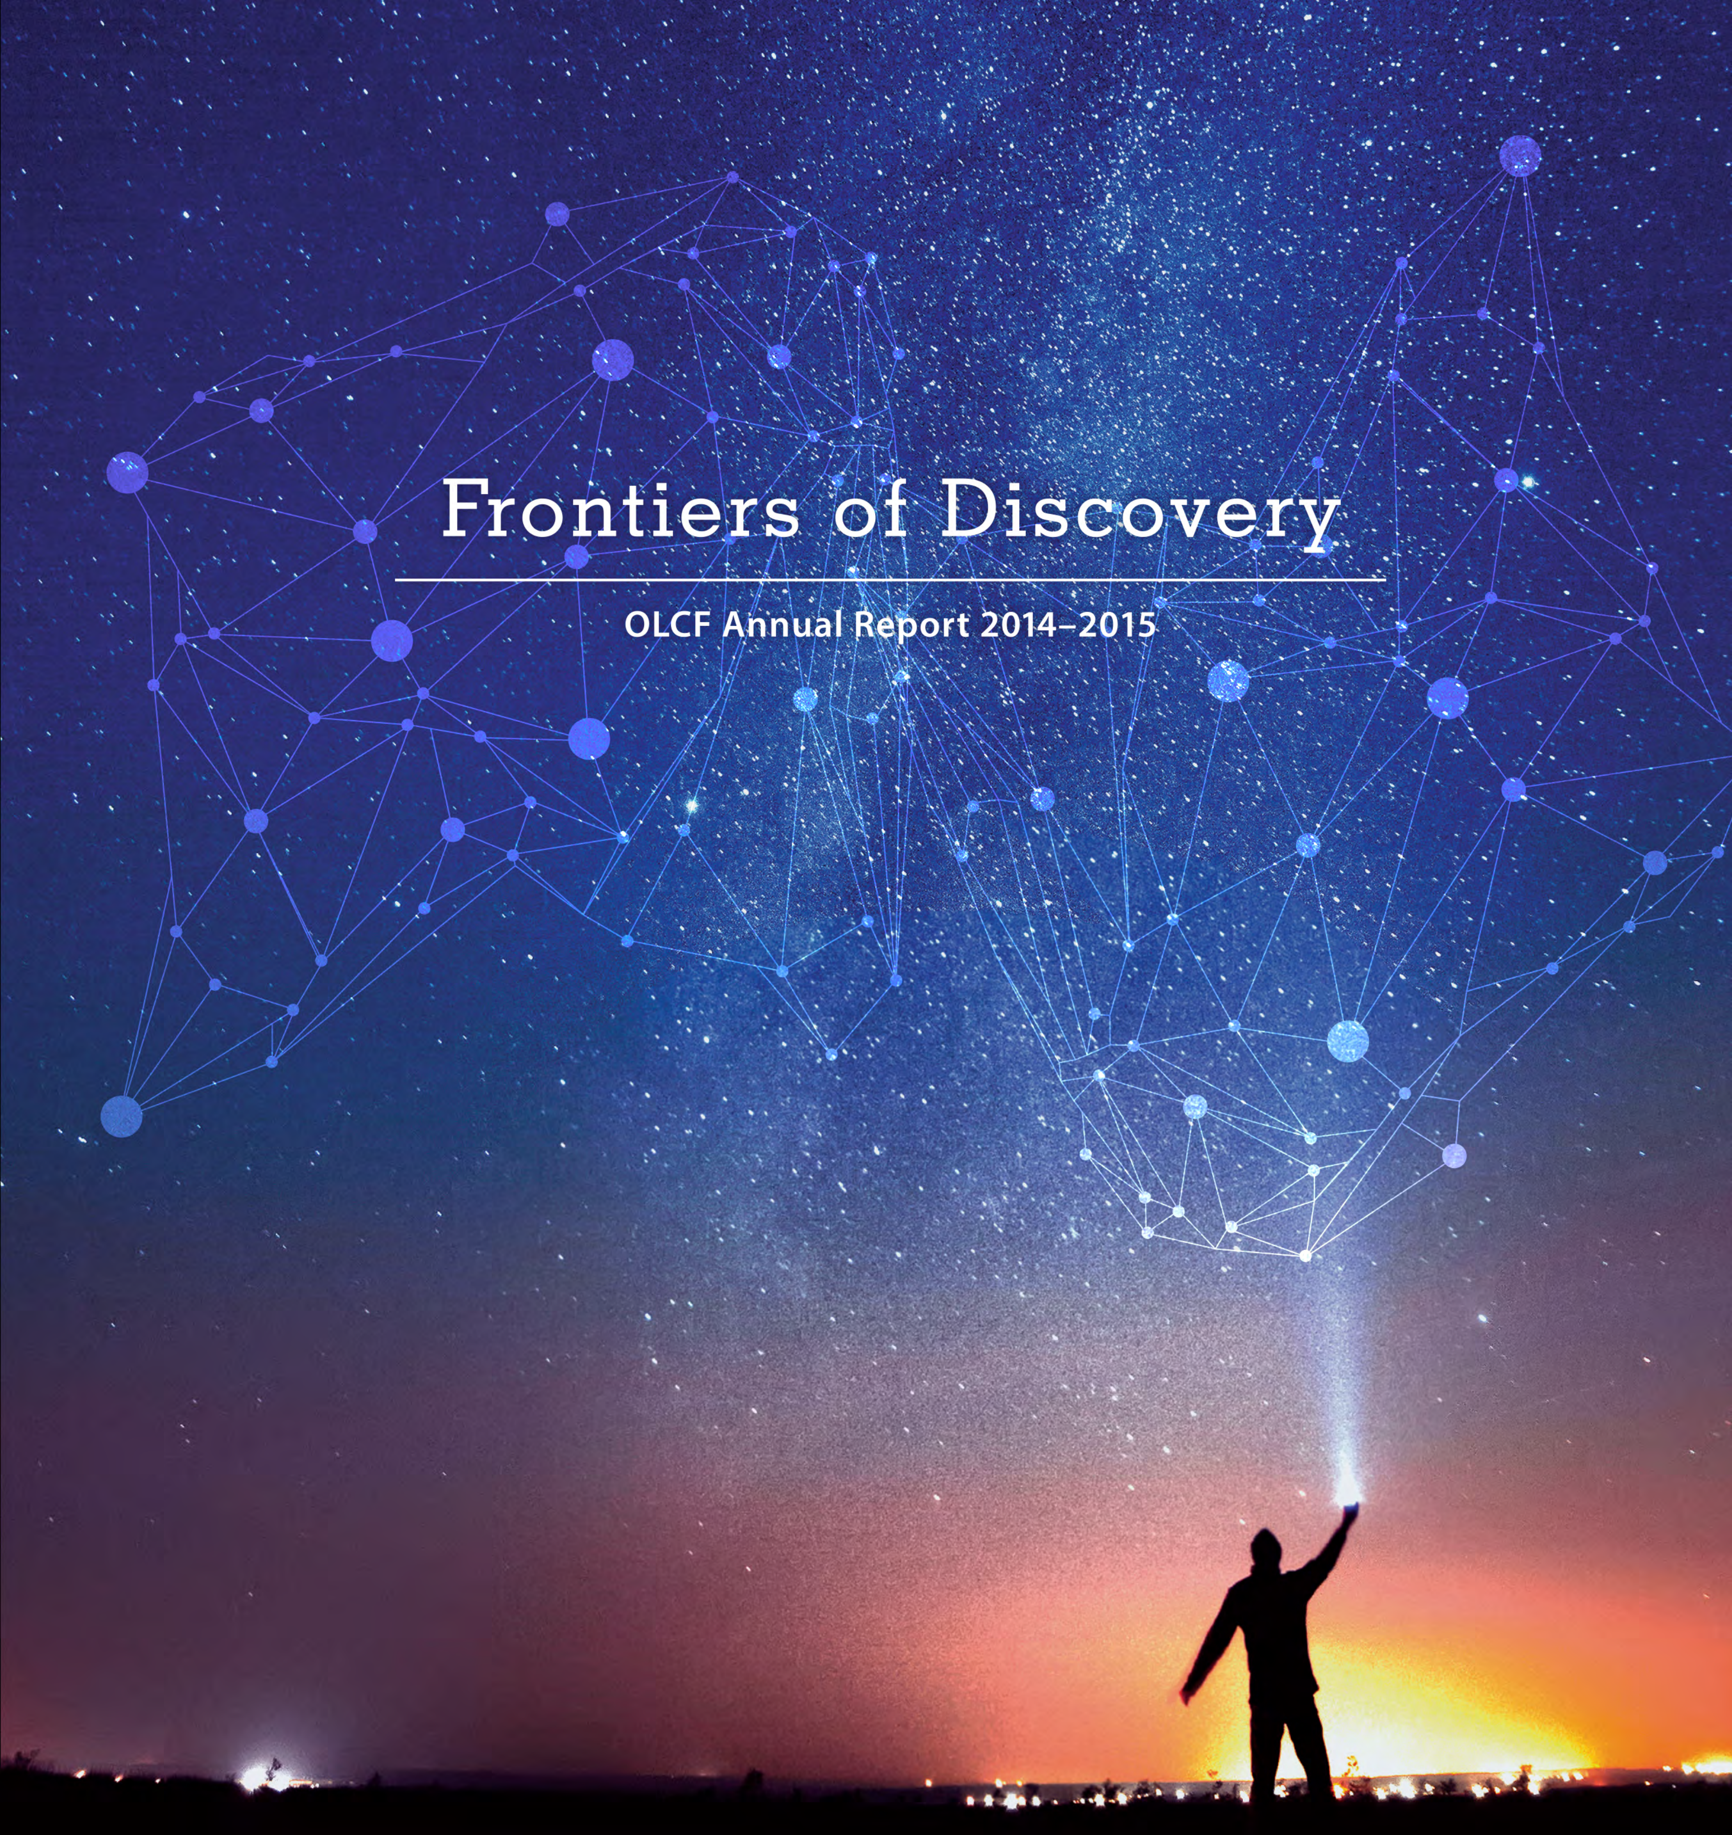

Frontiers of Discovery

OLCF Annual Report 2014–2015

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Contents

2 OLCF: Advancing the Frontiers of Discovery

4 The Complexities of Combustion

6 Titan Takes on the Universe

8 Researchers Get Warmer in Understanding High-Temperature Superconductors

10 Procter & Gamble and Temple University Scientists Model Skin's Makeup

12 Iron-based Superconductor Simulations Spin Out New Possibilities on Titan

14 Researchers Dig Up Biological Data from Microbes in Mines

16 Dependability, Delivery Hallmarks of a Successful Year for Titan

22 DATA Analysis and Workflow Demos Highlight DOE Booth at SC14

24 Ascending to Summit

26 Inside the OLCF

27 Training at the Threshold of Accelerated HPC

29 Education, Outreach and Training

31 OLCF Executive Board and User Groups

32 INCITE, ALCC, and Director's Discretionary

33 Specialized User Support for Accelerated HPC

35 OLCF Web Traffic Statistics

36 OLCF Compute Systems

37 Data Storage, Analysis, and Visualization Resources

38 2014 INCITE Projects

40 2014 ALCC Projects

42 2014 Publications

53 Acronyms

OLCF annual report team:

Outreach Project Manager: Jeff Gary

Writers: Eric Gedenk, Katie Elyce Jones, Dawn Levy, Jeremy Rumsey, Christie Thiessen

Graphic Designer: Jason Smith

Editors: Wendy Hames, Priscilla Henson

Photography: Jason Richards

Contact info:

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

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

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

Twitter: @OLCFGOV

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

OLCF: Advancing the Frontiers of Discovery

The hallmark of any successful organization is its people. The Oak Ridge Leadership Computing Facility (OLCF) at Oak Ridge National Laboratory (ORNL) is no exception.

The OLCF completed a very successful year in 2014. As has been the case in the past, our caring and dedicated staff members and our creative and talented scientific user community were at the heart of these successes.

I am happy to report that we have been able to deliver on our promises to our user programs as 2014 marked the first full production year for Titan.

We delivered 100 percent of INCITE allocations for Titan in 2014 and 100 percent of the ALCC and Director's Discretionary allocations for 2013–14 as well. Furthermore, Titan's scheduled availability proved to be very high, at 99.74 percent, with an overall availability of 95.85 percent. This is excellent when compared with Jaguar's most stable year, when scheduled availability was 98.11 percent and overall availability was 91.45 percent. Usage of available hours for Titan in 2014 was also very high at 89 percent compared with Jaguar's last year of service, where this figure was 84.39 percent. With only seven unscheduled downtimes in 2014, Titan has demonstrated it is a stable and reliable instrument for supporting the computational needs of our capability-driven user community.

We have documented a large number of outstanding scientific achievements through our INCITE program. In one such example, a team from Princeton Plasma Physics Laboratory and ORNL was able to effectively model from first principles the turbulence along the edge of the plasma in the DIII-D tokamak reactor at the General Atomics facility in San Diego. This same first-principles approach also revealed the divertor heat load footprint—the part of the reactor that extracts heat and helium ash from the plasma—acting as a vacuum system and ensuring a stable plasma and fusion reaction.

One of our ALCC projects provided a team from Rutgers University the opportunity to develop a computational model for predicting superconductivity. These researchers, studying iron-based superconductors, combined novel electronic structure algorithms with Titan's high-performance computing power to predict spin dynamics, or the ways electrons orient and correlate their spins in a material. These results were important contributions to the search for high-temperature superconducting materials because it is believed that spin dynamics create the conditions needed for superconductivity. This approach could expedite the search for new or modified materials that conduct electricity with little or no resistance.

The OLCF also supports a robust industrial partnership program. As one example from this program, Ramgen Power Systems—a small research and development company in Bellevue, Washington, that specializes in energy-related applications of supersonic aircraft technology—used Titan to optimize novel designs based on aerospace shock wave compression technology for gas compression systems, such as carbon dioxide compressors. This work gets them closer to their goal to commercialize new turbo compressor technology that will dramatically reduce the cost of carbon dioxide sequestration. In this case, the Ramgen CO_2 compressor is projected to reduce the capital costs of CO_2 compression by 50 percent and produce a minimum of 25 percent in operating cost savings.

In July 2014, the company announced it was being purchased by Dresser-Rand, one of the world's largest suppliers of custom-engineered rotating equipment solutions for the oil, gas, chemical, petrochemical, process, and power generation industries. Ramgen indicated that it was purchased in part because of the advanced modeling and simulation capabilities Ramgen engineers have acquired using OLCF resources, and the firm's resultant ability to apply large-scale computational fluid dynamics simulations to accelerate development of turbo-machinery technology. That new capability helped make Ramgen an attractive acquisition for a world-leading manufacturer.

To support our very active and diverse user community, we continue to evolve our user training program. In 2014, we provided more than 200 hours of training for the OLCF User Community. This year our training programs included our inaugural Hackathon in October, as well as our 2014

OLCF User Meeting in late July. Also in the summer, we sponsored a series of HPC fundamentals classes in addition to our monthly OLCF Users Conference Calls. In short, we believe our regimen of training continues to provide the best possible support for our user community and is targeted to the specific needs of that group.

The strength demonstrated by our center in 2014 also played an important role in informing our process to identify the OLCF's next generation machine.

The Collaboration of Oak Ridge, Argonne, and Lawrence Livermore National Laboratories—or CORAL—process began in earnest with the development of a joint Request for Information and the release of a Request for Proposals. Getting that RFP on the streets was a satisfying accomplishment but was only a step along the journey.

Once responses to the RFP were secured, a team representing each of the laboratories began evaluating the proposals. That complex selection process led to a successful Independent Project Review to proceed with awarding a contract with the successful vendor to acquire and deploy the OLCF-4 machine in 2018.

In November, we were able to announce our new system, Summit, and our key vendor partners of IBM, NVIDIA, and Mellanox.

The announcement, made in Washington, DC, by Department of Energy Secretary Ernest Moniz, Senator Lamar Alexander, and Congressman Chuck Fleischmann, was well received. Meanwhile, our OLCF outreach team members, working with members of the ORNL communications team, executed a successful launch of the announcement. More than 100 media stories have been generated to date, and more than 20,000 web hits occurred by the end of November.

The effort to bring Summit to our users in 2018 represents the next step along the path to exascale and continues to allow us the opportunity to provide the best high performance computing tools in the nation to help solve the most complex and dynamic scientific problems.

Over the course of the next 3 years, we will work to prepare our user community to be ready when the machine is delivered and applications prepared in a way to ensure their portability among the differing architectures.

At the OLCF, realize we have a technically-outstanding, hardworking, and dedicated staff. It's always gratifying, however, when our views are validated by the assessment of others. This year, a number of our OLCF staff members were formally recognized for their outstanding achievements. Several teams were recognized with Significant Event Awards for their efforts on upgrades to our visualization facilities and for the development and deployment of systems to dramatically improve efficiency, security, and flexibility of applications associated with writing, testing, reviewing, and deploying new code on OLCF systems.

The outstanding work of other OLCF staff members was recognized at the laboratory's annual awards night event. These awards included research accomplishments related to modeling and simulation of nuclear reactors, significant savings to OLCF operations, and science communication.

These achievements are emblematic of the excellent team we have at the OLCF. It is the hard work of this team that makes it possible for us to support the tremendous science discoveries facilitated by Titan and its unique computational ecosystem.

As we look to the year of 2015, we know this great work will continue, and the future of our facility is, indeed, a bright one.

James J. Hack
Director
National Center for Computational Sciences (NCCS)

The Complexities of Combustion

Despite the rush to green energy technologies, traditional combustion engines and devices will continue to be an economic reality for years to come, in part because they have the potential to employ low-carbon and renewable fuels. However, we can improve the efficiency of combustion devices to help reduce our reliance on fossil fuels and the amount of carbon dioxide released into the atmosphere.

Consider that Americans use two-thirds of their petroleum for transportation and one-third for heating buildings and generating electricity. "If low-temperature compression ignition concepts are widely adopted in next-generation automobiles, fuel efficiency could increase by as much as 25 to 50 percent," said Sandia National Laboratories' Jackie Chen, who uses the OLCF's Titan to study the combustion of a wide variety of fuels.

Recently, Ankit Bhagatwala, a member of Chen's team, employed the direct numerical simulation code known as S3D on Titan to simulate a jet flame burning dimethyl ether to probe fundamental turbulent flame physics associated with local extinction (where parts of the flame burn out) and reignition. If researchers can find strategies to minimize flame extinction, this will greatly enhance efficiency and minimize undesired emissions in combustion devices such as engines.

But even with computers as powerful as Titan, performing a direct numerical simulation of a diesel jet flame matching all of the aero-thermo-chemical conditions and chemical reactions is out of the question because of the multiscale complexity. Chemical properties manifest on small scales from billionths up to thousandths of a meter, whereas the motion of an engine valve can exert effects at scales from hundredths down to millionths of a meter. Researchers must therefore simulate a jet flame configuration at a lower Reynolds number (a measure of the mixing intensity and dynamic range of turbulence) that matches combustion's critical thermo-chemical conditions.

The latest S3D simulations on Titan delivered the highest Reynolds number ever achieved by Chen's team, 13,050. In fact, Titan and S3D pair so well together that the application is six times faster on Titan than it was on Jaguar, due almost entirely to the code's porting to Titan's GPUs with the incorporation of OpenACC, a programming standard that makes it easier for users to adapt their applications to GPU technology.

The Titan simulations included approximately 30 chemical molecules known as species to model the combustion of dimethyl ether. This enhanced number of species advances the team toward simulating the behavior of more realistic fuels, including biofuels. The increased Reynolds value also allows the team to resolve a wider range of turbulence scales in space and time—a major breakthrough when trying to match experimental conditions and evaluate turbulent mixing and combustion models.

Specifically, the team wanted to know the dependence of reignition on the local mixing rate, or the rate of fuel and air mixing during the combustion process. "We found that oxygenated fuels such as DME [dimethyl ether] generate considerably more stable intermediates such as formaldehyde, rendering the flame more robust to local extinction than conventional hydrocarbons such as methane," Bhagatwala said.

Ultimately, more accurate combustion models, analyzed by direct numerical simulations and experiments, will be used to optimize engine design, bringing us one step closer to more efficient combustion engines and devices.
—Gregory Scott Jones

Publication:

Ankit Bhagatwala, Zhaoyu Luo, Han Shen, Jeffrey A. Sutton, Tianfeng Lu, and Jacqueline H. Chen. 2014. "Numerical and experimental investigation of turbulent DME jet flames," Proceedings of the Combustion Institute 35(2): 1157–1166. doi:10.1016/j.proci.2014.05.147.

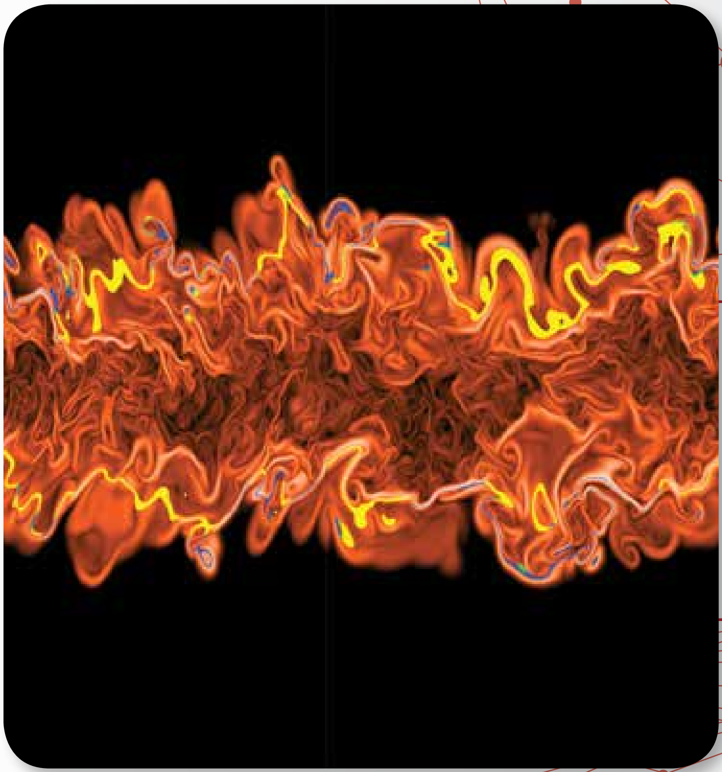

This image depicts the local mixing rate where white denotes high mixing rates and red, lower mixing rates.
Image credit: Ankit Bhagatwala.

2013 OLCF Early Science Project

Allocation: 79M hours

Jackie Chen
Sandia National Laboratories

Usage: 87.3M hours

Titan Takes on the Universe

It's a new era for cosmology. After decades of data collection from observational sky surveys and the progression of high-performance computing (HPC) power into the petascale, researchers can study the evolution of the universe through simulation, then statistically compare the virtual universe with real observations to refute or validate cosmological models.

In 2014, the largest cosmology simulation performed at resolutions required for modern-day galactic surveys ran on Titan. Scaling to almost 90 percent of Titan's GPUs, the "Q Continuum" simulation calculated the evolution of more than half a trillion particles, each representing roughly a thousandth of the mass of an average galaxy, to fill a virtual universe containing billions of galaxies.

The research team for this and related projects includes Salman Habib, Katrin Heitmann, Hal Finkel, Nicholas Frontiere, and Adrian Pope of DOE's Argonne National Laboratory. The team models structure formation in the universe using their Hardware/Hybrid Accelerated Cosmology Code (HACC), which is designed to exploit multiple supercomputing architectures by dividing the code into short- and long-range modules. Depending on the system architecture, the short-range, particle–particle level runs different algorithms optimized for either CPU or GPU architectures.

Over the next few years, the team will use this scalable, modular code to test the Standard Model of cosmology and explore the nature of dark energy and dark matter in a virtual universe on DOE leadership class computers.

"If we can understand the role dark energy and dark matter played in forming the structure of the universe, we can better interpret the astronomical phenomenon we study with our satellites and telescopes," Habib said.

The team will make results from HACC simulations available to astronomers and cosmologists who are seeking out distant galaxies and astrophysical phenomenon and cataloging their properties in large-scale sky surveys. Several surveys have members collaborating closely with the HACC team, including the Baryon Oscillation Spectroscopic Survey, the Dark Energy Survey, and the Large Synoptic Survey Telescope now under construction.

Using HACC data, researchers can compare different cosmological models to observational data; produce synthetic (or virtual) surveys to guide the design of galaxy surveys; and use the statistics generated from simulations to develop cosmological emulators, or precision prediction tools, that will help survey teams interpret large amounts of data.

The recent Q Continuum simulation on Titan mined for halos, or dense clumps of dark matter in which galaxy formation takes place. The computational power of Titan's NVIDIA GPUs was used for both the simulation and data analysis.

First, the simulation evolved half a trillion particles to represent universe expansion and mass distribution throughout about 13 billion years. The final stage of the simulation is a statistical correlation to the structure of our observable universe, although not an exact representation.

During data analysis, the team's halo finder program was able to identify clusters of particles representing galaxy formation. A few thousand particles alone can represent a dark matter structure (or halo) that hosts a galaxy. The halo finder recorded more than 167 million such halos (equaling 46 percent of all the particles in the simulation) in just under 36 minutes.

By identifying halos in the simulation, researchers can then seed the halos with galaxies to create mock observational catalogs of galaxy distribution, which will help optimize sky survey resources and guide the analysis of observed data.

"These simulations are helping surveys get ready for the massive amounts of data that come out of their own projects, and the information we have is changing rapidly," Habib said.—Katie Elyce Jones

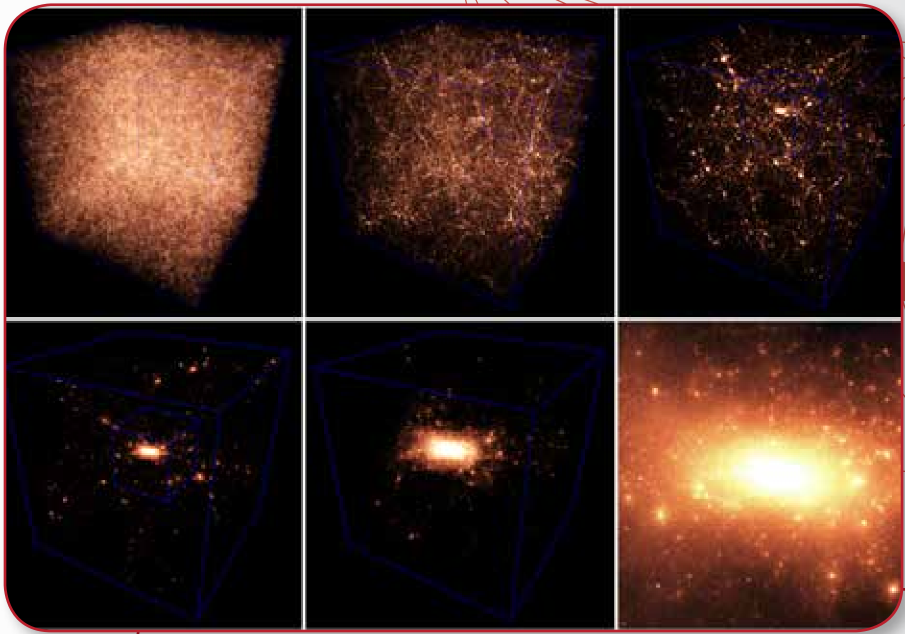

Halo particles (members of clumps of matter where galaxy formation takes place) in the Q Continuum simulation performed on Titan. The series of images (from upper left to bottom right) "zooms in" on one of the largest galaxy cluster-scale halos in the Q Continuum run, giving some impression of the enormous dynamic range of the simulation. Image credit: Heitmann et al.

2014 INCITE

Allocation: 100M hours

Usage: 190M hours

Salman Habib
Argonne National Laboratory

Researchers Get Warmer in Understanding High-Temperature Superconductors

A group at the US Department of Energy's (DOE's) Oak Ridge National Laboratory wanted to better understand superconductors' complex subatomic interactions and needed help from one of the world's fastest supercomputers.

DOE has long prioritized research surrounding materials at the subatomic level and considers better understanding electron manipulation in novel materials one of its grand challenges. More specifically, by being able to better predict superconductors' behaviors, scientists can expand research into a wide range of applications, including energy storage, catalysis, energy production, and metals that can be used as structural materials.

A group led by Paul Kent advanced its understanding of superconductivity by performing the first ab initio simulation of a high-temperature superconducting material, or cuprate.

"The goal of this research was to calculate the so-called exchange coupling, or interaction between adjacent copper molecules, from first principles," Kent said. "The challenge is that metal oxides are very difficult to predict, and our method had never been applied to this class of materials before."

Many materials scientists use density functional theory (DFT)—a method using mathematical functions to computationally model electrons—to try to accurately model subatomic interactions for transition metal oxides. Though DFT can offer accurate simulations, it has a difficult time simulating "d electrons," which are more localized around atoms. To simulate d electron-rich copper oxides accurately, the team needed another method.

It decided to use the quantum Monte Carlo (QMC) method for its simulations. Rather than using mathematical functions, QMC uses statistical data and random numbers to plot electrons in a simulation. Though QMC can be more computationally intensive work, it also allows the researchers to perform ab initio simulations far more accurately than using DFT.

"When you get to transition metal simulations using DFT, they can start to make big qualitative errors," Kent said. "Of course there are empirical fixes that you can do, but then of course you're putting your hand into the simulation to fix the result based on knowing what you want to get. That can be useful if you want to look at something else, but it isn't good if you want to make an ab initio prediction of magnetic couplings." By using Titan, the Kent team was able to perform the first truly ab initio simulation of a cuprate.

The team's initial runs using the QMCPACK code were computationally expensive, but through the end of 2014, the Kent team was able to reduce its memory usage by a factor of 8. This reduction allowed the team to do more computation on Titan's GPUs and enabled computation to be done on CPUs that previously required too much memory.

OLCF computational scientist Ying Wai Li helped the Kent team use memory more efficiently and created a new matrix inversion, allowing the team to delete hundreds of line of source code, which streamlined the computational demands of running QMCPACK. Before Kent's work, the researchers were limited to around 730 electrons in a simulation. After she updated the code, the team was able to run simulations with more than 1,000 electrons.

As computing power increases, the team will be able to simulate even heavier materials that could be used for large-scale superconducting applications.—Eric Gedenk

Publication:

K. Foyevtsova, et al. "Ab initio quantum Monte Carlo calculations of spin superexchange in cuprates: The benchmarking case for Ca2CuO3," Physical Review X 4: 031003 (2014). doi: http://dx.doi.org/10.1103/PhysRevX.4.031003

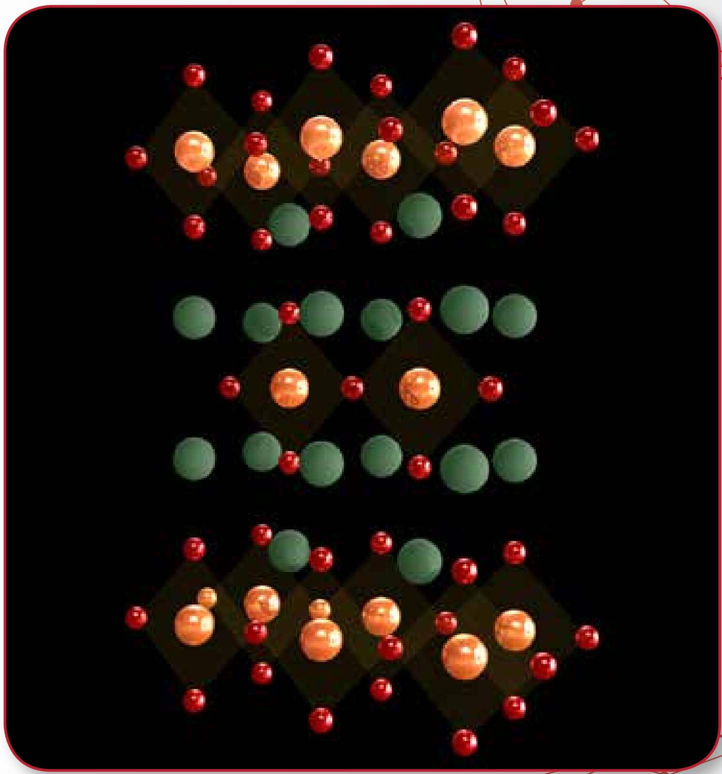

Using the quantum Monte Carlo method, the team was able to perform the first ab initio simulation of the cuprate $(Ca/Sr)_2CuO_3$. Cuprates show a remarkable property known as high temperature superconductivity. Superconductivity is a phenomenon where materials conduct electricity with zero resistance, and occurs in some materials at extremely low and high temperatures.

2014 INCITE

Paul Kent
Oak Ridge National Laboratory

Allocation: 50M hours

Usage: 53M hours

Procter & Gamble and Temple University Scientists Model Skin's Makeup

Skin is the body's largest organ. It is a protective barrier, keeping microbes out and moisture in. It also regulates temperature, enables sensation, and makes vitamin D. But researchers don't fully understand at the molecular level how our skin performs its functions.

Companies developing products applied to skin, such as lotions, body washes, and cosmetics, need a deep understanding of how ingredients molecularly interact with this barrier. They need to know which chemical compounds the skin blocks, which are absorbed, and which compromise its structural integrity. Moisturizers are meant to permeate skin, for instance; cosmetics are formulated mainly for the surface.

So in 2013, researchers at consumer-products giant Procter & Gamble (P&G) and Temple University began using the Titan supercomputer at the Department of Energy's (DOE's) Oak Ridge National Laboratory (ORNL), to better understand the three-dimensional structure of skin's outermost barrier, the stratum corneum. This barrier is composed of dead cells that are impermeable to some compounds and embedded in lipids like bricks set in mortar. No quantitative model existed that could predict with certainty the amount of a specific molecule that could cross the skin.

With Titan, the researchers achieved in a matter of weeks three milestones that would have taken years with their in-house high-performance computers. First, they simulated a 1-million-atom section of the stratum corneum's lipid mortar that revealed how skin can be simultaneously permeable and resilient when exposed to agents that transform it. Second, they used their molecular dynamics simulation to measure the quantity, rate, and pathway of chemicals permeating lipids of the stratum corneum to validate high-throughput models. Third, they gained important new insights into the mechanisms by which some compounds disrupt skin.

"The multilamellar, multicomponent structure of skin is one of the most complicated lipid systems one can study," said Temple's Michael L. Klein, a National Academy of Sciences member who led the project and has collaborated with P&G for 15 years. Russell DeVane of P&G and Giacomo Fiorin of Temple were his project partners.

"Skin lipids are complex mixtures in 3D arrangements that no one fully understands," Klein said. "Through modeling and simulation, we aim to improve our molecular-level understanding of skin penetration."

Improved understanding of skin permeation is important to P&G, an $85 billion firm established in 1837 that started off in soap. (Ivory was among its first brands.) Today its product lines encompass beauty and grooming (e.g., CoverGirl and Crest), health and wellbeing (e.g., Pepto-Bismol and Vicks), and household care (e.g., Dawn and Tide). It has 25 brands with net sales exceeding $1 billion, including Olay, Pantene, Wella, and Head & Shoulders. "It's hard to think of a product that we have that doesn't come into contact with your skin at some point," said P&G's DeVane.

The company has long relied on research to innovate, but the process of testing its products is expensive and time-consuming, using skin from human cadavers, surgical patients, and pigs, as well as synthetic skin. Replacing these physical tests with computational tests that model human skin could significantly reduce costs and accelerate screening of ingredients, bringing new and improved products to market faster.

"If a better understanding of skin brings improvements that allow us to capture even 5 percent more of a market, that's a huge return on investment," DeVane said. "That's why P&G has skin in the game."—Dawn Levy

Publications:

M. Paloncyova, R. H. DeVane, B. P. Murch, K. Berka K., and M. Otyepka. 2014. "Rationalization of reduced penetration of drugs through ceramide gel phase membrane," Langmuir 30(46): 13942–13948. doi: 10.1021/la503289v.

C. M. MacDermaid, R. H. DeVane, M. L. Klein, and G. Fiorin. 2014. "Dehydration of multilamellar fatty acid membranes: Towards a computational model of the stratum corneum," Journal of Chemical Physics 141(22): 22D526. doi: 10.1063/1.4902363.

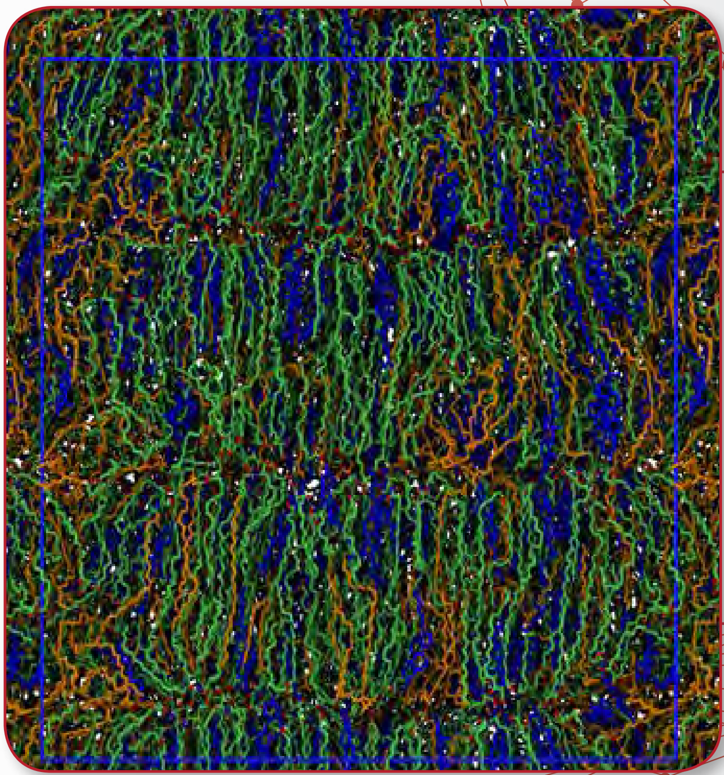

Access to Titan enabled a 2-microsecond simulation of 1 million atoms of skin lipid matrix—made of ceramides (waxy compounds, shown in green), fatty acids (orange), and cholesterol (blue) in water.

2013 INCITE

Michael Klein
Temple University

Allocation: 65M hours

Usage: 87.3M hours

Iron-based Superconductor Simulations Spin Out New Possibilities on Titan

Researchers studying iron-based superconductors are combining novel electronic structure algorithms with the high-performance computing power of the OLCF's Titan to predict spin dynamics, or the ways electrons orient and correlate their spins in a material. Because researchers have suggested that spin dynamics create the conditions needed for superconductivity, this approach could expedite the search for new or modified materials that conduct electricity with little or no resistance at higher temperatures, unlike current commercial superconductors, which must be expensively cooled to exhibit superconducting properties.

In a Nature Physics paper published in October, Zhiping Yin, Kristjan Haule, and Gabriel Kotliar of Rutgers University compute the dynamic spin structure factors—or the measure of how the spins of electrons align relative to each other at a given distance at different times—of 15 iron-based materials, including several high-temperature superconductors, in unprecedented detail.

"Our computational results are in good agreement with experimental results for experiments that have been performed, and we have several predictions for compounds that have not yet been measured," Kotliar said. "Once we validate the theory that our computational models are based on with experiments, then we can investigate materials computationally that are not being studied experimentally."

Computation offers a way for researchers to better understand spin dynamics and other material properties under many conditions, such as temperature change, rather than the singular condition present during a given experiment. Computation also allows researchers to simulate many of these materials at once, and the number of potential materials to explore rapidly increases as scientists introduce modifications to improve performance.

With the computational power at hand on the Titan system managed by the OLCF, the team was able to compare and contrast spin dynamics for all 15 materials simulated to identify tell-tale superconducting properties.

"By comparing simulations and experiments, we learned about which types of spin fluctuations actually promote superconductivity and which ones do not," Kotliar said.

In their model, the team used a technique called Dynamical Mean Field Theory to reduce the vast number of interactions involving electrons in a unit cell (the simplest repeating unit of a crystal material) and averaged these interactions in a mean field environment across the rest of the solid. The team used the Monte Carlo method to statistically select the best solutions for these techniques, achieving a new level of predictive accuracy for spin dynamics in these kinds of materials.

"We find these complex problems, as in superconductors, where you have to solve many degrees of freedom or a large number of variables, require supercomputing rather than computing on smaller clusters," Haule said. "Our algorithms are designed to work very efficiently on Titan's massively parallel architecture."

Using 20 million processor-hours on Titan, the team also discovered through simulation a new superconducting state, or electron pairing, found in the lithium-iron-arsenic compound, LiFeAs, which is consistent with experimental results.

In the future, they plan to simulate spin dynamics in other classes of superconductors and in non-superconducting materials that are exceptionally difficult to study experimentally, such as radioactive materials.

"Using computation as a substitute for experiment is an important step forward for designing new materials," Kotliar said. "The next time someone comes to us with potential materials for an application and asks, 'Should I work on this?' we hope to simulate that material through computation to select the most promising ones."—Katie Elyce Jones

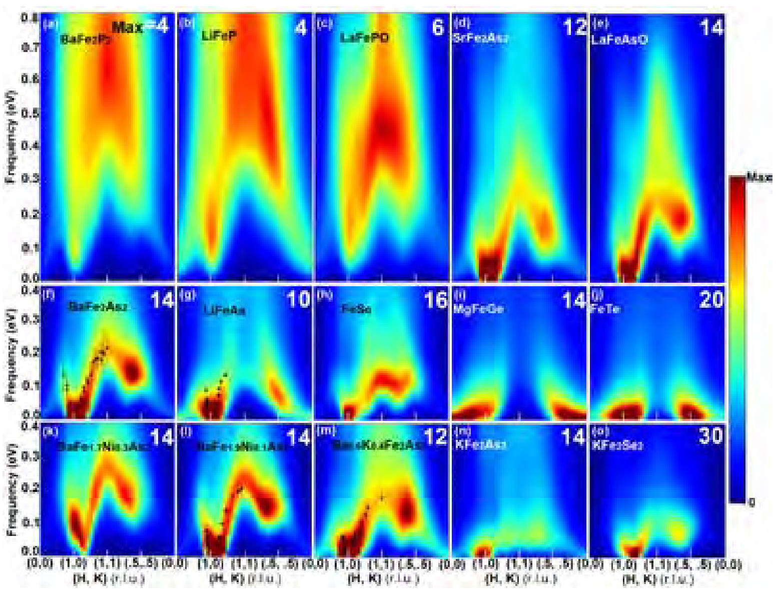

The 15 boxes in this image show the simulated intensity of spin excitations in 15 iron-based materials, including high-temperature superconductors (images d-h). Image credit: Kristjan Haule and Gabriel Kotliar, Rutgers University.

2013 Director's Discretionary

Allocation: 1M hours

Kristjan Haule
Rutgers University

Usage: 5M hours

Researchers Dig Up Biological Data from Microbes in Mines

Understanding microbial activities—and their response to climate change—is essential for understanding how climate shifts may affect large-scale systems, such as the carbon cycle. The DOE's Office of Biological and Environmental Research has invested heavily in better understanding complex microbial processes.

This investment led a group of researchers at ORNL to explore the biological functions of microbial communities with high-performance computing. The team, led by ORNL researcher Chongle Pan, studies data from analyzing tiny organisms using mass spectrometry, an analytical chemistry technique. The team recently published the results of its computational analyses in Nature Communications.

"We combine high-performance computing with high-throughput biological measurements in our experiments," Pan said. Specifically, University of California, Berkeley (UC Berkeley) biologists collected unique biological samples from the Richmond Mine in Iron Mountain, California. The ORNL researchers then measured these samples with mass spectrometry.

Mass spectrometry can generate millions of fragmentation patterns of positively or negatively charged molecules in a biological sample. Every fragmentation pattern contains thousands of unique data points. The ORNL researchers analyzed such complex, massive data using the Titan supercomputer at the OLCF. The interdisciplinary team from ORNL and UC Berkeley extracted unique insights into protein chemical changes in natural communities at an unprecedented scale and resolution from the computational results.

Microbes drive a large portion of the carbon cycle on Earth by fixing carbon dioxide in oceans and degrading terrestrial biomass such as falling leaves into carbon dioxide. The team used a proteogenomics approach for studying microbes in this ecosystem. The genetic sequence of a microbe, or its genome, reveals the microbe's potential functions. The changing protein makeup of a microbe, or its proteome, reveals its actual functions under specific growth conditions.

Often, proteins go through chemical changes called posttranslational modifications (PTMs) to diversify their functions. PTMs can alter, activate, or suppress the activities of proteins in response to different environmental conditions. However, because of the complexity of microbial communities, very little is known about how microbes use PTM in natural environments. The team developed a new approach to address this problem; using high-performance computing, they searched many types of potential PTMs on all proteins encoded by microbes in their genomes.

This new proteomic approach revealed a wide variety of PTMs on many proteins. The identified PTMs greatly increased proteins' structural and functional diversity. In addition, the researchers found PTMs were substantially different between the early and late growth stages of the community, which suggests PTMs play a significant role on many microbes' metabolic processes. Ultimately, the Pan team's findings underscore the importance of PTMs in the physiology of microbial communities.

After receiving a Director's Discretionary (DD) allocation in 2013, the team then received an allocation through the Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC) program in 2014. "We focused on our computational tools and development using the DD allocation, and then we used our ALCC allocation for full-scale analysis," Pan said.

Pan credits his collaboration with OLCF support staff for getting the Sipros computational code scaled efficiently to Titan. OLCF postdoc Juanjuan Chai played an integral role in scaling up Sipros using both OpenMP and MPI to effectively use 1,000 of Titan's nodes—a very large number for proteomics applications. As computer architecture continues to change, Pan is confident that next-generation systems—like the OLCF's Summit supercomputer, set to be installed in 2018—will have the computational power for even larger-scale analyses.
—Eric Gedenk

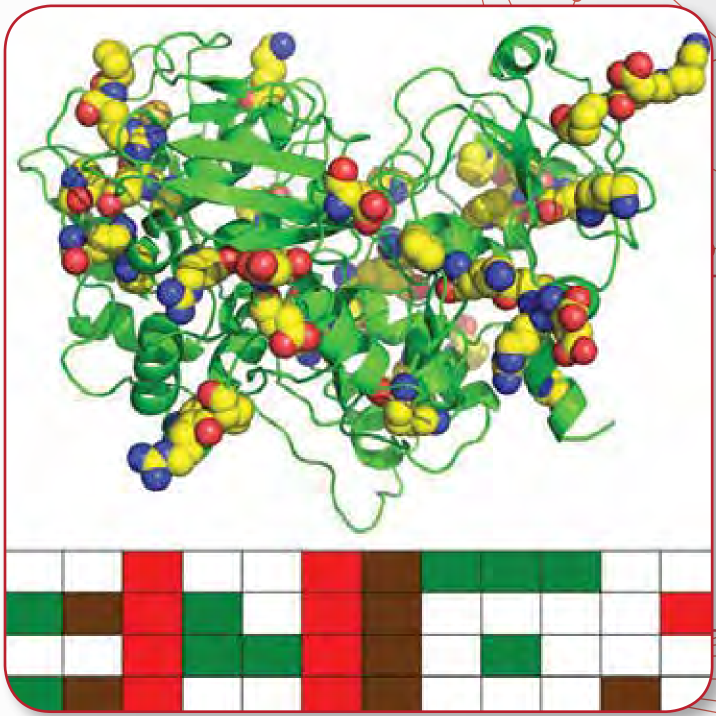

A variety of chemical changes were identified on a protein. These chemical changes were observed on thousands of proteins in a microbial community for the first time.

2014 ALCC

Chongle Pan
Oak Ridge National Laboratory

Allocation: 25M hours

Usage: 8M hours

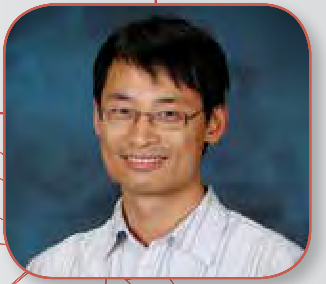

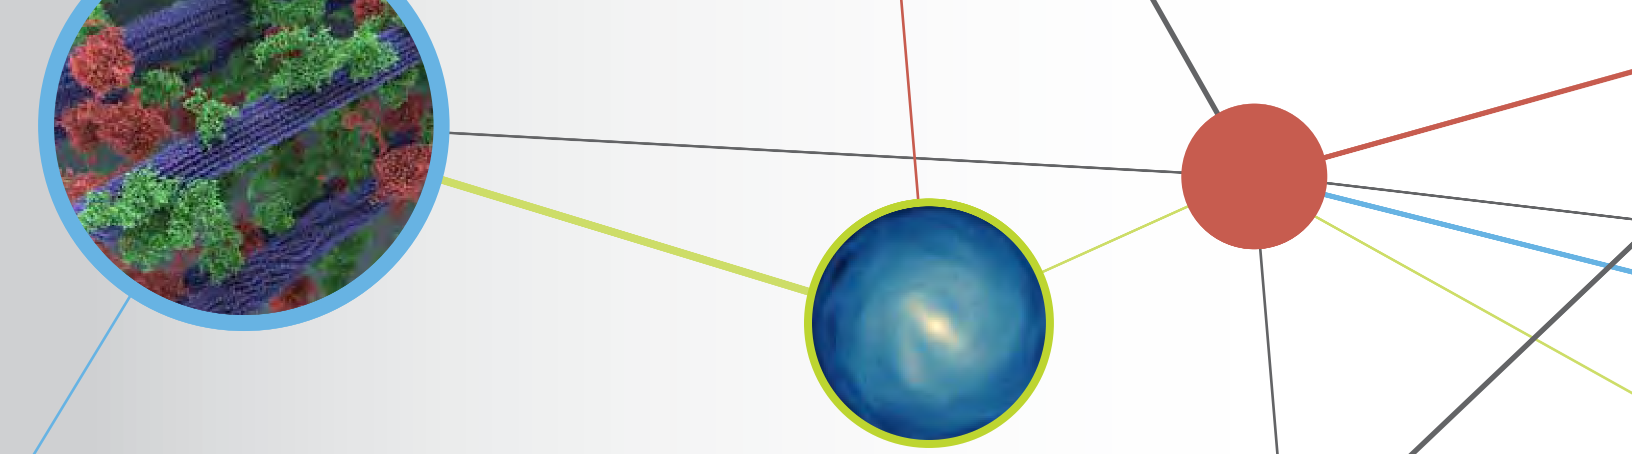

Dependability, Delivery Hallmarks of a Successful Year for Titan

After a challenging launch, the OLCF's Titan supercomputer is now running better than any previous OLCF system. Titan—the second most powerful computer in the world—is a Cray XK7 supercomputer capable of 27 petaflops, or 27 quadrillion calculations per second. First deployed in 2013, Titan had some early manufacturing issues, typical of a large serial number 1 system.

In 2013, a hardware issue was identified that could potentially lead to processor failure issues throughout the life of the machine. This defect was addressed through a wholesale processor change, under maintenance, rather than tolerating intermittent processor failures. A second round of maintenance, beginning in September and ending in December of that year, allowed Cray engineers to correct a heat-related issue in a staged activity that retained at least 80 percent of the system at any one time. Since that maintenance, the machine has been very stable and heavily utilized. "The machine is running very well," said Don Maxwell, task lead in the OLCF's HPC Operations Group. "Node failures are on par with what we'd expect. The system is very stable."

Titan's stability was very pleasing for users in 2014. In fact, the most improved score on the OLCF's 2014 user survey was the satisfaction with the number of unscheduled outages on Titan. The marks rose from a previous score of 3.8 to a 4.3 on the 2014 survey.

Adam Jacobs, team member on the INCITE project of Stony Brook University's Michael Zingale titled Approaching Exascale Models of Astrophysical Explosions, noted the smooth operation. (INCITE, which stands for Innovative and Novel Computational Impact on Theory and Experiment, is jointly managed by DOE's Leadership Computing Facilities at Argonne and Oak Ridge National Laboratories.)

"I ran simulations requiring tens of millions of CPU hours and hundreds of terabytes of storage in 2014," Jacobs said. "I had very few problems in carrying these simulations out, and the OLCF responded rapidly when I had any concerns."

Maxwell leads a team of four Cray employees and two ORNL staffers who keep Titan functioning. Their work includes planning scheduled downtimes for software upgrades, troubleshooting system problems, and doing around-the-clock duty to keep the machine running, with each team member taking a turn being on call. The beauty of this system, Maxwell said, is that every member of the team has a hand in Titan's successful operation.

Chris Fuson of the User Assistance and Outreach Group (UAO) said his communication with users confirms their appreciation of Titan's stability. "That just shows the maturity of the machine," he said, noting that jobs are interrupted less frequently. "There is more uptime, so more jobs can get through the queue and run to completion." Users reported in the 2014 annual survey an increase in overall satisfaction with Titan from a score of 4.2 in 2013 to a score of 4.5 in 2014, on a 5-point scale.

In 2014, the OLCF supported 1,071 users on more than 250 projects. Titan's users come from all over the world and represent many academic institutions, laboratories, and businesses. In addition, the research performed on Titan spans many disciplines such as chemistry, materials science, astrophysics, computer science, earth science, and biological science. Despite such a heavy and varied workload, Titan only had seven unscheduled outages during 2014.

Titan remained so stable largely because of hardware advancements that have allowed OLCF support staff to find, diagnose, and work on problems without making the machine unavailable. Cray created the Gemini interconnect, which has proven to be a valuable asset for keeping Titan

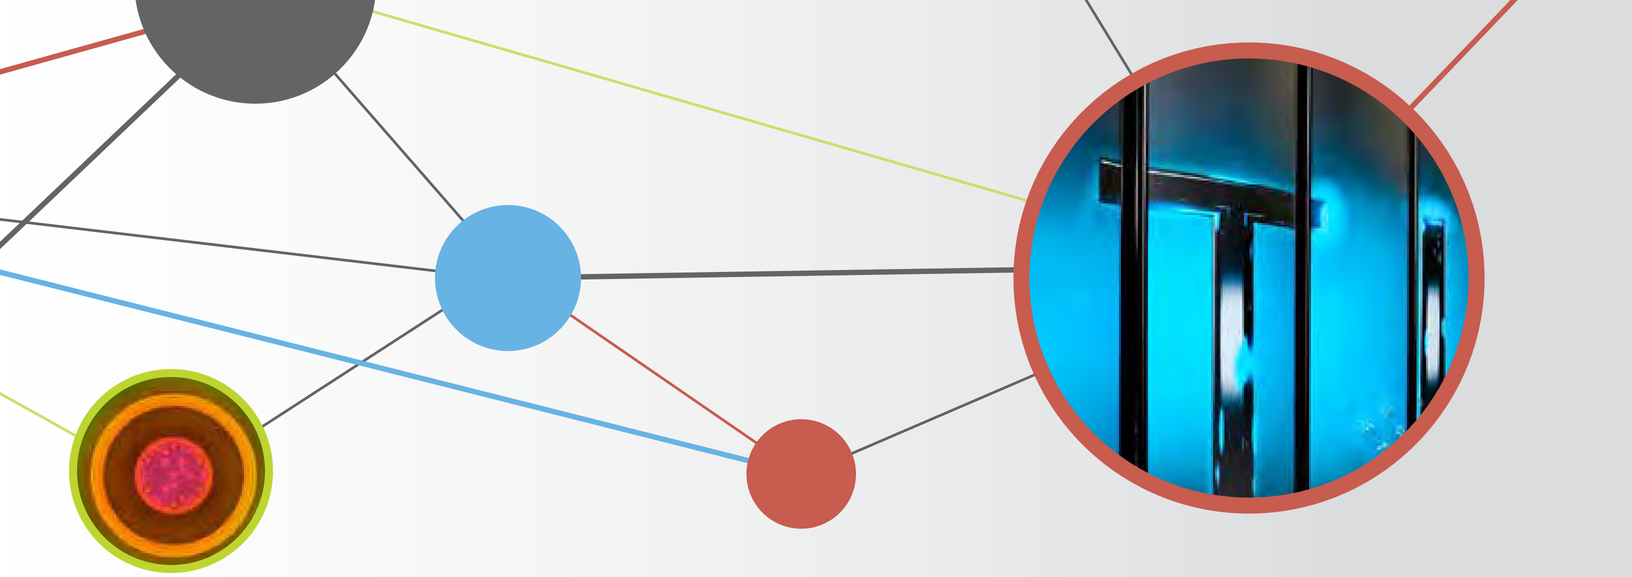

up and running. The Gemini interconnect allows staff to perform “warm swaps” of individual compute blades in the event of hardware problems.

Warm swaps allow hardware maintenance personnel to change out or repair hardware without turning off the machine. “While the machine is running, the Cray hardware team evaluates failed hardware and puts in a request to work on certain pieces of hardware, so we drain jobs from modules containing the failed hardware, and they’re able to take it out, fix it, and put it back in the system,” Maxwell said. “The Gemini network that Titan runs allows us to do that. The previous networks on the Cray did not allow us to do that, so it’s a very nice feature.”

Maxwell also pointed to advances in software and rigorous acceptance testing for making Titan more stable. The OLCF has worked closely with software vendors and has collaborated with various vendors to make their products scale to high-performance computing standards. In addition, staff members for both Cray and the OLCF work hard to make sure acceptance tests are as exhaustive as possible. “Our acceptance tests exercise both hardware and software in a variety of ways to attempt to pull out problems,” Maxwell said. “Those tests have contributed to the stability of the machine because we find bugs during acceptance, and the vendor is obligated to fix them before the machine can be accepted and then released to users for production.”

Other Center Innovations

The Titan system provides the largest extant heterogeneous architecture for computing and computational science. Titan usage is high, delivering on the promise of a system well-suited for capability simulations. This success is due in part to innovations in tracking and reporting the activity of the nodes and measuring their high degree of reliability and to the creation of novel ways to balance the workload over the entire node.

GPU Statistics and Reporting

In 2013, the OLCF reported the amount of time used by GPU-enabled jobs on Titan by using information about GPU-centric software libraries. This initial approach counted a job as GPU-enabled if any part of the job that executed had been linked against a GPU-centric software library. This technique, while enlightening, had two shortcomings:

- Jobs that did not link against any GPU-centric libraries but ran code on the GPUs via other mechanisms were not counted as GPU-enabled.
- Jobs that linked against GPU-centric libraries but did not actually run code on the GPUs were still counted as GPU-enabled.

In light of this, software developers within the OLCF User Support Group were tasked with defining and computing improved statistics for GPU activity on Titan.

By March, Cray’s Resource Utilization Reporting (RUR) tool was successfully installed on Titan and began to record GPU statistics logs for all each compute job (i.e., apruns on Cray machines). RUR reports the number of seconds that each GPU was active during a job using a GPU hardware timer. Each GPU exposes exactly one of these timers. By August, OLCF developers had programmed a new web application for storing and reporting on this new information. By correlating the RUR logs with other logs of interest, the developers were able to generate complete compute job records that were then archived into the new reporting application for future retrieval. Although the statistics using this timer provide better utilization information than the automated library tracking database (ALTD) tool, it does lose data if a job reaches its time limit and ends abnormally. We estimate that up to 25 percent of the run time on the system is attributed to such jobs. The OLCF is working with Cray and NVIDIA to find ways to capture all of the information.

With these new tools in place, the following two statistics can now be computed on Titan for any time span and any subset of compute jobs:

- Percentage of GPU-enabled runtime. This is defined as the percentage of runtime, within a queried timeframe, that employed GPUs. This statistic is comparable to the previous year's calculated GPU-enabled statistic, but staff members now know that a job used a GPU instead of inferring that it did via its associated GPU libraries.

- Percentage of GPU activity. This is defined as the percentage of compute job runtime during which a GPU was actually in use (applies only to GPU-enabled jobs).

The GPU activity statistics for different batch jobs can vary significantly based on the computational code employed and its algorithmic requirements. So although not completely appropriate for drawing generalized systemwide conclusions, these statistics are useful to OLCF staff when identifying projects that could benefit from extra assistance in GPU performance optimization tasks.

Currently, the OLCF is developing a customized RUR output plugin to streamline the collection of compute job data currently spread across many different logs. This plugin would eliminate the need for complex correlation of log files. OLCF developers are also working to provide more GPU-centric reports and to expand the application's data exploration capabilities.

Understanding, Quantifying, and Analyzing GPU Reliability on Titan

Using data gathered from RUR and ALTD, the OLCF is beginning to understand the performance efficiency of GPUs. However, the reliability characteristics of GPUs in a large-scale computing system have received less attention. Assessing, understanding, and ultimately further optimizing GPU reliability will lead to greater scientific productivity and higher operational efficiency. Therefore, the OLCF has made a long-term investment in analyzing the reliability characteristics of GPUs and how they affect system operations and applications. Staff members recently conducted a large-scale field study on GPU error characterization, quantification, and impact. Results appeared in the Proceedings of the 21st International Symposium on High-Performance Computer Architecture.

The study shows that Titan GPUs experience failures at a very low rate—lower, in fact, than vendor estimates. High-standard acceptance tests and rigorous testing of error-prone cards on another cluster during the production phase have enabled the OLCF to identify and eliminate bad cards, increasing the mean time to application interruption significantly (more than 40 hours system-wide for calendar year 2014).

The rate of occurrence of soft errors that cause GPU applications to terminate (i.e., double-bit memory errors) is low as well (one per week). This has again been possible because of careful and proactive management of GPU cards that exhibit reoccurring bit errors. The OLCF works with the

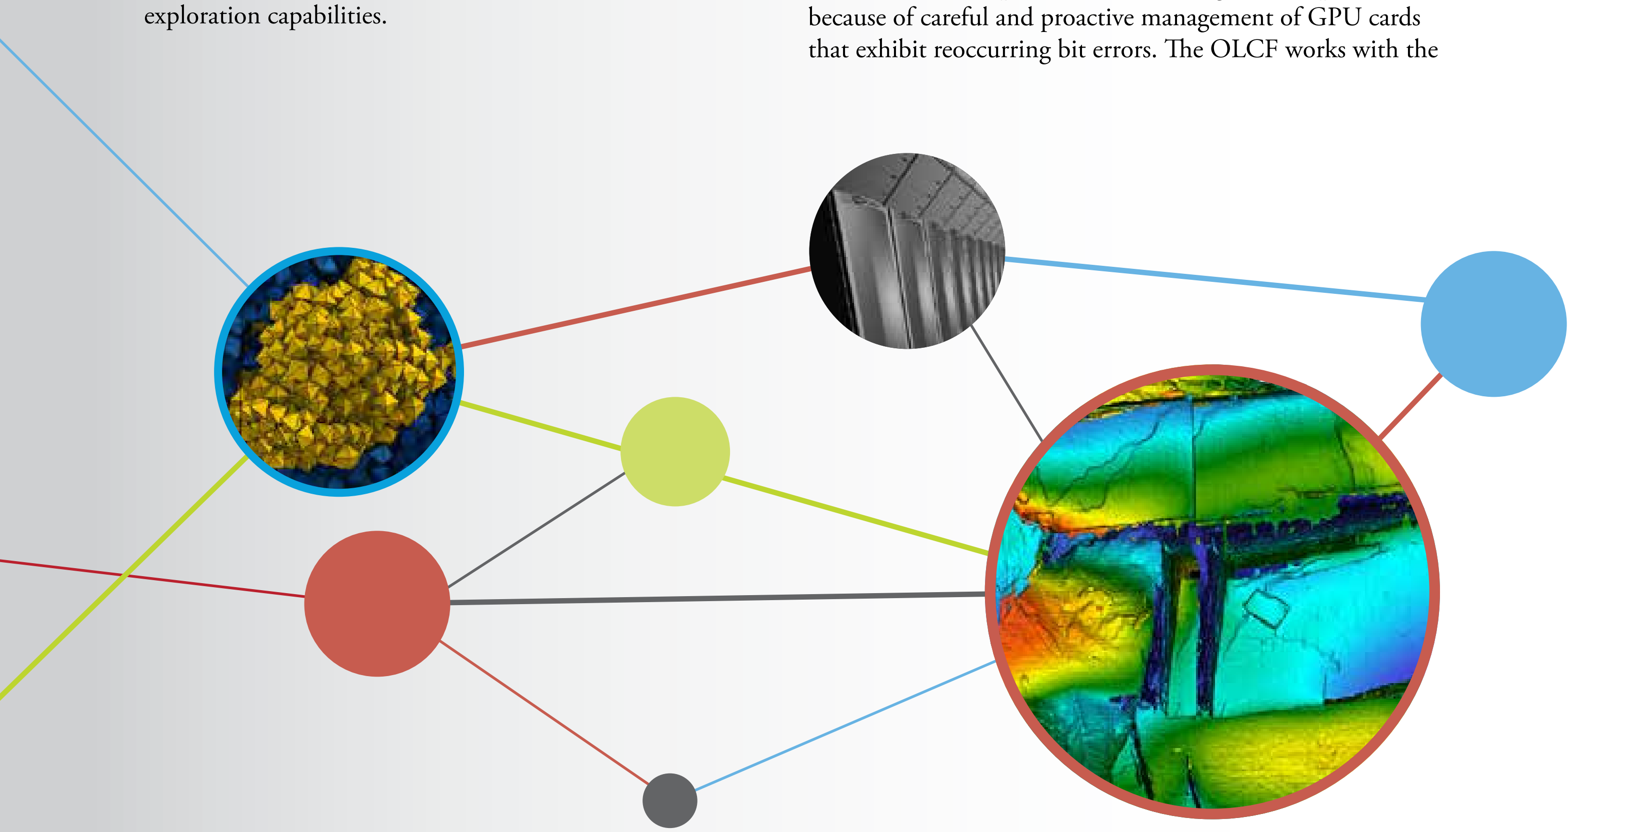

Cray staff to perform stress testing on GPU cards that exhibit memory errors. We do this on a separate cluster to verify they are working reliably before putting them back into Titan. The center also studied the impact of single-bit errors on the GPU cards. Single-bit errors do not affect an application, as they are automatically corrected by the ECC logic deployed in the memory controller. It is interesting that almost 98 percent of the single-bit errors have occurred in only 10 GPU cards (0.05 percent of the whole system). This suggests that a few cards may be significantly more prone to the recurrence of single-bit errors. The center also demonstrated that the GPU cards showing high single-bit errors can be identified early by checking the memory structure in which the errors are occurring. In addition, 98 percent of the single-bit errors occur on the GPU L2 cache, not the GDDR5 memory. This finding is useful for both GPU architects and the system operations team in identifying such cards early in the production phase.

OLCF staff members also have identified the impact of temperature on GPU system failures. A certain type of GPU system failure is relatively more frequent in cards running at higher temperatures. This has led the operations team to experiment with scheduling large-scale GPU jobs in the lower parts of the system cabinet (relatively cooler in temperature than in the upper cabinets) so the large-scale jobs are less likely to experience temperature-related failure.

Staff members have also conducted high-energy neutron-beam tests in collaboration with the Federal University of Rio Grande do Sul to study the resilience characteristics of the GPUs deployed on Titan (NVIDIA K20x GPUs). A neutron strike may flip one or more unprotected bits. The probability that a bit will actually be flipped depends on several factors, including the cell design and chip area. If a cell design is highly resilient, there is less likelihood a bit will be flipped. These radiation tests confirmed the finding about the stability of Titan GPUs and demonstrate that the GPU used in Titan is significantly more soft-error resilient than previous generations of GPUs designed by the same vendor.

The OLCF continues to work with vendors to improve the current and future generation of GPUs by helping the vendors identify which memory structures in the GPU architecture need better protection in the event of soft errors, and showing how to identify more soft-error prone GPU cards early in production. The OLCF has also found multiple inconsistencies related to vendor-provided GPU error logging tools, and are working with the vendor to resolve those. Staff members are also engaged with vendors in achieving more fine-grained and low-overhead methods to collect GPU-related errors. This is expected to have significant impact on managing GPUs, their error characteristics, and impact on production jobs. This study also is helping vendors in estimating the impact of GPU errors on future generations of supercomputers, such as Summit.

Functional Partitioning for Efficient End-to-End Computing on Heterogeneous Nodes

Smart resource utilization is critical to both optimal HPC operations and improved end-to-end performance of applications. Titan's heterogeneous node architecture presents a unique opportunity to address both of these issues. Typically, the accelerators are the primary compute engines, while the CPUs act as a master, offloading work onto the accelerators and moving the result/output data back to the main memory. An ideal balance isn't always achieved, and CPU usage may be proportionally greater than GPU usage or vice-versa. Or, due to scaling constraints or resource bottlenecks, scientific applications may not utilize all of the compute cores present on the node.

The OLCF has developed a Functional Partitioning (FP) runtime framework that helps applications use more of the resources on a compute node to expedite an application's own end-to-end workflow. It does this by concurrently scheduling tasks on a node from different parts of the job's workflow to take advantage of unused or underutilized cores and memory.

The FP runtime framework is initiated by the main application simulation to co-execute a set of tasks on a pre-specified number of CPU cores. The runtime framework creates an FP-agent per node. The FP-agent is responsible for several resource management activities such as setting up a communication infrastructure for the simulation processes and the FP-tasks to interact with each other, launching the FP-tasks, provisioning resources for the tasks, monitoring their progress, identifying unused resources, and running interference control.

The communication between the simulation and FP-tasks is accomplished via a shared memory transport atop the Common Communication Interface (CCI). CCI provides a network abstraction allowing better performance, scalability, and low latency in large-scale deployments. This is essential for communicating simulation output to the FP-task and for getting the intermediate results back from the FP-task. If the FP-agent detects that GPU resources are partially or fully available, and a GPU version of the same FP-task is available, the FP-task can be scheduled on the GPU to maintain high GPU utilization. Therefore, the FP-agent is responsible for intelligent scheduling of tasks while minimizing the performance impact on the main scientific application through interference control.

In addition to the ability to launch data analytics as FP-tasks, the framework also allows performance of I/O aggregation via an FP-task, thereby streamlining I/O, instead of multiple processes each doing I/O. To this end, staff used the Mercury I/O forwarding library that captures portable operating system interface (POSIX) calls and forwards them to a process, and developed a CCI plugin for the same.

One can also imagine future tasks such as feature extraction for visualization, health checks of a running simulation, and reliability tasks that can be co-executed with the main application simulation, using the FP framework.

The Community Earth System Model (CESM) is one of the widely used applications at the OLCF. In addition to performing the prognostic calculations, CESM also computes a set of diagnostics used for further analysis at the conclusion of the experiment. The software used for post processing and data analytics often would need to reinterpolate the model output to a standard structured grid before any analytics could be performed. Using the FP framework at large scale with CESM, researchers are able to run the regridding tasks on the output files from the ongoing simulation on the underutilized CPU cores on the same node in a pipelined fashion. This avoids the queue wait time and computation time spent in post-processing the output files. OLCF performance evaluation shows that FP can enable higher utilization of node resources, which leads to more efficient computing operations and overall higher scientific productivity.

The OLCF approach is different from extant techniques on in situ analysis in that it provides a framework for on-the-fly analysis on-node and dynamically uses underutilized resources.

Balanced Placement of I/O

The Spider file system offers a high peak throughput. However, applications do not always achieve such high rates due to a variety of factors such as contention for the shared storage, diverse access patterns, and delays at multiple levels from the compute node to the storage. Staff members have observed that a large I/O subsystem such as Spider suffers from severe contention, and there exists a significant load imbalance among the different components in the storage system.

To address the I/O load imbalance and contention issues, OLCF staff members have developed a topology-aware, balanced I/O placement strategy (libPIO) that is based on a site-defined, tunable, weighted cost function of selectable resources. The strategy is topology-aware because it assigns weight factors to separate resource components depending upon which storage layer they belong to, because different storage layers have different degrees of impact on the end-to-end performance. It is balanced as it keeps track of usage of all storage components and balances the load along the end-to-end I/O path. For evaluation purposes, the algorithm was implemented as an easy-to-use user-space library. The OLCF performed extensive experiments of libPIO on the Titan supercomputer and the Lustre-based Spider II file system.

Results demonstrate, with both synthetic benchmarks and a large-scale, scientific application, S3D, that the proposed strategy can significantly improve the I/O throughput regardless of the layout of the compute node allocation. Synthetic benchmark results indicate that the proposed strategy can improve the I/O performance by up to 50 percent and even more in some cases. For example, libPIO was easily integrated with S3D, a high-fidelity turbulent reacting flow solver. S3D writes the state of the simulation to the file system, which is later used for analysis. Only 30 lines of code needed to be added/modified in the checkpoint subroutine of the application to integrate the placement library. Evaluation results show up to 20 percent improvement in I/O performance at fairly large-scale runs (3,750 nodes—20 percent of the total compute nodes on Titan). Although tests were conducted in a production environment, substantial gains in I/O performance were observed. The ease of implementation with minimal code changes suggests that libPIO can be widely adopted by scientific users and middleware I/O libraries. OLCF personnel are actively seeking other applications to embrace this tool.

IOSI: Automatic I/O Signature Extraction from Noisy Server-Side Traces

Competing workloads on a shared storage system like Spider cause I/O resource contention and application performance vagaries. This problem is already evident in Spider and is likely to become acute at exascale. More interaction between application I/O requirements and system software tools will help alleviate the I/O bottleneck and move toward I/O-aware job scheduling. However, this requires rich techniques to capture application I/O characteristics, which remain evasive in production systems.

Traditionally, I/O characteristics have been obtained using client-side tracing tools, with drawbacks such as non-trivial instrumentation/development costs, large trace traffic, and inconsistent adoption. The OLCF has developed a novel approach, I/O Signature Identifier (IOSI), to characterize the I/O behavior of data-intensive applications. IOSI extracts signatures from noisy, zero-overhead server-side I/O throughput logs that are already collected on Spider, without interfering with the compiling/execution of applications. IOSI was evaluated using the Spider storage system, the S3D turbulence application, and benchmark-based pseudo-applications. Experiments confirmed that IOSI effectively extracts an application's I/O signature despite significant server-side noise. Compared with client-side tracing tools, IOSI is transparent and interface-agnostic, and it incurs no overhead. Compared with alternative data alignment techniques (e.g., dynamic time warping), it offers higher signature accuracy and shorter processing time.

The team was the first to study storage-server-side logs for signature extraction in the HPC domain. The extracted signature can be used in the future to perform I/O-aware scheduling such as staggering applications to prevent interference between the different applications' I/O.

Lazy Checkpointing

Checkpointing is a process in which a long-running application will periodically save information about the state of its progress to permanent storage. The perception among users about how frequently they should checkpoint varies widely and subjectively. Excessive checkpointing means less time spent solving the problem at hand as well as more contention for I/O on a shared file system. Alternative, too few checkpoints will ultimately result in larger amounts of rerun applications when there are faults that cause a job to fail. The balance is delicate because system efficiency is compromised by unnecessary I/O activity. As computing resources approach exascale—a hundred-fold increase over current supercomputers—increasing or even maintaining efficiency during checkpointing will be paramount.

The OLCF team has designed a new checkpointing technique that specifically addresses this issue. The checkpointing technique, lazy checkpointing, reduces I/O overhead and, hence, lowers the barriers for OLCF users for adopting checkpointing for their large-scale runs. The lazy checkpointing technique takes advantage of the key trait that failures have a strong temporal locality. The probability of a failure is high soon after a failure has occurred (more failures occur on the heels of a failure). The temporal locality in failures indicates that a significant fraction of failures are likely to occur within a relatively short period after the initial failure occurs, compared with the mean time between failures (MTBF) of the system. Therefore, as the time interval after the last observed failure increases, the necessity to checkpoint may decrease. Of course, there is a limit to every good thing—the lazy checkpointing technique puts upper bounds on how relaxed one can be before disrupting the balance of checkpointing versus an eventual hardware failure.

The OLCF developed detailed analytical models and simulations to study the efficacy of lazy checkpointing. Results indicate that lazy checkpointing reduces the I/O overhead significantly, by up to 30 percent, as need for checkpointing decreases over time between two failures. A prototype implementation of lazy checkpointing takes as input different system- and job-related parameters such as the job size, the checkpoint size, the OLCF file system performance, and the MTBF for the Titan system.

I/O Harness: New Testing Procedures and Introduction of a Systems Evaluation Team

I/O efficiency can be materially affected in unexpected ways. In this case, the OLCF's innovative response resulted in a new operational procedure and team, which has been shared as a best practice.

Titan is unique in its scale and computing ecosystem, which include the Spider center-wide file system that serves the high-speed storage needs of the OLCF for seamless data access between computing and visualization in the scientific workflow. The Spider file system went through a major hardware and Lustre software upgrade to enable improved performance, user functionality, and operational improvements, and was deployed to users in January 2014. Although significant testing was performed using synthetic tests and benchmarks on Spider before deployment, bugs not seen in earlier testing manifested themselves because of the size of the workload and diversity of applications running on Titan during production. User applications were impacted adversely through severe performance degradation.

The OLCF staff rolled back the Lustre software to a known, stable version to restore functionality to Titan. To mitigate future system degradations due to new hardware and/or software to the Titan ecosystem and to find a stable path to upgrade Spider as desired, the OLCF created a project team to provide a more robust systems testing approach that more closely mimicked user workload and behavior on Titan. The team consisted of members from all groups at the OLCF: Scientific Computing, User Assistance and Outreach, HPC Operations, and Technology Integration. A new application-centric test suite was created, hardened, and used repeatedly during a nine-month period to identify the issues with the Lustre software, work with the software developers, and identify a stable and performant version that was successfully deployed in October 2014. Because system maintenance and testing is competing with the need to provide production hours to the OLCF user community, testing windows are compressed to one day a week and often require the OLCF integrated project team to work late hours to make optimal use of testing time. OLCF systems evaluation team members have used this method and have incorporated it into new testing procedures to prepare for significant software upgrades to the operating system and other programming environment changes intended to improve user productivity and overall experience on Titan. Overall user satisfaction in 2014 reached its highest level ever, in no small part due to the responsiveness and subsequent proactive policy carried out by the OLCF to identify and resolve unanticipated challenges such as this.

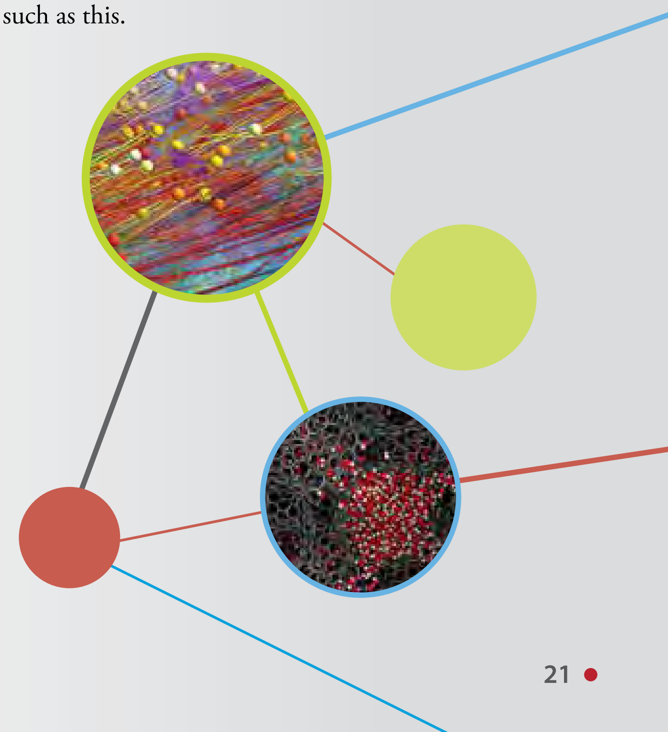

SC14

New Orleans, LA | hpc matters.

DATA Analysis and Workflow Demos Highlight DOE Booth at SC14

Eleven DOE science data pilot projects were featured at the DOE lab booth at the SC14 supercomputing conference. Of these, seven leveraged capabilities of the OLCF. While each data pilot project had unique components of its workflow, common requirements emerged, many of which could be met only by building upon the scalable compute and data storage capabilities of the OLCF, which continues to invest in innovations to promote data workflow.

Near Real-Time Analysis of Experiment

Alexander Hexemer, staff scientist at Lawrence Berkeley National Laboratory (LBNL) and Craig Tull, group leader of the Science Software Systems Group at LBNL, demonstrated the use of Titan to facilitate near real-time analysis of organic photovoltaics using x-ray scattering at the Advanced Light Source (ALS). As data was collected at the ALS, data movement to and subsequent analysis was triggered on Titan using more than 8,000 compute nodes running HipGISAXS, a massively-parallel high performance x-ray-scattering data analysis code. The reason for moving this data from LBNL to ORNL is that only Titan's GPU-accelerated architecture could run HipGISAXS with sufficient speed to process this data in real time.

Materials Science Data Analysis

Thomas Proffen, division director of Neutron Data Analysis and Visualization at ORNL, demonstrated the use of Titan to explore and classify features in large volumes of neutron scattering data from the Spallation Neutron Source. Diffuse scattering contains information about disorder in materials, a critical component in understanding material function. Advances in diffuse scattering instruments have enabled movement from 2D images to 3D volumes of measured scattering, making analysis of this data a significant challenge for scientists. This demonstration required advanced workflow capabilities coupled with the Titan supercomputer to enable near real-time feedback.

Near Real-Time Analysis

Rick Archibald of the Center for Applied Mathematics at ORNL and Sergei Kalinin of the Center for Nanophase Materials Science (CNMS) at ORNL demonstrated near real-time analysis of atomically resolved images utilizing advanced mathematics for feature detection and classification. Using high-throughput scanning transmission electron microscopy, the ferric oxide LaCoO3—a material of relevance to next-generation battery technology—was imaged with atomic resolution, resulting in more than one thousand images of the material. Current state-of-the-art at CNMS required many hours and hands-on analysis of these images to capture essential properties of the material, such as lattice structure and unit cell changes over time.

Analysis Workloads

Kenneth Read of the Physics Division at ORNL demonstrated the use of the Titan supercomputer coupled with the worldwide Large Hadron Collider computing grid to contribute to analysis workloads that will improve measurements and understanding of the recently discovered Higgs boson and its decay modes as well as to facilitate searches for supersymmetry. This coupling was made possible by recent advances in the ASCR-funded BigPanDA workflow management system.

Cosmology Simulations

Peter Nugent of the Computational Research Division at LBNL demonstrated the use of the Titan supercomputer for one of the highest resolution cosmology simulations to date to understand and constrain systematic uncertainties in cosmological observations. These and other simulation-based analysis approaches to understanding cosmological observations are of increasing importance as the ability to generate massive observation datasets is far outstripping the ability to effectively analyze and understand them. Instruments and observation surveys such as the Dark Energy Survey, the Sloan Digital Sky Survey, and the Large Synoptic Survey telescope are generating catalogs that will range in size from one to several hundred petabytes.

Complex Nanostructures

Simon Billinge of the Condensed Matter Physics & Materials Science Department at Brookhaven National Laboratory demonstrated the use of the Titan supercomputer at the OLCF coupled with scattering data from the National Synchrotron Light Source, Spallation Neutron Source, and Advanced Photon Source to elucidate the fundamental structure of complex nanostructures. Traditional crystallography provides a mechanism to identify material structure based on x-ray, neutron, or electron scattering by solving the inverse problem for a simulated scattering pattern. This type of complex modeling requires significant computational capabilities available only at Leadership Computing Facilities. In this demonstration configuration, three weeks of computation time on a modern desktop computer is required to generate and evaluate a single candidate solution.

PYRAMID: Building System Support for Complex Application Workflows

The goal of PYRAMID is to provide a lighter weight alternative to the native Cray/ALPS task management system, to support complex workflows with more concurrent tasks. It is based on the Scalable runTime Component Infrastructure (STCI) runtime system, which has been under development with OLCF support for some time with the primary goal of providing a more resilient runtime to support a variety of tool and communication middleware needs.

The PYRAMID project has extended STCI to support persistence of key portions of the runtime infrastructure across multiple computational tasks and to provide two-level resource management, which separates the assignment of job resources to computational tasks and deployment of the task to the assigned resources.

Ascending to Summit

HPC is more vital than ever. As America adapts to global changes and demands, so too must our supercomputers. Since its deployment in late 2012, the OLCF's Titan supercomputer has been the nation's most powerful system for advancing scientific discovery.

"Supercomputing is a national priority," said Jack Wells, ORNL's director of science. "Most every area of science and engineering today depends on it. And to be excellent in science and engineering in the 21st century means to be excellent in supercomputing."

On November 14, 2014, the coming of Summit—a system expected to push HPC capabilities farther than ever before—was announced.

"Summit will provide at least a fivefold increase in computational power. It's going to provide significantly greater memory, greater interconnection opportunities, and greater performance at the node level," said Jim Hack, director of the National Center for Computational Sciences at ORNL. "What that means is, scientists are going to be able to do much more realistic work on the new machine than they can do right now on Titan. And Titan is already enabling some pretty amazing things."

Summit is the result of a cooperative endeavor consisting of hundreds of laboratory decision makers, HPC experts, and vendor partners. This initiative, dubbed CORAL, is the collaboration between the two DOE Leadership Computing Facilities—the OLCF and Argonne Leadership Computing Facility (ALCF)—along with the National Nuclear Security Administration's (NNSA's) Lawrence Livermore National Laboratory (LLNL).

The establishment of CORAL (an acronym for Collaboration of Oak Ridge, Argonne, and Lawrence Livermore National Laboratories) was based on a joint memorandum of understanding between DOE and NNSA, which outlined goals for increasing coordination in HPC research and development and future HPC acquisitions.

This collaboration made it possible to share a broader range of technical expertise among the laboratories, improving the quality of technical requirements and improving efficiency across three DOE procurements. The result of the collaboration was a streamlined proposal process that culminated in the selection of IBM systems for both the OLCF and LLNL—subsequently named Summit and Sierra, respectively.

Immediately following the selection announcement by US Energy Secretary Ernest Moniz, the OLCF launched a website for Summit (https://www.olcf.ornl.gov/summit/), providing future users and the public with technical details and other important information on the new system. The website also offers an opportunity for potential users to get a head start using Summit through the Center for Accelerated Application Readiness (CAAR) partnerships project's early call for proposals, in which users will receive support in code refactoring and portability from ORNL's IBM/NVIDIA Center of Excellence and OLCF staff. Users will have access to computational resources like Titan, as well as access to systems with diversified architectures such as Mira at the ALCF, Edison at the National Energy Research Scientific Computing Center, and other early delivery systems as they become available.

A new set of CAAR applications were selected and announced in the Spring of 2015, and the project teams will be able to immediately start working to prepare their codes for Summit.

Between now and Summit's arrival in 2018, the OLCF will host a series of workshops, webinars, and tutorials to help prepare users for getting the most out of Summit and the facility's data and storage resources.

Until then, Titan will continue to be the workhorse of the OLCF, generating overwhelming amounts of data and shedding new light on many areas of science such as combustion, climate, energy storage, nuclear power, and many more.

"The OLCF is delighted to once again be delivering a new leadership computer for the DOE and science community," said Buddy Bland, the project director for Summit. "It will be one of the world's most powerful computers when delivered and will provide a natural follow-on system for the users of Titan."—Jeremy Rumsey

Introduction

For many OLCF users, Titan's hybrid architecture is their first experience with GPU accelerators. Because parallel computing significantly increases the amount of work users can perform on Titan and because many future HPC systems, including the OLCF's next-generation supercomputer Summit, will exploit parallelism, users are investing time and energy in modifying their applications to excel on accelerated architectures. To help users make the most of Titan, the OLCF is constantly seeking user input, developing new ways to inform and train users, and upgrading OLCF resources to meet the evolving needs of the HPC user community.

Training at the Threshold of Accelerated HPC

'Getting Started' rings in new projects on Titan

Entering 2014 with a new round of INCITE and ALCC projects, the OLCF offered two sessions of the three-hour "Getting Started" workshop geared at relaying essentials to new users. The first "Getting Started" workshop also was delivered as a webinar for users unable to visit ORNL.

"'Getting Started' was intended to give users a quick start to the center and a good overview of common pitfalls new users face when using leadership computing resources," said Fernanda Foertter, OLCF HPC user assistance specialist.

User assistance staff presented best practices for installing and updating software; resolving common runtime errors; using data management resources for storing and transferring data; downloading and running debugging tools like Allinea's Distributed Debugging Tool and Cray's performance analysis tool, CrayPAT-lite; and staying informed about OLCF computing resources via email lists, online status indicators, the @OLCFStatus Twitter feed, and OLCF Status smartphone apps.

OLCF SciComp Team Leads Portability effort

Members of the OLCF Scientific Computing Group have been leading efforts to enhance portability among ASCR facilities. The discussions will help inform efforts at each institution to train and assist users with code preparation for systems being deployed in the next 2–4 years. ASCR—which manages the OLCF, ALCF, National Energy Research Scientific Computing Center (NERSC) and ESNet—considers application portability and performance portability to be critical components of application readiness to better accommodate users who run their applications at multiple computing facilities.

First OpenACC Hackathon cuts coding time for users

Writing codes that are part science and part computer programming can be an all-around challenge. The OLCF invited users and GPU developers to spend an intense week code hacking at the OLCF's inaugural Hackathon October 27–31 in Knoxville, Tennessee.

During OLCF code-scaling training sessions, Fernanda Foertter, who leads training events and workshops, began identifying where users were likely to hit snags in scaling their scientific applications to Titan's GPUs. Users who experienced setbacks were trying to port their code but were becoming stuck, lacked the staff to properly conform code for hybrid systems, or were porting code but not getting the desired performance.

"I thought, 'Why not get the scientific developers together with the GPU developers?'" Foertter said.

After a two-week open call, Foertter chose six teams representing diverse scientific disciplines from climate to machine learning to participate in the event, which focused on porting scientific applications to Titan's GPUs using OpenACC, a compiler directive used to program GPUs.

OpenACC has been a vital tool for Titan and its users and has allowed researchers to attain a high level of parallelism and performance from their code, resulting in improved time to solutions and higher fidelity in simulations.

The event hosted 45 people over 5 days. Six teams of up to 12 users each were paired with two expert mentors from Cray, NVIDIA, The Portland Group, The Swiss National Computing Centre, and the OLCF's Scientific Computing and UAO groups.

"Without our mentors, it would have taken us several months to do what we were able to achieve in a week," said research scientist Rangan Sukumar of ORNL's Computational Data Analytics Group.

Since the Hackathon, vendor partners who participated have gathered compiler bugs observed during the event that might not have been picked up or submitted by users through traditional means, and three teams have requested projects with the OLCF to continue their application development on Titan.

Joint Facilities User Forum brings HPC staff together on data

In partnership with seven national laboratories, the OLCF co-organized a 3 day workshop to discuss approaches for handling data and the future of data-driven scientific discovery. The Joint Facilities User Forum on Data-Intensive Computing brought together users and HPC center staff June 16–18 in Oakland, California, to discuss advances in managing, analyzing, and visualizing data and successful versus unsuccessful methods.

"This is a continuation of what we started in Oak Ridge in 2013 with the workshop on analysis of very large datasets. National Energy Research Scientific Computing Center staff attended and expressed interest in producing a joint meeting the following year. Before we knew it, six laboratories were interested in participating," Fernanda Foertter, OLCF HPC user assistant specialist, said.

Several OLCF staff members took part in the workshop. Foertter and Ashley Barker, UAO group leader, served as session chairs during first and second day meetings, respectively. The OLCF's Chris Fuson presented on transferring large data sets over the wide area network, and Norbert Podhorszki gave a presentation on using ORNL's award-winning Adaptable I/O System, ADIOS. ORNL's Rangan Sukumar delivered a presentation on machine learning for data-driven discovery, and Stuart Campbell gave a presentation titled "Accelerating Scientific Discovery at the Spallation Neutron Source," a DOE neutron scattering facility located at ORNL that generates large amounts of data that can be more rapidly analyzed with leadership computing resources. Other ORNL participants included Clay England, Doug Fuller, and John Harney.

"It is interesting to hear how other centers do things, and it was a good experience to talk about these topics," Fuson said. "Our goal is to help users benefit from this collaboration."

OLCF introduces new use for user metrics at SC training workshop

At the SC14 supercomputing conference in New Orleans, several HPC centers joined the OLCF's lead in holding a workshop for those who develop and deliver HPC workshops and training programs. Foertter thought HPC centers such as the OLCF, NERSC, and others could benefit from sharing their training programs, so she and a group of HPC center colleagues led the first "Best Practices in HPC Training" workshop at SC14. About 50 people attended.

"We shared what's working and what's not working," Foertter said.

The format included a dozen 10-minute talks on workshop development, offering training courses, developing survey options to poll user response, and more. Foertter, who has begun gathering metrics on workshop attendees, presented on a new approach to tracking training impact.

"Until this year, if you had asked me 'How many users are attending our workshops?' I couldn't have told you very much," Foertter said. "We had head counts, and we knew who was coming to the workshops, but we weren't keeping up with the ratio of users to postdocs to graduate students. We weren't tracking what cross-section of users we were drawing: where are they from? Are they returning users? Did they go on to become users after attending the workshop?"

Foertter realized that by enrolling people in workshops, the OLCF already collected user data—but OLCF staff just weren't using that data to analyze the outcome of workshop training on different user groups. "Now we will," Foertter said.

Education, Outreach and Training

Tiny Titan teaches basics of parallel computing

When OLCF staff members Adam Simpson, Robert French, and Anthony DiGirolamo, built a portable nine-core parallel computer to illustrate the basics of parallel computing to students, they didn't expect the buzz "Tiny Titan" would generate.

"In the past there hasn't been a good way to explain to people, in a few minutes, how a supercomputer works," Simpson said. "Tiny Titan allows us to do that."

Since its development in December 2013, Tiny Titan has represented the OLCF at regional schools; through local and national news outlets such as Popular Science; at the US House of Representatives' National User Facility Organization meeting in June 2014; and at DOE's National Laboratory Day in Washington, D.C., in September 2014. A Tiny Titan exhibit is open at the American Museum of Science and Energy in Oak Ridge, and a Tiny Titan is displayed on the OLCF's observation deck, where staff members frequently give demonstrations for visitors and school groups.

Like Titan and Tiny Titan, most electronics operate on multicore systems. However, US schools are still primarily teaching serial processing, or computer processing on just one core.

"Students need to learn to program in parallel," French said. "Computers will continue to get more cores and become more complicated."

Tiny Titan visually teaches how multicore computing works. Each of the nine nodes includes a different colored light, and images on the connected monitor use the same colors to show what each processor is doing. The more colors that light up on the computers, the faster the program will run.

At its affordable cost of just under $1,000, Tiny Titan is designed for students to build and operate in a classroom setting. The OLCF team also manages a GitHub account, www.github.com/TinyTitan, where anyone can find instructions for building their own Tiny Titan as well as beginner parallel programming exercises.

For their outstanding efforts, the team received a Science Communicator Award at the 2014 ORNL Awards Night event.

HPC Fundamentals course offered lab-wide

Domain scientists today are encountering new and unexpected opportunities to use supercomputing in their own research, but it can be a difficult topic to tackle without a foundational understanding of HPC. That's why, for the fourth consecutive year, the OLCF in collaboration with the National Institute for Computational Sciences (NICS) offered an 8 week "HPC Fundamentals" course at ORNL.

Open to laboratory staff and interns, the course offered basic HPC concepts, tools, and terminology. The classes covered topics such as the role of HPC in science today, an introduction to UNIX/LINUX operating systems, and programming in Fortran, the general-purpose programming language for scientific applications.

Bobby Whitten, group leader for the NICS User Assistance Group, led classes with assistance from the OLCF's Suzanne Parete-Koon, Robert French, and Adam Simpson. Whitten said the goal of the course was "to give people an introduction to what supercomputing is all about—to demystify it just a little bit."

Participants in the course agreed.

"I really enjoyed that the sessions allowed us to learn through programing," said Tosin Alabi, a graduate student at Purdue University and a summer intern in 2014 through the HERE Program. "I found the examples simple and well scoped. This allowed me to gain knowledge of the parallel library functions more easily. I would defiantly recommend the course to anyone looking to get started with HPC!"

Handing the supercomputer over to the students

A group of graduate students from the University of Tennessee's (UT's) Department of Physics and Astronomy had the unique opportunity to perform research on Titan. Graduate students enrolled in Computational Physics (Physics 643) are required to complete an independent, semester-long project related to current physics research appearing in a recently published journal article.

Through use of an OLCF Director's Discretionary project led by UT Physics Professor Ken Read and distinguished research staff member in the Physics Division at ORNL, nearly a dozen Physics 643 students interested in using Titan were allocated time on the system.

OLCF staff members Fernanda Foertter, Suzanne Parete-Koon, and Adam Simpson supported the students and taught classroom sessions designed to deliver HPC basics. The students involved were researching a range of scientific disciplines; some were using codes already developed for Titan, while others chose to write their own. To meet different needs, Foertter, Parete-Koon, and Simpson provided technical support throughout the duration of their projects.

"We want to foster a future generation of scientists who will want to compute on Titan," Foertter said. "A lot of students don't know these supercomputers exist or that they are accessible to them or their professors. It's not just a lab resources—it's a worldwide resource."

While most strong physics departments offer a computational physics class, Read said he's not aware of any leading supercomputing center making their resources available to a class of college students.

"They usually don't award hundreds of thousands of hours for those projects and contribute multiple staff to help co-teach and design the class," he said.

The experience was especially meaningful for the students.

"These seminars were given during class and consisted of an introduction to high performance computing as well as interactive lessons on how to program for these systems. Class access to Titan was provided, and we began learning how to use the OpenACC standard to take advantage of the GPU processors that many HPC systems have begun utilizing in recent years," said graduate student Nick McNutt. "The result of these seminars inspired individual student projects, and the OLCF graciously provided one hundred thousand hours of compute time on Titan for this purpose.

"My project in particular used LAMMPS to perform an atomistic simulation of a carbon system. Adam Simpson assisted me with this, and I was able to discuss ideas for the project with him and ask questions on how to best implement the simulation on Titan," McNutt said. "I consider this experience to have been highly valuable, as the skills I obtained from the seminars have had direct applicability to my computational research as a graduate student."

OLCF summer interns share their experiences

Every summer, hundreds of college students join the ranks at ORNL as interns, and in 2014, the OLCF attracted 29 interns interested in gaining hands-on exposure to HPC tasks. Here are a few examples of projects on which some of the interns focused.

Sean McDaniel, who plans to pursue a PhD in computer science at the University of Delaware in 2015 following his OLCF internship, created smaller, faster-running kernels that mimic real-world applications used to test supercomputers and ensure the hardware and software are always robust. McDaniel worked specifically on the HACC framework.

"HACC is this big scientific application that has to do with simulating the birth of the universe, but it's huge," McDaniel explained. "So we're trying to extract the key parts and shorten the testing time while still capturing how the application interacts with the system."

Daniel Wherry, a sophomore in computer science at Austin Peay State University, spent his summer developing a tool to improve users' everyday research experiences. His project, "Estimating Lustre Striping Impact on Scientific Application Checkpointing," gives users a better understanding of how data is saved as well as insight into different strategies they can use to improve their speed.

"Basically the tool will mimic the part of the programs they make that actually write to the Luster File system and apply different striping configurations to those writes. Striping is just one way of improving the speed of writing to a file," Wherry said. "The end goal is that the user will be able to boost parts of their program that need to be as quick as possible so they can get back to doing the real science."

Chris Martin, a sophomore at the University of Tennessee who decided to change his major from economics to computer science after his OLCF internship, combined graphic design experience with an interest in computer science to present energy data for the Titan system.

"I've been working on a project that is basically a web front end for power and water usage effectiveness coupled with humidity and temperature," Martin said. "It's a radial graph that shows how those values fluctuate within 24 hours, how they fluctuate over 7 days, 30 days, 90 days, a year. It's interesting data because it jumps around a lot."

Martin said he gained broad communication skills to low-level UNIX system knowledge during his internship—skills sure to help him in his future as a computer science major and beyond.

OLCF Executive Board and User Groups

Transition to elected, rather than appointed, user board increases engagement

For the first time, OLCF users elected representatives to the OLCF User Group (OUG) User Executive Board (UEB). Formerly the OLCF User Council (which comprised interested and active users appointed by the OLCF), the UEB eventually will be entirely user-appointed. The board's functions are to advise and provide feedback to the OLCF regarding policies, resources, and requirements and to lead the OUG. Balint Joo of J-Lab is chair of the OUG.

User assistance specialist Suzanne Parete-Koon wrote a charter, with input from the existing User Council, to establish the UEB. The charter states that 10 members including a chair and vice-chair will serve 3 year terms. With three members elected in June 2014—vice-chair Mike Zingale of Stony Brook University and Rangan Sukumar and John Turner of ORNL—the entire board should be elected members by 2016.

The move hasn't just changed board operations; it also has encouraged a greater number of users to engage more often. More than 100 users submitted votes for the election, and the monthly user meeting has seen a rise in participants.

"The election publicized the fact that the board was present, and we now have 20–50 people participating in user group calls every month," Parete-Koon said. "One reason is that members have to run on a platform explaining why they want to be UEB members, which leads to more user-driven discussion. Users are more likely to call in and participate if they feel they are represented."

One issue already being discussed by the new board reflects increased user interest: how to facilitate user-driven content. Users are discussing how to share content, such as custom software tools developed for their own projects, that have broader applications other users might find helpful.

OLCF User Meeting introduces successful research angle

The annual OLCF user meeting drew 104 Titan users, including 26 principal investigators, to ORNL in July. The 2014 meeting included a significant formatting change that received positive feedback from participants and OLCF staff.

"In the past we've done a user meeting that was very training oriented—how to get on the system, basic debugging, how to compile your code—that sort of thing," said Judy Hill, an OLCF computational scientist who was on the planning committee for the user meeting. “This year we really wanted to attract the PIs and the software developers, especially since we’re preparing for software readiness for OLCF-4.”

While users still received training and skills preparation on the first day, as in years past, the rest of the event was much more research oriented. The format included a keynote speaker in the morning followed by an afternoon session focused on the scientific progress users had made on Titan.

“We had a science PI paired with a software developer,” Hill said. “So the PI set the stage for the overall science, and the code developer talked about the overall challenges that they had to overcome.”

The OLCF has always administered surveys and other methods to collect user comments—both positive and negative. However, Hill believes that in-person conversations about user experiences, such as those that came about during the research-oriented sessions, are much more effective.

“This was a much bigger success than we thought—a wild success,” Hill said. “So now I think the challenge is to meet the expectation that we set and exceed it for next year—try to build on what we’ve got.”

INCITE, ALCC, and Director’s Discretionary

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

For large-scale, computationally intensive research projects, DOE’s Leadership Computing Facilities lead the INCITE program. INCITE awards are sizeable allocations (typically, tens to hundreds of millions of processor-hours per project) to address grand challenges in science and engineering. In 2014, 30 INCITE projects were allocated a collective 2.25 billion hours on Titan. For more information or to apply for an INCITE project, visit www.doeleadershipcomputing.org.

The ALCC program is open to scientists from the research community in national laboratories, academia, and industry. The ALCC program allocates computational resources at the OLCF and other computing facilities for special situations of interest to DOE programmatic needs with an emphasis on high-risk, high-payoff simulations in areas directly related to DOE’s energy mission. For more information or to submit a proposal, please visit www.science.energy.gov/ascr/facilities/alcc/.

Director’s Discretionary projects are dedicated to leadership computing preparation, INCITE and ALCC scaling, and application performance to maximize scientific application efficiency and productivity on leadership computing platforms. The OLCF Resource Utilization Council, as well as independent referees, review and approve all DD requests. Applications are accepted year round at www.olcf.ornl.gov/support/getting-started/olcf-director-discretion-project-application/.

There’s not just one path for one project

A General Electric (GE) Global Research team led by Vittorio Michelassi first took advantage of support resources at the OLCF through a DD project that ultimately led to a winning 2015 INCITE award. Michelassi’s project uses a high performance solver for turbulence and aeroacoustic research (HiPSTAR) that enables direct numerical simulations of flow under real aircraft engine conditions. The GE Global Research team hopes that with a more accurate computational model, they ultimately will be able to increase gas turbine efficiency by 2 to 4 percent.

Because HiPSTAR requires a lot of computational power, GE applied for a DD project in 2012 to port the solver to Titan’s accelerated architecture. Support staff from the OLCF’s Cray Center of Excellence and the OLCF helped the team scale HiPSTAR on Titan, and the GE team plans to leverage the advances from this DD project in its 2015 INCITE project.

Specialized User Support for Accelerated HPC

New system upgrades improve data storage, transfer, and analysis for users

In 2014, users began a new year with a new center-wide file system, Atlas. At 32 petabytes of storage capacity, Atlas provides a location to temporarily store large amounts of data produced on Titan, Eos, and Rhea (a 196-node Linux cluster that provides a conduit for large-scale scientific discovery via pre- and post-processing of simulation data generated on Titan). Atlas also has a new directory structure that organizes files by project so that users with multiple projects now have multiple work areas.

"Previously, a user's project and their work areas were separate, making it extremely difficult to share data between the two," said user support specialist Chris Fuson. "As a result of the new directory structure, users are no longer required to change permissions on their project directory to share data.

Atlas is divided into two separate file systems, Atlas1 and Atlas2, so that one file system is available at all times for users to manage their project data.

"The increased capacity and greater performance will allow users even greater depth in their research, with less restrictions on data, both how fast it's generated and the overall amount," said Dustin Leverman, deployment lead for Atlas.

Following the Atlas transition, the OLCF also is in the process of upgrading the Network File System (NFS), which serves as user and project home directories. Plans to quadruple NFS storage from around 45 terabytes to approximately 140 terabytes will give users more long-term storage for source code, batch script, and other data requiring routine access.

In March 2014, Eos, the OLCF's 744-node Cray XC30 cluster, became available to all users to provide additional computing resources for the OLCF users. Eos now also supports cross-system job submission. Titan, Eos, and Rhea can automatically direct analysis and storage of large data sets via batch scripts.

"Now users can run compute jobs on Eos and then automatically visualize the results on Rhea, for example," said Robert French, user assistant specialist.

Transferring data to the OLCF's High-Performance Storage System has also improved thanks to testing of the archival HSI/HTAR data transfer nodes.

OLCF adds Globus Endpoint and updates guides to make multiple-streaming transfers easier

OLCF staff members extended and updated the data management user guide after a thoughtful review and an examination of what challenges users faced when conducting large data transfers.

Users routinely produce tens or hundreds of terabytes of data, and many users predict their needs will multiply significantly in the next 5 years. While the OLCF provides scratch and archival storage, users often need to move terabytes of data to other facilities for analysis and long-term use.

"We're adding data at an amazing rate," said Suzanne Parete-Koon, user assistance specialist.

In 2013, the OLCF revisited its data transfer capabilities and dedicated two nodes for outgoing batch transfers. At the beginning of 2014, that capability increased to 10 nodes. These dedicated batch nodes provide a scheduled data transfer source determined by user requests.

However, OLCF staff Parete-Koon, Hai Ah Nam, and Jason Hill discovered researchers were often underutilizing resources by using single-stream transfer methods such as scp rather than multiple-stream methods that break data into multiple, simultaneous streams, resulting in higher transfer rates.

"Using scp when multiple-stream methods are available is akin to drivers using a single lane of a 10-lane highway," Nam said.

Staff tested three multiple-stream tools optimized for OLCF systems—bbcp, GridFTP, and Globus Online—for ease of use and reliable availability. Results showed that bbcp and GridFTP transferred data up to 15 times faster than scp in performance tests.

Through a survey and conversations with users, the team identified why researchers were stuck on scp. Users were willing to sit on standby during long stretches of data transfer not only because they were typically more familiar with scp, but also because some multiple-stream methods require a lengthy setup process that can take several days, creating a high barrier to entry.

"Because of a high level of cybersecurity at the OLCF and other computing facilities, transferring files requires additional steps that seem inconvenient and overly complicated to users," Parete-Koon said. "For example, to use GridFTP, OLCF users need an Open Science Grid Certificate, which is like a passport that users 'own' and are responsible for maintaining."

Despite the improved transfer capabilities of GridFTP, only 37 Open Science Grid Certificates had been activated on OLCF systems as of November 2013.

Based on the performance test results for bbcp, GridFTP, and Globus Online, the OLCF is recommending users with large data transfers apply for the certificate and take advantage of the specialized data transfer nodes.

To support Globus Online, the OLCF set up a Globus Endpoint that allows 12-hour automated transfers.

"For a long time, the OLCF had an unofficial Globus Endpoint, and we needed to make sure that supporting an official endpoint wouldn't create security issues," Parete-Koon said.

To help users navigate the complicated setup process, Parete-Koon wrote documentation and presented information on multiple-stream methods, including how to set up a proxy certificate and access Globus's online services. During a user conference call, she shared a sample, automated workflow that uses three scripts: the first to schedule a data transfer node, the second to launch an application, and the third to transfer files to longer-term data storage or a remote site for analysis.

"After the conference call, one user thanked us because he was able to transfer 2 years of data—almost 19 terabytes—in five streams with an average rate of 1,290 megabits per second," Nam said. "That is about five times faster than an optimized scp transfer."

In addition to providing users with the information to carry out data transfers, UAO staff assisted multiple projects with creating workflows that used Globus Online and ORNL's Compute and Data Environment for Science) data transfer nodes for serving data to the broader scientific community. Among these projects were large-scale cosmology and climate simulation projects that plan to share data with an outside community of researchers.

To further increase support for data transfers, the OLCF has set up and is utilizing a center-specific PerfSONAR system—a network monitoring system run by an international collaboration including DOE's ESNet. PerfSONAR allows system administrators to troubleshoot data transfers out of the OLCF.

New accelerator guide continues to grow in 2014

Launched in November 2013, the OLCF's Accelerated Computing User Guide was updated throughout 2014 to expand GPU-related documentation for users.

"Through an annual survey, users indicated a desire for more online GPU-related documentation," said Adam Simpson, user support specialist. "I used that feedback and GPU-related user tickets to determine what would be useful to cover in the guide."

With input from other staff, Simpson wrote the guide that provides background, basics, techniques, and additional resources on GPU computing.

"It contains a lot of information that is specific to accelerated computing on Titan," Simpson said.

The accelerator guide was the fourth most frequently accessed user guide on the web in 2014, followed by the Titan, Eos, and Rhea system guides.

OLCF Web Traffic Statistics

General Statistics	
OLCF website unique visitors	122,287
Unique page views	278,434
Top page views	
Titan	61,239 unique views
Landing page	24,178 unique views
Titan User Guide	11,191 unique views
Titan Computing Resources	6,763 unique views
Getting Started Overview	4,453 unique views
Tutorials top views	
OpenCL Vector Addition	4,620 unique views
CUDA Vector Addition	4,403 unique views
Serial to Parallel: Monte Carlo Operation	1,123 unique views
GPUDirect: CUDA aware MPI	1,005 unique views
OpenACC Vector Addition	984 unique views
User guide views	
Titan User Guide	19,904 unique views
Landing page	11,195 unique views
Eos User Guide	1,582 unique views
Rhea User Guide	1,279 unique views
Accelerated Computing User Guide	1,134 unique views




OLCF Compute Systems

Titan

Titan is the nation's premier supercomputer for scientific discovery. The Cray XK7 supercomputer is the second fastest in the world as of November 2014. It has demonstrated 17.59 petaflops, or 17.59 quadrillion calculations per second, and is theoretically capable of 27 petaflops.

Titan combines GPUs with traditional CPUs.

Occupying about 5,300 square feet, Titan comprises 200 cabinets that house 18,688 NVIDIA Tesla K20X GPU accelerators and 18,688 16-core AMD Opterons (299,008 cores total), with 710 terabytes of memory.

Titan had 805 users in 2014, and they kept the computer in constant use—researchers consumed 4.2 billion core-hours and completed 326,516 batch jobs. Despite such a heavy workload, Titan's availability was better than last year. Titan's scheduled availability, which does not take scheduled downtime into account, was 99.6 percent in 2014, and Titan was available 95.8 percent of the time. Titan's overall availability increased 2 percent from 2013 to 2014.

Titan generates a lot of heat, and its components must be cooled by heat sinks—devices that dissipate heat into the surrounding air. In addition, a large liquid cooling system developed by Cray sends chilled refrigerant over Titan through an arrangement of pipes affixed to each cabinet. That heat is then extracted directly to the chilled water system running beneath the computer room floor.

Titan represents a new, innovative course for supercomputing developments. Titan has, in many ways, set a new bar for supercomputing standards and will continue to solve some of the world's most pressing problems for the next several years.

Eos

Eos is a smaller CPU-only system that supports the OLCF's users. The OLCF made Eos available in October 2013 as a compute system to augment on Titan, analyzing results, and general processing. The 736-node Cray XC30 cluster has 47 terabytes of memory and uses the Intel Xeon E5-2670 processor.

The Eos compute nodes are organized in blades. Each blade contains four nodes connected to a single Aries interconnect. Every node has 64 GB of DDR3 SDRAM and two sockets with eight physical cores each. Intel's Hyper-Threading (HT) Technology allows each physical core to work as two logical cores, meaning each node can functions as if it has 32 cores. In total, the Eos compute partition contains 11,776 traditional processor cores (23,552 logical cores with HT enabled).

Rhea

Rhea is a 512-node Linux cluster. OLCF staff upgraded Rhea in 2014, increasing its node count from 196 to 512. The primary purpose of Rhea is to provide a conduit for large-scale scientific discovery via pre- and post-processing of simulation data generated on Titan.

Each of Rhea's nodes contains two 8-core 2.0 GHz Intel Xeon processors with HT and 64 GB of main memory. Rhea is connected to the OLCF's 32 PB high performance Lustre file system Atlas.

Data Storage, Analysis, and Visualization Resources

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 InfiniBand high-performance storage area network to quickly move data between systems in the OLCF and 100 gigabit Ethernet to move 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; InfiniBand keeps Titan connected to each of those platforms.

Spider II

The Lustre-based file system, Spider, serves as the centerpiece of the OLCF's technological integration. With more than 20,000 compute nodes mounting it, Spider allows OLCF users to organize data from all of the OLCF computing platforms into a unified file system. The original Lustre-based file system, Spider I, began in 2009, and Spider II transitioned to the primary operating file system in October 2013 connecting Titan, Eos, Rhea, and the center's dedicated GridFTP servers—at an aggregate bandwidth of over 1 terabyte per second. Spider has increased both the capacity available to OLCF users and the bandwidth that was designed carefully to match the performance of Titan and the demanding I/O workloads of OLCF users' applications.

High-Performance Storage System

HPSS uses high-speed data movers to write data into a high speed disk cache, and then the data ultimately is transferred to tape storage for long-term archival. Because of increasing demand for the archival storage resource, staff members are constantly adding more disk space and tape. HPSS was upgraded from version 7.3.2 to 7.4.2p1 this year, offering benefits such as faster retrieval through better aggregation, better responsiveness due to code improvements, DB2 performance optimizations, and the implementation of Redundant Arrays of Independent Tape (RAIT), providing for increased redundancy, reliability, and performance. In addition, more than 50 servers were upgraded to RHEL 6.5, and staff added 32 T10K-D tape drives and 40 gigabit Ethernet for the tape movers. The six SL8500 tape libraries that make up HPSS contain 160 tape drives capable of just over 26 GB per second of performance.

Sith

Sith is an Opteron-based InfiniBand cluster running Linux. The system is provided as an end-to-end resource for center users, meaning users can input data, design their workflow, and expect a result without having to do any intermediary mechanics or further data manipulation 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 and 64 GB of memory. The system is configured with an 86 terabyte Lustre file system for scratch space.

EVEREST

The Exploratory Visualization Environment for Research in Science and Technology, better known as EVEREST, serves as the centerpiece of data analysis at the OLCF. Eighteen 1,920 × 1,080 stereoscopic Barco projection displays create a 37 megapixel main display wall that is 30 × 8.5 ft, while an adjacent, 16-panel, 33 megapixel display measures 7.6 × 13.5 ft. A dedicated Linux cluster accompanied by two nodes controls the Barco wall. The cluster is a 9-node, 16-core Intel machine with 96 GB of memory and an NVIDIA Quadro 5000 graphics card for each node. The cluster is directly attached to the center-wide Lustre file system so that simulation data from other OLCF resources can be accessed without duplication of files. The new ARTTRACK 3 tracking system allows interactivity with large datasets using Titan.

Data analysis

Using supercomputers for scientific discovery is no easy feat. With any job on Titan there can be a seemingly insurmountable amount of data that users have to process to understand their results. When it's time to share results, a visual representation is often best. Therefore, each Titan user is granted access to the powerful data visualization resources of the OLCF. By using software such as VisIt, CEI EnSight, POV-Ray, AVS/Express, ParaView, and IDL, the OLCF data analysis team helps to provide researchers with powerful insight into their data.

2014 INCITE Projects

Biology

Assembling and Sustaining the "Acid Mantle" of the Human Skin Barrier
Michael Klein, Temple University
75,000,000 hours

Designing O_2 Tolerant Hydrogenases
Vijay Pande, Stanford University
16,000,000 hours

Petascale Computing of Biomolecular Systems
Klaus Schulten, University of Illinois at Urbana-Champaign
150,000,000 hours

Chemistry

DNS of Turbulent Combustion Towards Fuel-Flexible Gas Turbines and IC Engines
Jacqueline Chen, Sandia National Laboratories
106,000,000 hours

Computational Actinide Chemistry: Reliable Predictions and New Concepts
David Dixon, University of Alabama
150,000,000 hours

Large-Scale Coupled-Cluster Calculations of Supramolecular Wires
Poul Jørgensen, Aarhus University
24,000,000 hours

Precision Many-Body Quantum Simulations of Functionalized Structures
Shiwei Zhang, College of William and Mary
30,000,000 hours

Computer Science

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

Collaborative Research into Exascale Systemware, Tools, and Applications (CRESTA)
Lorna Smith, The University of Edinburgh
42,000,000 hours

Earth Sciences

High Frequency Physics-based Earthquake System Simulations
Thomas Jordan, University of Southern California
48,000,000 hours

Advancing Models for Multiphase Flow and Transport in Porous Medium System
James McClure, Virginia Tech
60,000,000 hours

High Resolution Simulation for Climate Means, Variability and Extreme
Mark Taylor, Sandia National Laboratories
100,000,000 hours

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

CESM Century-Scale Climate Experiments with a High-Resolution Atmosphere
Warren Washington, UCAR
102,800,000 hours

Materials Sciences

Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo
Dario Alfe, University College London
70,000,000 hours

Scalable First Principles Calculations for Materials at Finite Temperature
Markus Eisenbach, Oak Ridge National Laboratory
150,000,000 hours

Predictive and Insightful Calculations of Energy Materials
Paul Kent, Oak Ridge National Laboratory
50,000,000 hours

QMC Simulations DataBase for Predictive Modeling and Theory
Jeongnim Kim, Oak Ridge National Laboratory
100,000,000 hours

Safety in Numbers: Discovery of New Solid Li-Ion Electrolytes
Boris Kozinsky, Bosch in the USA
60,000,000 hours

Innovative Simulations of High-Temperature Superconductors
Thomas Maier, Oak Ridge National Laboratory
60,000,000 hours

Quantum Monte Carlo Simulations of Hydrogen and Water Ice
Richard Needs, University of Cambridge
55,000,000 hours

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

Physics

High-Fidelity Simulation of Tokamak Edge Plasma Transport
Choong-Seock Chang, Princeton Plasma Physics Laboratory
129,000,000 hours

Linkages between Turbulence and Reconnection Kinetic Plasmas
William Daughton, Los Alamos National Laboratory
86,000,000 hours

Cosmological Simulations for Large-Scale Sky Surveys
Salman Habib, Argonne National Laboratory
100,000,000 hours

Lattice QCD
Paul Mackenzie, Fermilab
100,000,000 hours

Three-Dimensional Simulations of Core-Collapse Supernovae with Chimera
Eric Lentz, University of Tennessee
85,000,000 hours

Nuclear Structure and Nuclear Reactions
James Vary, Iowa State University
104,000,000 hours

Probing Dark Matter at Extreme Scales
Michael Warren, Los Alamos National Laboratory
80,000,000 hours

Petascale Simulations of Type 1a Supernovae
Stan Woosley, University of California, Santa Cruz
50,000,000 hours

2014 ALCC Projects

Biology

Molecular Simulation in Bioenergy
Jeremy Smith, Oak Ridge National Laboratory
59,000,000 hours

Energy coupling in membrane protein function: Mechanisms of Na+-coupled transporters and effects of their environment
Harel Weinstein, Cornell University
10,000,000 hours

Computational Analysis of Complex Proteogenomics Data for Characterization of Terrestrial Carbon Turnover by Soil Microbial Communities
Chongle Pan, Oak Ridge National Laboratory
25,000,000 hours

Earth Science

Multi-scale Water Cycle Processes in Climate Change: Sensitivity to Modeling Frameworks
Ruby Leung, Pacific Northwest National Laboratory
18,000,000 hours

Delivering the Department of Energy's Next-Generation High-Resolution Earth System Model
Peter Thornton, Oak Ridge National Laboratory
30,000,000 hours

Chombo-Crunch: Modeling Pore Scale Reactive Transport Processes Associated with Carbon Sequestration
David Trebotich, Lawrence Berkeley National Laboratory
50,000,000 hours

Computer Science

Hobbes: Operating System and Runtime Research for Extreme Scale
Ron Brightwell, Sandia National Laboratories
30,000,000 hours

Fusion Energy

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

Understanding Helium Plasma Mediated Tungsten Surface Response to Better Predict Fusion Plasma Facing Component Performance in ITER
Brian Wirth, University of Tennessee
30,000,000 hours

Engineering

Multi-hole Injector Optimization for Spark-Ignited Direct-Injection Gasoline Engines
Tang-Wei Kuo, General Motors
15,000,000 hours

High Fidelity Simulations of Combustion Approaches with Increased Efficiency and Reduced Emissions
Peter Cocks, United Technologies Research Center
20,000,000 hours

Simulating Cyclic Variability in Dilute Internal Combustion Engine Operation
Charles Finney, Oak Ridge National Laboratory
17,500,000 hours

Development of High-Fidelity Multiphase Combustion Models for Large Eddy Simulation of Advanced Engine Systems
Joseph Oefelein, Sandia National Laboratories
75,000,000 hours

Supersonic Shockwave Compression & Engine Technology: Time-Varying Phenomena & Geometric Optimization
Ravi Srinivasan, Ramgen Power Systems, LLC
41,000,000 hours

Large Scale Turbulent Clean Coal Combustion
Martin Berzins, University of Utah
30,000,000 hours

Simulating Multiphase Heat Transfer in a Novel Receiver for Concentrating Solar Power (CSP) Plants
Christina Hrenya, University of Colorado
15,000,000 hours

Physics

Calculation of Neutron Scattering Cross Section of Plutonium and its Compounds
Gabriel Kotliar, Rutgers University
90,000,000 hours

Quark and Glue Structure of the Nucleon with Lattice QCD
Keh-Fei Liu, University of Kentucky
68,800,000 hours

Laser-Driven Relativistic Electron Beam Filamentation in Solids
Andreas Kemp, Lawrence Livermore National Laboratory
30,000,000 hours

The Spectrum and Properties of Exotic Mesons in Quantum Chromodynamics
Robert Edwards, Jefferson Science Associates, LLC
250,000,000 hours

Quantum Computational Science
Itay Hen, University of Southern California
45,000,000 hours

Hypernuclei and Charmed Nuclei
Martin Savage, University of Washington
65,100,000 hours

Nuclear Energy

Delivering Advanced Modeling & Simulation for Nuclear Energy Applications
John Turner, Oak Ridge National Laboratory
80,000,000 hours

Advanced Simulation of HFIR for LEU Conversion
Gregory Davidson, Oak Ridge National Laboratory
50,000,000 hours

2014 Publications

1 Liu, Z.; Lu, Z.; et al. (2014) “Evolution and Forcing Mechanisms of El Nino Over the Past 2,100 Years.” Nature Volume: 515 Issue: 7528 DOI: 10.1038/nature13963

2. Otto-Bliesner, Bette L.; et al. (2014) “Coherent Changes of Southeastern Equatorial and Northern African Rainfall During the Last Deglaciation” Science Volume: 346 Issue: 6214 DOI: 10.1126/science.1259531

3. Liao, H.; Zherebetskyy, D.; et al. (2014) “Facet Development During Platinum Nanocube Growth” Science Volume: 345 Issue:6199 DOI: 10.1126/science.1253149

4. Zherebetskyy, D.; Scheele, M.; Zhang, Y.; et al. (2014) “Hydroxylation of the Surface of PbS Nanocrystals Passivated with Oleic acid” Science Volume: 344 Issue:6190 DOI: 10.1126/science.1252727

5. Buizert, C.; Gkinis, V.; Severinghaus, J.; et al. (2014) “Greenland Temperature Response to Climate Forcing During the Last Deglaciation” Science Volume: 345 Issue: 6201 DOI: 10.1126/science.1254961

6. Sterpone, F.; Melchionna, S.; et al. (2014) “The OPEP Protein Model: from Single Molecules, Amyloid Formation, Crowding and Hydrodynamics to DNA/RNA Systems” Chemical Society Reviews Volume: 43 Issue: 13 DOI: 10.1039/c4cs00048j

7. Sumpter, B.; Liang, L.; Nicolai, A.; Meunier, V. (2014) “Interfacial Properties and Design of Functional Energy Materials” Accounts of Chemical Research Volume: 47 Issue: 11 DOI: 10.1021/ar500180h

8. Yin, Z.P.; Haule, K.; Kotliar, G. (2014) “Spin Dynamics and Orbital-Antiphase Pairing Symmetry in Iron Based Superconductors” Nature Physics Volume: 10 Issue: 11 DOI: 10.1038/NPHYS3116

9. Cardall, C.; Budiardja, R.; Endeve, E.; Mezzacappa, A. (2014) “Genasis: General Astrophysical Simulation System 1. Refinable Mesh and Nonrelativistic Hydrodynamics” Astrophysical Journal Supplement Series Volume: 210 Issue: 2 DOI: 10.1088/0067-0049/210/2/17

10. Bryan, G.; Norma, M.; O’Shea, B.; et al. (2014) “Enzo: An Adaptive Mesh Refinement Code for Astrophysics” Astrophysical Journal Supplement Series Volume:211 Issue: 19 DOI: 10.1088/0067-0049/211/2/19

11. Davidson, E.; Klimes, J.; Alfe, D.; Michaelides, A. (2014) “Cooperative Interplay of van der Waals Forces and Quantum Nuclear Effects on Adsoprtion: H at Graphene and at Coronene” ACS Nano Volume: 8 Issue: 10 DOI: 10.1021/nn505578x

12. Tarafder, K.; Surendranath, Y.; et al. (2014) “Hole Transfer Dynamics from a CdSe/CdS Quantum Rod to a Tethered Ferrocene Derivative” Journal of the American Chemical Society Volume: 136 Issue: 13 DOI: 10.1021/ja500936n

13. Li, Z.; Wang, Y.; et al. (2014) “Diverse and Divergent Protein Post-Translational Modifications in Two Growth Stages of a Natural Microbial Community” Nature Communications Volume:5 DOI: 10.1038/ncomms5405

14. Wang, Y.; Wang, M.; Zhang, R.; et al. (2014) “Assessing the Effects of Anthropogenic Aerosols on Pacific Storm Track Using a Multiscale Global Climate Model” Proceedings of the National Academy of Sciences of the United States of America Volume: 111 Issue: 19 DOI: 10.1073/pnas.1403364111

15. Tessum, C.; Hill, J.; Marshall, J. (2014) “Life Cycle Air Quality Impacts of Conventional and Alternative Light-Duty Transportation in the United States” Proceedings of the National Academy of Sciences of the United States of America Volume: 111 Issue: 52 DOI: 10.1073/pnas.1406853111

16. Foyevtsova, K.; Krogel, J.; et al. (2014) “Ab initio Quantum Monte Carlo Calculations of Spin Superexchange in Cuprates: The Benchmarking Case of Ca2CuO3” Physical Review X Volume: 4 Issue: 3 DOI: 10.1103/PhysRevX.4.031003

17. Guo, F.; Li, H.; Daughton, W.; Liu, Y. (2014) “Formation of Hard Power Laws in the Energetic Particle Spectra Resulting from Relativistic Magnetic Reconnection” Physical Review Letters Volume:113 Issue: 15 DOI: 10.1103/PhysRevLett.113.155005

18. Bulgac, A.; Forbes, M.; et al. (2014) “Quantized Superfluid Vortex Rings in the Unitary Fermi Gas” Physical Review Letters Volume: 112 Issue: 2 DOI: 10.1103/PhysRevLett.112.025301

19. Zhang, C.; Harriger, L.; Yin, Z.; et al. (2014) “Effect of Pnictogen Height on Spin Waves in Iron Pnictides” Phyiscal Review Letters Volume:112 Issue: 21 DOI: 10.1103/PhysRevLett.112.217202

20. Hong, L.; Smolin, N.; Smith, J. (2014) “de Gennes Narrowing Describes the Relative Motion of Protein Domains” Physical Review Letters Volume: 112 Issue: 15 DOI: 10.1103/PhysRevLett.112.158102

21. Romero-Redondo, C.; Quaglioni, S.; Navratil, P.; Hupin, G. (2014) "He-4 + n plus Continuum within an Ab initio Framework" Physical Review Letters Volume: 114 Issue: 3 DOI: 10.1103/PhysRevLett.113.032503

22. Virgus, y.; Purwanto, W.; Krakauer, H.; Zhang, S. (2014) "Stability Energetics and Magnetic States of Cobalt Adatoms on Graphene" Physical Review Letters Volume: 113 Issue: 17 DOI: 10.1103/PhysRevLett.113.175502

23. Jansen, G. R.; Engel, J.; Hagen, G.; Navratil, P.; Signoracci, A. (2014) "Ab initio Coupled-Cluster Effective Interactions for the Shell Model: Applicatio to Neutron-Rich Oxyen and Carbon Isotopes" Physical Review Letters Volume: 113 Issue: 14 DOI: 10.1103/PhysRevLett.113.142502

24. Fiskel, G.; Fox, W.; et al. (2014) "Magnetic Reconnection between Colliding Magentized Lazer-Produced Plasma Plumes" Phyiscal Review Letters Volume: 113 Issue: 10 DOI: 10.1103/PhysRevLett.113.105003

25. Alhassid, Y.; Gilbreth, C.N.; Bertsch, G.F. (2014) "Nuclear Deformation at Finite Temperature" Physical Review Letters Volume: 113 Issue: 26 DOI: 10.1103/PhysRevLett.113.262503

26. Dudek, J.; Edwards, R.; Thomas, C.; Wilson, D. (2014) "Resonances in Coupled πK–ηK Scattering from Quantum Chromodynamics" Physical Review Letters Volume: 113 Issue: 18 DOI: 10.1103/PhysRevLett.113.182001

27. Wu, J.; Zhou, Y.; Gao, Y.; et al. (2014) "Estimation and Uncertainty Analysis of Impacts of Future Heat Waves on Mortality in the Eastern United States" Environmental Health Perspectives Volume: 122 Issue: 1 DOI: 10.1289/ehp.1306670

28. Lopez-Benzanilla, A.; Ganesh, P.; Kent, P.; Sumpter, B. (2014) "Spin-Resolved Self-Doping Tunes the Intrinsic Half-Metallicity of AIN Nanoribbons" Nano Research Volume: 7 Issue: 1 DOI: 10.1007/s12274-013-0371-1

29. Nguyen, T.; Carrillo, J.; Brown, W.M.; Matheson, M. (2014) "Rupture Mechanism of Liquid Crystal Thin Films Realized by Large-Scale Molecular Simulations" Nanoscale Volume: 6 Issue: 6 DOI: 10.1039/C3NR05413F

30. Raskin, C.; Kasen, D.; Moll, R.; Schwab, J.; Woosley, S. (2014) "Type Ia Supernovae From Merging White Dwarfs. II. Post-Merger Detonations" Astrophysical Journal Volume: 788 Issue: 75 DOI: 10.1088/0004-637X/788/1/75

31. Papatheodore, T.; Messer, O.E. (2014) "On Numerical Considerations for Modeling Reactive Astrophysical Shocks" Astrophysical Journal Volume: 782 Issue: 1 DOI: 10.1088/0004-637X/782/1/12

32. Fragile, P. C.; Olejar, A.; Anninos, P. (2014) "Numerical Simulations of Optically Thick Accretion Onto A Black. II. Rotating Flow" Astrophysical Journal Volume: 796 Issue: 1 DOI: 10.1088/0004-637X/796/1/22

33. Moll, R.; Raskin, C.; Kasen, D.; Woosley, S.E. (2014) "Type 1a Supernovae from Merging White Dwarfs. I. Prompt Detonations" Astrophysical Journal Volume: 785 Issue: 2 DOI: 10.1088/0004-637X/785/2/105

34. Malone, C.M.; Nonaka, A.; et al. (2014) "The Deflagration Stage of Chandrasekhar Mass Models for Type 1a Supernovae. I. Early Evolution" Astrophyiscal Journal Volume: 782 Issue: 1 DOI: 10.1088/0004-637X/782/1/11

35. Sengupta, B.; Friston, K.J.; Penny, W.D. (2014) "Efficent Gradient Computation for Dynamical Models" Neuroimage Volume:98 DOI: 10.1016/j.neuroimage.2014.04.040

36. Sengupta, B.; Stemmler, M. (2014) "Power Consumption During Neuronal Computation" Proceedings of the IEEE Volume: 102 Issue: 5 DOI: 10.1109/JPROC.2014.2307755

37. Long, M.S.; Keene, W.C.; et al. (2014) "Sensitivity of Tropospheric Chemical Composition to Halogen-Radical Chemistry Using a Fully Coupled Size-Resolved Multiphase Chemistry-Global Climate System: Halogen Distributions Aerosol Composition and Sensitivity of Climate-Relevant Gases" Atmospheric Chemistry and Physics Volume:14 Issue:7 DOI: 10.5194/acp-14-3397-2014

38. Pickel, D.; Kilbey, S. M.; Uhrig, D.; et al. (2014) "Structural Evolution of Polylactide Molecular Bottlebrushes: Kinetics Study by Size Exclusion Chromatography, Small Angle Neutron Scattering, and Simulations" ACS Macro Letters Volume: 3 Issue: 9 DOI: 10.1021/mz5003454

39. Kusne, A. G.; Gao, T.; Mehta, A.; et al. (2014) "On-the-Fly Machine-Learning for High-Throughput Experiments: Search for Rare-Earth-Free Permanent Magnets" Scientific Reports Volume: 4 DOI: 10.1038/srep06367

40. Vizcaino, M.; Lipscomb, W.H.; Sacks, W.; van den Broeke, M. (2014) "Greenland Surface Mass Balance as Simulated by the Community Earth System Model. Part II: Twenty-First-Century Changes" Journal of Climate Volume: 27 Issue: 1 DOI: 10.1175/JCLI-D-12-00588.1

41. Bacmeister, J.; Wehner, M.F.; et al. (2014) "Exploratory High-Resolution Climate Simulations Using the Community Atmosphere Model (CAM)" Journal of Climate Volume: 27 Issue: 9 DOI: 10.1175/JCLI-D-13-00387.1

42. Padmanath, M.; Edwards, R.; Mathur, N.; Peardon, M. (2014) "Spectroscopy of Triply Charmed Baryons from Lattice QCD" Physical Review D Volume: 90 Issue: 7 DOI: 10.1103/PhysRevD.90.074504

43. Upadhye, A.; Biswas, R.; et al. (2014) "Large-Scale Structure Formation with Massive Neutrinos and Dynamical Dark Energy" Physical Review D Volume: 89 Issue: 10 DOI: 10.1103/PhysRevD.89.103515

44. Sandeep, S.; Stordal, F.; Sardeshmukh, P.D.; Compo, G. (2014) "Pacific Walker Circulation Variability in Coupled and Uncoupled Climate Models" Climate Dynamics Volume: 43 Issue: 1-2 DOI: 10.1007/s00382-014-2135-3

45. Hagos, S.; Leung, L.; Gustafson, W.; Singh, B. (2014) "Eddy Fluxes and Sensitivity of the Water Cycle to Spatial Resolution in Idealized Regional Aquaplanet Model Simulations" Climate Dynamics Volume: 42 Issue: 3-4 DOI: 10.1007/s00382-013-1857-y

46. Hagos, S.; Leung, L. et al. (2014) "Assessment of Uncertainties in the Response of the African Monsoon Precipitation to Land Use Change Simulated by A Regional Model" Climate Dynamics Volume: 43 Issue: 9-10 DOI: 10.1007/s00382-014-2092-x

47. Wang, D.; Xu, Y.; Thornton, P.; et al. (2014) "A Functional Test Platform for the Community Land Model" Environmental Modelling and Software Volume: 55 DOI: 10.1016/j.envsoft.2014.01.015

48. Nakamura, T.K.M.; Daughton, W. (2014) "Turbulent Plasma Transport Across the Earth's Low-Latitude Boundary Layer" Geophysical Research Letters Volume: 41 Issue: 24 DOI: 10.1002/2014GL061952

49. Karra, S.; Painter, S.L.; Lichtner, P.C. (2014) "Three-Phase Numerical Model for Subsurface Hydrology in Permafrost-Affected Regions (PFLOTRAN-ICE v1.0)" Cryosphere Volume: 8 Issue: 5 DOI: 10.5194/tc-8-1935-2014

50. Edwards, T.; Fettweis, X.; Gagliardini, O.; et al. (2014) "Effect of Uncertainty in Surface Mass Balance–Elevation Feedback on Projections of the Future Sea Level Contribution of the Greenland Ice Sheet" Cryosphere Volume: 8 DOI: 10.5194/tc-8-195-2014

51. Dicaire, N.Mm.; Omand, C.; Navratil, P. (2014) "Alternative Similarity Renormalization Group Generators in Nuclear Structure Calculations" Physical Review C Volume: 90 Issue: 3 DOI: 10.1103/PhysRevC.90.034302

52. Pei, J.C.; Fann, G.I.; et al. (2014) "Adaptive Multi-Resolution 3D Hartree-Fock-Bogoliubov Solver for Nuclear Structure" Physical Review C Volume: 90 Issue: 2 DOI: 10.1103/PhysRevC.90.024317

53. Shi, Y.; Ward, D.E.; Carlsson, B.G.; et al. (2014) "Structure of Superheavy Nuclei Along Decay Chains of Element 115" Physical Review C Volume: 90 Issue: 1 DOI: 10.1103/PhysRevC.90.014308

54. Pais, H.; Newton, W.G.; Stone, J. (2014) "Phase Transitions in Core-Collapse Supernova Matter at Sub-Saturation Densities" Physical Review C Volume: 90 Issue: 6 DOI: 10.1103/PhysRevC.90.065802

55. Wong, C.; Sen, A.; Gerhard, J.; Torrieri, G.; Read, K. (2014) "Analytical Solutions of Landau (1+1)-Dimensional Hydrodynamics" Physical Review C Volume: 90 Issue: 6 DOI: 10.1103/PhysRevC.90.064907

56. Sadhukhan, J.; Dobaczewski, J.; Nazarewicz, W.; Sheikh, J.; Baran, A. (2014) "Pairing-Induced Speedup of Nuclear Spontaneous Fission" Physical Review C Volume: 90 Issue: 6 DOI: 10.1103/PhysRevC.90.061304

57. Jaganathen, Y.; Michel, N.; Ploszajczak, M. (2014) "Gamow Shell Model Description of Proton Scattering on Ne-18" Physical Review C Volume: 89 Issue: 3 DOI: 10.1103/PhysRevC.89.034624

58. Vajta, Zs.; Stanoiu, M.; Sohler, D.; et al (2014) "Excited States in the Neutron-Rich Nucleus F25" Physical Review C Volume: 89 Issue: 5 DOI: 10.1103/PhysRevC.89.054323

59. Schunck, N.; Duke, D.; Carr, H.; Knoll, A. (2014) "Description of Induced Nuclear Fission with Skyrme Energy Functionals: Static Potential Energy Surfaces and Fission Fragment Properties" Physical Review C Volume: 90 Issue: 5 DOI: 10.1103/PhysRevC.90.054305

60. Maris, P.; Vary, J.P.; Calci, A.; et al. (2014) "C12 Properties with Evolved Chiral Three-Nucleon Interactions" Physical Review C Volume: 90 Issue: 1 DOI: 10.1103/PhysRevC.90.014314

61. Hagen, G.; Papenbrock, T.; Ekstroem, A.; et al. (2014) "Coupled-Cluster Calculations of Nucleonic Matter" Physical Review C Volume: 89 Issue: 1 DOI: 10.1103/PhysRevC.89.014319

62. Kruppa, A.T.; Papadimitriou, G.; Nazarewicz, W.; Michel, N. (2014) "Nuclear Three-Body Problem in the Complex Energy Plane: Complex-Scaling Slater Method" Physical Review C Volume: 89 Issue: 1 DOI: 10.1103/PhysRevC.89.014330

63. McDonnell, J.D.; Nazarewicz, W.; et al. (2014) "Excitation-Energy Dependence of Fission in the Mercury Region" Physical Review C Volume: 90 Issue: 2 DOI: 10.1103/PhysRevC.90.021302

64. Sheikh, J.A.; Hinohara, N.; et al. (2014) "Isospin-Invariant Skyrme Energy-Density Functional Approach with Axial Symmetry" Physical Review C Volume: 89 Issue: 5 DOI: 10.1103/PhysRevC.89.054317

65. Kortelainen, M.; McDonnell, J.; et al. (2014) "Nuclear Energy Density Optimization: Shell Structure" Physical Review C Volume: 89 Issue: 5 DOI: 10.1103/PhysRevC.89.054314

66. Chandler, D.; Struempfer, J.; et al (2014) "Light Harvesting by Lamellar Chromatophores in Rhodospirillum Photometricum" Biophysical Journal Volume: 106 Issue: 11 DOI: 10.1016/j.bpj.2014.04.030

67. Bhagatwala, A.; Shen, H.; Sutton, J.; Lu, T.; Chen, J.H. (2014) "Numerical and Experimental Investigation of Turbulent DME Jet Flames" Proceedings of the Combustion Institute Volume: 35 Issue: 2 DOI: 10.1016/j.proci.2014.05.147

68. Hammond, G.E.; Lichtner, P.C.; Mills, R.T. (2014) "Evaluating the Performance of Parallel Subsurface Simulators: An Illustrative Example with PFLOTRAN" Water Resources Research Volume: 50 Issue: 1 DOI: 10.1002/2012WR013483

69. Chen, Z.; Li, J.; Li, S.; Wang, L. (2014) "Approximate Hessian for Accelerating Ab Initio Structure Relaxation by Force Fitting" Physical Review B Volume: 89 Issue: 14 DOI: 10.1103/PhysRevB.89.144110

70. Shi, H.; Jimenez-Hoyos, C.; et al (2014) "Symmetry-Projected Wave Functions in Quantum Monte Carlo Calculations" Physical Review B Volume: 89 Issue: 12 DOI: 10.1103/PhysRevB.89.125129

71. Rodriguez-Guzman, R.; Jimenez-Hoyos, C.A.; Scuseria, G.E. (2014) "Multireference Symmetry-Projected Variational Approximation for the Ground State of Doped One-Dimensional Hubbard Model" Physical Review B Volume: 89 Issue: 19 DOI: 10.1103/PhysRevB.89.195109

72. Starr, P.; Maier, T.; Schulthess, T.C. (2014) "Two-Particle Correlations in A Dynamic Cluster Approximation with Continuous Momentum Dependence: Superconductivity in the Two-Dimensional Hubbard Model" Physical Review B Volume: 89 Issue: 19 DOI: 10.1103/PhysRevB.89.195133

73. Lin, Y.; Cohen, R.E.; et al. (2014) "Equations of State and Stability of MgSiO3 Perovskite and Post-Perovskite Phases from Quantum Monte Carlo Simulations" Physical Review B Volume: 90 Issue: 18 DOI: 10.1103/PhysRevB.90.184103

74. Rodriguez-Guzman, R.; Jimenez-Hoyos, C.A.; Scuseria, G. (2014) "Variational Description of the Ground State of the Repulsive Two-Dimensional Hubbard Model in Terms of Nonorthogonal Symmetry-Projected Slater Determinants" Physical Review B Volume: 90 Issue: 19 DOI: 10.1103/PhysRevB.90.195110

75. Alfe, D.; Bartok, A.P.; Csanyi, G.; Gillan, M.J. (2014) "Analyzing the Errors of DFT Approximations for Compressed Water Systems" Journal of Chemical Physics Volume: 141 Issue: 1 DOI: 10.1063/1.4885440

76. Quigley, D.; Alfe, D.; Slater, B. (2014) "Communication: On the Stability of Ice 0, Ice I, and I h" Journal of Chemical Physics Volume: 141 Issue: 16 DOI: 10.1063/1.4900772

77. Cox, S.; Towler, M.D.; Alfe, D.; Michaelides, A. (2014) "Benchmarking the Performance of Density Functional Theory and Point Charge Force Fields in Their Description of sI Methane Hydrate Against Diffusion Monte Carlo" Journal of Chemical Physics Volume: 140 Issue: 17 DOI: 10.1063/1.4871873

78. Kennedy, M.R.; McDonald, A.R.; et al. (2014) "Communication: Resolving the Three-Body Contribution to the Lattice Energy of Crystalline Benzene: Benchmark Results from Coupled-Cluster Theory" Journal of Chemical Physics Volume: 140 Issue: 12 DOI: 10.1063/1.4869686

79. Menzeleev, A.R.; Bell, F.; Miller, T.F. (2014) "Kinetically Constrained Ring-Polymer Molecular Dynamics for Non-Adiabatic Chemical Reactions" Journal of Chemical Physics Volume: 140 Issue: 6 DOI: 10.1063/1.4863919

80. Karim, E.T.; Shugaev, M.; et al. (2014) "Atomistic Simulation Study of Short Pulse Laser Interactions With A Metal Target Under Conditions of Spatial Confinement by a Transparent Overlayer" Journal of Applied Physics Volume: 115 Issue: 18 DOI: 10.1063/1.4872245

81. Carrillo, J.Y.; Sumpter, B.G. (2014) "Structure and Dynamics of Confined Flexible and Unentangled Polymer Melts in Highly Adsorbing Cylindrical Pores" Journal of Chemical Physics Volume: 141 DOI: 10.1063/1.4893055

82. Al-Hamdani, Y.S.; Alfe, D.; von Linienfeld, O.A.; Michaelides, A. (2014) "Water on BN Doped Benzene: A Hard Test for Exchange-Correlation Functionals and the Impact of Exact Exchange on Weak Binding" Journal of Chemical Physics Volume: 141 DOI: 10.1063/1.4898356

83. Tubman, N.M.; Kylanpaa, I.; Hammes-Schiffer, S.; Ceperley, D.M. (2014) "Beyond the Born-Oppenheimer Approximation With Quantum Monte Carlo Methods" Physical Review A Volume: 90 Issue: 4 DOI: 10.1103/PhysRevA.90.042507

84. Pindzola, M.S. (2014) “Electron-Impact Ionization of the Inner Subshells of Uranium” Physical Review A Volume: 90 Issue: 2 DOI: 10.1103/PhysRevA.90.022708

85. Nguyen, T.D.; Fuentes-Cabrera, M.; Fowlkes, J.D.; Rack, P.D. (2014) “Coexistence of Spinodal Instability and Thermal Nucleation in Thin-Film Rupture: Insights from Molecular Levels” Physical Review E Volume: 89 Issue: 3 DOI: 10.1103/PhysRevE.89.032403

86. Vogel, T.; Li, Y.W.; Wuest, T.; Landau, D.P. (2014) “Scalable Replica-Exchange Framework for Wang-Landau Sampling” Physical Review E Volume: 90 Issue: 2 DOI: 10.1103/PhysRevE.90.023302

87. Marino, R.; Mininni, P.D.; et al. (2014) “Large-Scale Anisotropy in Stably Stratified Rotating Flows” Physical Review E Volume: 90 DOI: 10.1103/PhysRevE.90.023018

88. Muhammad, A.; Cencerrado, A.; et al. (2014) “Impact of I/O and Data Management in Ensemble Large Scale Climate Forecasting Using EC-Earth3” 2014 International Conference on Computational Science Volume: 29 DOI: 10.1016/j.procs.2014.05.221

89. Evans, K.J.; Mahajan, S.; et al. (2014) “A Spectral Transform Dynamical Core Option Within the Community Atmosphere Model (CAM4)” Journal of Advances in Modeling Earth Systems Volume: 6 Issue: 3 DOI: 10.1002/2014MS000329

90. Chan, W.L.; Kolla, H.; Chen, J.H.; Ihme, M. (2014) “Assessment of Model Assumptions and Budget Terms of the Unsteady Flamelet Equations for A Turbulent Reacting Jet-In-Cross-Flow” Combustion and Flame Volume: 161 Issue: 10 DOI: 10.1016/j.combustflame.2014.04.007

91. Coon, E.T.; Porter, M.L.; Kang, Q. (2014) “Taxila LBM: A Parallel Modular Lattice Boltzmann Framework for Simulating Pore-Scale Flow in Porous Media” Computational Geosciences Volume: 18 Issue: 1 DOI: 10.1007/s10596-013-9379-6

92. Dathar, G.K.P.; Tsai, Y.; et al. (2014) “Identifying Active Functionalities on Few-Layered Graphene Catalysts for Oxidative Dehydrogenation of Isobutane” Chemsuschem Volume: 7 Issue: 2 DOI: 10.1002/cssc.201301006

93. Jiang, T.; Evans, K.J; Deng, Y.; Dong, X. (2014) “Intermediate Frequency Atmospheric Disturbances: A Dynamical Bridge Connecting Western US Extreme Precipitation With East Asian Cold Surges” Journal of Geophysical Research-Atmospheres Volume: 119 Issue: 7 DOI: 10.1002/2013JD021209

94. Heitmann, K.; Habib, S.; et al. (2014) “Large-Scale Simulations of Sky Surveys” Computing in Science and Engineering Volume: 16 Issue: 5 DOI: 10.1109/MCSE.2014.49

95. Paolucci, S.; Zikoski, Z.J.; Grenga, T. (2014) “WAMR: An Adaptive Wavelet Method for the Simulation of Compressible Reacting Flow. Part II. The Parallel Algorithm” Journal of Computational Physics Volume: 272 DOI: 10.1016/j.jcp.2014.03.059

96. Orlandini, G.; Bacca, S.; et al. (2014) “Coupling the Lorentz Integral Transform (LIT) and the Coupled Cluster (CC) Methods: A Way Towards Continuum Spectra of ‘Not-So-Few-Body’ Systems” Few-Body Systems Volume: 55 Issue: 8-10 DOI: 10.1007/s00601-013-0772-4

97. Hagos, S.; Feng, Z.; Landu, K.; Long, C.N. (2014) “Advection moistening and shallow-to-deep convection transitions during the initiation and propagation of Madden-Julian Oscillation” Journal of Advances in Modeling Earth Systems Volume: 6 Issue: 3 DOI: 10.1002/2014MS000335

98. Krisman, A.; Tang, J.C.K.; et al. (2014) “A DNS Evaluation of Mixing Models for Transported PDF Modelling of Turbulent Nonpremixed Flames” Combustion and Flame Volume: 161 Issue: 8 DOI: 10.1016/j.combustflame.2014.01.009

99. Justice, N.B.; Li, Z.; et al. (2014) “N-15- and H-2 Proteomic Stable Isotope Probing Links Nitrogen Flow to Archaeal Heterotrophic Activity” Environmental Microbiology Volume: 16 Issue: 10 DOI: 10.1111/1462-2920.12488

100. McInnes, L.; Smith, B.; Zhang, H.; et al. (2014) “Hierarchical Krylov and Nested Krylov Methods for Extreme-Scale Computing.” Parallel Computing: Volume: 40 Issue: 1 DOI: 10.1016/j.parco.2013.10.001

101. Ibrahim, A.; Wawan, M.; Mosher, S.; et al. (2014) “Assessment of Fusion Facility Dose Rate Map Using Mesh Adaptivity Enhancements of Hybrid Monte Carlo/Deterministic Techniques.” Fusion Engineering and Design Volume: 89 Issue: 9-10 DOI: 10.1016/j.fusengdes.2014.02.046

102. Norman, M. (2014) “A WENO-Limited ADER-DT Finite-Volume Scheme for Efficient Robust and Communication-Avoiding Multi-Dimensional Transport.” Journal of Computational Physics Volume: 274 DOI: 10.1016/j.jcp.2014.05.043

103. Perumalla, K.; Park, A. (2014) “Reverse Computation for Rollback-Based Fault Tolerance in Large Parallel Systems.” Cluster Computing Volume: 17 Issue: 2 DOI: 10.1007/s10586-013-0277

104. Perumalla, K.; Park, A.; Tipparaju, V. (2014) “Discrete Event Execution with One-Sided and Two-Sided GVT Algorithms on 216 000 Processor Cores.” ACM Transactions on Modeling and Computer Simulation Volume: 24 Issue: 3 DOI: 10.1145/2611561

105. Emerson, S.; Zhu, T.; Davis, Timothy; et al. (2014) “Liquid Phase Reforming of Woody Biomass to Hydrogen.” International Journal of Hydrogen Energy Volume: 39 Issue: 1 DOI: 10.1016/j.ijhydene.2013.09.041

106. Salaway, R.; Zhigilei, L. (2014) “Molecular Dynamics Simulations of Thermal Conductivity of Carbon Nanotubes: Resolving the Effects of Computational Parameters.” International Journal of Heat and Mass Transfer Volume: 70 Pages: 954-964 DOI: 10.1016/j.ijheatmasstransfer.2013.11.065

107. Petra, C.; Schenk, O.; Anitescu, M. (2014) “Real-Time Stochastic Optimization of Complex Energy Systems on High-Performance Computers.” Computing in Science & Engineering Volume: 16 Issue: 5 Pages: 32-42

108. Nowicki, S.; Bindschadler, R.; Abe-Ouchi, A.; et al. (2014) “Density-Fitted Singles and Doubles Coupled Cluster on Graphics Processing Units.” Molecular Physics Volume: 112 Issue: 5-6 Pages: 844-852 DOI: 10.1080/00268976.2013.874599

109. Nowicki, S.; Bindschadler, R.; Abe-Ouchi, A.; et al. (2014) “Simulation at Extreme-Scale: Co-Design Thinking and Practices.” Archives of Computational Methods in Engineering Volume: 21 Issue: 1 Pages 39-58 DOI: 10.1007/s11831-014-9095-y

110. Wang, D.; Schuchart, J.; Janjusic, T.; et al. (2014) “Toward Better Understanding of the Community Land Model within the Earth System Modeling Framework.” 2014 International Conference on Computational Science Volume: 29 Pages: 1515-1524 DOI: 10.1016/j.procs.2014.05.137

111. Wu, C.; Zhigilei, L. (2014) “Microscopic Mechanisms of Laser Spallation and Ablation of Metal Targets from Large-Scale Molecular Dynamics Simulations.” Applied Physics A - Materials Science & Processing Volume: 114 Issue: 1 Pages: 11-32 DOI: 10.1007/s00339-013-8086-4

112. Yu, Y.; Bittencourt, M.; Karniadakis, G. (2014) “A Semi-Local Spectral/Hp Element Solver For Linear Elasticity Problems.” International Journal for Numerical Methods in Engineering Volume: 100 Issue: 5 DOI: 10.1002/nme.4739

113. Dubey, A.; Almgren, A.; Bell, J.; et al. (2014) “A Survey Of High Level Frameworks In Block-Structured Adaptive Mesh Refinement Packages.” Journal of Parallel And Distributed Computing Volume: 74 Issue: 12 Pages: 3217-3227 DOI: 10.1016/j.jpdc.2014.07.001

114. Lakshminarasimhan, S.; Zou, X.; Boyuka, D.; et al. (2014) “DIRAQ: Scalable In Situ Data - And Resource-Aware Indexing For Optimized Query Performance.” Cluster Computing – The Journal of Networks Software Tools and Applications Volume: 17 Issue: 4 DOI: 10.1007/s10586-014-0358-z

115. Huebl, A.; Pugmire, D.; Schmitt, F.; et al. (2014) “Visualizing the Radiation of the Kelvin-Helmholtz Instability.” IEEE Transactions on Plasma Science Volume: 42 Issue: 10 Pages: 2638-2639 DOI: 10.1109/TPS.2014.2327392

116. Hagos, S.; Feng, Z.; Burleyson, C.; et al. (2014) “Evaluation Of Convection-Permitting Model Simulations Of Cloud Populations Associated With The Madden-Julian Oscillation Using Data Collected During The AMIE/DYNAMO Field Campaign.” Journal of Geophysical Research-Atmospheres Volume: 119 Issue: 21 DOI: 10.1002/2014JD022143

117. Samsonidze, G.; Park, C.; Kozinsky, B. (2014) “Insights And Challenges Of Applying The GW Method To Transition Metal Oxides.” Journal of Physics – Condensed Matter Volume: 26 Issue: 47 DOI: 10.1088/0953-8984/26/47/475501

118. Ilsche, T.; Schuchart, J.; Cope, J.; et al. (2014) “Optimizing I/O Forwarding Techniques For Extreme-Scale Event Tracing.” Cluster Computing Volume: 17 Issue: 1 DOI: 10.1007/s10586-013-0272-9

119. Poyraz, E.; Xu, H.; Cui, Y.; et al. (2014) “Application-specific I/O Optimizations on Petascale Supercomputers.” 2014 International Conference on Computational Science Volume: 29 Pages: 910-923 DOI: 10.1016/j.procs.2014.05.082

120. Tang, W.; Wang, B.; Ethier, S. (2014) “Scientific Discovery in Fusion Plasma Turbulence Simulations at Extreme Scale.” Computing in Science & Engineering Volume: 16 Issue: 5 Pages: 44-52 DOI: 10.1109/MCSE.2014.54

121. Kim, Y.; Lee, J.; Oral, S.; et al. (2014) “Coordinating Garbage Collection for Arrays of Solid-State Drives.” IEEE Transactions on Computers Volume: 63 Issue: 4 Pages: 888-901 DOI: 10.1109/TC.2012.256

122. Smith, C.; Compo, G.; Hooper, D. (2014) “Web-Based Reanalysis Intercomparison Tools (WRIT) for Analysis and Comparison of Reanalyses and Other Datasets.” Bulletim of the American Meteorological Society Volume: 95 Issue: 11 Pages: 1671-1678 DOI: 10.1175/BAMS-D-13-00192.1

123. Teodoro, G.; Pan, T.; Kurc, T.; et al. (2014) “Region Templates: Data Representation And Management For High-Throughput Image Analysis.” Parallel Computing Volume: 40 Issue: 10 DOI: 10.1016/j.parco.2014.09.003

124. Mueller, E.; Scheichl, R. (2014) “Massively Parallel Solvers For Elliptic Partial Differential Equations In Numerical Weather And Climate Prediction.” Quarterly Journal of the Royal Meteorological Society Volume: 140 Issue: 685 Pages: 2608-2624 DOI: 10.1002/qj.2327

125. Pindzola, M.S.; Lee, T.G.; Colgan, J. (2014) "Antiproton-Impact Ionization Of H-2 At Low Incident Energies." Journal of Physics B – Atomic molecular and Optical Physics Volume: 47 Issue: 18 DOI: 10.1088/0953-4075/47/18/185202

126. Nicholson, D.M.; Odbadrakh, Kh.; Shassere, B.A.; et al (2014) "Modeling And Characterization Of The Magnetocaloric Effect In Ni2MnGa Materials." International Journal of Refrigeration-Revenue Inernationale du Froid Volume: 37 Pages: 289-296 DOI: 10.1016/j.ijrefrig.2013.10.001

127. Potter, H.D.; Fischer, S.; Maris, P.; et al. (2014) "Ab Initio Study Of Neutron Drops With Chiral Hamiltonians." Physics Letters B Volume: 739 Pages: 445-450 DOI: 10.1016/j.physletb.2014.10.020

128. Tang, Y.H.; Karniadakis, G. (2014) "Accelerating Dissipative Particle Dynamics Simulations On Gpus: Algorithms Numerics And Applications." Computer Physics Communications Volume: 185 Issue: 11 Pages: 2809-282 DOI: 10.1016/j.cpc.2014.06.015

129. McLaughlin, B.M.; Bizau, J.M.; Cubaynes, D.; et al. (2014) "K-Shell Photoionization Of B-Like Oxygen (O3+) Ions: Experiment And Theory." Journal of Physics B – Atomic Molecular and Optical Physics Volume: 47 Issue: 11 DOI: 10.1088/0953-4075/47/11/115201

130. Bao, J.; Kuley, A.; Lu, Z.X. (2014) "Particle Simulation Of Lower Hybrid Wave Propagation In Fusion Plasmas." Plasma Physics and Controlled Fusion Volume: 56 Issue: 9 DOI: 10.1088/0741-3335/56/9/095020

131. Green, D.L.; Berry, L.A. (2014) "Iterative Addition Of Parallel Temperature Effects To Finite-Difference Simulation Of Radio-Frequency Wave Propagation In Plasmas." Computer Physics Communications Volume: 185 Issue: 3 Pages: 736-743 DOI: 10.1016/j.cpc.2013.10.032

132. Kim, Y.; Gupta, A.; Urgaonkar; et al. (2014) "Hybridplan: A Capacity Planning Technique For Projecting Storage Requirements In Hybrid Storage Systems." Journal of Supercomputing Volume: 67 Issue: 1 Pages: 277-303 DOI: 10.1007/s11227-013-0999-3

133. Romero-Redondo, C.; navratil, P.; Quaglioni, S.; et al. (2014) "Ab Initio NCSM/RGM for Three-Body Cluster Systems and Application to He-4+N+N." Few-Body Systems Volume: 55 Issue: 42226 Pages: 927-930 DOI: 10.1007/s00601-014-0876-5

134. Loehner, R. (2014) "Recent Advances in Parallel Advancing Front Grid Generation." Archives of Computational Methods in Engineering Volume: 21 Issue: 2 Pages: 127-140 DOI: 10.1007/s11831-014-9098-8

135. Philip, B.; Wang, Z.; Berrill, M.A.; et al. (2014) "Dynamic Implicit 3D Adaptive Mesh Refinement For Non-Equilibrium Radiation Diffusion." Journal of Computational Physics Volume: 262 Pages: 17-37 DOI: 10.1016/j.jcp.2013.12.058

136. Trossman, D.S.; Thopson, L.; Mecking, S.; et al. (2014) "Causes and Implications of Persistent Atmospheric Carbon Dioxide Biases in Earth System Models." Ocean Modeling Volume: 74 DOI: 10.1016/j.ocemod.2013.11.001

137. Hoffman, F.M.; Randerson, J.T.; Arora, V.K.; et al. (2014) "Causes and Implications of Persistent Atmospheric Carbon Dioxide Biases in Earth System Models." Journal of Geophysical Research-Biogeosciences Volume: 119 Issue: 2 Pages: 141-162 DOI: 10.1002/2013JG002381

138. Pindzola, M.S.; Lee, T.G.; Abdel-Baby, Sh A.; et al. (2014) "Neutron-Impact Ionization of He." Journal of Physics B – Atomic Molecular and Optical Physics Volume: 47 Issue: 19 DOI: 10.1088/0953-4075/47/19/195202

139. Kolla, H.; Hawkes, E.R.; Kerstein, A.R.; et al. (2014) "On Velocity and Reactive Scalar Spectra in Turbulent Premixed Flames." Journal of Fluid Mechanics Volume: 754 Pages: 456-487 DOI: 10.1017/jfm.2014.392

140. Shankar, S.K. (2014) "Numerical Investigation of Turbulence in Reshocked Richtmyer-Meshkov Unstable Curtain of Dense Gas." Shock Waves Volume: 24 Issue: 1 Pages: 456-487 DOI: 10.1017/jfm.2014.392

141. Cheng, Y.Q.; Bi, Z.H.; Huq, A.; Feygenson, M.; et al. (2014) "An Integrated Approach for Structural Characterization of Complex Solid State Electrolytes: The Case of Lithium Lanthanum Titanate." Journal of Materials Chemistry A Volume: 2 Issue: 7 Pages: 2418-2426 DOI: 10.1039/c3ta14433j

142. Daughton, W.; Nakamura, T.K.M.; Karimabadi, H.; et al (2014) "Computing the Reconnection Rate in Turbulent Kinetic Layers by Using Electron Mixing to Identify Topology." Physics of Plasmas Volume: 21 Issue: 5 DOI: 10.1063/1.4875730

143. Schoenherr, M.; Slater, B.; Hutter, J.; et al. (2014) "Dielectric Properties of Water Ice the Ice Ih/XI Phase Transition and an Assessment of Density Functional Theory." Journal of Physical Chemistry B Volume: 118 Issue: 2 Pages: 590-596 DOI: 10.1021/jp4103355

144. Taylor, C.N.; Allain, J.P.; Luitjohan, K.E.; et al. (2014) "Differentiating the Role of Lithium and Oxygen in Retaining Deuterium on Lithiated Graphite Plasma-Facing Components." Physics of Plasmas Volume: 21 Issue: 5 DOI: 10.1063/1.4874340

145. Fulton, D/P.; Lin, Z.; Holod, I.; et al. (2014) "Microturbulence In DIII-D Tokamak Pedestal. I. Electrostatic Instabilities." Physics of Plasmas Volume: 21 Issue: 4 DOI: 10.1063/1.4871387

146. Rouleau, C.M.; Shih, C-Y; Wu, C.; et al. (2014) "Nanoparticle Generation and Transport Resulting from Femtosecond Laser Ablation of Ultrathin Metal Films: Time-Resolved Measurements and Molecular Dynamics Simulations." Appied Physics Letters Volume: 104 Issue: 19 DOI: 10.1063/1.4876601

147. Zhou, J. (2014) "Stacking Interactions of Nickel Bis(Dithiolene) with Graphene and Beyond." RSC Advances Volume: 4 Issue: 26 DOI: 10.1039/c3ra48058e

148. Iyer, K.P.; Yeung, P.K. (2014) "Structure Functions and Applicability of Yaglom's Relation in Passive-Scalar Turbulent Mixing at Low Schmidt Numbers with Uniform Mean Gradient." Physics of Fluids Volume: 26 Issue: 8 DOI: 10.1063/1.4892581

149. Ambrosetti, A.; Alfe, D.; DiStasio, R.; et al. (2014) "Hard Numbers for Large Molecules: Toward Exact Energetics for Supramolecular Systems." Journal of Physical Chemistry Letters Volume: 5 Issue: 5 Pages: 849-855 DOI: 10.1021/jz402663k

150. Kamenetskiy, D.; Bussoletti, J.; Hilmes, C.; et al. (2014) "Numerical Evidence of Multiple Solutions for the Reynolds-Averaged Navier-Stokes Equations." AIAA Journal Volume: 52 Issue: 8 Pages: 16-86 DOI: 10.2514/1.J052676

151. Liu, D.; Zhang, W.; McClenaghan, J.; et al. (2014) "Verification of Gyrokinetic Particle Simulation of Current-Driven Instability in Fusion Plasmas. II. Resistive Tearing Mode." Physics of Plasmas Volume: 21 Issue: 12 DOI: 10.1063/1.4905074

152. Hong, L.; Sharp, M.; Poblete, S.; et al. (2014) "Structure and Dynamics of a Compact State of a Multidomain Protein the Mercuric Ion Reductase." Biophysical Journal Volume: 107 Issue: 2 Pages: 393-400 DOI: 10.1016/j.bpj.2014.06.013

153. Davidson, G.; Evans, T.; Jarrell, J.; et al. (2014) "Massively Parallel Three-Dimensional Transport Solutions for the k-Eigenvalue Problem." Nuclear Science and EngineeringVolume: 177 Issue: 2 Pages: 111-125 http://darkknight.ornl.gov:9001/548614f360158e15a51eb02d.pdf

154. Ellingson, S.; Dakshanamurthy, S.; Brown, M.; et al. (2014) "Accelerating Virtual High-Throughput Ligand Docking: Current Technology and Case Study on a Petascale Supercomputer." Concurrency and Computation-Practice & Experience Volume: 26 Issue: 6 Pages: 1268-1277 DOI: 10.1002/cpe.3070

155. Chai, J.; Kora, F.; Ahn, T-H.; et al. (2014) "Functional Phylogenomics Analysis of Bacteria and Archaea Using Consistent Genome Annotation with Unifam." BMC Evolutionary Biology Volume: 14 Issue: 207 DOI: 10.1186/s12862-014-0207-y

156. Kumar, R.; Sumpter, B.; Muthukumar, M. (2014) "Enhanced Phase Segregation Induced by Dipolar Interactions in Polymer Blends." Macromolecules Volume: 47 Issue: 18 Pages: 6491-6502 DOI: 10.1021/ma501067r

157. Zhou, J.; Huang, J.; Sumpter, B.; et al. (2014) "Theoretical Predictions of Freestanding Honeycomb Sheets of Cadmium Chalcogenides." Journal of Physical Chemistry C Volume: 118 Issue: 29 Pages: 16236-16245 DOI: 10.1021/jp504299e

158. McClenaghan, J.; Lin, Z.; Holod, I.; et al. (2014) "Verification of Gyrokinetic Particle Simulation of Current-Driven Instability in Fusion Plasmas. I. Internal Kink Mode." Physics of Plasmas Volume: 21 Issue: 12 DOI: 10.1063/1.4905073

159. Liu, Y-H.; Daughton, W.; Karimabadi, H.; et al. (2014) "Do Dispersive Waves Play a Role in Collisionless Magnetic Reconnection?." Physics of Plasmas Volume: 21 Issue: 2 DOI: 10.1063/1.4865579

160. Lopez-Bezanilla, A. (2014) "Electronic and Quantum Transport Properties of Substitutionally Doped Double-Walled Carbon Nanotubes." Journal of Physical Chemistry C Volume: 118 Issue: 3 Pages: 1472-1477 DOI: 10.1021/jp410648p

161. Lopez-Bezanilla, A.; Zhou, W.; Idrobo, J-C. (2014) "Electronic and Quantum Transport Properties of Atomically Identified Si Point Defects in Graphene." Journal of Physical Chemistry Letters Volume: 5 Issue: 10 Pages: 1711-1718 DOI: 10.1021/jz500403h

162. den Toom, M.; Dijkstra, H.; Weijer, W.; et al. (2014) "Response of a Strongly Eddying Global Ocean to North Atlantic Freshwater Perturbations." Journal of Physical Oceanography Volume: 44 Issue: 2 Pages: 464-481 DOI: 10.1175/JPO-D-12-0155.1

163. Chen, R.; McClean, J.; Gille, S.; et al. (2014) "Isopycnal Eddy Diffusivities and Critical Layers in the Kuroshio Extension from an Eddying Ocean Model." Journal of Physical Oceanography Volume: 44 Issue: 8 Pages: 2191-2211 DOI: 10.1175/JPO-D-13-0258.1

164. Sanyal, J.; New, J.; Edwards, R.; et al. (2014) "Calibrating Building Energy Models Using Supercomputer Trained Machine Learning Agents." concurrency and Computation-Practice & Experience Volume: 26 Issue: 13 Pages: 2122-2133 DOI: 10.1002/cpe.3267

165. da Cruz, C.; Seabra, G. (2014) "Molecular Dynamics Simulations Reveal a Novel Mechanism for ATP Inhibition of Insulin Degrading Enzyme." Journal of Chemical Information and Modeling Volume: 54 Issue: 5 Pages: 1380-1390 DOI: 10.1021/ci400695m

166. Seo, J.; Chang, C.S.; Ku, S.; et al. (2014) "Intrinsic Momentum Generation by a Combined Neoclassical and Turbulence Mechanism in Diverted DIII-D Plasma Edge." Physics of Plasmas Volume: 21 Issue: 9 DOI: 10.1063/1.4894242

167. Wang, S.; Kavaipatti, B.; Kim, S-J.; et al. (2014) "Atomic and electronic structures of lattice mismatched Cu2O/TiO2 interfaces." Applied Physics Letters Volume: 104 Issue: 21 DOI: 10.1063/1.4880942

168. Zhao, X.; Shu, Q.; Manh Cuong Nguyen, n.; et al. (2014) "Interface Structure Prediction from First-Principles." Journal of Physical Chemistry C Volume: 118 Issue: 18 Pages: 9524-9530 DOI: 10.1021/jp5010852

169. Paloncyova, M.; DeVane, R.; Murch, B.; et al. (2014) "Rationalization of Reduced Penetration of Drugs through Ceramide Gel Phase Membrane." Langmuir Volume: 30 Issue: 46 Pages: 13942-13948 DOI: 10.1021/la503289v

170. Ganesh, P.; Kim, J.; Park, C.; et al. (2014) "Binding and Diffusion of Lithium in Graphite: Quantum Monte Carlo Benchmarks and Validation of van der Waals Density Functional Methods." Journal of Chemical Theory and Computation Volume: 10 Issue: 12 Pages: 5318-5323 DOI: 10.1021/ct500617z

171. Hix, W.; Lentz, E.; Endeve, E.; et al. (2014) "Essential Ingredients in Core-Collapse Supernovae." AIP Advances Volume: 4 Issue: 4 DOI: 10.1063/1.4870009

172. Baer, M.; Kuo, I-F.; Tobias, D.; et al. (2014) "Toward a Unified Picture of the Water Self-Ions at the Air-Water Interface: A Density Functional Theory Perspective." Journal of Physical Chemistry B Volume: 118 Issue: 28 Pages: 8364-8372 DOI: 10.1021/jp501854h

173. Warren, M. (2014) "2HOT: An Improved Parallel Hashed Oct-Tree N-Body Algorithm for Cosmological Simulation." Scientfic Programming Volume: 22 Issue: 2 Pages: 109-124 DOI: 10.3233/SPR-140385

174. Paloncyova, M.; DeVane, R. Murch, B.; et al. (2014) "Amphiphilic Drug-Like Molecules Accumulate in a Membrane below the Head Group Region." Journal of Physical Chemistry B Volume: 118 Issue: 4 Pages: 1030-1039 DOI: 10.1021/jp4112052

175. Xiao, S.; Klein, M.; LeBard, D.; et al. (2014) "Magnesium-Dependent RNA Binding to the PA Endonuclease Domain of the Avian Influenza Polymerase." Journal of Physical Chemistry B Volume: 118 Issue: 4 Pages: 873-889 DOI: 10.1021/jp408383g

176. Jiang, P.; Lin, Z.; Holod, I.; et al. (2014) "Effects of Magnetic Islands on Drift Wave Instability." Physics of Plasmas Volume: 21 Issue: 12 DOI: 10.1063/1.4903910

177. Mintz, T.; Hernandez, O.; Bernholdt, D. (2014) "A Global View Programming Abstraction for Transitioning MPI Codes to PGAS Languages." Lecture Notes in Computer Science Volume: 8356 Pages: 120-133 DOI: 10.1007/978-3-319-05215-1_9

178. Li, Y.; Vogel, T.; Wuest, T.; et al. (2014) "A New Paradigm for Petascale Monte Carlo Simulation: Replica Exchange Wang-Landau Sampling." Journal of Physics: Conference Series Volume: 510 Issue: 1 DOI: 10.1088/1742-6596/510/1/012012

179. Liu, Y.; Gunasekaran, R.; Ma, X.; et al. (2014) "Automatic Identification of Application I/O Signatures from Noisy Server-Side Traces." 12th Usenix Conference on File and Storage Technologies https://www.usenix.org/system/files/conference/fast14/fast14-paper_liu.pdf

180. Oral, S.; Simmons, J.; Hill J.; et al. (2014) "Best Practices and Lessons Learned from Deploying and Operating Large-Scale Data-Centric Parallel File Systems." SC14 International Conference for High-Performance Computing, Networking, Storage and Analysis Pages: 217-228 DOI: 10.1109/SC.2014.23

181. Sanyal, J.; New, J. (2014) "Building Simulation Modelers—Are We Big Data Ready?." Proceedings of the Ashrae/IBSA-USA Building Simulation Conference Pages: 449-456 https://www.ashrae.org/File%20Library/.../57_Sanyal.pdf

182. Gunasekaran, R.; Kim, Y. (2014) "Feedback Computing in Leadership Compute Systems." 9th International Workshop on Feedback Computing https://www.usenix.org/system/files/conference/feedbackcomputing14/feedback14-gunasekaran.pdf

183. Liu, G.; Logan, J.; Tian, Y.; et al. (2014) "Hello ADIOS: The Challenges and Lessons of Developing Leadership Class I/O Frameworks." Concurrency and Computation: Practice & Experience Volume: 26 Issue: 7 Pages: 1453-1473 DOI: 10.1002/cpe.3125

184. Kartsaklis, C.; Park, E.; Cavazos, J. (2014) "HSLOT: The HERCULES Scriptable Loop Transformations Engine." SC14 International Conference for High-Performance Computing, Networking, Storage and Analysis Pages: 31-41 DOI: 10.1109/WOLFHPC.2014.1

185. Tiwari, D.; Gupta, S.; Vazhkudai, S. (2014). "Lazy Checkpointing : Exploiting Temporal Locality in Failures to Mitigate Checkpointing Overheads on Extreme-Scale Systems." 2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Pages: 25-36 DOI: 10.1109/DSN.2014.101

186. Maxwell, D.; Ezell, M.A.; et al. (2014) "Monitoring Cray Cooling Systems" 2014 Cray User Group Volume: 5 HTTP Link: https://cug.org/proceedings/cug2014_proceedings/includes/files/pap181.pdf

187. Hjelm, N.T.; Gutierrez, S.K.; Gorentla Venkata, M. (2014) "On the Current State of Open MPI on Cray Systems" 2014 Cray User Group HTTP Link: http://cs.unm.edu/~samuel/publications/2014/cug2014-on-the-state.pdf

188. Salko, R.; Schmidt, R.; Avramova, M.N. (2014) "Optimization and Parallelization of the Thermal-Hydraulic Sub-Channel Code CTF for High-Fidelity Multi-Physics Applications" Annals of Nuclear Energy In Press DOI: 10.1016/j.anucene.2014.11.005

189. Dai, H.; Kartsaklis, et al. (2014) "RACB: Resource Aware Cache Bypass on GPUs" 2014 International Symposium on Computer Architecture and High-Performance Computing Workshop DOI: 10.1109/SBAC-PADW.2014.14

190. Boyuka, D.; Lakshminarasimhan, S.; Zou, X.; et al. (2014) "Transparent In Situ Data Transformations in ADIOS" 13th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing DOI: 10.1109/CCGrid.2014.73

191. Grosvenor, A.; Rixon, G.; et al. (2014) "High Resolution RANS Nonlinear Harmonic Study of Stage 67 Tip Injection Physics" Journal of Turbomachinery Volume: 137 Issue: 5 DOI: 10.1115/1.4028550

192. Wang, T.; Oral, S.; Wang, Y.; et al. (2014) "BurstMem: A High-Performance Burst Buffer System for Scientific Applications" IEEE Big Data 2014 HTTP Link: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7004215

193. Zhang, Y.; Liu, G.; Klasky, S.; et al. (2014) "Active Workflow System for Near Real-Time Extreme Scale Science" 2014 Proceedings of the Workshop on Parallel Programming for Analytical Applications" DOI: 10.1145/2567634.2567637

194. Boehm, S.; Elwasif, W.R.; Naughton III, T.J.; Vallee, G.R. (2014) "Toward Improved Support for Loosely Coupled Large Scale Simulation Workflows" 2014 Cray User Group HTTP Link: https://cug.org/proceedings/cug2014_proceedings/includes/files/pap147.pdf

195. Jones, T.R.; Settlemyer, B.W. (2014) "Fan-In Communications on a Cray Gemini Interconnect" 2014 Cray User Group HTTP Link: https://cug.org/proceedings/cug2014_proceedings/includes/files/pap138.pdf

196. Ezell, M.A.; Dillow, D.; et al. (2014) "I/O Router Placement and Fine-Grained Routing on Titan to Support Spider II" 2014 Cray User Group HTTP Link: https://cug.org/proceedings/cug2014_proceedings/includes/files/pap116.pdf

197. Touma, D.E.; Ashfaq, M.; et al. (2014) "A Multi-Model and Multi-Index Evaluation of Drought Characteristics in the 21st Century" Journal of Hydrology In Press DOI: 10.1016/j.jhydrol.2014.12.011

198. Patel, P.; Schmidt, D.; Chen, W.; Ostrouchov, G. (2014) "High-Level Analytics with R and pbdR on Cray Systems" 2014 Cray User Group

199. Grosvenor, A.; Rixon, G.; et al. (2014) "High Resolution RANS NLH Study of Stage 67 Tip Injection Particles" ASME Turbo Expo 2014: 2B Turbine Technical Conference and Exposition DOI: 10.1115/GT2014-27219

200. Park, E.J.; Kartsaklis, C.; Cavazos, J. (2014) "HERCULES: Strong Patterns Towards More Intelligent Predictive Modelling" 2014 International Conference on Parallel Processing DOI: 10.1109/ICPP.2014.26

201. Steed, C.; Evans, K.J.; et al. (2014) "Web-Based Visual Analytics for Extreme Scale Climate Science" IEEE Big Data 2014 DOI: 10.1109/BigData.2014.7004256

202. Langan, R.T.; Archibald, R.K.; et al. (2014) "Stochastic Parameterization to Represent Variability and Extremes in Climate Modeling" Procedia Computer Science Volume: 29 DOI: 10.1016/j.procs.2014.05.103

203. Liakh, D.I. (2014) "Scale-Adaptive Tensor Algebra for Local Many-Body Methods of Electronic Structure Theory" International Journal of Quantum Chemistry Volume: 114 Issue: 23 DOI: 10.1002/qua.24732

204. Kim, S.; Luong, M.; Chen, J.; Yoo, C. (2014) "A DNS Study of the Ignition of Lean PRF/Air Mixtures with Temperature Inhomogeneities Under High Pressure and Intermediate Temperature" Combustion and Flame in press DOI: 10.1016/j.combustflame.2014.09.001

205. Dong, T.; Dobrev, V.; et al. (2014) "A Step Towards Energy Efficient Computing: Redesigning A Hydrodynamic Application on CPU-GPU" IEEE 28th International Parallel and Distributed Processing Symposium DOI: 10.1109/IPDPS.2014.103

206. Maris, P. (2014) "Ab Initio Calculations of P-Shell Nuclei with JISP16" International Conference on Nuclear Theory in the Supercomputing Era (NTSE 2013) HTTP Link: http://www.ntse-2013.khb.ru/Proc/Maris.pdf

207. Potter, H.; Oryspayev, D.; et al. (2014) "Accelerating Ab Initio Nuclear Physics Calculations with GPUs" International Conference on Nuclear Theory in the Supercomputing Era (NTSE 2013) HTTP Link: http://www.ntse-2013.khb.ru/Proc/Sosonkina.pdf

208. Vary, J. Shirokov, A.M. & A.I. Mazur eds. (2014) "Computational Nuclear Physics: Key to Discovery Opportunities" International Conference on Nuclear Theory in the Supercomputing Era (NTSE 2013) HTTP Link: http://www.ntse-2013.khb.ru/Proc/JPVary.pdf

209. Nguyen, H.; Sho, H.; Xu, J.; Zhang, S. (2014) "CPMC-Lab: A Matlab Package for Constrained Path Monte Carlo Calculations" Computer Physics Communications Volume: 185 Issue: 12 DOI: 10.1016/j.cpc.2014.08.003

210. Zherebetskyy, D.; Wang, L. (2014) "In-Gap States in Electronic Structure of Nonpolar Surfaces of Insulating Metal Oxides" Advanced Material Interfaces Volume: 1 Issue: 8 DOI: 10.1002/admi.201300131

211. Abe, T.; Maris, P.; et al. (2014) "Monte Carlo Shell Model for Ab Initio Nuclear Structure" EPJ Web Conferences Volume: 66 DOI: 10.1051/epjconf/20146602001

212. Gaudin, W.; Mallinson, A.; et al. (2014) "Optimising Hydrodynamics Applications for the Cray XC30 with the Application Tool Suite" 2014 Cray User Group HTTP Link: https://cug.org/proceedings/cug2014_proceedings/includes/files/pap124.pdf

213. McClure, J.; Wang, H.; et al. (2014) "Petascale Application of A Coupled CPU-GPU Algorithm for Simulation and Analysis of Multiphase Flow Solutions in Porous Medium Systems" IEEE 28th International Parallel and Distributed Processing Symposium DOI: 10.1109/IPDPS.2014.67

214. Vary, J. (2014) "Quantum Hamiltonian Physics with Supercomputers" Nuclear Physics B-Proceedings and Supplements Volume: 251-252 DOI: 10.1016/j.nuclphysbps.2014.05.003

215. Karimabadi, H.; Roytershteyn, V.; Daughton, W.; Liu, Y.H. (2014) "Recent Evolution in the Theory of Magnetic Reconnection and Its Connection with Turbulence" Microphysics of Cosmic Plasmas: Space Sciences Series of ISSI Volume: 47 DOI: 10.1007/978-1-4899-7413-6_9

216. Almgren, A.; Bell, J.; Nonaka, A.; Zingale, M. (2014) "Low Mach Number Modeling of Stratified Flows" Finite Volumes for Complex Applications VII-Methods and Theoretical Aspects Springer Proceedings in Mathematics and Statistics Volume: 77 DOI: 10.1007/978-3-319-05684-5_1

217. Phillips, J.; Sun, Y.; et al. (2014) "Mapping to Irregular Torus Topologies and Other Techniques for Petascale Biomolecular Simulation" SC14 International Conference for High-Performance Computing, Networking, Storage and Analysis DOI: 10.1007/978-3-319-05684-5_1

218. Lyra, S.; Wilde, B.; et al. (2014) "Structure of Hydrogen-Rich Transverse Jets in A Vitiated Turbulent Flow" Combustion and Flame In press DOI: 10.1109/SC.2014.12

219. Gahvari, H.; Gropp, W.; et al (2014) "Algebraic Multigrid on A Dragon y Network: First Experiences on a Cray XC30" SC14 International Conference for High-Performance Computing, Networking, Storage and Analysis HTTP Link: http://www.dcs.warwick.ac.uk/~sdh/pmbs14/PMBS14/Workshop_Schedule_files/1-AlgebraicMultigrid.pdf

220. Mudalige, G.; Reguly, I.; et al. (2014) "Performance Analysis of A High-Level Abstractions-Based Hydrocode on Future Computing Systems" SC14 International Conference for High-Performance Computing, Networking, Storage and Analysis HTTP Link: http://www.dcs.warwick.ac.uk/~sdh/pmbs14/PMBS14/Workshop_Schedule_files/11-PerformanceHydrocode.pdf

221. Bedorf, J.; Gaburov, E.; et al. (2014) "24.77 Pflops on a Gravitational Tree-Code to Simulate the Milky Way Galaxy with 18,600 GPUs" SC14 International Conference for High-Performance Computing, Networking, Storage and Analysis DOI: 10.1109/SC.2014.10

222. Landge, A.; Pasucci, V.; et al. (2014) "In-Situ Feature Extraction of Large-Scale Combustion Simulations Using Segmented Merge Trees" SC14 International Conference for High-Performance Computing, Networking, Storage and Analysis DOI: 10.1109/SC.2014.88

223. Gamell, M.; Katz, D.; Kolla, H.; et al. (2014) "Exploring Automatic Online Failure Recovery for Scientific Applications at Extreme Scales" SC14 International Conference for High-Performance Computing, Networking, Storage and Analysis DOI: 10.1109/SC.2014.78

224. Dong, R.; Calzolari, A.; et al. (2014) "Optical Enhancement in Heteroleptic Ru(II) Polypyridyl Complexes Using Electrong-Donor Ancillary Ligands" Journal of Physical Chemistry C Volume: 118 Issue: 17 DOI: 10.1021/jp409733a

225. Vose, A.; Mitchell, B.; Levesque, J. (2014) "Tri-Hybrid Computational Fluid Dynamics on DOE's Cray XK7 Titan" 2014 Cray User Group

226. Zhang, Z.; Carrillo, J.M.; et al. (2014) "Atomistic Structure of Bottlebrush Polymers" Simulations and Neutron Scattering Studies" Macromolecules Volume: 47 Issue: 16 DOI: 10.1021/ma500613c

227. Parete-Koon, S.; Caldwell, B.; et al. (2014) "HPC's Pivot to Data" 2014 Cray User Group

228. Janjusic, T.; Kartsaklis, C.; Wang, D. (2014) "Scalability Analysis of Gleipnir: A Memory Tracing and Profiling Tool, on Titan" 2014 Cray User Group

229. Carlyle, A.; French, R.; Renaud, W. (2014) "Designing Service-Oriented Tools for HPC Account Management and Reporting" 2014 Cray User Group

Acronyms

ALCC–ASCR Leadership Computing Challenge

ALCF–Argonne Leadership Computing Facility

ALS–Advanced Light Source

ALTD–Automatic Library Tracking Database

ASCR–Advanced Scientific Computing Research

CAAR–Center for Accelerated Application Readiness

CCI–Common Communication Interface

CESM–Community Earth System Model

CNMS–Center for Nanophase Materials Science

CORAL–Collaboration of Oak Ridge, Argonne, and Lawrence Livermore National Laboratories

CPU–central processing unit

DD–Director's Discretionary

DFT–density functional theory

DME–dimethyl ether

DOE–US Department of Energy

EVEREST–Exploratory Visualization Environment for Research in Science and Technology

FP–Functional Partitioning

GE–General Electric

GPU–graphics processing unit

HACC–Hardware/Hybrid Accelerated Cosmology Code

HiPSTAR–high performance solver for turbulence and aeroacoustic research

HPC–high-performance computing

HPSS–High-Performance Storage System

HT–Hyper-Threading

I/O–input/output

INCITE–Innovative and Novel Computational Impact on Theory and Experiment

IOSI–I/O Signature Identifier

LBNL–Lawrence Berkeley National Laboratory

LLNL–Lawrence Livermore National Laboratory

LSST–Large Synoptic Survey Telescope

MTBF–Mean Time Between Failures

NERSC–National Energy Research Scientific Computing Center

NFS–Network File System

NICS–National Institute for Computational Sciences

NNSA–National Nuclear Security Administration

OLCF–Oak Ridge Leadership Computing Facility

ORNL–Oak Ridge National Laboratory

OUG–OLCF User Group

P&G–Procter & Gamble

PI–principal investigator

PTM–posttranslational modification

QMC–quantum Monte Carlo

RAIT–Redundant Arrays of Inexpensive Tape

RUR–Resource Utilization Reporting

SDRAM–synchronous dynamic random access memory

STCI–Scalable runTime Component Infrastructure

UAO–User Assistance and Outreach Group

UC Berkeley–University of California, Berkeley

UEB–User Executive Board

UT–University of Tennessee

Managed by UT-Battelle for the US Department of Energy