, and K. Roche, "Cooper Pairing Above the Critical Temperature in a Unitary Fermi Gas," Physical Review Letters, **110** (9), 90401 (2013)

P. Krstic, J. Allain, C. Taylor, J. Dadras, S. Maeda, K. Morokuma, J. Jakowski, A. Allouche, and C. Skinner, "Deuterium Uptake in Magnetic-Fusion Devices with Lithium-Conditioned Carbon Walls," Physical Review Letters, **110** (10), 105201 (2013)

Fusion

Y. Liu, W. Daughton, H. Karimabadi, H. Li, and V. Roytershteyn, "Bifurcated Structure of the Electron Diffusion Region in Three-Dimensional Magnetic Reconnection," Physical Review Letters, **110** (26), 265004 (2013)

W. Fox, G. Fiksel, A. Bhattacharjee, P. Chang, K. Germaschewski, K., S. Hu, and Nilson, "Filamentation Instability of Counterstreaming Laser-Driven Plasmas," Physical Review Letters, **111** (22), 225002 (2013)

V. Roytershteyn, S. Dorfman, W. Daughton, H. Ji, M. Yamada, and H. Karimabadi, "Electromagnetic instability of thin reconnection layers: Comparison of three-dimensional simulations with MRX observations," Physics of Plasmas, **20** (6), 61212 (2013)

Y. Huang, and A. Bhattacharjee, "Plasmoid instability in high-Lundquist-number magnetic reconnection," Physics of Plasmas, **20** (5), 55702 (2013)

C. Forssen, G. Hagen, M. Hjorth-Jensen, W. Nazarewicz, and J. Rotureau, "Living on the edge of stability, the limits of the nuclear landscape," Physics Scripta, **T152**, 14022 (2013)

Chemistry

Y. Yang, S. Pope, H. Wang, and J. Chen, "Large-eddy simulation/probability density function modeling of a non-premixed CO/H-2 temporally evolving jet flame," Proceedings of the Combustion Institute," **34**, 1241-1249 (2013)

J. Kang, S. Li, J. Li, J.B. Xia, and L. Wang, "Electronic structural Moire pattern effects on MoS2/MoSe2 2D heterostructures," Nano Letters, **13** (11), 5485-5490 (2013)