



# Oak Ridge Leadership Computing Facility

Annual Report 2012—2013

Oak Ridge Leadership Computing Facility  
Annual Report 2012—2013







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




- 4** Titan Proves Its Value
- 6** Researchers Harness ORNL Supercomputers to Understand Solar Storms
- 8** Researchers Use ORNL Supercomputer to Pinpoint Process Behind Last Deglaciation
- 10** Supernovas Acting as Nuclear Pasta Factories
- 12** Big Computing Cures Big Pharma
- 14** Simulation Tackles the Challenge of Keeping Fuel Rods in Place



- 16** Engine Bay Redesign Boosts Fuel Efficiency
- 18** Oceans Will Rise, but by How Much?
- 20** Ramgen Simulates Shock Waves, Makes Shock Waves Across Energy Spectrum
- 22** Jaguar Guides Demonstration of Novel Quantum State
- 24** Mapping the Nuclear Landscape
- 26** Titan Is the World's Most Powerful Supercomputer for Science
- 32** Top Rankings Give Experts Chance To Share Titan's Science Impact
- 34** OLCF Experts Ready Users for Titan
- 38** Putting the "User" in User Facility
- 44** Advancements
- 48** OLCF Systems Overview
- 53** High-Impact Publications
- 54** Education, Outreach, and Training
- 56** Computing Allocations for 2013





**James Hack**  
Director, National Center  
for Computational Sciences

## Titan Proves Its Value

*ORNL's accelerated behemoth lays the groundwork for years of scientific breakthrough*

Now that ORNL's Titan supercomputer has completed acceptance testing and is dedicated full time to scientific discovery, I'd like to look back on some of the experiences we have witnessed in standing up this ground-breaking system.

One of the most important decisions made in the plan to provide significantly enhanced computational capabilities was to upgrade the Jaguar system in place. We recognized that this kind of upgrade would present additional complexities, but it also appreciably lowered the cost of completing an order of magnitude augmentation to what Jaguar was able to provide to our user community. We fully appreciated the impact this upgrade strategy would have on the user community and are grateful for your patience as we completed the elaborate upgrade plan. We realize that the process was at times frustrating, but the end result is an unmatched capability available to the open science community.

One of the critical decisions in the definition of the Titan architecture was to pursue a relatively new technology—GPU accelerators. The technology that made Jaguar the world's most powerful system in 2009 and 2010—hundreds of thousands of CPU cores connected in parallel—had reached reasonable limits of space and power. For example, a CPU-only system to replace Jaguar would have required five times the space and electricity to obtain the same performance as Titan. It had become clear that the

path to exascale capabilities, 500 times the computational performance of Jaguar, would require new architectural and technological approaches.

So, we turned to the enormous potential of GPUs. GPUs were created to perform highly parallel graphic tasks associated with image processing, but their natural ability to exploit parallelism has enormous potential for highly structured scientific computing applications.

The upshot is that Titan is an order of magnitude more powerful than Jaguar while occupying the same physical space and demanding only modestly more power. We are extremely pleased with Titan's early results.

As long ago as November, Titan was ranked in first place on the Top500 list of the world's leading supercomputers, solving a synthetic benchmark faster than any other system at that time. But Titan's real value lies in its ability to tackle formidable science and engineering problems, and here it is succeeding as well. Titan's performance is nearly four times that of a comparable CPU-only system on two leadership codes—Denovo, which is used to simulate neutron transport in a nuclear reactor, and WL-LSMS, which analyzes the statistical mechanics of magnetic materials. The speedup is even greater—7½ times—on the high-performance molecular dynamics code LAMMPS. We are seeing other examples of success stories for many of our large user applications.

It is also important that we are not alone in developing tools and programming approaches that exploit the potential of this hybrid computer architecture. ORNL is part of a larger community of researchers and centers pushing the limits of accelerated supercomputing. Examples include researchers at San Diego State University and the San Diego Supercomputing Center who have used GPUs to double their potential performance for AWP-ODC, a leading code for simulating ground movement during an earthquake. Our colleague Thomas Schulthess, who spends most of his time as director of the Swiss National Supercomputing Center, has seen a speedup of more than 400 percent with DCA++, a leading code for condensed matter physics. As researchers spend more time with Titan and other GPU-accelerated systems, we expect to see these numbers continue to rise. We are, in fact, here to ensure our users get years of breakthrough results with Titan.

At the same time, though, we're looking forward to our next advance in supercomputing. ORNL has teamed with the Argonne and Lawrence Livermore national laboratories to take the next step toward the exascale. It's far too early to say what our next system will look like, but I can say one thing: the added levels of parallelism that our users are exposing for GPU accelerators will be critical to their success in the future. Everything we know about the future of supercomputing points to the importance of exploiting very fine-grained, local parallelism in their applications. The investments the user community makes to fully exploit the Titan architecture are going to be realized into the indefinite future. We look forward to traveling that path with you and facilitating the many scientific breakthroughs available through leadership-class computational approaches.

“Titan is an order of magnitude more powerful than Jaguar while occupying the same physical space and demanding only modestly more power.”





**Homa Karimabadi**  
University of California,  
San Diego  
PI: William Daughton  
Los Alamos National Laboratory  
INCITE, 75 million hours

## Researchers Harness ORNL Supercomputers to Understand Solar Storms

The sun provides invaluable services to life on Earth but also creates space weather, spewing massive amounts of radiation in our direction.

Fortunately for us, Earth's magnetic field creates a shield known as the magnetosphere that protects us from much of the sun's worst storms. Unfortunately, it's not perfect. The radiation unleashed by the sun as an ionized gas known as plasma can sneak past Earth's defenses in a process known as magnetic reconnection, creating the northern lights while wreaking havoc on electronics. These solar storms can knock out power grids, rendering regions without electricity. More than \$4 billion in satellite damage from space weather has been documented.

Radiation from a solar storm can take one to five days to reach Earth. If we knew a storm was coming, we could take preemptive measures, preventing power and communication disturbances. "When a storm goes off on the sun, we can't really predict the extent of damage that it will cause here on Earth. It is critical that we develop this predictive capability," said Homa Karimabadi, a space physicist at the University of California–San Diego. Karimabadi's team used about 30 million hours on Jaguar in 2012 to do just that.

Essentially, Karimabadi's team studied how plasma interacts with Earth's magnetosphere, what gets through and why.

Thanks to the power of Jaguar, the complex processes behind solar storms are finally becoming understandable. While Karimabadi has been at this research for 15 years, only with petascale computing have the simulations come close to reality. It's now possible to model magnetic reconnection, including kinetic effects—the individual particles that make up the plasma ejected by the sun. The particles are numerous, hence the need for systems such as Jaguar.

Back in 2006, the team's largest simulations contained around 1 billion individual particles. With Jaguar's capabilities, they modeled 3.2 trillion particles, revealing the reconnection phenomenon in a whole new light.

The simulations have motivated observational studies to search for the newly discovered effects. Currently, the team's research is being used in planning for the National Aeronautics and Space Administration's (NASA's) Magnetospheric Multiscale Mission. In 2014, NASA will launch four spacecraft into the magnetosphere to study the processes that play a role in space weather.


And now the researchers are ready to take advantage of the increase in power available with Titan. The team's most recent simulations on Jaguar used more than 200,000 of the machine's 224,000-plus compute cores, no small feat for even the most aggressive applications.

With Titan clocking in at more than 20 petaflops, space physicists like Karimabadi will soon have a chance to take advantage of even more computing muscle, a union that will produce an unprecedented picture of space weather's effect on Earth's magnetosphere and our lives. In the meantime we will continue to weather the sun's storms.

— By Gregory Scott Jones

For more: <https://www.olcf.ornl.gov/2012/02/06/when-worlds-collide/>

“When a storm goes off on the sun, we can't really predict the extent of damage that it will cause here on Earth. It is critical that we develop this predictive capability,” — Homa Karimabadi, University of California, San Diego



*The complex entanglement and mixing of the field lines results from the formation of a magnetic flux rope generated by the magnetic reconnection process. Image courtesy of H. Karimabadi and B. Loring*





**Feng He**

University of Wisconsin,  
Madison  
PI: Warren Washington  
National Center for Atmosphere  
Research  
INCITE, 48 million hours

## Researchers Use ORNL Supercomputer to Pinpoint Process Behind Last Deglaciation

While much is made of today's melting Arctic ice cap, the phenomenon is not a new one.

In fact, about 22,000 years ago the Earth's great ice sheets began to decline, slowly at first, but gradually more rapidly. Given the growing concerns about today's shrinking glaciers and polar ice caps, scientists are very interested in knowing what happened the last time the Earth shed much of its ice. A better understanding of natural climate change should greatly assist in advancing our understanding of man-made climate change.

Researchers agree that a rapid, natural release of carbon dioxide (CO<sub>2</sub>) about 17,000 years ago led to a rise in global temperatures that further encouraged deglaciation in both the northern and southern hemispheres, though at very different rates. What started the ball rolling 22,000 years ago, however, was until recently still a mystery.

Now researchers from the University of Wisconsin—Madison (UW—Madison), Harvard University, Oregon State University, and the National Center for Atmospheric Research (NCAR) have discovered the trigger for the beginning of the last great deglaciation. The team ran transient, or continuous, simulations on Oak Ridge National Laboratory's (ORNL's) Jaguar supercomputer over three years to create the first physics-based test of hemispheric deglaciation. Their culprit: a combination of increase in insolation (solar radiation that reaches the Earth's surface) caused by changes in the Earth's orbit and ocean circulation.

The simulations, conducted by Feng He and Zhengyu Liu of UW—Madison and Bette Otto-Bliesner of NCAR, help to recreate the climate during the first half of the last deglaciation period and identify why temperatures and deglaciation rates differed between the hemispheres. The research builds on earlier simulations performed at ORNL and featured in *Science* in 2009 and *Nature* in 2012. Their latest finding detailing ocean circulation as the primary cause of early deglacial warming in the Southern Hemisphere appears in the February 7 issue of *Nature*.

The research is part of a larger initiative that has succeeded in obtaining a mean global temperature for the past 21,000 years, giving scientists a much-needed tool with which to compare carbon dioxide levels and temperatures across the world. Results from the simulations are shedding light on our current climate conundrum.

And quite possibly our past. The team also believes that this same mechanism may apply to all of the major deglaciations for the last 450,000 years.

The team's weapon of choice in deciphering the climate of the last 20,000-plus years: the Community Climate System Model (CCSM). Since its creation in 1983, CCSM has evolved into the world's leading tool for climate simulation, regularly contributing to the United Nations' Intergovernmental Panel on Climate Change reports.

"The simulation reproduces the Southern Hemisphere proxy records beautifully. A good model is the result of many people's efforts," said He.


The Oak Ridge Leadership Computing Facility (OLCF) has given the project nearly four continuous years of access, allowing the team to run climate simulations over 22,000 years and produce nearly 300 terabytes of data. "We have the resources to stage all data online for analysis," said the OLCF's Valentine Anantharaj, who worked with the team to make sure they got the most from their time on Jaguar. Anantharaj now works with users on the tenfold-more-powerful Titan system, and according to him the OLCF represents a valuable end-to-end resource capability: "Our facility supports a scientific workflow that enables our users to run their simulations, do their analyses, and visualize and archive the results."

— By Gregory Scott Jones


For more: <https://www.olcf.ornl.gov/2013/02/08/lessons-from-the-past/>

"The simulation reproduces the Southern Hemisphere proxy records beautifully. A good model is the result of many people's efforts."

— Feng He, University of Wisconsin–Madison



*Global paleoclimate reconstruction of temperature and CO<sub>2</sub> and the Jaguar supercomputer simulation showing that CO<sub>2</sub> drove global warming at the end of the last ice age.*



**Jirina Stone**

University of Tennessee,  
Knoxville

PI: Anthony Mezzacappa  
Oak Ridge National Laboratory  
INCITE, 28 million hours

## Supernovas Acting as Nuclear Pasta Factories

The collapse and bounce of a giant star's iron core is an event of such jaw-dropping violence that it blows the star into space.

This is a good thing. Not only is the nebula produced by a core-collapse supernova beautiful; it also provides the seed material for new solar systems, including our own. You, in fact, are supernova flotsam.

The supernova illustrates all of nature's forces—from the gravity that extends across the universe to the nuclear force that binds and arranges the smallest building blocks of matter. It also illustrates what happens when these forces battle one another for dominance.

The University of Tennessee—Knoxville's Jirina Stone and Helena Pais have run the most detailed and accurate simulation to date of one such battle. For a split second as the star begins to blow outward, the Coulomb force and surface tension will be in balance, causing about 20 percent of the core's matter to form into cylinders, sheets, bubbles, and other odd shapes.

This is nuclear pasta—a form of frustrated matter—a rare ordering found within the supernova and in the neutron star that will be all that remains when it is over. As the star blows outward, the layer of nuclear pasta will be in the core about 100 kilometers from the center.



The Coulomb force governs charged particles, dictating, for instance, that positively charged particles will repel one another and attract negatively charged particles. Surface tension, on the other hand, is the internal pressures that pull liquids into the smallest possible area (think water droplets); in this case it is enticing the neutrons and protons to form coherent structures.

“Imagine what happens,” said Stone. “When the neutrons and protons are close enough that the strength of the Coulomb force and nuclear forces are comparable, the Coulomb force is telling all the nucleons to get apart, while the surface forces are saying, ‘Come on, I want to keep you together.’”

The team simulated the system with an approach called Skyrme-Hartree-Fock, with Pais handling the calculations. All told, the project used 30 million processor hours on Jaguar, with each 12-hour run making use of 45,000 processor cores.


“‘Hartree-Fock’ is the mathematical method that makes a description of systems made of many quantum particles possible to the best approximation,” Stone explained, “as exact treatment of such a system is currently beyond reach of the fastest computers in the world. The Skyrme interaction is a prescription of how nucleons talk to each other.”

The project simulated the conditions of nuclear pasta by simulating a series of tiny cells, each just 25 quadrillionths of a meter on a side. By modeling a range of these cubes and calculating what happens to the matter under different temperature and density conditions, the project provides all the information needed by a supernova simulation to accurately describe the pasta phase of nuclear matter.

Stone and Pais published their results in the October 12, 2012, issue of *Physical Review Letters*. — *By Leo Williams*

For more: <https://www.olcf.ornl.gov/2013/03/26/supernovas-as-nuclear-pasta-factories/>

The supernova illustrates all of nature’s forces—from the gravity that extends across the universe to the nuclear force that binds and arranges the smallest building blocks of matter. It also illustrates what happens when these forces battle one another for dominance.



*Example of the neutron density distribution in supernova matter forming cylindrical holes, a pasta structure, where the red areas represent the maximum density values and the blue ones the minimum. Image courtesy of Helena Pais*



**Jerome Baudry**  
University of Tennessee,  
Knoxville, and ORNL  
PI: Jeremy Smith  
Oak Ridge National Laboratory  
INCITE, 8 million hours

## Big Computing Cures Big Pharma

Drugmakers spend \$50 billion each year on research and development in hopes of finding the next big money maker, but the failure rate of new drug candidates is staggering. In fact, less than 1 percent of drug candidates starting in the lab make it to market.

According to Jerome Baudry, an assistant professor at the University of Tennessee (UT) and member of the Center for Molecular Biophysics at ORNL, the high cost of drugs is due in part to the fact that manufacturers have to recoup the cost of the drugs that fail.

The industry responsible for keeping Americans healthy is itself in need of a prescription. UT and ORNL researchers believe that injecting the pharmaceutical industry with a strong dose of supercomputing might just be the remedy.

"We are working toward delivering new drugs to the market at a fraction of the cost and time that it takes now," says Baudry.

The cost of developing and bringing one drug to market typically ranges anywhere from a few hundred million dollars to more than a billion. And it's not out of the ordinary for drug development to take from 10 to 15 years before patients get the meds they need. The simple reason is that it takes a really long time to test pharmaceuticals.

## The computational elixir

The cure-all calls for tens of thousands of processing cores churning through 1 quadrillion calculations per second. In fact, the Jaguar supercomputer at ORNL has already tested 2 million different compounds against a targeted receptor—and it did so amazingly quickly. Comparatively speaking, what would have taken smaller computing systems months to accomplish, and conventional test-tube methods even longer, Jaguar did in less than 72 hours.

Pharmaceuticals work because they bind to specific molecular receptors. However, when compounds attach themselves to receptors other than the ones causing the problem, the drug becomes less efficient.

By parallelizing the publicly licensed virtual high-throughput softwares Autodock4 and AutodockVina from the Scripps Research Institute, Baudry and his colleagues were able to re-create the molecular binding events with three-dimensional biological simulations.

With so many combinations, though, “It’s like trying to find a needle in a haystack,” says Baudry.

Fortunately, being able to run simulations greatly reduces the size of the haystack, making it easier to find the right combination. As poor drug candidates get eliminated, more specifically binding, therefore more efficient, drugs start to emerge.


The next time Baudry and his collaborator, Jeremy Smith, conduct virtual screenings for new drug candidates, it will not be on Jaguar but on Titan, Jaguar’s successor.

“In the field of new drug discovery, Titan will be the equivalent of the Saturn V rocket. We are going to be able to fly to the moon,” Baudry says. “We will be able to simulate going inside patients and inside their cells instead of the test tube, and that’s a revolution.”

—By *Jeremy Rumsey*

For more: <https://www.olcf.ornl.gov/2012/10/18/big-computing-cures-big-pharma/>

“We are working toward delivering new drugs to the market at a fraction of the cost and time that it takes now.” —Jerome Baudry, University of Tennessee



*Computational approaches are used to describe how molecular compounds of a drug candidate (displayed in colored spheres) bind to its specific protein receptor target. Image courtesy of Jerome Baudry, ORNL*





*Energy Secretary Ernest Moniz takes in the latest findings from ORNL's Consortium for the Advanced Simulation of Light-Water Reactors.*



**Mark Christon**  
Los Alamos National Laboratory  
PI: John Turner  
Oak Ridge National Laboratory  
ALCC, 21.9 million hours

## Simulation Tackles the Challenge of Keeping Fuel Rods in Place

It's tough to keep a nuclear fuel rod from rattling in its cage.

At 12 feet long, the rod is barely wider than a crayon as it is buffeted by hot, swirling, pressurized water and bombarded by neutrons. It and more than 200 others are held in place by as many as a dozen grids containing strong, sturdy springs, yet the frames and the springs sit in the same extreme environment as the rods.

Eventually the springs lose their grip, the rods get banged around, and the cladding around the rods gets worn away. If the rod is left in place long enough, products of the fission process taking place in the fuel can leak into the water. This is not good.

Calming the water is not the answer. The water's turbulence is so important that the grids have mixing vanes, little wings that jut out to keep the water stirred up.

One major challenge for the reactor's operator, then, is to keep the rods safely immobilized inside the reactor as long as possible without risking excessive fuel wear. Researchers with the Consortium for the Advanced Simulation of Light-Water Reactors (CASL) are working to understand this challenge better by simulating it on Titan.

The technical term for damage done to rattling fuel rods is grid-to-rod fretting (GTRF). CASL researchers re-create the problem on Titan with a computational fluid dynamics code called Hydra-TH.

“GTRF is the primary failure mechanism for fuel,” said CASL collaborator Mark Christon of Los Alamos National Laboratory (LANL). “The mixing vanes tend to generate a lot of turbulence, which leads to fluctuations in the pressure that make the rods vibrate.

“If they can design it out with better material or spacers, they can get a longer life from the fuel, as well as increased power output.”

The researchers at this point are focused on pressurized water reactors (PWRs), the most common design for American power plants. By using Hydra-TH, Christon and collaborators within CASL are able to model the thermal hydraulics of a PWR by dividing it into as many as 250 million separate cells.

On Titan, the project has scaled up to 36,000 of the system’s nearly 300,000 central processing unit (CPU) cores. Jozsef Bakosi, another LANL researcher, said the team is working with graphics processing unit (GPU) maker NVIDIA to take full advantage of Titan’s GPU accelerators. Specifically, he said, they’re working with the company to adapt its linear algebra library for Hydra-TH.


As it moves forward Hydra-TH will be used for problems other than GTRF, noted John Turner of ORNL, CASL’s lead for virtual reactor integration.

“There are multiple things that we need Hydra for,” he said. “You need a detailed flow to understand the heat transfer, the convection of the heat out of the rods in the fluid. Up to now there’s been much more crude or course flow models, and as long as you could validate to experiments, that was fine. But we’re getting into regimes where we need a more detailed description of the flow.” — *By Leo Williams*

For more: <https://www.olcf.ornl.gov/2013/06/19/hold-it-right-there/>

“If they can design it out with better material or spacers, they can get a longer life from the fuel, as well as increased power output.”

—Mark Christon,  
Los Alamos National  
Laboratory



*Two-phase flow near the spacer grid for a 3 x 3 rod bundle. The isosurfaces show the volume-fraction of the liquid fluid. Image courtesy of Mark Christon, LANL*



**Alex Akkerman**

Ford Motor Company  
PI: Brad VanDerWege  
Ford Motor Company  
Director's Discretion,  
1 million hours

## Engine Bay Redesign Boosts Fuel Efficiency

Designing cars that satisfy consumer expectations is hard enough as it is. Car buyers want it all—safety, speed, power, beauty, fuel efficiency—without sacrificing amenities. Furthermore, the White House has mandated that new cars average 35.5 miles per gallon by 2016 and 54.5 mpg by 2025.

No stranger to challenges, the Ford Motor Company is focusing its creative energy on reaching these milestones.

“Delivering a more energy-efficient vehicle while maintaining design leadership and high-package efficiency, and fulfilling future safety standards, is a very complex task requiring a new engineering approach,” said Burkhard Hupertz, a thermal and aerosystems engineer at Ford of Europe.

That new approach required new design methods, in Hupertz's case to optimize the underhood package, or engine bay. The challenge: air flowing in through the front grill plays a critical role in cooling the engine, but the process creates air drag that increases the engine's workload.

Also influencing airflow are parameters like the size and location of the front-end openings, the position of heat exchangers, grill shutters, and speed flaps, and even fan power. Additionally, operating conditions such as external temperatures, speed, terrain, and towing also complicate the problem.



Instead of testing select parameters and operating conditions individually, Ford researchers wanted to simultaneously test multiple parameters with multiple operating conditions. This called for a computational fluid dynamics code (CFD) along with a high-performance computing (HPC) system powerful enough to handle that much data.

CFD simulates the flow of liquids and gases around solid surfaces. Ford's UH3D code simulates the flow of air within the engine compartment in order to determine the most aerodynamic model for an underhood engine bay design.

The project required thousands of simulations and an HPC system much more powerful than the one Ford had in-house.

### Power equals performance


Ford's research needs led Hupertz and senior HPC technical specialist and lead investigator Alex Akkerman to ORNL. Through the lab's HPC Industrial Partnerships Program, the team was awarded time on ORNL's Jaguar supercomputer, one of the most powerful in the world.

After scaling UH3D to run on Jaguar, Ford's team used approximately 1 million processor hours to test 11 geometric and nongeometric parameters against four different operating conditions. In total they ran 1,600 simulation cases. It was the first time Ford designers were able to collect this much data so quickly. Jaguar was so precise in fact, that researchers found less than a 1 percent deviation in the results between the computational and physical road tests.

"Being able to understand the interactions between the performance of the cooling system and the large number of design parameters in such great detail has provided us with a level of insight we have not been able to gain before," said Akkerman.

As a result, not only will the country benefit from the increased fuel efficiency, but Ford customers will be driving around with a little more cash in their pockets. —By *Jeremy Rumsey*

"Delivering a more energy-efficient vehicle while maintaining design leadership and high-package efficiency, and fulfilling future safety standards, is a very complex task requiring a new engineering approach,"  
—Burkhard Hupertz,  
Ford Motor Company



*Visualization of a vehicle cooling airflow that has been optimized using advanced Design of Experiment-based computer-aided engineering processes.*



**Gerald Meehl**

National Center for Atmospheric  
Research

PI: Warren Washington

National Center for Atmospheric  
Research

INCITE, 48 million hours

## Oceans Will Rise, but by How Much?

The thought of rising sea levels often conjures images of glaciers melting and ice sheets collapsing into the sea. However, these dramatic events account for less than half of the rising ocean levels we can expect in the coming years.

There is another cause, one that is just as important but far less dramatic: thermal expansion. As water warms, it expands, and when it warms at the scale of the world's oceans, the outcome is sea-level rise. Researchers led by Gerald Meehl of the National Center for Atmospheric Research used Jaguar to explore just how much sea level is likely to rise under various circumstances.

The scientists found that sea level will continue to rise substantially, even if the global community takes aggressive action to slow climate change. The ocean has a long memory, and climate changes to date ensure that sea level will rise, at least to some degree. The question, then, becomes "how much?" The answer depends greatly on the actions we take.

The group released its findings in the July 2012 edition of *Nature Climate Change*.

The researchers ran simulations for four greenhouse gas mitigation scenarios. In the most aggressive, more carbon dioxide would be removed from the atmosphere than added by the year 2070. The least aggressive



scenario would be a hands-off approach under which essentially nothing is done to curb carbon emissions. The team concluded that while aggressive measures could slow the rise in sea level, it could not halt it altogether.

Under the most aggressive approach, the global average temperature would cool, but thermal expansion combined with water from melting glaciers and land ice would nevertheless result in a 1½-foot rise in sea level by 2300. The least aggressive mitigation scenario led to a sea-level rise of over 8 feet by 2300. Such an outcome would be devastating, making it costly or impossible to protect major cities from flooding and creating severe environmental problems.

Using the four emission mitigation scenarios, the researchers ran climate simulations with the Community Climate System Model, which integrates simulations of atmosphere, ocean, land, and ice. To draw their conclusions, the scientists needed models to go out to 2300 and include five simulations for each of the four mitigation scenarios. All told, they simulated a total of 4,500 years, which would have been impossible without a powerful supercomputer like Jaguar. The simulation used 2,064 processors, which allowed the researchers to obtain 10 years worth of data in a day of simulations.


According to the *Nature Climate Change* paper, sea level will continue to rise regardless of the aggressiveness of mitigation, but it will rise much more if less stringent measures are taken. Not only that, but it will rise faster, leaving less time for coastal areas to adjust. The researchers hope that their findings will help in urging governments to adopt more aggressive greenhouse gas mitigation techniques.

“If we are able to mitigate and stabilize the climate change, it would buy us more time to actually adapt to this ongoing rise of sea level,” says Meehl.—*By Leah Moore*

For more: <https://www.olcf.ornl.gov/2013/01/24/sea-level-rise-will-continue-even-with-aggressive-emission-mitigation/>

“If we are able to mitigate and stabilize the climate change, it would buy us more time to actually adapt to this ongoing rise of sea level.”

—Gerald Meehl, National Center for Atmospheric Research



*Sea level will continue to rise into 2300 because of thermal expansion under the most aggressive mitigation scenario (cooling), but the rise will be slowed and give us more time to adapt. With less aggressive mitigation (stable) and (warming), there would be less time for adaptation.*





**PI: Allan Grosvenor**  
Ramgen Power Systems  
ALCC, 35 million hours

## Ramgen Simulates Shock Waves, Makes Shock Waves Across Energy Spectrum

One of the most pressing scientific challenges facing the United States and the world is reducing greenhouse gas emissions. Compounding that challenge is the fact that power plants burning fossil fuels account for more than 40 percent of the world's energy-related CO<sub>2</sub> emissions and will continue to dominate the supply of electricity until the middle of the century. There is an urgent need for cost-effective methods to capture and store their carbon emissions.

The Department of Energy (DOE) is currently sponsoring large-scale demonstration projects to prove the viability of carbon capture and sequestration (CCS). The principal barrier to widespread application of CCS is its cost. Once the CO<sub>2</sub> is captured, compressing it to the required 100 atmospheres represents approximately 33 percent of the total cost of CCS.

Ramgen Power Systems, a small, Seattle-based energy research and development (R&D) firm, is developing a novel gas compressor system based on shock-wave technology used in supersonic flight applications. This technology holds important promise for the turbomachinery industry of engines and compressors. Ramgen is a world leader in applying this shock-wave-based compression technique to gases, including CO<sub>2</sub>—a more challenging application than air because of CO<sub>2</sub>'s larger molecular weight.

The traditional process to design and optimize new turbomachinery, or machines that transfer energy between a rotor and a fluid, involves the testing of multiple physical prototypes, which is expensive and takes more time than DOE's demonstration schedule permits.

Thanks to Jaguar and Titan, the company has been making dramatic progress in designing its advanced turbomachinery for compressing CO<sub>2</sub> and for generating electricity with a breakthrough engine capable of using dilute methane for fuel, work that appeared in the Proceedings of the 4th European Conference for Aerospace Sciences.

Ramgen, its software vendor NUMECA International, and the OCLF have transformed the workflow of this turbomachinery design project in a way that exploits the strengths of both Jaguar and Titan, involving performance improvements in the simulation code and memory reductions per compute core to fully use all of Jaguar/Titan's nodes.

All of these improvements have enabled the use of intelligent optimization in which ensemble simulations of varying design parameters are combined into a single run on Jaguar/Titan, capable of using more than 240,000 cores. The analysis of these ensembles drives the Ramgen designs toward optimal configurations, leading to accelerated timelines for product development and deployment.

Accelerating this timeline is critical to the nation's effort to curb the amount of CO<sub>2</sub> emitted into the atmosphere. Ramgen's prototype compressor is projected to reduce the capital costs of CO<sub>2</sub> compression by 50 percent and to produce a minimum of 25 percent savings in operating costs. Applying these cost savings to a new 400-megawatt clean coal plant would result in capital cost savings of approximately \$22 million and an annual operating cost savings of approximately \$5 million.

However, the immense amount of data generated by the Ramgen project has created the need for improved analysis capabilities. The OLCF has stepped in with expertise in visualization in areas such as shock wave volume rendering, development of new boundary layer detection techniques, and statistical analysis.


"The use of Jaguar has cut the projected time from concept to a commercial product by at least 2 years and the cost by over \$2 million," said Ramgen's CEO and Director Doug Jewett.

With the help of Ramgen's dedicated engineers and the increasing importance and power of supercomputers such as Jaguar and Titan, these technologies can enter the marketplace sooner, proving that together industry and high-performance computing can change the world in profound ways.—By Gregory Scott Jones


For more: <https://www.olcf.ornl.gov/2012/08/14/ramgen-simulates-shock-waves-makes-shock-waves-across-energy-spectrum/>

**"The use of Jaguar has cut the projected time from concept to a commercial product by at least 2 years and the cost by over \$2 million."**

**—Doug Jewett, Ramgen Power Systems**



***Ramgen Power Systems is using Jaguar to simulate equipment that will achieve carbon sequestration at a significantly lower cost than that offered by conventional equipment. This image is the high-resolution result of a billion-cell two-body simulation showing the complex reflected structures colored by mach number. Visualization by Mike Matheson, ORNL***



**PI: Tommaso Roscilde**  
ENS de Lyon  
INCITE (2009), 1.3 million hours

## Jaguar Guides Demonstration of Novel Quantum State

Researchers combining the supercomputing muscle of ORNL's Jaguar with the experimental abilities of powerful research magnets have confirmed an exotic quantum state known as Bose glass.

In it, particles condense into separate regions within a material, with the particles in each region sharing the same wave function. This is the closest that particles in a quantum mechanical system can get to being in the same place at the same time.

The researchers, affiliated with institutions spread across three continents, published their findings in the September 2012 issue of the journal *Nature*.

In this collaboration, the calculations performed on Jaguar guided experimental researchers both in creating Bose glass and in confirming that their creation was indeed Bose glass. The computational work was led by Tommaso Roscilde of France's École Normale Supérieure de Lyon, who worked with colleagues Rong Yu of Rice University and Stephan Haas of the University of Southern California.

"If the theoretical and the experimental data match up for the macroscopic properties, one gains from the theoretical study the fundamental information on the microscopic physics at the basis of the observed behavior," Roscilde explained.



The researchers focused their exploration on a very cold magnet placed in a powerful magnetic field. The project is concerned not just with the magnet as a whole, but with the magnetic moment of each atom, which can be thought of as a small compass.

Under the influence of the external magnetic field, these atomic moments are partly aligned along the field. The projections of the atomic moments along the field then behave as if they were separate particles, called “quasiparticles.” It is these quasiparticles that condense into a single quantum state.

If the magnet is made of a single material, the quasiparticles gather in a condensate spread across the entire magnet. This is Bose-Einstein condensation. To get them to condense into smaller, separate areas within the magnet—i.e., Bose glass—the researchers carefully added impurities to the magnet in a process called “doping.”

Doping works because this is a quantum system and the quasiparticles behave as waves as well as particles. Just as a rock poking out of a pond scatters a wave in the water, an impurity in the magnetic material will scatter the waves of the quasiparticles.

As a result, separate condensates gather in areas that are relatively free of impurities. While particles will be in the same state as nearby particles, they will not be in the same state as the other islands of condensate.

Although Bose-Einstein condensates had been verified experimentally, this is the first experimental confirmation of Bose glass.

Roscilde and his computational colleagues used a Quantum Monte Carlo technique to predict the proper doping of the material, as well as the ideal temperatures and magnetic field for producing Bose glass. The simulations averaged around 1,000 quasiparticles. They also predicted the identifying features, or signatures, that would demonstrate that an experiment had produced a Bose glass.

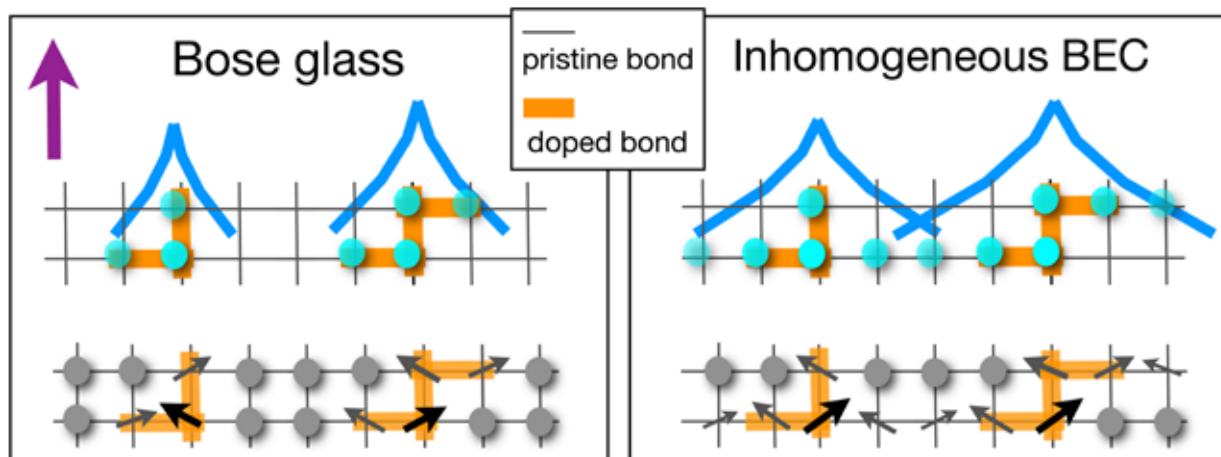
Moving forward, Roscilde and his teammates—both computational and experimental—will be exploring this system in more depth.

“We are planning to study the dynamics of the system,” he noted, “namely the nature of the elementary excitations in the Bose glass and the way that such elementary excitations would show up in a neutron scattering experiment.”—By *Leo Williams*

For more: <https://www.olcf.ornl.gov/2013/06/05/jaguar-guides-demonstration-of-novel-quantum-state/>

“If the theoretical and the experimental data match up for the macroscopic properties, one gains from the theoretical study the fundamental information on the microscopic physics at the basis of the observed behavior.” —Tommaso Roscilde, École Normale Supérieure de Lyon, France

*Below: Sketch of Bose glass and inhomogeneous Bose-Einstein condensate. The transition between these two phases is induced by an applied magnetic field (indicated by the violet arrow in the left panel). The upper lines show the phases in terms of quasiparticles (the cyan dots); the lower level in terms of spins (arrows corresponding to atomic magnetic moments). Grey dots symbolize spins that remain unaffected by the applied field. Illustrations courtesy of Tommaso Roscilde*





**Witek Nazarewicz**  
University of Tennessee,  
Knoxville  
PI: James Vary  
Iowa State University  
INCITE, 51 million hours

## Mapping the Nuclear Landscape

How many isotopes are there—theoretically?

We know we haven't seen them all. Many would not be found except in giant stars at the moment of their self-destruction. But how many combinations of protons and neutrons could there be in the nucleus of an atom?

A team from ORNL and the University of Tennessee has answered that question, using Jaguar to calculate the number of isotopes allowed by the laws of physics.

The team, led by Witek Nazarewicz, used a quantum approach known as density functional theory, applying it independently to six leading models of the nuclear interaction to determine that about 7,000 isotopes would be allowed in bound nuclei with up to 120 protons (a hypothetical element called “unbinilium”). Of that total, only about 3,000 have been seen in nature or produced in nuclear physics laboratories.

The team's results were presented in the June 28, 2012, issue of the journal *Nature*.

The computations on Jaguar allowed the team to identify the nuclear drip lines that mark the borders of nuclear existence. For each number of protons in a nucleus, there is a limit to how many neutrons are allowed, and vice versa.

The closer an isotope is to one of these drip lines the faster it decays into more stable forms. In fact, said Nazarewicz, all radioactive isotopes decay until they are transformed into one of 288 isotopes that form the so-called “valley of stability.”

Nazarewicz noted that results from the six separate models were surprisingly consistent.

“This is not a young field,” he said. “Over the years we’ve tried to improve the models of the nucleus to include more and more knowledge and insights. We are building a nuclear model based on the best theoretical input guided by the best experimental data.”

The calculations themselves were massive, with each set of nuclei taking about two hours to calculate on the 244,256-processor Jaguar system. Each of these runs included about 250,000 possible nuclear configurations.

“Such calculation would not be possible two to three years ago,” Nazarewicz said. “Jaguar has provided a unique opportunity for nuclear theory.”

Nazarewicz noted that this work has both existential value, helping us to better understand the evolution of the universe, and potential practical applications.


“We are not doing nuclear physics just to see whether you can get 7,000 species,” he explained. “There are various nuclei that we can use to our advantage, eventually. Those we call ‘designer nuclei.’”

Among these valuable nuclei are iron-45, a collection of 26 protons and 19 neutrons, which may help us understand superconductivity between protons; a pear-shaped radium-225, with 88 protons and 137 neutrons, which will help us understand why there is more matter than antimatter in the universe; and terbium-149, with 65 protons and 84 neutrons, which has shown an ability to attach to antibodies and irradiate cancer cells without affecting healthy cells.

“Applications will certainly follow from the basic knowledge,” Nazarewicz said. — *By Leo Williams*

For more: <https://www.olcf.ornl.gov/2012/06/27/ornlutk-team-maps-the-nuclear-landscape/>

“Such calculation would not be possible two to three years ago. Jaguar has provided a unique opportunity for nuclear theory.” — Witek Nazarewicz, University of Tennessee







## Titan

# Titan Is the World's Most Powerful Supercomputer for Science

The Titan supercomputer is a green powerhouse. In October 2012 it debuted at ORNL as the largest deployment of accelerated, energy-efficient computing the world has ever seen. Its creation marks a crucial step toward exascale computing, or machines capable of executing 1,000 quadrillion calculations each second. A month later Titan took the #1 spot on the TOP500 list, which ranks computers by performance, and #3 on the Green500 for energy-efficient performance (i.e., calculations executed per watt). It is now the research community's most powerful computational tool for exploring solutions to grand challenges in science and engineering.

The media coverage that Titan's debut and these rankings generated demonstrated that the public is excited about technological achievements and the science they enable. The stories gave top technologists platforms for educating the taxpayers, who support national labs, about the research that highly specialized, shared facilities make possible.

Titan delivered science even before its completion at the OLCF, which supports national science priorities through deployment and operation of advanced supercomputers. While the supercomputer was still under construction, scientists began using portions of the machine, and OLCF staffers began readying users with ramped-up training sessions. These pioneering users blazed trails for entire communities of scientists by running early simulation experiments on Titan, which uses energy-efficient GPUs as well as state-of-the-art CPUs. Titan's hybrid architecture allows application software codes to run at unprecedented speeds and thereby enables greater realism in simulations, giving OLCF science teams the ability to ask questions that could not be addressed before.

Prior to Titan's acceptance in the spring of 2013, its early users ran applications to explore magnetic properties of advanced materials; combustion of commercially important fuels; life-cycle extension of aging nuclear reactors; adaptation and mitigation in the face of climate change; molecular behavior of complex films for solar cells and liquid crystals for sensors; and radiation transport, a key process in astrophysics, medical imaging, and nuclear energy. Researchers expect that Titan will accelerate transformational discoveries and innovations in diverse fields.

## Super among supercomputers

In Salt Lake City in November 2012 at the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC12), Titan topped a roster of the world's fastest supercomputers by running Linpack, a benchmark application, at a speed of 17.59 quadrillion calculations per second (petaflops). The accomplishment demonstrated the US commitment to advancing high-performance computing and built on the OLCF's reputation for deploying powerful supercomputers (its Intel Paragon XP/S 150 was the fastest machine in the world in 1995, as was Titan's predecessor, Jaguar, from November 2009 to November 2010). Titan's theoretical peak performance is 27 petaflops.

Titan was created by upgrading Jaguar, a Cray XT5 system that occupied 200 cabinets and filled a space the size of a basketball court. That upgrade was Plan B. Plan A had been to have two separate machines: Jaguar would live on, and users would be able to run simulations on it while a new machine, Titan, was being built. Plan A would have required a lot of floor space. While the OLCF had the space, using it would have required installing an additional 10 megawatts of power and expanded cooling infrastructure to enable Jaguar and Titan to run at the same time as well as buying new cabinets, interconnects, and other equipment needed to support two systems. "When we priced all that out, it was going to cost about \$25 million more to be able to have Titan be a separate machine than if we did an upgrade of the existing Jaguar machine," said Titan project manager Buddy Bland. "In a period of limited budgets, that meant that Titan would be a substantially less capable system for our science teams than we were able to achieve by doing an upgrade."

In consultation with program managers at DOE, the OLCF managers opted for Plan B. The supercomputing center experienced substantial downtime as Jaguar's boards were removed, replaced, tested, repaired to remediate any problems, retested, and released to users. "Because our leadership computing mission compels us to field powerful supercomputers for the scientific community, we decided to take Jaguar out of production and in essence pay its resources forward to create Titan—an even more powerful system for scientists," Bland said.

Recycling Jaguar's components provided ample benefits. "That saved a basketball-court-sized volume of equipment from going to a landfill," Bland said. "It also saved us probably close to 10 percent on the cost of the machine, to

replace all those cabinets, power supply, and cables. And from the standpoint of getting the machine running more quickly, installing all of the infrastructure and making sure that all of it is working well is a major effort in one of these big system installations. The fact that all of it was already there—installed, debugged, and working—meant that we didn't have to do all that installation and testing."



*Jaguar Cray XT5*

The upgrade turned all-CPU Jaguar into a hybrid system of GPUs and CPUs. Each of Titan's 18,688 nodes sports a 16-core AMD Opteron 6274 processor and an NVIDIA Tesla K20X GPU accelerator. The resulting XK7 system has 710 terabytes of memory. Titan is 10 times as fast as Jaguar was but consumes only 20 percent more electricity.

Because GPUs handle hundreds of calculations simultaneously, they can carry out more calculations than can CPUs in a given period. Titan's 299,008 CPU cores guide simulations while allowing its 18,688 GPUs to do the heavy lifting. As a result, scientific calculations run with unmatched speed and fidelity.

## Countdown to Titan

Large supercomputing deployments are years in the making. The OLCF has successfully deployed supercomputers through projects named OLCF-1 (which produced the Phoenix Cray X1E machine) and OLCF-2 (Jaguar XT3, XT4, and XT5 systems). OLCF-3 recently delivered Titan.

Such deployments require many critical decisions (CDs), or approval steps. The signing of a mission need statement (CD0) by Ray Orbach, director of DOE's Office of Science, in January 2009, set the OLCF-3 project in



motion. It acknowledged that Office of Science needed tens of petaflops of computing capability to fulfill its mission to foster, formulate, and support forefront basic and applied research programs that advance the science and technology foundations necessary to accomplish DOE missions: efficiency in energy use, diverse and reliable energy sources, improved health and environmental quality, and fundamental understanding of matter and energy. CD0 stipulated that the leadership computing facilities at Argonne and Oak Ridge should provide a combined computing capability between 20 and 40 petaflops in 2013.

After a review in July 2009 to define how the project would proceed, DOE approved CD1, which provided a strategy for acquiring the requisite 10 to 20 petaflops of computing capability at the OLCF and an analysis of alternatives that could satisfy that need. The review raised concerns about plans to use a system employing application accelerators. How would software developers program a machine featuring both CPUs and GPUs? In December 2009, OLCF managers selected a group of applications representative of the facility's workload. The Advanced Scientific Computing Research program manager at DOE's Office of Science approved the applications. OLCF staffers spent the next several months developing strategies about how to effectively accelerate applications on a GPU-enabled system.

In July 2010, an external panel of experts reviewed the selected applications for their readiness to run on a hybrid, accelerated architecture. It ascertained that the proposed supercomputer could indeed meet the needs of the Innovative & Novel Computational Impact on Theory and Experiment (INCITE) program, the major means through which the scientific community gains access to leadership-class supercomputing. "The reviewers were impressed with our designs, our analysis of the applications and how well they would run on the GPUs, and our projections for what performance we would get," Bland said. With that sign-off the OLCF managers proceeded to the next stage, a CD3a review in December 2010 that approved issuance of a request for proposals. DOE approved other CDs in August 2011—CD2 to set the project scope, schedule, and budget and CD3b to approve the machine's acquisition.


ORNL issued the contract to upgrade Jaguar in autumn 2011. Goals included doubling memory and installing the XK7 node boards with 16-core CPU processors and the Gemini network for interconnecting processors. In a 10-cabinet partition called TitanDev, a system for application-development work, Cray also installed 960 NVIDIA Fermi GPUs. Users began learning how to work with GPUs. In September 2012, the OLCF received even faster NVIDIA Kepler GPUs and installed them in all 18,688 nodes of the machine. After rigorous testing and acceptance of Titan, in spring 2013 came CD4, the decision to approve the start of operations. After four years of preparation, Titan was officially open for business.

"Our users have really embraced the GPUs," Bland emphasized. "One of our biggest concerns was that users would say, 'This is too hard,' or, 'We don't want to use the GPUs,' but that has turned out to not be the case at all.






**S3D**




**Denovo**




**CAM-SE**



**NRDF**



**WL-LSMS**



**LAMMPS**

*The six applications chosen by the OLCF demonstrate the potential of Titan's GPUs.*

We have offered a lot of training on how to use GPUs, which has been very well attended. In addition the GPUs are available in many laptop personal computers, and university clusters are starting to see GPUs. So many of our users had started getting experience or had started porting their codes to GPUs even prior to Titan's arrival. We're seeing excitement about the availability of a very large machine with GPUs."

Bland stressed that CPU-only architectures are not practical or affordable as machines further scale up in number of processors. It would take a nuclear plant to power an all-CPU exascale supercomputer built using today's processors. "Titan is America's first openly available hybrid multicore system, which we think is an important direction in getting ready for exascale computing," he said. Titan's energy-efficient GPUs in concert with its CPUs will enable scientific leadership by providing unprecedented computing power for energy, climate change, materials, and other research fields. "Titan is the first of these big machines that is generally available to the scientific community, so users are starting to convert their codes to run on this type of architecture. Our science teams are already using Titan to do great science, and that will continue throughout its lifetime. But its legacy will be all of the applications that people have converted to run on this hybrid, multicore environment. They will be the first wave of applications to run on the next generations of energy-efficient, extreme-scale, parallel supercomputers needed to support the science demands of the nation over the next decade."

## Early science

Since the first day the OLCF made GPUs available, early users with codes capable of employing GPUs at scale have made substantial use of the accelerators. The demand for GPUs continues to rise.

"Early science refers to that vanguard set of apps and problems that we get going on the machine as quickly as we possibly can—codes capable of exploiting Titan and problems that require that much computational power," explained Center for Accelerated Application Readiness (CAAR) lead Bronson Messer, a computational astrophysicist at ORNL and member of the Scientific Computing Group at the National Center for Computational Sciences (NCCS). "We want them to have access to as much of the machine as possible, as early as possible, both to let them achieve those scientific goals and to take the machine on a shakedown cruise."

At the Accelerating Computational Science Symposium (ACSS) March 28–30, 2012, in Washington, DC, top experts in science, engineering, and computing gathered from around the world to discuss research advances made possible with extreme-scale hybrid supercomputers. The OLCF, National Center for Supercomputing Applications, and Swiss National Supercomputing Center cohosted the conference. Attendees compared the performance of scientific applications on computing systems that had two CPUs per compute node (i.e., a Cray XE6 architecture) versus one CPU and one GPU per compute node (i.e.,



***The accelerator-based compute blade used in Titan.***

a Cray XK6 architecture) to ascertain which architecture was a better investment for science. The comparisons allowed them to find out if buying a GPU instead of another CPU to populate the second socket of the node conferred an advantage. The results indicated that hybrid supercomputers could speed discoveries, such as deeper understanding of phenomena from earthquakes to supernovas, and innovations, such as next-generation catalysts, materials, energy-efficient engines, and nuclear reactors. Presenters, including many OLCF users, shared advances enabled by hybrid supercomputers in fields including chemistry, combustion, biology, and physics. Vanguard researchers using TitanDev, a Cray XK6, presented code speedups of 1.5- to 3-fold when comparing the performance of its one Interlagos CPU and one Fermi GPU per compute node to the performance of an XE6 machine having two Interlagos CPUs per node.

Among the ACSS pioneers were OLCF experts working with application teams at ORNL's CAAR. Born with the OLCF-3 project, CAAR was established to develop the techniques needed to expose more parallelism in applications so that the GPUs could be effectively used on Titan. The result would be a handful of codes that from the outset could take advantage of Titan's new hybrid architecture by exploiting hierarchical parallelism. "There had never been a bigger hybrid GPU machine built, so we had added to the usual problems that we have getting people to scale up to many, many processors. And, 'Oh, and by the way, you have to be able to marshal the GPU as well,'" Messer said.

The codes selected for development at CAAR were S3D (for direct numerical simulation of turbulent combustion), CAM-SE (high-resolution climate-scale simulation of the atmosphere over decades to centuries), Denovo (neutron transport in nuclear reactors), 2009 Gordon Bell winner WL-LSMS (statistical mechanics of magnetic systems), and LAMMPS (molecular dynamics). The researchers also investigated NRDF (radiation transport and advanced algorithms), a proof-of-concept code for adaptive mesh refinement. Speedups for CAAR codes on the XK6 versus XE6 ranged from 1.4- to 3.3-fold. Other key, non-CAAR, community codes saw 1.5- to 6-fold speedups.

Using Kepler GPUs instead of Fermis when the former were installed to create Titan's Cray XK7 architecture, the CAAR researchers obtained even bigger speedups. "Kepler is a significantly different and more capable GPU architecture relative to Fermi," Messer said. "It's faster and has more cores ... but the real difference between those two GPUs is the programming model. Keplers can execute multiple independent tasks simultaneously, and the Fermis couldn't. From a user's perspective, the Keplers are even more parallel than the Fermis ever were."

Researchers achieved science results from the first days of access to Titan's Kepler GPUs. Mike Brown, a member of NCCS's Scientific Computing Group, obtained LAMMPS speedups of up to 7.4 (with performance depending on the number of atoms each node calculates and the employment of lower [mixed] precision on GPUs and higher [double] precision on CPUs).

One LAMMPS project simulated liquid crystals that can be used as biomedical sensors to detect bacteria, antibodies, or other signs of illness. Liquid crystals—which exist in orderly phases akin to solid crystals—are common in nature (e.g., the proteins embedded in cell membranes) and technology (e.g., electronic displays). Optical microscopy can detect a biomolecule in a liquid crystal layer without tagging of the biomolecule. Simulations show how thin films of liquid crystal molecules "dewet," or rupture and then self-assemble into droplets, patterns, or complex structures on top of a substrate. "We have shown on Titan that our model is capable of reproducing dewetting patterns observed in experiments," Brown said. "Simulations may aid the design of liquid crystal devices with desired properties."

Another early science application that benefited from Kepler GPUs uses the Wang-Landau locally self-consistent multiple scattering (WL-LSMS) code to calculate the Curie point, the temperature above which a material loses its magnetism. The results illuminate the behavior of nanoscale systems and provide understanding needed to advance energy-efficient industrial materials, electric motors, and generators. ORNL's Markus Eisenbach led development of the code to run on GPUs. The speedup of WL-LSMS with Kepler GPUs was 3.8 as compared to the Cray XE6. With early access to Titan at the end of 2012, Eisenbach ran the code at above 10 petaflops—about five times as fast as the world-record-holding 2011 Gordon Bell winner—to calculate magnetization and free energy for a 250-atom iron "supercell" using 10 million Titan core-hours. He will further

the studies with an award of 105 million core-hours on Titan in 2013 through the INCITE program. This initial free-energy calculation was not executed using all of Titan's nodes. Eisenbach more recently ran WL-LSMS on the entire Titan machine, simulating an even larger supercell of 1,024 iron atoms at a speed of 14.5 petaflops.

"Our early science projects will allow all of the work that we did as part of CAAR and the work that others have done outside of CAAR to produce the first wave of GPU-accelerated science results," said Messer. "We expect that every one of the early science results will be important. We also think that every one of those results will lead to yet another result, and so on. An important goal of CAAR was to expose parallelism in a way that is portable to many architectures, not only GPUs. We expect that many future machines will be able to take advantage of these code architectures that we've set up."

## Lessons learned

The lessons learned in helping half a dozen codes run efficiently on Titan were both technological and sociological. "We learned that 80 percent or more of the work that is required to take a code that ran well on Jaguar and make it run well on Titan isn't caught up in writing new code to engage the GPUs," Messer said. "Most of the work is dedicated to restructuring all the rest of the code so that the data is ready to be consumed by the GPU."

That lesson had two corollaries, he added. One was that no matter the code, porting it to Titan's hybrid architecture took two person-years. "That could either be one person over two years or two people over one year or four people over six months, whatever, but it takes two people-years, basically, to get it done," Messer said. "Importantly, the time to do the conversion now is likely less, as the tools and compilers available have become much better relative to when we started CAAR."

The other corollary was that if users restructured their code to run well on Titan's GPUs but then turned off all the GPUs and just ran the code on CPUs, the code would run twice as fast. "One might think it should run slower because the individual CPU cores on Titan have a lower clock speed than Jaguar. There are more cores, but they actually run slower," Messer said. "The reason the code's faster is the changes you had to make for the GPU also help getting operands out of memory. They make the CPU code more efficient. It does a better job of cache reuse. It's just better CPU code. Programming for the GPUs forced you to do things that you probably should have done to your code a long time ago to make it run faster."

Of the sociological lessons learned, community code developers need to keep abreast of changes others make to the code. "Just because we're working on a GPU port of S3D doesn't mean that the rest of the S3D development team isn't doing stuff to that code. The hallmark of most codes at the OLCF is that they're under constant development," Messer said. "The ground is constantly moving under your feet, and a change you make today might be a complete non sequitur in the code tomorrow."

We had to ensure that we synched with the main line of code development constantly."

To make the most of GPUs, developers have to uncover a deeper level of parallelism in codes so more processors can be put to work on solving problems. By way of analogy, Messer explained different strategies at play when Jaguar got new CPUs versus when it received GPUs to turn it into Titan. "Let's say I have a field that I have to excavate down to 6 feet," he said. The Jaguar strategy increases the number of parallel processors: "It's like giving 16 people each a shovel and giving them some portion of that field to dig down to 6 feet. The limiting speed is how fast the slowest guy can shovel." The TitanDev strategy with Fermi GPUs is analogous to replacing each of those workers with multiple-armed mutants that can hold more shovels, Messer said. The Titan strategy with Kepler GPUs, however, is like giving workers motorized buckets instead of shovels—digging and dumping of dirt is quick and continuous. Domains of a simulation are decomposed into threads of calculations that are fed into the GPUs. Increasing the number of threads accelerates the execution of concurrent tasks. "Now the limiting step is how much time they lose between picking up the dirt and putting it back down." That is, the limiting step is data caching and data movement.

"We expect that every one of the early science results will be important. We also think that every one of those results will lead to yet another result, and so on." —Bronson Messer, ORNL

Amdahl's Law, famous in parallel computing, says that any computational task has some work that is serial. A processor has to wait for the result of the last step before it can start the next step. A lot of code restructuring aims to remove serial bottlenecks. "The GPUs run so unimaginably faster than your CPU that if you can keep them busy, you always win," Messer said.

The critical task of GPUs is to increase node-level performance, Messer said. "There's no computational scientist who's ever lived who, if given the choice between running a given problem and it taking a day or an hour, would choose a day. In that sense every code could use some acceleration." Five years from now it won't be realistic for users of leadership-computing platforms to run CPU-only code. "The accelerator may take the form of an NVIDIA GPU or an Intel Xeon Phi or something that hasn't been invented yet," Messer said, "but they'll be using an accelerator." —By Dawn Levy





## Titan

# Top Rankings Give Computing Experts Chance To Share Titan's Science Impact

When Titan took the TOP500 list's #1 spot among the world's most powerful supercomputers in November 2012, the win gave ORNL's top techies the opportunity to talk with the global media about research projects the HPC system would make possible and the importance of these projects to fulfilling the DOE's science and energy missions. The ensuing surfeit of stories—nearly 1,000 were published and broadcast—proved that the popular press is excited about science and technology and generated visibility and anticipation among the wider public about future Titan-enabled science discoveries and engineering innovations. Moreover, Titan's subsequent ranking as #3 on the Green500 list, which reorders the supercomputers on the Top500 list according to how many calculations they can execute per watt of electricity, created the chance to discuss how Titan's energy-efficient hybrid architecture of GPUs as well as CPUs reduced long-term operating costs and resulted in a more capable machine than its all-CPU predecessor.

"We were trying to tell people that Titan would be a unique national resource for science and that from a technology point of view it was a solution to a new set of problems that supercomputing facilities have to address: energy consumption," said Jack Wells, director of science at the OLCF. "With Titan, in the same footprint and with a marginal 10 or 20 percent increase in peak electrical power consumed, we enabled a factor-10 increase in compute capability."

While those messages can be conveyed without winning international prizes, awards garner attention. "Our mission requirements as a leadership computing facility compel us to field the most powerful computer we can and it still be a production supercomputer—general-purpose, production, day-to-day workhorse—but at the limits of what's possible to build," Wells said. "The TOP500 and Green500 rankings, competitive metrics for

overall performance or efficiency of performance, give us measures that we are satisfying our mission in that we are at the limits of what's possible to build.”

To make the public aware of Titan's technological achievements and the science they enable, ORNL with industry partners NVIDIA and Cray began a media launch well in advance of Titan's unveiling. This global introduction of Titan in October and November 2012 included press releases; videos; journalist interviews with researchers and computing experts; journalist site visits to ORNL; and visits by ORNL, Cray, and NVIDIA executives to journalists in their offices. Reporters from AnandTech and the BBC came to ORNL on October 9 for tapings and interviews with ORNL Associate Laboratory Director for Computing and Computational Sciences Jeff Nichols, Wells, and OLCF Project Director Buddy Bland. Between October 10 and 25, the executives engaged in back-to-back interviews that brought the story to the journalists at their places of business, including the offices at the *New York Times*, the *Washington Post*, and *National Geographic*. Nichols met with reporters in New York whereas Wells handled interviews in San Francisco and Washington. Steve Scott, chief technology officer of the Tesla business unit at NVIDIA, did all of the interviews that Nichols and Wells did, plus others on his own. From Cray, Barry Bolding, vice president of storage and data management and corporate marketing, participated in the San Francisco interviews, whereas Bill Blake, Cray's senior vice president and chief technology officer, joined conversations in New York. Sumit Gupta, general manager of the Tesla Accelerated Computing Business at NVIDIA, also did some interviews. ORNL, NVIDIA, and Cray executives participated in phone interviews with, among others, *The Economist*, *Bio-IT World*, *Fortune*, NPR, *New Scientist*, *LeFigaro*, *Scientific Computing*, *Drug Discovery and Development*, and *Venture Beat*. Similarly, Nichols's phone interviews included *Popular Science*, while Wells's included the *Wall Street Journal*.

The media launch was wildly successful in getting the word out to the business, financial, science, technology, HPC, and consumer press, who in turn produced numerous online, print, and broadcast stories. Besides appearing in such HPC stalwarts as *HPCwire* and *insideHPC*, Titan stories also turned up in hundreds of media outlets worldwide, including *All Things Digital*, *Bio-IT World*, *CBS News*, *CNET*, *CNN*, *eWeek*, *Forbes*, *Fortune*, *IDG*, *LiveScience*, *Financial Times*, *NBC*, *Reuters*, *Scientific American*, *Slashgear*, *Time*, *VentureBeat*, *WIRED*, *WirtschaftsWoche*, *Yahoo*, and *ZDNet*.

The stories generated visibility for ORNL supercomputing resources and scientific leadership and highlighted US leadership in technological innovation and scientific discovery. They educated business, science, and consumer audiences about the benefits and possibilities of supercomputer-enabled research.

From industrial competitiveness to the importance of specialized user facilities available through national labs, the story angles were as varied as the reporters who wrote them. Headlines focused on speed (“US Lab's

‘Titan’ Named World's Fastest Supercomputer”—*National Geographic*) or science (“Interview: Buddy Bland on How the #1 Titan Supercomputer Will Power Science”—*insideHPC* video) or both (“The Fastest Science Machine in the World”—*Popular Science*). The technology was described in terms from geeky (“Inside the Titan Supercomputer: 299K AMD x86 Cores and 18.6K NVIDIA GPUs”—*AnandTech* photo gallery) to gushing (“Titan Crowned World's Most Awesome Supercomputer”—*WIRED*) to grandiose (“Oak Ridge Lab: Behold, I Am TITAN, Hear My 20 Petaflop ROAR—One Giant Leap for a GPU, One Small Step for Exascale”—*The Register*).

Some stories concentrated on competition among technologists (“Titan Knocks Off Sequoia as Top Supercomputer”—*HPCwire*) and nations (“The 5 Fastest Supercomputers”—*CNN Money*, in an article accusing America of ceding supercomputing ground to China, Japan, and Germany in recent years). Others emphasized the role of HPC in helping the US economy remain globally competitive by bringing great minds to America (“Supercomputers Act Like Talent Magnets”—NPR). Further, some articles explored the benefits of such collaborations. A *National Geographic* article described a partnership between researchers from Procter & Gamble and Temple University to use Titan to develop the first molecular-based model for understanding how lotions or drugs are delivered through skin.

Articles described the GPU heart of accelerated computing in terms consumers could quickly grasp (“‘Giant Playstation’ Titan Supercomputer to Solve Biggest Riddles”—*The Australian*). They explored “green” angles such as Jaguar's recycling to create Titan (“Cray's Jaguar Supercomputer Upgraded with NVIDIA Tesla GPUs, Renamed Titan”—*Engadget*). Headlines were clear (“US Supercomputer at Tenn. Lab Is World's Fastest”—*USA Today*) and clever (“Out to Crunch: US Energy Department Unleashes Its Titan Supercomputer”—*Scientific American*).

“Titan's #1 TOP500 and #3 Green500 rankings gave us an opportunity to talk about the science that could be performed on Titan that could not have been performed on previous generations of leadership computing, for example, Jaguar,” Wells said. “In combustion science, on Titan we can study real liquid transportation fuels rather than proxies for transportation fuels that are studied in the laboratory. In nuclear reactor simulation, 3D physical reactors are now within our reach with Titan. We can put that simulation capability in the hands of an engineer at Westinghouse or Tennessee Valley Authority. In the past, these engineers have needed to be satisfied with approximations.”

Added Wells: “In part we had a platform for telling the science story because of the interest in the competition around being #1 on the TOP500 list. It provides a focus in terms of not only talking about broader scientific goals in terms of discovery or innovation, but we were also able to talk about what science can Titan do, what science is sized for Titan. It built anticipation about what is going to be done in the next few years.”—By Dawn Levy



## Titan

# OLCF Experts Prepare Users for Titan

In 2012 the OLCF began offering a smorgasbord of training to prepare approximately 1,100 supercomputer users to achieve science from day one of accessing Titan, an HPC system at ORNL that harnesses brawny, energy-efficient application accelerators called GPUs, as well as brainy but traditional number-crunching CPUs. Workshops, seminars, and tutorials sated the intellectual appetites of current supercomputer users, the thirst for knowledge of a new generation of users, and the hunger for GPU specifics of code developers preparing applications to run on the world's fastest supercomputer.

Late in 2012 the final configuration of Titan's GPU/CPU architecture became known, and OpenACC, a directives-based interface that simplifies programming of accelerators, was announced as a standard. By then CUDA, an extension of the C programming language, had become entrenched as the means of facilitating the programming of GPUs. Before March 11, 2013, only those pioneering the development of a select group of codes had been allowed to run on Titan. But on that day researchers from a wide spectrum of scientific disciplines were granted access to the world's largest deployment of GPU computing. Having been trained for success, all reaped immediate benefits from recent upgrades to CPUs. But in the race toward scientific discoveries, engineering innovations, and game-changing industrial improvements, those able to harness the GPUs that give Titan its theoretical peak performance of 27 quadrillion calculations per second, or petaflops, may own the future.

Users hit the ground running thanks to the dramatic ramp-up of training in 2012. An OLCF survey conducted in 2011 had clarified the needs of users uncertain about if or how they'd adjust to running codes on the new GPU/CPU hybrid architecture. "[Users] were very



clear that they wanted more OpenACC, more GPU programming, so they wanted CUDA and they wanted OpenACC talks,” said OLCF User Assistance and Outreach Group Leader Ashley Barker. Since January 2012 the OLCF has delivered training to more than 300 current and prospective users as well as users of other supercomputing centers that do not have their own training programs.

## Training days

OLCF staffers literally went the extra mile to bring training to users. After delivering a basic overview class in fall 2012 at Oak Ridge, in January of 2013, they traveled to Santa Clara, California, to deliver a three-day workshop for Titan users and application developers on the West Coast. The following month in Knoxville, about half an hour from ORNL, they repeated the class.

New users received guidance about how, for example, to submit a simulation to run on Titan or add collaborators to a project. Experienced users benefited from “best practices” for, say, transferring data in and out of the OLCF. Regardless of the programming paradigm used (e.g., CUDA or OpenACC), developers learned a key lesson: “You have to rearrange your code to take advantage of the calculations,” said HPC user assistance specialist Fernanda Foertter, who since March 2012 has determined the OLCF’s training needs, goals, curriculum, delivery method, and schedule. “It’s good for the science.”

The Titan workshops, which covered topics from parallelization to debugging, were free and open to the public. “Access to supercomputers is very limited,” Foertter said. “The way most of us come to supercomputing or HPC in general is we get a project that needs to use it.” The OLCF training is an important way to increase participation in the field, she added. Recent training has attracted “hopefuls,” she said—people who want to run codes on Titan but are afraid to apply because they don’t think they will get an allocation. Foertter recalled running into someone at a conference who had participated in OLCF training. “[That person] said, ‘Thank you so much for the training because it made me see that it was doable, that I could do this [simulation experiment] with this code, and I’m getting the code ready so I can apply [for time on Titan].’ That’s a success story.”

Both West Coast and East Coast classes were webcast, increasing participation by at least 30 percent, according to Barker. “We knew our users have a hard time traveling, especially with the new budget restrictions right now, so making sure that our webcasting was of good quality was very important in opening up training to more of our users,” she said.

Users can access online training as well. Recent improvements to the OLCF website enable searches of a single reference guide rather than separate articles, provide tutorials about GPU programming, and consolidate OLCF policies in one location. A new section provides articles about each software package in use at the OLCF. Publishing of these articles is integrated with the tools the center uses to maintain the software, ensuring that software changes are automatically reflected on the website.

Abundant and varied, 2012 and early 2013 training offerings included workshops about Titan’s new architecture, software performance analysis tools, programming languages (e.g., R and CUDA), message passing interfaces (e.g., MPI 3.0), and



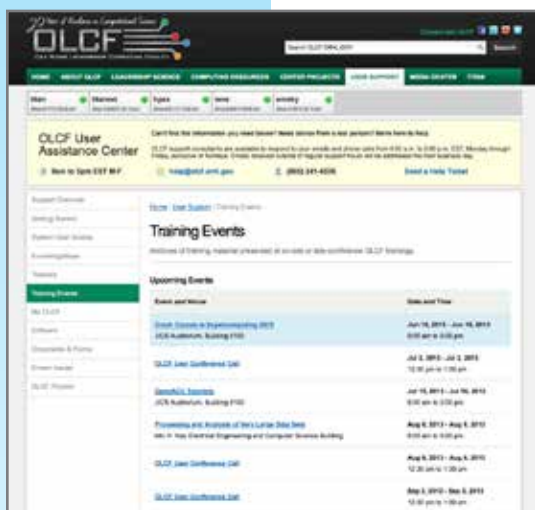
**Ashley Barker, User Assistance and Outreach group leader.**

applying for supercomputing time. A workshop cohosted by the OLCF and the University of Tennessee's Joint Institute for Computational Sciences covered innovative accelerator hardware and multicore chips that enable unprecedented simulations of the atomic structure and electronic properties of materials. Internships also introduced a new generation to HPC as a career. For example, 10 Appalachian students and teachers gathered at the OLCF in summer 2012 to write, compile, and execute a parallel program on a Beowulf cluster they had built out of Mac minis. Further bolstering the user community were a crash course in supercomputing; a series of workshops about HPC fundamentals; and meetings and workshops for new users, existing users, and developers.

Some training took the form of outreach. Cohosted by the OLCF, the National Center for Supercomputing Applications, and the Swiss National Supercomputing Center, the Accelerating Computational Science Symposium 2012 in Washington, DC, gave pioneering researchers the opportunity to share early simulation results obtained using accelerated codes for combustion, climate, nuclear reactors, magnetic systems, and biophysical systems explored at ORNL's Center for Accelerated Application Readiness (CAAR) as well as non-CAAR codes for chemistry, seismology, medicine, and more. Moreover, OLCF Director of Science Jack Wells and computational astrophysicist Bronson Messer spoke to students about HPC at the DOE Computational Science Graduate Fellowship's annual conference, and OLCF nuclear physicist Hai Ah Nam and Foertter participated in a session on women in computing at SC12, an international supercomputing meeting.

Barker is most proud of the fact that OLCF training allows the center's specialists to quickly share lessons learned through two years of assessing the performance of CAAR applications. "One of the good things about the CAAR effort is the application developers all went about programming GPUs a different way," she said. "It showed that you could come at the problem in different ways. They all encountered some of the same issues that they all had to work through, but maybe they worked through them using different tools. Having those different perspectives was really important because what they showed is there's not just one magic bullet. It's not just CUDA or Open ACC or some of the earlier things they were trying out. They were able to tell people, 'Here's what you really want to look at before you get started.' That will help people down the line by making them really stop and think."

Added Barker, "The CAAR talks showed people some of the lessons learned when porting [application codes] to accelerators. Regardless of the path used to accelerate an application, be it CUDA or Open ACC or CAPS directives, the codes require work in exposing parallelism and smarter data access. The downside is most codes need significant restructuring, but the payoff is that these adapted codes are more portable to new architectures, even if they still require tuning to the specific hardware."



*The OLCF website informs users of upcoming training opportunities.*




*Cray's John Levesque speaks to the OLCF's Spring Training and Users' Meeting.*

### Unusual homework

The workshops offered a peerless hands-on lab. Instead of merely listening to a speaker, trainees followed along and put together a code. Participants at the January workshop then ran their codes on Chester, a machine of just one cabinet. But so many users signed up for the February workshop that the hands-on exercises required a bigger machine. Foertter asked the Resource Utilization Council (RUC)—the body governing use of OLCF supercomputers, clusters, and storage systems—for time on Titan. The RUC granted Foertter 100 of Titan's 18,688 nodes for the training. The trainees put into practice what they learned in the classroom by doing “homework” on the world's most powerful supercomputer. “New users were very excited about that,” she said.

Other highlights were training about Open ACC and CUDA delivered by vendors. Users appreciated the insights of professionals from NVIDIA, Cray, PGI, and CAPS, as GPUs add a level of complexity; programmers must code for two pieces of hardware (i.e., both GPUs and CPUs). Said Foertter, “That may also be difficult for the user to think about, but the payoff is *big* and it's worth it, and by payoff I mean doing science in half the time or in depth, doing higher-resolution science problems, or new science that we didn't think before to do because we just didn't have the ability to calculate at 27 petaflops.” —By Dawn Levy





## Users

### Putting the “User” in User Facility

OLCF users represent a wide spectrum of research interests, from industry to academia to the US government.

And those interests represent an even wider spectrum of science, from more aerodynamic tractor trailers to climate change to novel materials to organic solar power, to name a few.


While former OLCF flagship Jaguar and its current incarnation Titan are impressive machines in their own right, it's the marriage of Titan's hardware and perfectly tailored applications that produce revolutionary scientific breakthroughs while keeping America competitive.

For this reason, OLCF users occupy a high priority at the nation's leading computing facility. Be it General Electric, Ford, liquid crystals, or the explosion of core-collapse supernovas, science is front and center, and that means taking care of the user community.


In 2012 alone Jaguar hosted 157 projects consuming a total of nearly 1.3 billion processor hours. And the demand is only growing. The number of INCITE proposals submitted for 2013 surpassed the previous year by 20 percent, much of which reflects completely new areas of science that have never before sought out a computer the size of Titan. This, in turn, reflects the increasing importance of HPC in R&D. No matter how a project arrives at the OLCF, once researchers sign on there is no time to waste. In scientific computing, time is of the essence, and the faster user teams can get up and running the better.

## OLCF Stats and Numbers


**Allocation by Domain 2012**




**System Usage by Domain 2012**




**User by Affiliation 2012**



**2012 HelpDesk Tickets Issued**



**2012 HelpDesk Tickets by Category**



The OLCF employs a multi-pronged approach to meeting its users' needs. Be it improving their codes to take advantage of Titan's GPUs, regular training and seminars to keep users up-to-date on hardware and software changes, or centers of excellence manned by OLCF vendor partners on-site to keep Titan at the height of its game, the OLCF ensures users they have the tools they need to accomplish the breakthroughs they're seeking.

With Titan's hybrid architecture now up and running, users will lean more heavily on OLCF resources than ever before. In porting or scaling their codes, programming to take advantage of the GPUs, or cutting-edge data analysis and visualization, the OLCF has never been better prepared to meet the needs of its users head-on.



## Communication is Key

**Balint Joo, Jefferson Lab**

With users from around the world spread out across multiple scientific domains, getting a handle on their individual problems and concerns is no small task.

Like the OLCF's user community itself, the issues and recommendations that arise from the hundreds of researchers on Jaguar and Titan vary greatly, yet they each need to be heard and addressed.

To facilitate communication between the OLCF and the researchers it hosts, the center formed a User Council (UC) soon after its inception, an institution that continues to this day. With Titan's revolutionary hybrid architecture coming online, the UC is certain to be more important than ever.

The UC provides a forum for the exchange of ideas and development of recommendations to the OLCF regarding the center's current and future operation and usage policies and is made up of researchers who have active accounts on OLCF compute resources. Essentially, the UC serves as a liaison between the OLCF and its user community, operating independently to provide feedback and advice to the OLCF on the community's needs and concerns.

A Steering Committee with five members leads the UC. This committee consists of a chairperson (two-year term), a chairperson pro tem (two-year term), an ad hoc member of the OLCF Resource Utilization Council (one-year term), and up to two members at large (one-year terms). The chair serves as a liaison between the group and OLCF management, a role that includes addressing feedback and comments from users to raise concerns to a higher level within the OLCF organization.

The current chair is Balint Joo, a computer scientist at Jefferson Lab in Newport News, Virginia, with extensive experience in algorithm development and lattice quantum chromodynamics.

In 2012–2013, the UC contributed to the annual survey and helped increase participation by OLCF users. The UC also provided input into the OLCF User Support website redesign, and members of the UC participated and gave talks at various OLCF workshops and training events, including the 2012 Accelerating Computational Science Symposium in Washington, DC.

"The User Council serves several roles. Our primary role is user advocacy, either directly or indirectly," said Joo. "Direct representation would involve advocating on behalf of users with concerns they bring to us, should they feel that the OLCF is not listening to their needs. The fact that during my time as chairman I have not had to engage in such direct representation is a sign that the more regular channels of communication between the center and its users are functioning very well."

—By Gregory Scott Jones

## The Road to Success

Neils Bohr once said, "An expert is a person who has found out by his own painful experience all the mistakes one can make in a very narrow field."

Thanks to computers such as Jaguar and Titan, researchers can make those mistakes cheaper and faster than ever before. And this means scientific breakthroughs, well, cheaper and faster than ever before.

Where experiments are costly, time-consuming, and sometimes dangerous, petascale computing systems are forging a path forward via simulation, providing researchers with an unprecedented tool for achieving scientific breakthroughs in record time.

Just ask Ramgen Power Systems, a small, Seattle-based energy R&D firm. Ramgen is currently developing a novel gas compressor system based on shock-wave technology used in supersonic flight applications, research that shows serious promise in helping the United States tackle one of its most pressing environmental challenges: the management of carbon dioxide.

Power plants burning fossil fuels account for more than 40 percent of the world's energy-related CO<sub>2</sub> emissions and will continue to dominate the supply of electricity until the middle of the century. There is an urgent need for cost-effective methods to capture and store their carbon emissions.

DOE is currently sponsoring large-scale demonstration projects to prove the viability of CCS, or the capture of CO<sub>2</sub>, and its subsequent storage underground. The principal barrier to widespread application of CCS is its cost; once the CO<sub>2</sub> is captured, compressing it to the required 100 atmospheres represents approximately 33 percent of the total cost.

Enter Ramgen. Their novel compression system has the potential to achieve CCS safely and economically, when perfected. And without Jaguar and Titan, achieving that perfection could take years, or even decades, far more time than DOE's demonstration schedule permits.

That's because the traditional process to design and optimize new turbomachinery, or machines that transfer energy between a rotor and a fluid, involves the testing of multiple physical prototypes, testing that can be both expensive and time-consuming. Ramgen modified this conventional development process by using more extensive computer simulations validated by test results. DOE leadership determined that applying the most capable, HPC systems and modern CFD analysis would further accelerate optimizing the turbomachinery's performance.



Ramgen runs hundreds of design combinations with up to 50 design parameters simultaneously (“ensembles”) to find the optimal design solution for their compression technology. Such a process requires systems like Jaguar and Titan with hundreds of thousands of processors, along with sophisticated mathematical algorithms to analyze and predict the optimal solutions. Other firms have tried similar experiments but have not been successful.

In fact, some of Ramgen’s most recent simulations involving the time-varying nature of a flow field have run at one billion and four billion grid cells on Titan, sometimes solving ten times the number of partial differential equations at any given time as previous simulations.

“It’s only because of the world-class nature of Jaguar and Titan that it’s possible to run these jobs,” said Ramgen’s CFD Lead and Principal Investigator Allan Grosvenor. “If we had to run them on lesser systems, these would take a number of years.”

### **Crawl before you walk**


Ramgen’s success is a testament to the symbiotic relationship the OLCF shares with its users.

In just a few years, the company has revolutionized its R&D with the help of Jaguar, and now, Titan.

“In 2009, we started speaking in detail with ORNL people and ASCR [DOE Office of Advanced Scientific Computing Research] and DOE,” said Grosvenor. At the urging of then-Secretary of Energy Steven Chu, Ramgen was awarded a director’s discretionary allocation on Jaguar. “Secretary Chu really moved the ball forward,” said Grosvenor, “when he took the time to understand the technology and decided it was important to the nation.”

Since those initial conversations Ramgen has made enormous leaps in its ability to apply the formidable capability of HPC to solve complex engineering problems. For the last three years it has received its allocation time from the ASCR Leadership Computing Challenge (ALCC) program, a testament to the firm’s rapid progress. ASCR gives allocations solely to researchers capable of using the fastest computers in the world, and Ramgen’s award three years running speaks volumes for their success.

But success like Ramgen’s is the culmination of hard work not only on the part of researchers, but also on the OLCF. Back in 2009, when Ramgen first began working with Jaguar, Ramgen’s ISV’s FINE/Turbo code was not ready for a system with Jaguar’s power. “We needed to close a very large technological gap in order to reach the



***Doug Jewett, CEO of Ramgen Power Systems, stands by an engine and turbine his company uses to compress carbon dioxide. Image courtesy of The Seattle Times***

HPC level we're now at, and we've achieved this by organizing a series of development projects. Our success was due to efforts of highly talented people at Numeca and with the expert support of some amazing people at the OLCF," said Grosvenor.

Through a strong collaboration between Ramgen, Numeca (their software vendor), and OLCF staff, a series of major rewrites vastly improved the code's parallelism, so much so that it would soon be able to utilize as much as 80 percent of Jaguar's peak performance. Ramgen can now run their huge, ensemble-based optimization computations across more than 240,000 cores (for instance 240 candidates in an optimization database each distributed over 1,000 cores), bringing the optimized design parameters within reach faster than anyone ever thought possible. This long-lasting collaboration has also led to a series of different development projects and a major rewrite of Numeca's FINE/Turbo code that now runs on Titan.

One of the more valuable contributions to come out of the collaboration is Ramgen's implementation of ORNL's Adaptable I/O System (ADIOS). ADIOS provides a simple, flexible way for scientists to describe the data in their code that may need to be written, read, or processed outside of the running simulation.

"ADIOS allowed us to radically change the I/O so that the code could run at scale," said Grosvenor. Originally Ramgen's I/O strategy was serial, tedious to say the least, "and it used a file format that highly limited the possible achievable acceleration of reading and writing files." In fact, ADIOS has been so helpful that Numeca is the first commercial company to adopt it into a commercial code.

The rewrite resulted in a 100-fold speedup in time to solution and a 2-fold decrease in memory usage per core, which allows higher-resolution computational meshes to be used, resulting in more accurate answers for specific aerodynamic performance issues.

But the simulations are only half of the story, especially when optimization is the ultimate goal, as is Ramgen's. "When you're running high-resolution simulations to study specific phenomena, you need high-level post-processing and visualization," said Grosvenor. "Without it, you probably shouldn't run these kinds of simulations at all." Here, too, the OLCF has been critical in helping Ramgen to analyze data rapidly, thus expediting their design to meet DOE's strict timeline. "Visualization has been game-changing for us and we would love to continue this interaction," said Grosvenor.

In the beginning, Ramgen relied heavily on the popular visualization software known as VisIT. However, they soon encountered issues similar to those they faced before adopting ADIOS: VisIT proved challenging when paired with an industrial code. OLCF staff helped Ramgen evolve their visualizations in line with their simulations, enabling Grosvenor and his colleagues to capture sophisticated aerodynamics, thus making the optimization process much more effective. In fact, Ramgen now has a team of engineers developing internal codes for data mining.

And these advances may just be the tip of the iceberg. As Ramgen ramps up and prepares to tackle Titan, OLCF staff have located additional promise in the area of "smart core pinning." Because Ramgen requires more memory process, they can't yet use all of the cores on each node. "Even with half the cores, they will have used up all the memory, leaving the other cores idle," said Sudharshan Vazhkudai, technology integration group leader at the OLCF.

On Titan, each compute node has eight "Bulldozer" modules, each with two cores that share a floating point unit (FPU). In the OLCF's internal testing with applications that do not use all the cores, staff are noticing a speedup if they only use every alternate core for their processes, which eliminates the contention for the FPU. The system does not automatically schedule processes in this fashion and leaves it to the user. "We hope to work with Ramgen, and similar such cases, to understand their usage patterns and enable smart placement and pinning of processes to core," said Vazhkudai.

The success of the OLCF/Ramgen collaboration was recognized throughout the industry. In fact, this partnership enabled Ramgen to win an International Data Corporation HPC Innovation Excellence Award at SC12 in Salt Lake City, Utah.

## A hard day's work rewarded

Ramgen's intelligently driven optimization process reduced what used to be months of work to a mere 8 hours. Grosvenor noted: "Two years ago it was impossible for us to run an intelligently driven ensemble of simulations for design optimization like this. The sophisticated design analyses that we are running now are having a significant impact on our turbomachine development work."

Ramgen's CEO and Director Doug Jewett added: "Jaguar provides, in a remarkably short time, data that we use to predict optimal designs. It's enabling us to advance the design of our equipment in a timeframe that simply would not have been possible without Jaguar, Titan, and the assistance of the OLCF's Scientific Computing Group."

And there is even more good news. There is enormous synergy in the work done on shock-wave-based compression for CO<sub>2</sub> and Ramgen's Integrated Supersonic Component Engine. This technology will generate electricity—efficiently and cost effectively—using dilute methane gases released during coal mining operations and from landfills. Methane, per volume, traps 21 times more heat in the atmosphere than CO<sub>2</sub> but has a shorter atmospheric lifetime compared with it.

Engineers from Jim Walter Resources, a mining company, and Ramgen staff have developed an approach which they projected can use up to 75 percent of the methane now being emitted into the atmosphere worldwide as fuel to generate electricity. Currently approximately 90 percent of this methane is simply vented to the atmosphere.

This compression process has additional exciting possible applications with other gases and for other products. "What all these technologies have in common," said

Jewett, "is that they'll boost the gas pressure significantly at lower cost than alternatives."

"I believe that by applying the optimization enabled by Jaguar and Titan, Ramgen will accomplish levels of aerodynamic refinement with shock-wave-based technology in 5 years that took 50 years with gas turbines," said Ramgen's Chief Technology Officer Shawn Lawlor.

That dramatic increase in time-to-production will be necessary if Ramgen is to meet DOE's deadline for the development of efficient, cost-effective CCS. And thanks to Jaguar and Titan, it just might happen.

—By Gregory Scott Jones

## Three Paths to Discovery

---

Users arrive at the OLCF through three different avenues: the INCITE program, which allocates the majority of computer time on Titan, the ALCC program, or via Director's Discretion (DD), which represents projects chosen by the center for a particular research angle or to prepare for INCITE or ALCC awards.

### INCITE

Open to researchers from academia, government labs, and industry, the INCITE program is the major means by which the scientific community gains access to some of the fastest supercomputers. The program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering.

### ALCC

The mission of the ALCC is to provide an allocation program for projects of interest to DOE with an emphasis on high-risk, high-payoff simulations in areas directly related to the DOE mission and for broadening the community of researchers capable of using leadership computing resources.

Open to scientists from the research community in industry, academia, and national laboratories, the ALCC program allocates up to 30 percent of the computational resources at ASCR's supercomputing facilities. ASCR supercomputing facilities include NERSC at Lawrence Berkeley National Laboratory and the Leadership Computing Facilities at Argonne and Oak Ridge national laboratories. These resources represent some of the world's fastest and most powerful supercomputers.

### Director's Discretion

Since its inception in 2006, the OLCF's DD program has granted allocations in virtually all areas of science identified by DOE as strategic for the nation. Additional allocations have been made to promote science education and outreach, and requests and awards have grown steadily each year.

The goals of the DD program are threefold: development of strategic partnerships, preparation for leadership computing competitions (i.e., INCITE and ALCC), and application performance development and measurement. These goals are aligned with particular strategic goals for the OLCF, namely the expansion of the leadership computing science community and enhancement of the pervasive use of leadership computing in a variety of scientific fields.

The DD program is also accessible by the general HPC community to carry out prototyping and development exercises for nascent and less-efficient applications. The program also supports a variety of "data-only" projects that require data storage and bandwidth capabilities but few compute resources.

The Resource Utilization Council makes the final decision on DD applications, using written reviews from subject matter experts. The actual DD project lifetime is specified upon award: allocations are for 1 year or less. The typical size of DD awards can range from tens of thousands of hours to 4 million hours or more.





## Inside the OLCF

# Advancements

## Introduction

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

The OLCF strives to develop new tools to make scientific discovery at the facility more efficient. Here are some of the facility's 2012 innovations.

## Innovations in supercomputer optimization

ORNL is a leader in the development of features for batch schedulers used on Cray and other systems around the world. Batch schedulers schedule jobs in the queue to optimize performance. OLCF staff members have worked on Adaptive Computing, Inc.'s, Moab and TORQUE, which are used on supercomputers around the world. The intelligent management of Moab saves time and money by optimizing HPC, and TORQUE is an advanced open-source product that incorporates significant advances in scalability, reliability, and functionality. The first port of Moab and TORQUE products for the Cray X-series was performed at the OLCF.

In 2012 OLCF staff collaborated with Adaptive Computing to alleviate synchronization issues between the native Cray application-level placement scheduler and the Moab batch scheduler. The new design allows the Moab and TORQUE servers to easily exist outside the Cray machine on external service nodes, which enables users to submit



***Because of their expertise in large-scale data management, OLCF personnel are helping to manage and analyze the large data sets produced by ORNL's Spallation Neutron Source.***

and manipulate jobs when the Cray compute partition is unavailable. The OLCF hosted the beta testing of this new design at scale on Jaguar and provided numerous patches to Adaptive Computing for the redesigned code. This redesign has now been incorporated in Adaptive Computing's product line.

### **Innovations in scientific data management at SNS**

In addition to its computing expertise, ORNL also operates the world's brightest neutron source, the Spallation Neutron Source (SNS). Funded by the DOE Office of Basic Energy Sciences, this national user facility hosts hundreds of scientists from around the world, providing a platform to enable breakthrough research in materials science, sustainable energy, and basic science. Because of their expertise in large-scale data management, OLCF personnel are helping to manage and analyze the large data sets—over 1 terabyte per data set—generated by the intense pulses of neutrons.

OLCF collaborated with SNS staff to complete the Accelerating Data Acquisition, Reduction, and Analysis (ADARA) Lab-Directed Research and Development project. As a result of the ADARA project, a new data infrastructure was created that enhances users' ability to collect, reduce, and analyze data as it is taken; create data files immediately after acquisition, regardless of size; reduce a data set in seconds after acquisition; and provide the resources for any user to do postacquisition reduction, analysis, visualization, and modeling without requiring users to be on-site at the SNS facility.

ADARA provides a streaming data backplane, allowing scientists to go from experiment to data reduction to obtaining an energy spectrum or diffraction pattern

nearly instantaneously and while the experiment is still running. Rather than the previous approach—saving data in “buckets” and, once the bucket is full, handing the bucket off to the next process—ADARA uses a streaming approach. As data is being captured, translation is done concurrently. Every event coming off a detector is translated to a common data format as the experiment progresses. While performing translation, ADARA also does live data reduction; as neutron events are coming off the detectors, that same data is reduced live into an energy spectrum or diffraction pattern.

### **Innovations in high-performance parallel file systems**

At the OLCF and many other computing facilities, the Lustre parallel file system provides support to client systems, data storage, and fast throughput. The OLCF cofounded the OpenSFS organization, which allows the Lustre code to be improved through debugging, maintenance, and feature development. OLCF staff members have overseen the development of features that significantly benefit the OLCF and the broader Lustre community. In 2012 these enhanced features included a metadata server (MDS) survey tool, imperative recovery, and general metadata improvements. Metadata is “data about data,” which provides information about files within a system, and a good survey tool is needed for a file system to run smoothly.

One of the greatest challenges to applications on the HPC systems is MDS performance. An MDS survey tool was developed through a contract with Whamcloud, Inc., a division within Intel that develops and supports Lustre features. The tool simulates standard workloads, permitting performance testing without requiring any clients to create metadata operations. Several metadata

handling defects were identified and addressed during 2012. One such defect was the single lock on a single directory. This lock would serialize access to a directory if several processes simultaneously attempted to access this location in the file system. The solution was to introduce more granular locking on directory structures, thus allowing more concurrent operations.

Today's servers are normally multicore, but Lustre had not taken advantage of these resources. To address this issue, another 2012 performance improvement was the minimization of penalties that occur when processes migrate from one core to another.

In addition to the work leveraged by OpenSFS, internal work at the OLCF has also helped enhance Lustre. Whamcloud supports only Red Hat-based servers. The OLCF has helped bridge the gap for support in this unique environment. The OLCF has made Lustre compatible with newer kernels, or the underlying structures that allow the hardware and software to communicate, with various operating systems. One of these is SuSE, which was not previously supported. This work has been completed and will be a part of the Lustre 2.4 release, which is important for the Titan deployment.

The Gemini interconnect is a new networking technology used in the Cray XE and XK systems. While the underlying Gemini technology was supported on Lustre 1.8, newer Lustre versions do not support this technology, limiting the OLCF from taking advantage of its new features and performance improvements. To overcome this limitation, OLCF staff have ported and improved the Gemini Lustre Networking driver for Lustre. The base work has been completed and will be merged in Lustre 2.4 support. The goal is to have that work also merged into the Lustre 2.4 version to support Titan.

Monitoring and reporting of disk usage on Lustre file systems is a resource-intensive task and can affect metadata performance if not done in a centralized and scalable way. LustreDU is a nonintrusive tool developed by OLCF personnel to address this issue. It provides an end-user utility that queries a database of file- and directory-size information. This database is updated daily from a separate process that runs on the Lustre servers. This approach is significantly more efficient than going through the Lustre client-side application programming interface (API) that the normal "du" utility would have to use. This tool has had a significant positive impact on OLCF operations as system administrators use it to trim the scratch space usage of the Spider Lustre parallel file system.

### **Innovations in the HPSS archival storage**

High Performance Storage System (HPSS) is the software used by the OLCF for managing more than 30 petabytes of data on disk and robotic tape. HPSS also has a large installed base outside of the OLCF. It was developed and is maintained through an industry and national laboratory

collaboration and provides highly flexible and scalable hierarchical storage management that keeps recently used data on disk and less recently used data on tape.

For the last several years, OLCF staff members have had the primary responsibility for developing several HPSS subsystems, including the Storage System Manager (SSM), or the graphical and command-line interface for monitoring, configuring, and controlling the system; the Bitfile Server (BFS), which makes up one-third of the Core Server; the logging subsystem; and the accounting subsystem.

During 2012 the OLCF's efforts were devoted to HPSS release 7.4, HPSS release 7.p, and quality improvements. The capabilities implemented in release 7.4 had significant contributions from OLCF HPSS developers. Developers from the facility contributed to dynamic drive updates, which build upon the dynamic drive's add-and-delete functionality and allow device configuration updates without system downtime. OLCF staff members were responsible for enhancements to "hpssadm," which is the command-line interface to SSM. Thanks to this improvement, lengthy system configuration changes can now be automated in a batch script, reducing or eliminating downtime. A complete system can now be configured from a script, enabling quick setup. Work done by OLCF developers provided logging enhancements to change log files from binary to text format, which dramatically improves real-time debugging. Log archiving was also improved to be more flexible and avoid potential loss of logging data during times of high activity; previous systems could lose some log data when a log file could not be archived quickly enough. OLCF staff members made contributions by improving the redundant array of inexpensive tape, or RAID, which will improve reliability of the data that the OLCF stores and retrieves in production.

Final testing of the 7.4 release was conducted in 2012, and release general availability is targeted for March 2013.

The first step in the planning for moving HPSS to exascale is release 7.p, a performance-centered release. The strategy is to partition the metadata database across multiple storage devices and adapt HPSS to exploit the parallelism that this partitioning makes available. Work on release 7.p began in mid-2012 and is ongoing.

A major focus of HPSS development for 2012 was on improving quality. By redesigning the collaboration software development process (SDP), the HPSS Technical Committee is addressing the challenge of maintaining software quality and agility in the face of declining resources. The OLCF developers have taken a major leadership role in the redesign of the SDP. In addition, in the past year the OLCF developers have significantly expanded, improved, and automated the Logging and BFS test suites—and are in the process of doing the same to the SSM test suite—to ensure greater reliability and higher quality in each release.



## **Innovations in networking with the Common Communication Interface**

The OLCF poses scalability issues for everything from storage to debugging tools. In addition to Titan, the OLCF includes many different types of hardware as well as multiple types of network infrastructures. Each network provides at least two APIs: Berkeley sockets and the network's native interface, which provides better performance through direct access to the network hardware.

For each new generation of hardware, various groups within the OLCF modify applications to use each network's native API to obtain the best performance, but recently OLCF staff members have been working on a new programming interface that will provide a common API for applications, allowing them to take advantage of current networking hardware and next-generation hardware as it is acquired. This new API, known as the Common Communication Interface (CCI), is being jointly developed by ORNL, the University of Tennessee, Myricom, and Cisco. CCI is designed for portability, scalability, and performance. The new API will create one interface for all applications, which makes doing work at the OLCF easier. The software will be ready for adoption in 2013.

## **Innovations in hierarchical collective communication library with Cheetah**

Collective operations are used in parallel computing to synchronize processes and for other operations in which all processes participate, such as broadcasts and reductions. Collective operations are often among the most sensitive parts of scientific simulations with regard to performance and scalability. Unfortunately, collective operations that have been designed for older supercomputers can present scalability and performance bottlenecks on today's multicore compute nodes and when used on different network fabrics. To solve this problem, the OLCF's Computer Science Research Group built Cheetah, a framework for implementing collective operations to address these problems from the ground up.

The goal of Cheetah is to provide an efficient collective-operations implementation for modern supercomputers and various programming models, including message passing and global address space models. Cheetah achieves this goal by designing the collective operation as a combination of simple collective primitives, which are each optimized for a different homogeneous block in the heterogeneous architecture. The reference implementation outperforms the native implementations on Cray XE/XK and InfiniBand systems and scales to over 100,000 cores on Cray XE/XK systems.

Cheetah is a joint effort of ORNL and Mellanox Technologies, and the Cheetah research shaped the features and capabilities of Mellanox's CORE-Direct technology. Cheetah 1.0.0, the first version to be officially released, will appear in forthcoming releases of Open MPI, a widely used open-source implementation of the MPI-2 standard.

## **Innovations in on-site support for the DDT debugger and HMPP compiler**

The OLCF has long-standing partnerships with both Allinea Software, Ltd., and CAPS Enterprise, and the facility has now implemented subcontracts with the two companies to enhance and extend their products to meet the scale and functionality needs of OLCF users. Allinea develops the distributed debugging tool (DDT), which helps programmers to spot problems with their codes quickly. CAPS Enterprise is the developer of the Hybrid Multicore Parallel Programming (HMPP) compiler suite for GPUs. A compiler translates programming language into machine language. In 2012 the OLCF expanded the partnerships with the companies to include the full-time placement of partner employees on-site at ORNL to provide in-depth support for these tools to OLCF staff and users.

On-site CAPS and Allinea staff actively work with application teams to address problems that go beyond typical tool support, often assisting in use of the tools in challenging situations. For example, Allinea on-site staff member Dirk Schubert recently helped an application team use DDT to isolate and fix a bug that occurred in its application only when it was run at full scale on Titan. When coupled with a long-term partnership such as OLCF's with Allinea and CAPS, placing support staff on location at the center is a best practice that enhances the partnership, leads to improvements in the partner's products, and facilitates center operations.



## Inside the OLCF

# OLCF Systems Overview

### Introduction to OLCF

Home to the fastest supercomputer in the nation, Titan, the OLCF is a leading HPC center, striving to help researchers unlock the many secrets held by our universe. Titan is accompanied by one of the largest data storage systems in the world.

Working for the OLCF are some of the brightest minds in HPC. Working in five separate groups—Scientific Computing, User Assistance and Outreach, Technology Integration, High-Performance Computing Operations, and Application Performance Tools—the OLCF staff is responsible for ensuring that Titan users get the best experience possible.

ORNL places a premium on efficiency. The Computational Sciences Building where Titan resides was one of the first computing centers in the country to be Leadership in Energy and Environmental Design (LEED)-certified. It was deemed so due to the vapor barrier that keeps the inside of the building at a higher pressure than the outside. The setup allows hot air to escape out, while preventing outside air from coming in, which helps to guarantee Titan's longevity.

Due to the effort put forth by the OLCF, its vendors and its users, Jaguar, Titan's predecessor, became the first supercomputer to run working scientific applications at the petascale. In the days ahead, Titan will be poised to take us one step closer to the exascale—computers capable of a million trillion calculations a second.

## Titan

Upgraded from Jaguar in October 2012, Titan is now the nation's premier supercomputer for scientific discovery. One month after its upgrade, the Top500 list—a semiannual ranking of the most powerful computing systems in the world—announced Titan as the new number one. The Cray XK7 supercomputer reportedly reached 17.59 petaflops but is theoretically capable of 27 petaflops, i.e., 27 quadrillion calculations per second.


Titan combines GPUs with traditional CPUs. As always, energy efficiency is a must, and Titan does require an enormous amount of power—roughly 8.2 megawatts (1.2 megawatts more than Jaguar). However, that increase in power consumption is well worth it due to the fact that the newly added GPUs make Titan 10 times more powerful than its predecessor. Consequently, Titan landed in third place on the list of the most energy-efficient supercomputers in the world, the Green500.

Occupying less than 4,500 square feet, Titan comprises 200 cabinets that house 18,688 NVIDIA Tesla K20 GPU accelerators and 18,688 16-core AMD Opterons (299,008 cores total), with a total of 710 terabytes of memory.

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

Titan generates a lot of heat, and 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, rushing water over Titan through an arrangement of pipes affixed to each cabinet.

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







## Data management

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

## High-Performance Storage System

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



## Spider

The Lustre-based file system Spider lies at the center of the OLCF's technological integration. With more than 26,000 nodes mounting it, Spider is one of the largest file systems in the world. Spider organizes data from the multiple computing platforms into a unified file system. The project began in 2005, and now Spider is the main operational file system connecting Titan, the Lens visualization cluster, the Smoky development cluster, and the center's dedicated GridFTP servers—at a bandwidth of 240 gigabytes per second. Since Titan's arrival has put added strain on Spider's I/O demands and capabilities, preparations are being made for a 2013 upgrade to Spider2. Spider2, a Lustre parallel file system, will be more than capable of handling Titan's added workload, offering over 18 petabytes to open science users of the OLCF, and with an increased, blazing-fast bandwidth of 1,000 gigabytes per second.

## 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 comes time to share those results, a visual representation is often the most effective. 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.



## EVEREST

For years the Exploratory Visualization Environment for Research in Science and Technology, better known as EVEREST, has given researchers an in-depth look at their results. With EVEREST's recent upgrade, researchers can now watch a frame-by-frame supernova or fusion reactor in stunning, three-dimensional detail.

Eighteen  $1,920 \times 1,080$  stereoscopic Barco projection displays create a 37 megapixel main display wall that is 30 by 8.5 feet, while an adjacent, 16-panel, 33 megapixel display measures 7.6  $\times$  13.5 feet. A touch panel at the control desk can access both display walls, as well as individual panels. Access can also be routed to other devices, and the conference table provides pop-out boxes for additional peripheral video inputs.

A dedicated Linux cluster accompanied by two fat nodes controls both display walls. The cluster is a 9-node, 16-core Intel machine with 96 gigabytes of memory containing an NVIDIA 5000 and an NVIDIA 600 graphics card.

EVEREST uses the Lens system—a 77 fat-node cluster dedicated to data analysis and visualization. The cluster is cross mounted with the center-wide Lustre file system so that simulation data from other OLCF resources can be accessed without duplication of files.

Now fully upgraded, EVEREST offers state-of-the-art visualization and rendering facilities to the Titan user community.

## Lens

Lens is a 77-node Linux cluster dedicated to data analysis and high-end visualization. Each of the 77 nodes contains four 2.3 GHz AMD Opteron processors. Of the nodes, 32 are configured with 64 gigabytes of main memory, an NVIDIA 8800 GTX GPU with 768 MB of memory, and an NVIDIA Tesla general-purpose GPU with 4 gigabytes of memory. The other 45 nodes do not contain GPUs but are configured with 128 gigabytes of memory. The primary purpose of Lens is to enable data analysis and visualization of simulation data generated on Titan so as to provide a conduit for large-scale scientific discovery. Members of allocated Titan projects will automatically be given accounts on Lens.







## Developmental Systems

In order to ensure that researchers using Titan are spending their allocations solving scientific problems instead of troubleshooting applications, the OLCF uses two safeguard systems, Sith and Smoky. These systems help researchers scale up algorithms and identify problems before they work on Titan. Both systems can also operate independently to solve ancillary calculations for other projects outside Titan.

### 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 and expect a result without having to do any intermediary mechanics, or further data manipulation in order to get the output. It is used for workflow automation for jobs running from Titan and for advanced data analysis. The system contains 40 compute nodes. Each compute node contains four 2.3 GHz eight-core AMD Opteron processors and 64 gigabytes of memory. The system is configured with an 86 terabyte Lustre file system for scratch space.



### Smoky

Smoky is a development resource provided to users needing a comparable system to the larger resources for application development. More simply stated, it can be used as a test bed for smaller jobs so that workflow doesn't affect heavier jobs being done on Titan. Smoky is configured as an 80-node Linux cluster consisting of four quad-core 2.0 GHz AMD Opteron processors per node, 32 gigabytes of memory (2 gigabytes per core), a gigabit Ethernet network with Infiniband interconnect, and access to Spider, the center-wide Lustre-based file system. Its primary purpose is application development, specifically for petascale applications. Its programming environment mirrors the environment available on the Titan system. Only a limited number of center users are allowed access to this resource.



# High-Impact Publications

---

## Biology

B. Lindner and J.C. Smith, 2012, "Sassena — X-ray and neutron scattering calculated from molecular dynamics trajectories using massively parallel computers", *Comp. Phys. Comm.*, **183** (7), 1491 (2012).

## Chemistry

J. Brabec *et al.*, "Parallel Implementation of Multireference Coupled-Cluster Theories Based on the Reference-Level Parallelism," *J.Chem. Theory Comput.* **8** (2), 487 (2012).

Q. Li *et al.*, "Supramolecular Self-Assembly of  $\pi$ -Conjugated Hydrocarbons via 2D Cooperative CH/ $\pi$  Interaction," *ACS Nano* **6** (1), 566 (2012).

M. O'Hagen *et al.*, "Proton Delivery and Removal in [Ni(P2N2R)-N-R']<sub>2</sub> Hydrogen Production and Oxidation Catalysts," *J. Am. Chem. Soc.* **134** (47), 19409 (2012).

## Climate

C. K. G. Castillo, S. Levis, and P. Thornton, "Evaluation of the New CNDV Option of the Community Land Model: Effects of Dynamic Vegetation and Interactive Nitrogen on CLM4 Means and Variability," *J. Climate* **25** (11), 3702 (2012).

D. M. Lawrence *et al.*, "The CCSM4 Land Simulation, 1850-2005: Assessment of Surface Climate and New Capabilities," *J. Climate* **25** (7), 2240 (2012).

J. D. Shakun *et al.*, "Global Warming Preceded by Increasing Carbon Dioxide Concentrations during the Last Deglaciation," *Nature* **484** (7392), 49 (2012).

G.M. Meehl, A.Hu, C. Tebaldi, J.M Arblaster, W.M. Wahsington, H. Teng, B.M. Sanderson, T. Ault, W.G. Strand, J.B. White III, "Relative outcomes of climate change mitigation related to global temperature versus sea-level rise," *Nature Clim. Change* **2**, 576 (2012).

G. A. Meehl *et al.*, "Climate System Response to External Forcings and Climate Change Projections in CCSM4," *J. Climate* **25** (11), 3661 (2012).

## Fusion

H. S. Zhang, Z. Lin, and I. Holod, "Nonlinear Frequency Oscillation of Alfvén Eigenmodes in Fusion Plasmas," *Phys. Rev. Lett.* **109** (2), 025001 (2012).

## Materials

D. P. Hashim *et al.*, "Covalently Bonded Three-Dimensional Carbon Nanotube Solids via Boron Induced Nanojunctions," *Sci. Rep.* **2** (363), (2012).

A. Lopez-Bezanilla *et al.*, "Boron Nitride Nanoribbons Become Metallic," *Nano Lett.* **11** (8), 3267 (2012).

A. Morelos-Gomez *et al.*, "Clean Nanotube Unzipping by Abrupt Thermal Expansion of Molecular Nitrogen: Graphene Nanoribbons with Atomically Smooth Edges," *ACS Nano* **6** (3), 2261 (2012).

A. N. Volkov and L. V. Zhigilei, "Heat Conduction in Carbon Nanotube Materials: Strong Effect of Intrinsic Thermal Conductivity of Carbon Nanotubes," *Appl. Phys. Lett.* **101** (4) 043113, (2012).

R. Yu *et al.*, "Bose Glass and Mott Glass of Quasiparticles in a Doped Quantum Magnet," *Nature* **489** (7416), 379 (2012).

D. P. Hashim *et al.*, "Covalently bonded three-dimensional carbon nanotube solids via boron induced nanojunctions," *Scientific Reports* **2**, 363 (2012).

## Physics

J. Erler *et al.*, "The Limits of the Nuclear Landscape," *Nature* **486** (7404), 509 (2012).

G. Hagen *et al.*, "Continuum Effects and Three-Nucleon Forces in Neutron-Rich Oxygen Isotopes," *Phys. Rev. Lett.* **108** (24), 242501 (2012).

G. Hagen *et al.*, "Evolution of Shell Structure in Neutron-Rich Calcium Isotopes," *Phys. Rev. Lett.* **109** (3), 032502 (2012).

S. Palaniyappan *et al.*, "Dynamics of Relativistic Transparency and Optical Shuttering in Expanding Overdense Plasmas," *Nature Phys.* **8** (10), 763 (2012).

V. Roytershteyn, W. Daughton, H. Karimabadi, and F. S. Mozer, "Influence of the Lower-Hybrid Drift Instability on Magnetic Reconnection in Asymmetric Configurations," *Phys. Rev. Lett.* **108** (18), 185001 (2012).

M. Wan, W.H. Matthaeus, H. Karimabadi, V. Roytershteyn, M. Shay, P. Wu, W. Daughton, B. Loring, S.C. Chapman, "Intermittent dissipation at kinetic scales in collisionless plasma turbulence," *Phys. Rev. Lett.* **109**, 195001 (2012).

E. J. Lentz *et al.*, "On The Requirements For Realistic Modeling of Neutrino Transport in Simulations of Core-Collapse Supernovae," *Astrophys. J.* **747** (1) 73, (2012).

H. Pais and J. Stone, "Exploring the Nuclear Pasta Phase in Core-Collapse Supernova Matter," *Phys. Rev. Lett.* **109** (15), 151101 (2012).



## Inside the OLCF

# Education, Outreach, and Training

The OLCF is committed to educating the next generation of scientists, reaching out to the community, and training users to ensure their success. In 2012, the OLCF hosted a number of programs to accomplish these goals.

## Education at the OLCF

---

HPC education is critical for scientists and students alike. Supercomputer simulations allow scientists to explore areas that don't lend themselves to experiment, such as global climate change and astrophysics. The OLCF also helps the next generation of scientists through summer internship programs and courses for high school students. Allowing students to do hands-on work with high-performance computing fosters an enthusiasm for computational science that can last a lifetime.

### Students gain supercomputing knowledge in crash course

In June, the OLCF offered a crash course in supercomputing for ORNL staff and summer interns, covering the basics of HPC. The two-day workshop in June exposed interns and non-technical staff to the basics of the tools and techniques used in parallel processing. The 90 attendees of the workshop programmed, compiled, ran, and debugged their own codes on Smoky, an 80-node Linux cluster normally used for application development.

<https://www.olcf.ornl.gov/2012/07/16/students-gain-supercomputing-knowledge-in-crash-course/>

### Educating future generations in HPC basics

The OLCF offered an eight-week course over the summer in HPC for interested summer interns and employees. The course was designed to provide a glimpse into the OLCF's mission and to introduce participants to the basic concepts of HPC. Attendees learned

about Unix command lines and parallel computing. At the end of the course, they assembled a cluster of Mac Minis to run their own programs.

<https://www.olcf.ornl.gov/2012/10/18/educating-future-generations-in-supercomputing-basics/>

### OLCF summer programs educate next generation

Over the summer, the OLCF hosted education programs for students in high school and college. As part of the Appalachian Regional Commission's Summer Math-Science-Technology Institute, high school students from across the Appalachian region spent two weeks learning the basics of parallel computing by assembling a cluster of Mac Mini computers to work in parallel. The OLCF also hosted a number of summer interns, who spent time gaining experience that will help them in their future careers as programmers and scientists.

<https://www.olcf.ornl.gov/2012/09/17/oak-ridge-leadership-computing-facility-summer-programs-educate-next-generation/>

### Outreach at the OLCF: tours

---

OLCF tours focus on work being done at the facility. Visitors are able to gain insight into the possibilities for research that exist when using the resources at the OLCF.

### Women in Physics attendees tour ORNL

Attendees of the Southeast Conference for Undergraduate Women in Physics toured the ORNL campus in January, including the OLCF. The conference introduced the women to a national laboratory setting and exposed them to the many career possibilities stemming from their education in physics.

<https://www.olcf.ornl.gov/2012/02/06/women-in-physics-attendees-tour-oak-ridge-national-laboratory/>

### Chu visits ORNL

In early 2012, former Energy Secretary Steven Chu visited ORNL, where he got a briefing on the advanced computer simulations of nuclear energy being done by CASL. Chu was updated on the progress being made on Jaguar's upgrade to Titan and took a turn experiencing a three-dimensional visualization of a virtual reactor's nuclear core.

<https://www.olcf.ornl.gov/2012/02/27/chu-visits-ornl/>

### Outreach at the OLCF: seminars and lectures

---

The OLCF invites scientists to ORNL and sends representatives to lecture around the country to make researchers aware of the resources that are available at the facility. Spreading the word about the OLCF allows high-impact science to continue at the center.

### Wells takes OLCF message to industry

OLCF Director Jack Wells spoke at the United Technologies Engineering Fellows Lecture series at Pratt & Whitney in East Hartford, Connecticut, on July 26. Wells lectured about the range of computational research performed at the OLCF and about the resources available at the facility. He emphasized the ways that modeling and simulations can help transform research and development for industry. Pratt & Whitney is an OLCF user and a world leader in the design, manufacture, and service of aircraft engines, industrial gas turbines, and space propulsion systems.

<https://www.olcf.ornl.gov/2012/08/13/wells-takes-olcf-message-to-industry/>

### Astrophysicist Messer speaks to next-generation computational scientists about HPC

The OLCF's Bronson Messer spoke to recipients of the Computational Science Graduate Fellowship at their annual conference July 26–28 in Arlington, Virginia. Messer spoke about the OLCF's resources and described how the fellows could obtain time at the OLCF for their own research through a Director's Discretion allocation. Messer said he wants the next generation of computational scientists to emphasize the high-performance when it comes to high-performance computing and to think big when it comes to research opportunities.

<https://www.olcf.ornl.gov/2012/08/13/astrophysicist-messer-speaks-to-next-generation-computational-scientists-about-high-performance/>

### Training at the OLCF

---

As the OLCF upgrades its resources to keep them world class, users must constantly learn how to use new systems, tools, and techniques. Training sessions from the OLCF experts help with that process. With the installation of Titan, training sessions became especially useful for prospective users and old OLCF users alike.

### Workshop prepares HPC users for Titan

Before Titan was up and running, the OLCF began preparing users for the changes that the new system would bring. For example, at a January workshop in 2012, OLCF specialists taught about how the use of debuggers, compilers, and performance analysis tools would change and prepared the users to integrate their codes into Titan's CPU-GPU architecture.

<https://www.olcf.ornl.gov/2012/02/27/workshop-prepares-hpc-users-for-titan/>





Inside the  
OLCF

## OLCF Computing Allocations for 2013

**INCITE** | [doeleadershipcomputing.org](http://doeleadershipcomputing.org)

### **Advanced Modeling of the Human Skin Barrier**

Michael Klein, Temple University

65,000,000 hours

### **Petascale Computing of Biomolecular Systems**

Klaus Schulten, University of Illinois, Urbana-Champaign

110,000,000 hours

### **Simulations of Ribosome Biogenesis and Cellular Processes**

Zan Luthey-Schulten, University of Illinois, Urbana-Champaign

51,410,000 hours

### **Cellulosic Ethanol: Simulation of Multicomponent Biomass Systems**

Jeremy Smith, Oak Ridge National Laboratory

78,000,000 hours

**High-Fidelity Simulations for Advanced Engine Combustion Research**

Joseph Oefelein, Sandia National Laboratories  
100,000,000 hours

**Precision Many-Body Quantum Simulations of Functionalized Structures**

Shiwei Zhang, College of William and Mary  
27,000,000 hours

**Non-Perturbative QED Study for Matter and Anti-Matter Collisions**

Michael Pindzola, University of Auburn  
30,000,000 hours

**Performance Evaluation and Analysis Consortium (PEAC) End Station**

Leonid Oliker, Lawrence Berkeley National Laboratory  
45,000,000 hours

**Collaborative Research into Exascale Systemware, Tools, and Applications (CRESTA)**

Lorna Smith, University of Edinburgh  
21,000,000 hours

**Climate-Science Computational Development Team: The Climate End Station II**

Warren Washington, University Corporation for Atmospheric Research  
60,000,000 hours

**CyberShake 3.0: Physics-Based Probabilistic Seismic Hazard Analysis**

Thomas Jordan, University of Southern California  
45,000,000 hours

**Global Seismic Tomography Based on Spectral-Element and Adjoint Methods**

Jeroen Tromp, Princeton University  
100,000,000 hours

**Multiscale Blood Flow Simulations**

George Karniadakis, Brown University  
51,000,000 hours

**Explosive Hazard Predictions with the Uintah Framework**

Martin Berzins, University of Utah  
45,000,000 hours

**Parameter Studies of Boussinesq Flows**

Susan Kurien, Los Alamos National Laboratory  
34,000,000 hours

**The Solution of Three-Dimensional PWR Neutronics Benchmark Problems for CASL**

Thomas Evans, Oak Ridge National Laboratory  
21,000,000 hours

**Computational Prediction and Discovery of Magnet Materials**

Bruce Harmon, Ames Laboratory  
45,000,000 hours

**Scalable First Principles Calculations for Materials at Finite Temperature**

Markus Eisenbach, Oak Ridge National Laboratory  
105,000,000 hours

**Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo**

Dario Alfe, University College London  
55,000,000 hours

**Quantum Monte Carlo Simulations of Hydrogen and Water Ice**

Richard Needs, University of Cambridge  
75,000,000 hours

**Predictive and Insightful Calculations of Energy Materials**

Paul Kent, Oak Ridge National Laboratory  
45,000,000 hours

**Ab Initio Simulations of Carrier Transports in Organic and Inorganic Nanosystems**

Lin-Wang Wang, Lawrence Berkeley National Laboratory  
25,000,000 hours

### **Safety in Numbers: Discovery of New Solid Li-Ion Electrolytes**

Boris Kozinsky, Bosch

48,000,000 hours

### **Petascale Simulations of Type Ia Supernovae**

Stan Woosley, University of California, Santa Cruz

55,000,000 hours

### **Three Dimensional Simulations of Core Collapse Supernovae**

Anthony Mezzacappa, ORNL

35,000,000 hours

### **Simulating Reionization of the Local Universe: Witnessing Our Own Cosmic Dawn**

Paul Shapiro, University of Texas

40,000,000 hours

### **Lattice QCD**

Paul Mackenzie, Fermilab

140,000,000 hours

### **Nuclear Structure and Nuclear Reactions**

James Vary, Iowa State University

74,000,000 hours

### **Transformative Advances In Plasma-Based Acceleration**

Warren Mori, University of California, Los Angeles

30,000,000 hours

### **Unraveling the Physics of Magnetic Reconnection with 3D Kinetic Simulations**

William Daughton, Los Alamos National Laboratory

55,000,000 hours

### **Magnetic Reconnection in High-Energy-Density-Laser-Produced Plasmas**

Amitava Bhattacharjee, University of New Hampshire

35,000,000 hours

### **High-Fidelity Simulation of Tokamak Edge Plasma Transport**

C.S. Chang, Princeton Plasma Physics Laboratory

100,000,000 hours



**Gyrokinetic Simulation of Energetic Particle Turbulence and Transport**

Zhihong Lin, University of California, Irvine  
20,000,000 hours

**HPC Colony: Adaptive System Software for Improved Resiliency and Performance**

Terry Jones, ORNL  
3,000,000 hours

**Coupled electronic and nuclear dynamics in solar photocatalytic water splitting**

Tom Miller, Caltech  
5,000,000 hours

**Predictive Simulations of Cuprate High-Temperature Superconductors**

Thomas Maier, ORNL  
50,000,000 hours

**Massively Parallel High-Fidelity Simulation of Spray Atomization**

Xiaoyi Li, United Technologies Research Center  
20,000,000 hours

**High-Level Studies of Excited States in Light Harvesting Systems and Complex Emergent Phenomena**

Karol Kowalski, Pacific Northwest National Laboratory  
5,000,000 hours

**VUQ Assessment of a Large Eddy Simulation Tool for Clean-Coal Technology**

Jeremy Thornock, University of Utah  
20,000,000 hours

**Exploring the Nature of the Lightest Massive Particles in the Universe**

Karen Heitmann, Argonne National Laboratory  
6,000,000 hours

**Multiscale Modeling of CO<sub>2</sub> Sequestration in Carboxysomes**

Greg Voth, University of Chicago  
8,000,000 hours

**Impact of the Inlet Boundary Condition on High-Pressure Turbine Temperature Predictions**

Anne Dord, GE Global Research  
34,000,000 hours

**Petascale Atomistic Simulations of Ultra Scaled Transistors**

Behtash Behin-Aein, GlobalFoundries  
8,000,000 hours

**Protein Folding and Computational Models**

Ken Dill, Stony Brook University  
2,000,000 hours

**Supercomputer-Enabled Accelerated Development of Revolutionary Supersonic Shock Wave-Based Turbomachines: Achieving DOE Goals for Compressing Carbon Dioxide and Achieving High Energy Efficiency via High Resolution CFD**

Allan Grosvenor, Ramgen  
40,000,000 hours

**Generation of Intrinsic Toroidal Rotation in Tokamak Plasma to Enable Stable Fusion Energy Production**

C. S. Chang, Princeton Plasma Physics Laboratory  
20,000,000 hours

**Time-Dependent Density Functional Theory (TDDFT) Approach to Nuclear Reactions**

Aurel Bulgac, University of Washington  
12,000,000 hours

**Designing Bio-inspired Catalysts for Energy Harvesting & Renewable Energy**

Pratul Agarwal, ORNL  
5,000,000 hours



