PETASCALE SCIENCE DELIVERED

Oak Ridge Leadership Computing Facility | Annual Report 2009

Oak Ridge Leadership Computing Facility
Annual Report 2009

Creative Director: Jayson Hines
Writers: Gregory Scott Jones, Dawn Levy, Leo Williams, Caitlin Rockett, Wes Wade, and Agatha Bardoel
Editor: Priscilla Henson and Deborah Counce
Graphic Design and Layout: Jason B. Smith
Graphic Support: Jane Parrott and Andy Sproles
Photography: Jason Richards
Additional images: iStockphoto

James J. Hack, Director
National Center for Computational Sciences

Arthur S. Bland, Project Director
Oak Ridge Leadership Computing Facility

Oak Ridge National Laboratory
P.O. Box 2008, Oak Ridge, TN 37831-6008
TEL: 865-241-6536
FAX: 865-241-2850
EMAIL: help@nccs.gov
URL: www.olcf.ornl.gov

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

About the Cover

Neutron power distribution in a 17x17 assembly, PWR900 reactor core. Each assembly contains a 17x17 array of fuel pins that are illustrated in the figure. Three different assembly-level fuel loadings are modeled. The calculation was performed on 17,424 cores on the Jaguar XT5 supercomputer and required the solution of 78.5 billion unknowns.
- Simulation by Tom Evans, ORNL
  Visualization by Dave Pugmire, ORNL

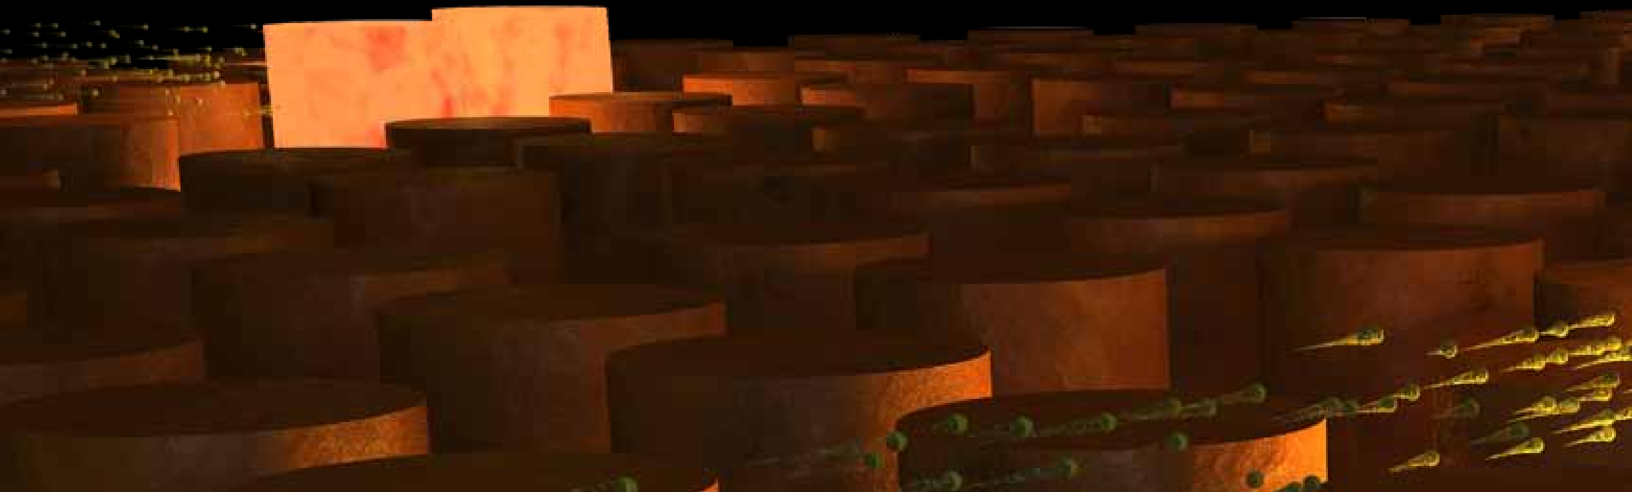

Contents ● ● ● ●

ORNL's supercomputing program grows from humble beginnings to deliver the most powerful system ever seen. On the way, it helps researchers deliver practical breakthroughs and new scientific knowledge in climate, materials, nuclear science, and a wide range of other disciplines.

OLCF: A HISTORY OF MAKING A DIFFERENCE

For Scientific Discovery and Innovation, the Pace Picks Up at the Petascale

Computing architecture shifts to hybrid multicore as supercomputers evolve for exascale

Jeff Nichols, ORNL, associate laboratory director for computing and computational sciences

In 2008 the Oak Ridge Leadership Computing Facility (OLCF) project at the National Center for Computational Sciences gave civilian scientists their first petascale high-performance computing (HPC) system. Its Cray XT supercomputer, "Jaguar," became the first computer able to achieve performance at more than a sustained petaflop (a million billion floating point operations per second) on full science applications. In November 2009 Jaguar's processors were upgraded, increasing the system's total computing capability to 2.6 petaflops.

Built for science with a balanced architecture, Jaguar took first place at the 2009 HPC Challenge in solving a dense matrix of linear algebra equations, executing the global fast Fourier transform algorithm, and demonstrating sustained memory bandwidth. It took third place in executing the global random access measure. At the same time, scientists immediately began making new discoveries using Jaguar. Gordon Bell Prize winner Markus Eisenbach calculated novel superconductor properties, and finalist Edo Aprà explored the electronic structure of water. Others pioneered petascale runs in biology, chemistry, climate, combustion, nuclear fission and fusion, geoscience, materials science, physics, and turbulence—fulfilling the OLCF mission of generating scientific breakthroughs.

The OLCF resources, housed at Oak Ridge National Laboratory (ORNL), are accelerating the pace of discovery and innovation in areas of strategic interest to the nation, such as developing clean, renewable energy and addressing global climate change. They also speed progress in fundamental and applied research, such as elucidating the underpinnings of health and disease and advancing the frontiers of astrophysics and genetics. Models and simulations, which address topics from batteries and biofuels to photovoltaics and superconductors, explore complex systems and even systems of systems. The OLCF provides a peerless capability for delving into complexity, igniting insight, and arriving at answers faster than ever in the face of problems from Alzheimer's disease to global climate change.

Computation has joined theory and experiment as the third pillar of science, and Jaguar is speeding the translation from laboratory to marketplace of energy systems discoveries. For example, our carbon footprint can be reduced by expanding nuclear energy, but reducing the time to market of safer, more productive reactors will take simulations that help make the most of the current reactor fleet and optimize and certify the next generation of reactors.

We achieved a historical milestone in 2009. The accomplishment of deploying the world's fastest supercomputer was the culmination of a journey begun in 2004 with the inception of the OLCF. Stepwise, the OLCF grew in power through teamwork, strategy, a commitment to excellence, and service to science. Jaguar could not have ascended to #1 on the TOP500 list without close relationships with Cray and AMD. Simulations can generate a petabyte of data, so success in developing storage systems was crucial as well. The Spider file system now stores almost 11 petabytes of data and moves data on and off large platforms and into archival storage at a rate of up to 240 gigabytes per second.

Ascending the mountain from teraflops to petaflops could not have been accomplished in one giant leap. The incremental approach of developing and deploying applications on a series of ever-more-capable systems enabled scientists to appreciate the changing scenery of discoveries at each camp along the way. A common programming environment made the system usable for a wide range of scientific and engineering challenges. An expert project management team mitigated risks in setting up a world-class infrastructure. Likewise, liaisons to the scientific teams were vital to progress.

The OLCF's major accompishments include generating 40 percent of the climate data cited in the fourth assessment report of the Intergovernmental Panel on Climate Change, which unequivocally concluded that human activity contributed to warming the planet in the 20th century, and proving the 2D Hubbard Model depicts high-temperature superconductors better than does a conventional theory, a line of thought that may someday revolutionize electric-power generation, electronics, and transportation. Simulation results with the potential to stimulate the economy and save fuel include a Boeing investigation to optimize performance of aircraft wings.

This annual report outlines major accomplishments. Looking to the future, there are more mountains to climb. Simulating next-generation energy systems will take exascale computers capable of a million trillion calculations per second—a thousand times faster than petascale computers.

Reaching the exascale summit will not be easy. As computers are pushed to address increasing complexity, system architecture must evolve to meet technical challenges, such as providing sufficient bandwidth for information to flow into and out of memory and storage. Hybrid multicore systems with both central processing unit (CPU) chips and graphical processing unit (GPU) chips, common in gaming systems, will supplant current multicore systems, which are based solely on CPUs. Computing systems will have to become more energy efficient. Software will have to scale up to use millions of processors and will have to become more resilient to failure. Today's workforce will have to educate and train tomorrow's.

These challenges may be dwarfed by those application programmers will face with regard to concurrency, or parallelism. Jaguar has 225,000 cores, giving it a concurrency of 225,000 tasks. An exascale machine will have a concurrency of a billion. How will programmers deal with that?

Overcoming those difficulties and reaching the exascale will entail a 10-year journey. There will be other mountains to climb, other passes and valleys to cross on the way that will incrementally bring the community to the exascale, just as happened while achieving the petascale.

To this end ORNL, Lawrence Berkeley National Laboratory, and Los Alamos National Laboratory (LANL) have formed the Hybrid Multicore Consortium. These labs have invested substantially to deploy computing platforms based on current and emerging accelerator technologies, which speed processing on the leading large-scale computational resources they provide for the scientific community. Universities renowned for hybrid multicore architecture and software research, such as the Georgia Institute of Technology and the Swiss Federal Institute of Technology, are part of the team. While hybrid multicore technologies will be important in future high-end computing systems, most applications will require considerable re-engineering to take advantage of the systems. The consortium will address the migration of applications to hybrid systems to maximize the return on investment.

Reaching sustained exascale performance on science applications will involve the engagement of multiple agencies, including the Department of Energy as well as the international computational science community. Demand has been overwhelming not only for currently deployed resources but also for those to be deployed for various organizations to accomplish their scientific missions. For example, a $215 million agreement between the OLCF and the National Oceanic and Atmospheric Administration in 2009 will result in another petascale computer at ORNL focused on improving our understanding and prediction of the behavior of the atmosphere, the oceans, and climate. As the ORNL computing complex reaches maximum capacity for space, power, and cooling, planning is under way to expand the computing infrastructure to a 140,000-square-foot facility called the Extreme Scale Computing Facility.

Also in the works is OLCF-3, a project to deliver an HPC system of approximately 20 petaflops in 2012. This machine will be a hybrid multicore system. There is excitement and anticipation across the scientific community about scaling scientific applications to the OLCF-3 system and running them in a production manner. This hybrid multicore system will be the base camp to the exascale, with future installations and upgrades at Oak Ridge to achieve exascale computing by 2018.

The Oak Ridge Leadership Computing Facility:

A History of Scientific Discovery

Milestones on the road from terascale to petascale point the way to exascale

Today the Oak Ridge Leadership Computing Facility (OLCF) is the most powerful high-performance computing (HPC) complex in the world, solving problems in energy, climate, combustion, biology, nanotechnology, and other disciplines crucial to U.S. science leadership. Yet its beginnings were humble—even improbable. Its history shows great achievement made possible through consistent, strategic steps.

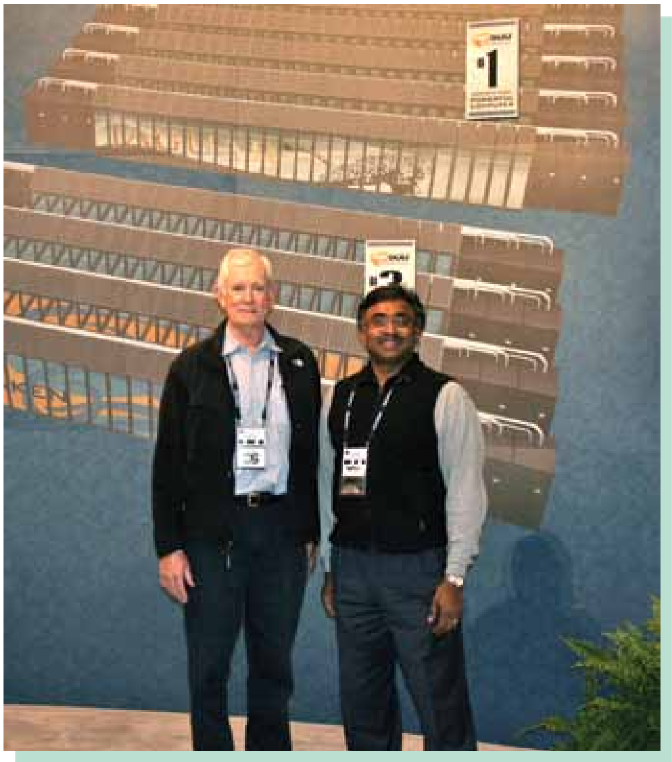

Former ORNL Director Al Trivelpiece, left, with ORNL's Deputy for Science and Technology Thomas Zacharia together at ORNL's booth during the 2009 Supercomputing conference. Both men were instrumental in bringing leadership computing to ORNL.

In 1991 the Department of Energy (DOE) announced a competition to select the site of an HPC research center that would be created by the High Performance Computing and Communications Act. At the time Oak Ridge National Laboratory (ORNL) had no large-scale scientific computing program. "Oak Ridge was not even on the map in computing," recalled Thomas Zacharia, who helped charter and found the OLCF. Nonetheless, a group led by Ed Oliver, head of scientific computing at ORNL, and Al Trivelpiece, the lab director, devised a plan. They developed a proposal, Partnership in Computatonal Sciences, that coupled key HPC science teams from around the country with ORNL computer scientists and hardware vendor Intel. Partnering with the science teams and vendor was a novel strategy for the time but showed ORNL's focus on scientific discovery. The ORNL proposal won, and in 1992 the Center for Computational Sciences (CCS) was born.

The CCS enabled users to run scientific applications at almost unimaginable speeds—billions of floating point operations, or calculations, each second (gigaflops)—on a 5-gigaflop Intel Paragon. The computer system grew through stepwise upgrades to 35 gigaflops, and later, 150 gigaflops. With each upgrade CCS experts worked closely with science teams nationwide to scale their applications to run on the increasingly parallel architecture. That partnership was productive, and simulations on the Paragon machine provided insight into how solids melt, pollutants travel through groundwater, materials behave in vehicular collisions, combustion occurs in engines, magnetic atoms interact, and air flows over an airplane's wing.

Ten years later, in 2002, the Japanese Earth Simulator eclipsed U.S. dominance in supercomputing. With the unprecedented speed of 36 teraflops, or trillions of calculations per second, it had the processing power of America's 14 fastest computers combined and was five times faster than its swiftest, the IBM ASCI White system at DOE's Lawrence Livermore National Laboratory. "The power of the Earth Simulator really took the world by surprise," OLCF project director Buddy Bland said.

To respond to the Earth Simulator, the DOE Office of Science issued a call for proposals in 2004 to develop a "leadership computing" capability in support of investigations in a broad spectrum of scientific domains. ORNL again partnered with science teams and a computer vendor, Cray, to put together a proposal. Its goal was ambitious—increase computing capability by a thousandfold in just over 4 years. Science applications would run at a petaflop, or a quadrillion calculations, per second. This speed and power in solving problems was necessary to achieve DOE's missions, contribute to the nation's security, and improve standards of living and health. With a petascale system, calculations that once took months would take only minutes.

DOE chose ORNL to deploy the nation's first leadership computing facility. The OLCF delivered a series of upgrades to a Cray XT3

1992

The High Performance Computing and Communications Act authorizes the Department of Energy to deploy an HPC research center.

1993

A 35-gigaflop Intel Paragon explores how solids melt, cars crash, groundwater transports pollutants, and more.

1996

Awards abound for advancements in concurrency, simulation, and networking.

called Jaguar with staged increases to 26, 54, 119, and then 263 teraflops. In 2008 a new component was added to Jaguar, a 1.382-petaflop XT5 system. Jaguar became the first system in the world to run a scientific application at a sustained petaflop.

Because OLCF users had a consistent Cray XT architecture through the years, scientists were able to quickly use the petascale system. The first petascale production runs dealt with materials science, climate, combustion, fusion, and seismology. Another key to the facility's success is the scientific liaisons, who understand the leadership computers as well as the science areas. In fact, other centers have adopted the OLCF liaison model.

As a national user facility, the OLCF is made available to scientists and engineers in industry, academia, and government through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. INCITE supports high-impact, time-critical, computationally intensive research with allocations of millions of processor hours. Major advancements in computational science used OLCF supercomputers to gain extraordinary insights into supernovas, dark matter, fuel combustion, nuclear fusion, superconductivity, mathematics, and more. In 2010 the OLCF will deliver more than one billion processor hours on Jaguar to users.

Early computers at the CCS were used by researchers for the analysis of materials performance in automotive applications. Computer modeling of materials and vehicles helps to yield solutions in less time and at a lower cost than prototypes do.

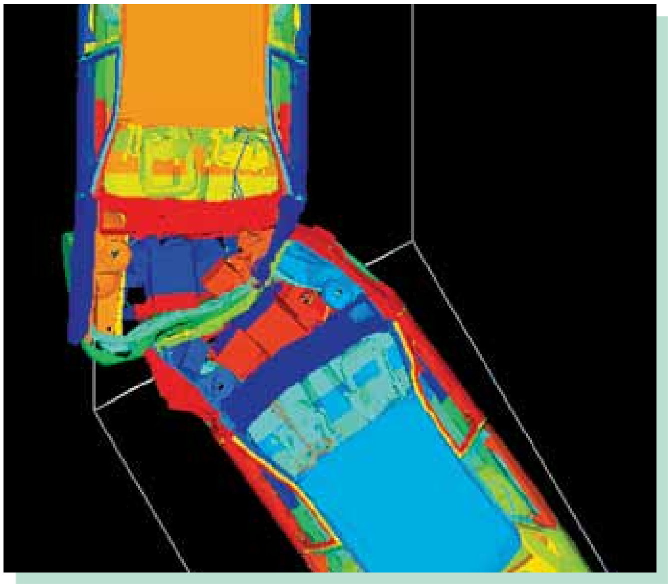

In November 2009 the XT5 component was further upgraded, to 2.332 petaflops, making it the world's most powerful supercomputer. It sports 300 terabytes of memory—three times more than any other computer. With more than 10 petabytes of file-system storage capacity, it can hold 1,000 times as much data as contained in the printed collection of the Library of Congress. Data moves in and out of its almost 225,000 processing cores at 240 gigabytes per second. Moreover, the OLCF houses leading-edge systems for storage, networking, and visualization. These computing resources serve the nation in exploring areas of strategic importance, such as battery technology, biofuels, fusion energy, and climate science.

At the 1992 inception of the CCS, ORNL had a staff of approximately 50 employees working in HPC (i.e., research to optimize the performance of supercomputers) or computational science (i.e., research in domains, such as chemistry or medicine, that employ supercomputers). Today the OLCF employs 71 professionals and ORNL has an additional 430 workers in HPC or computational sciences. This expertise has been critical to attracting research partners worldwide.

Other petascale supercomputers give ORNL the critical mass to make it the world's most powerful computing complex. In 2009 Kraken, a Cray XT5 supercomputer run by the University of Tennessee and ORNL and funded by the National Science Foundation (NSF), became the first petaflop computer system for academic use. A third petascale supercomputer is scheduled for deployment at ORNL in 2010.

Scientific discovery requires advances to software as well as hardware. In 2001 DOE launched the Scientific Discovery through

1998

An ORNL team wins the Gordon Bell Prize for the first science application to sustain a teraflop.

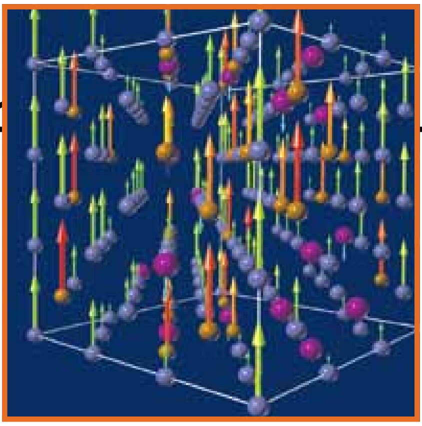

2000

In June, Eagle, an IBM Power3, becomes the 11th fastest computer in the world.

2001

DOE launches SciDAC. Cheetah, an IBM Power4 supercomputer, is installed and later runs some of the world's largest climate simulations.

Advanced Computing, or SciDAC, program to create a new generation of scalable scientific software, applications, middleware, libraries, and mathematical routines needed to make the most of parallel computing systems as they evolved from terascale to petascale and beyond. The SciDAC program is still going strong, scaling applications and tools to hundreds of thousands and even millions of threads of execution.

Supercomputing speeds science

"The nation that leads the clean energy economy will be the nation that leads the global economy," President Obama told Americans in his January 2010 State of the Union address. "America must be that nation." He envisioned continued investments in advanced technologies to make the cleanest use of carbon-based fuels, next-generation nuclear power plants as a step to move beyond carbon, and incentives to make clean energy profitable.

"As we address the important national problems that President Obama talked about in his State of the Union address, we know we're going to need much more computing capability," Bland said. "The appetite for high-performance computing from the science community is ever growing to find solutions to the difficult problems in energy, the environment, and national security."

Of immediate importance is obtaining a better understanding of what is happening to our planet's climate. Warren Washington of the National Center for Atmospheric Research leads a large community of scientists running the next-generation Community Climate System Model at the OLCF. The researchers aim to use the Climate-Science Computational End Station for predicting how climate might change if various energy policies were to be adopted, leading to different emission scenarios. Studies enabled by Jaguar will improve the level of certainty for future conclusions of the Intergovernmental Panel on Climate Change.

In 2006, ORNL's Cray XT3 and XT4 systems were combined and upgraded to quad-core processors. When the upgrades were complete, the new and improved Jaguar delivered a peak processing speed of 119 teraflops and 46 terabytes of aggregate memory.

Other OLCF research focuses on improving the efficiency of conventional fuels and the vehicles they power. For example, fuel efficiency could improve by at least 25 percent if next-generation vehicles burn lean fuel mixtures in low-temperature, compression-ignition engines. Toward this goal, Jacqueline Chen of Sandia National Laboratories led combustion simulations on Jaguar that generated 120 terabytes of data about ignition and stabilization of flames from burning fuels. Engineers use the detailed information in this data library to advance designs for fuel-efficient vehicle engines and industrial boilers.

Energy-assurance efforts also include a project to squeeze more energy out of biomass. Jeremy Smith of the University of Tennessee and ORNL leads an OLCF project to simulate cellulose, a complex

2003

The first cabinet is installed for Phoenix, a Cray X1 vector system venerated by researchers in computational fluid dynamics, climate, fusion, astrophysics, and materials science.

2003

ORNL completes construction of America's most advanced facility for unclassified scientific computing. Supernova shock wave instability is discovered.

2004

In May DOE announces the CCS will lead the OLCF project to build the world's most powerful supercomputer.

carbohydrate that gives leaves, stalks, and stems their rigidity. Its sugar subunits can be fermented to produce ethanol, enabling full use of plants for fuel. Revealing the structure, motion, and mechanics of cellulose in unmatched detail may lead to bioengineering of cell walls that are less resistant to decomposition and enzymes that are more efficient at breaking down cellulose.

Moreover, researchers use Jaguar to speed and optimize the design of advanced power plants to tap coal's oxidative potential and trap its pollutants. Instead of building an expensive prototype plant and seeing how effectively it burns coal, engineers can model the prototype, simulate its performance, and adjust operating conditions in hundreds of virtual experiments to learn what works best before they build.

ORNL's Tom Evans runs simulations on Jaguar to help design next-generation power plants employing nuclear fission. Improvements could dramatically shrink the volume of waste from such plants. Farther in the future, nuclear fusion, the process that powers the sun, could be commercially harnessed to provide clean, abundant energy, and numerous projects at the OLCF support that goal.

Petascale systems are invaluable in investigating nanomaterials for energy efficiency. A General Motors project uses Jaguar to simulate the atomic arrangement of a thermoelectric material. Pushing electrons to the side, the material turns the waste heat in a tailpipe into an electric generator. The extra power could help run the electric motors of hybrid vehicles, saving hundreds of millions of gallons of fuel annually.

Likewise, OLCF supercomputers aid development of electrochemical devices that surpass conventional batteries in performance; catalysts that improve industrial processes; magnets that drive motors in electric vehicles; and strong, lightweight materials for vehicles requiring less fuel to move. Scientists employ HPC to improve understanding of high-temperature superconductors, which transmit energy with no resistance and may someday lessen losses in transmission lines. In addition, supercomputers help characterize the atomic properties of materials that absorb light and transmit electrons—a key to improving the efficiency of solar panels.

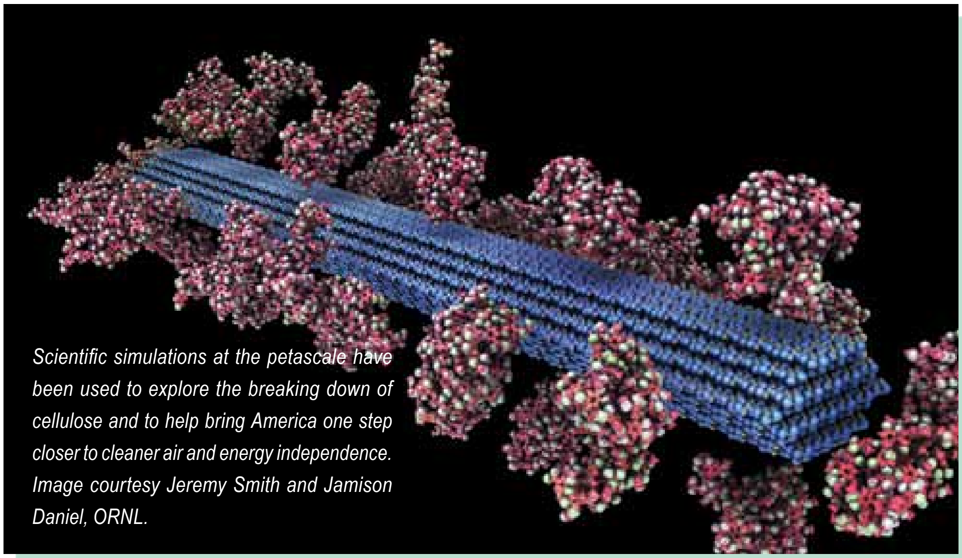

Scientific simulations at the petascale have been used to explore the breaking down of cellulose and to help bring America one step closer to cleaner air and energy independence. Image courtesy Jeremy Smith and Jamison Daniel, ORNL.

2005

The OLCF comes into full production, with time allocations on Jaguar, a Cray XT3, and Phoenix, a Cray X1. The OLCF holds its first users meeting.

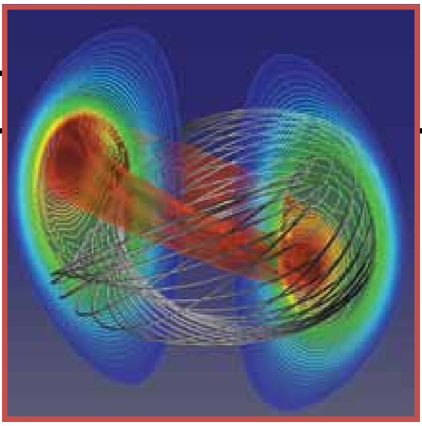

2006

The Jaguar XT3 is upgraded to 119 teraflops.

2008

Jaguar breaks the petaflop barrier with an upgrade to 1.38 petaflops and lands the #2 spot on the TOP500 list of the world's fastest supercomputers.

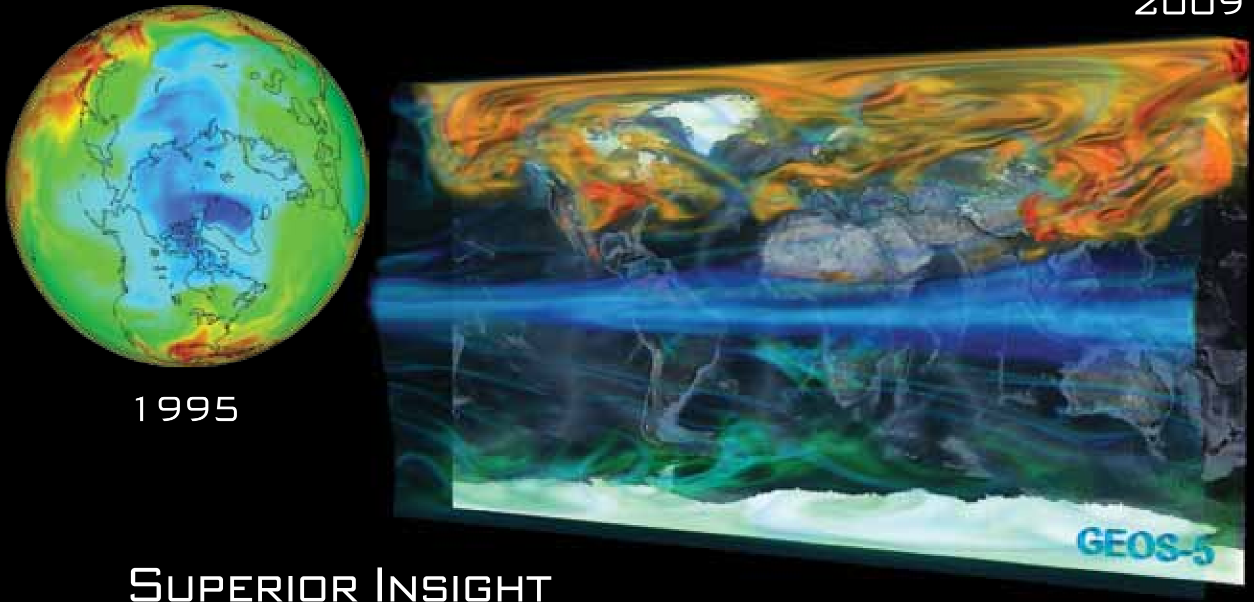

SUPERIOR INSIGHT

Today's state-of-the-art models and simulations provide scientists with unparalleled insights into the world's most daunting challenges, such as understanding the global carbon cycle. At right, a NASA model shows atmospheric carbon dioxide is higher in regions of anthropogenic emissions. Atmospheric carbon dioxide was simulated using the NASA Global Earth Observing System (GEOS-5) model running on Jaguar at the Oak Ridge Leadership Computing Facility. This research, a joint climate science project between David Erickson of ORNL and Steven Pawson of NASA's Global Modeling and Assimilation Office, was supported by the DOE Office of Science and the NASA Carbon Data Assimilation project based at NASA Goddard Space Flight Center. Image credit: Jamison Daniel, Oak Ridge Leadership Computing Facility. The image at left, in contrast, shows the state of the art in 1995, as depicted by a project using models to simulate a century of global climate. Image credit: Argonne and Oak Ridge national laboratories and the National Center for Atmospheric Research.

2009

Recovery Act funds enable an upgrade of the Jaguar XT5 from four-core to six-core processors. The XT5 runs at 1.759 petaflops and becomes #1 on the TOP500 list.

TOP500®
SUPERCOMPUTER SITES

#1

2010 and beyond

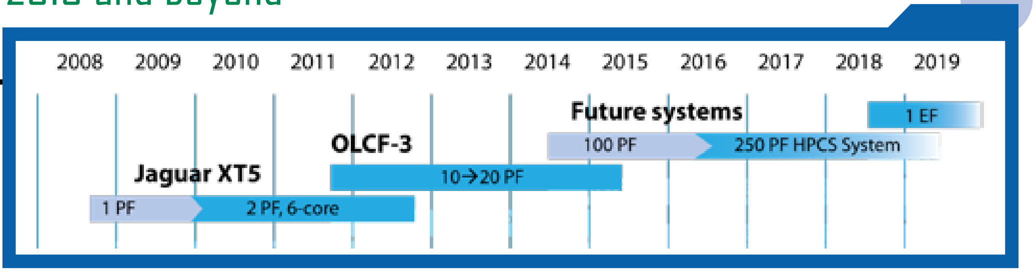

Advancing to the exascale

In 2008 OLCF science and technology teams laid out a roadmap to achieve exascale computing by 2018. The first step in the roadmap is the design and deployment of a 20-petaflop DOE system in 2012, followed by a series of ever larger systems, an approach similar to that used to successfully reach the petascale. "Our plan is to deliver a 100-petaflop machine by 2015, upgrade that to 250 petaflops, and by 2018 to deliver an exaflop computer system," Bland said.

Exascale HPC systems will be game changers in future scientific discoveries, allowing explorations at complexities and speeds not possible with today's systems. They will provide insights into the complex biogeochemical cycles that underpin global ecosystems and control the sustainability of life on Earth. Exascale systems will allow granularity fine enough to resolve cloud systems, validate models, and perhaps enable climate prediction over regions and decades—scales of interest to managers of infrastructures supplying water, food, and electricity.

Helping ensure energy for transportation and buildings, exascale computing will enable forefront research in areas from combustion of advanced fuels to catalysis of chemicals. It will help researchers improve efficiency, storage, transmission, and production of energy from fossil-fuel, nuclear-fission, nuclear-fusion, biomass, wind, solar, and hydroelectric sources.

Exascale supercomputers will accelerate the design, characterization, and manufacture of nanoscale materials for diverse applications. They will help build fundamental knowledge of how stars collapse, proteins fold, or elements of an ecosystem interact. Reducing the time to deployment of rationally designed drugs and economically viable industrial processes will be another exascale boon.

The OLCF's success as a world-class resource for discovery continues to attract research collaborators from other agencies, including NSF, NOAA, and the Department of Defense. Computational scientists and HPC professionals at the OLCF look forward to working together on grand challenges too tough for any one organization to solve.

As OLCF visionary Zacharia told ORNL computing employees in November 2009, after Jaguar had earned the number one spot on the TOP500 list of the world's fastest supercomputers, "When you're at the top of your game, that's when you have to lead and that's when you collaborate. Collaborate broadly."

There has never been a better time for computational science. The role of computation in accelerating breakthroughs is recognized so clearly that it has joined theory and experiment as the third pillar of scientific practice. Using models and simulations, OLCF users have achieved a greater understanding of the natural world and the substances we create to improve our lives. Scientists have farther to go in pushing past the frontiers of the unknown to arrive at important milestones and destinations. They will get there only if computational resources keep pace.—*by Dawn Levy*

The petascale Jaguar system is a dream realized. The world's most powerful supercomputer spent 2009 tackling previously intractable scientific problems while shining a powerful, new light on the world's most pressing challenges and the universe's most elusive mysteries.

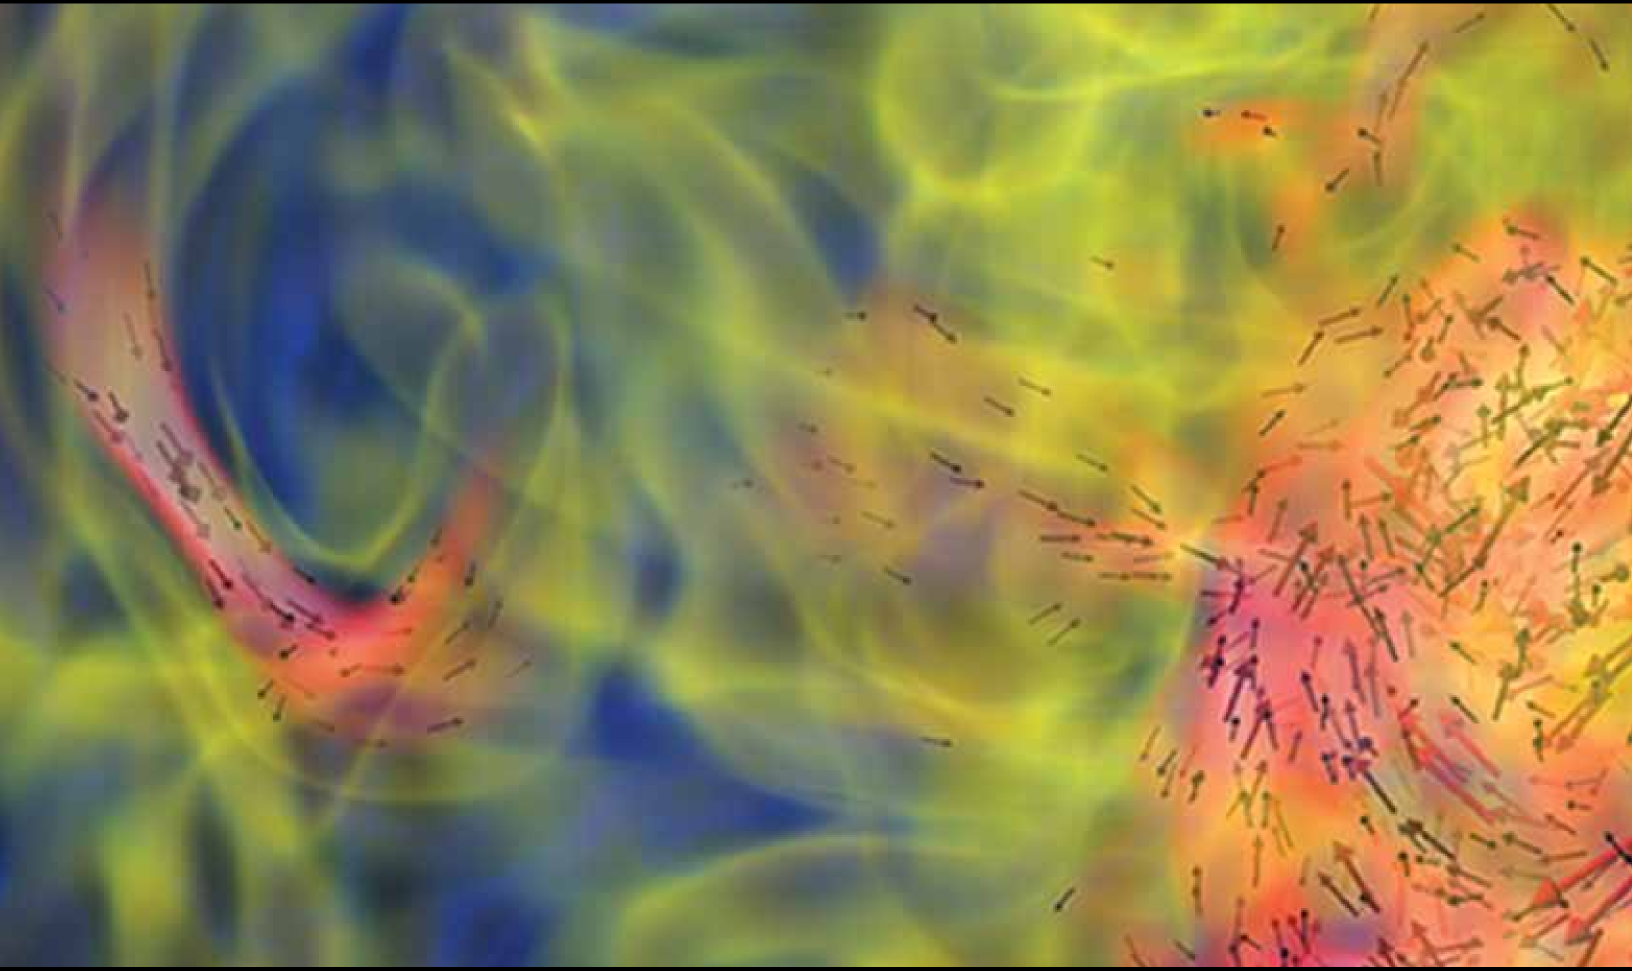

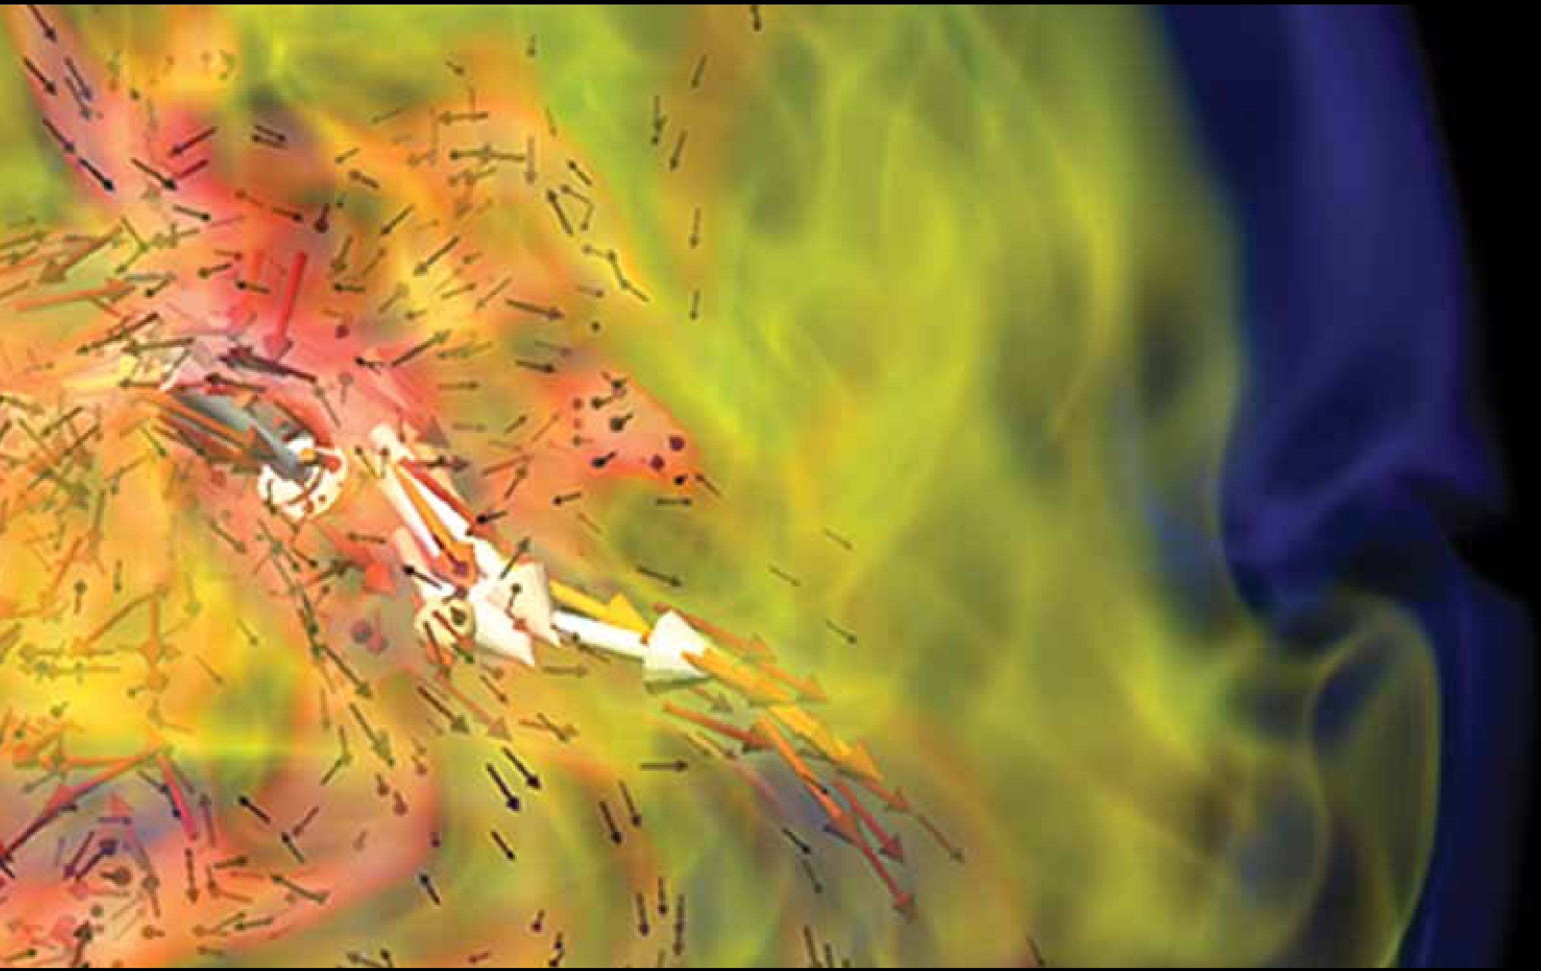

BREAKTHROUGH SCIENCE

Doug Kothe,
director of science

2009:

A YEAR OF PETASCALE SCIENCE

The year 2009 was, without a doubt, the most successful 365 days in the OLCF's six-year history. The center's numerous accolades, including Jaguar's ascension to the number one spot on the TOP500 List, were eclipsed only by the scientific accomplishments enabled by the new petascale Jaguar.

Last year revealed the promise of petascale computing—unprecedented resolution and physical fidelity that will help us to solve some of our most daunting energy challenges while keeping America scientifically competitive.

Two programs in particular highlight the potential of simulation at the petascale: The Petascale Early Science program, designed to test Jaguar's mettle early, and the INCITE program, designed to host a broad community of users across domains.

Twenty-eight research teams participated in the six-month Petascale Early Science program during the first half of 2009, using more than 360 million combined processor hours. These projects tackled some of the most pressing issues in climate science, chemistry, materials science, nuclear energy, physics, bioenergy, astrophysics, geosciences, fusion, and combustion.

The OLCF's primary goals for this period were threefold: to deliver important, high-impact science results and advancements; to harden the new Jaguar system for production; and to embrace a broad user community capable of using the system. Several research teams using Jaguar during this early petascale phase ran the largest calculations ever performed in their respective fields, and three codes achieved sustained performance of more than one petaflop during production simulations.

Petascale early science projects in climate included models of unprecedented resolution. One project carried out century-long simulations with a coupled atmospheric–oceanic climate model to explore the possibility of predicting climate on a decadal scale. The new, ultra-high resolution also allowed researchers to study fine-scale phenomena such as ocean eddies and storm patterns and determine the effect of a warming climate on tropical storm frequency and intensity.

The list of accomplishments also included the largest known simulations of a nuclear reactor. This work provided researchers with an accurate depiction of the distribution of power within a modern fission reactor, paving the way for the design of next-generation nuclear power devices.

Other projects included studies of high-temperature superconductors, cost-effective means for producing cellulosic ethanol from biomass, and calculations of the influx of uranium pollution from the Hanford 300 waste disposal site into the Columbia River basin.

With powerful applications and huge computational allocations on Jaguar, these researchers were able to punch through long-standing barriers and create new knowledge. We are confident that their achievements will stand over time as substantial contributions to science and technology.

The INCITE program, jointly managed by the Argonne and Oak Ridge leadership computing facilities, likewise produced plenty worth bragging about. Open to researchers from academia, government labs, and industry, INCITE is the major means by which the scientific community gains access to Jaguar. The program aims to accelerate scientific discoveries and technological innovations by awarding time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering.

For example, one INCITE allocation ran some of the biggest combustion simulations to date, studying the burning of various fuels and creating data libraries that will assist engineers in the development of models used to design the engines and boilers of tomorrow. The engine research could lead to a significant increase in national fuel efficiency if employed in next-generation autos.

Another project sought to determine if carbon sequestration, or storing carbon dioxide from sources such as smokestacks underground, is an effective way to keep CO_2 out of the atmosphere.

Also worth mentioning is the GEOS-5 project, a collaboration between NASA and the National Oceanic and Atmospheric Administration (NOAA) that tackled mesoscale climate simulations on a 14-kilometer grid and produced the highest-ever-resolution simulations of cloud systems, an important driver that has so far been grossly approximated in most climate simulations. As the world's fastest computing resource, Jaguar is well suited to handle the complexities and

Climate

Energy

High-Energy Physics

computational demands of the world's most advanced climate simulations, allowing GEOS-5 to provide researchers with increasingly accurate climate scenarios.

Other noteworthy projects included the mechanism behind core-collapse supernovas, stellar explosions that litter the universe with the elements necessary for life as we know it, and several fusion projects assisting with the design of ITER, a cooperative prototype fusion reactor being built to test fusion's potential for commercial power production.

Both the Petascale Early Science period and INCITE represent the primary source of computing cycles for the DOE's SciDAC program. The program was founded in 2001 to create a new generation of scalable scientific software, applications, middleware, libraries, and mathematical routines needed to make the most of parallel computing systems as they evolved from terascale to petascale and beyond. Realizing the promise of SciDAC—real scientific discovery through computation—relies on the availability and effective use of leadership computing, and the OLCF is the ideal proving ground for the techniques, implementations, and applications developed via the program. This synergistic relationship was made manifest on numerous occasions in 2009.

And 2010 looks equally promising. With Jaguar's two-plus petaflops in full swing, the system will continue to host three applications that all exceed a performance of more than one petaflop: WL-LSMS, the 2009 Gordon Bell Prize winner; NWChem, a 2009 Gordon Bell finalist; and DCA++, the 2008 Gordon Bell Prize winner.

WL-LSMS achieved a performance of 1.84 petaflops to become the fastest application in history and take the Gordon Bell in 2009. The application allows researchers to directly and accurately calculate the temperature above which a material loses its magnetism—known as the Curie temperature. Areas that stand to benefit from the research include the design of lighter, more resilient steel and the development of future refrigerators that use magnetic cooling.

NWChem, the 2009 Gordon Bell runner-up, clocked in at 1.39 petaflops in a first principles, quantum mechanical exploration of the energy contained in clusters of water molecules.

And the 2008 winner, DCA++, will continue to harness Jaguar's petascale power to simulate high-temperature superconductors, technology that could revolutionize power transmission and potentially several other arenas. The application ran at 1.35 petaflops in 2008 and 2.24 petaflops in 2009, helping researchers understand the underpinnings of superconducting materials that ultimately may prove important to electrical transmission.

These three leading applications aren't the only testament to the potential of computational science at the OLCF, however. Perhaps the most promising advances will be made in climate—few areas of science have ever been as pressing, and there is little debate surrounding its current importance. Prizes and petaflops aside, the 2010 climate projects on Jaguar promise to truly advance understanding.

For example, a climate end-station project will run the world's largest dataset for climate research. This dataset will be accessed directly by researchers working on the United Nations' Intergovernmental Panel on Climate Change report, the seminal document in climate change legislation. And Jaguar will host the highest-ever-resolution coupled simulation of the Community Climate System Model (CCSM). This enhanced CCSM project will study all four of the model's variables (atmosphere, ocean, ice, and land) at an equal resolution, a feat possible only with Jaguar's 2-plus petaflops of raw computing power.

Other equally impressive projects in the climate realm will explore the evolution of cloud systems and the effects of hurricanes and study past abrupt climate change to see if similar events could arise in our near future.

Few areas of science are as vital to the DOE mission as climate, and with the help of Jaguar the OLCF is better preparing policymakers to predict and deal with the ramifications of climate change.

Ultimately, climate research is an excellent example of the value of high-performance computing. While HPC is indeed still evolving—with researchers plotting the path to the exascale—2009 will stand out as a pivotal year in which supercomputing truly entered the petascale, with never-before-seen properties of nature revealed via powerful simulation. We hope to build on last year's success and make 2010 even more groundbreaking, pushing the limits of Jaguar and our knowledge of the world around us.

Resolving Global Climate Change

Jaguar helps coupled models explore limits of forecasting the statistical behavior of the atmosphere in a warming world

In 1969 scientists from the Geophysical Fluid Dynamics Laboratory (GFDL) in Princeton, New Jersey, published results from the world's first climate model. Though the model gave scientists their first look at how the ocean and atmosphere interact to influence climate, it covered only one-sixth of the surface of the Earth and did not take into account any human-made, or anthropogenic, causes of climate change.

Forty years later, using an ancestor of this pioneering model, a team led by GFDL climate scientist Venkatramani Balaji is using the Cray XT5 known as Jaguar—housed in the OLCF—to simulate and assess both natural and anthropogenic causes of climate change at exceptional resolutions.

"Resolving our models to 50 kilometers, you can see states and counties. That's what people are really asking for [from climate science]," said Balaji, head of the Modeling Systems Group at GFDL, a branch of NOAA devoted to developing and using mathematical models and computer simulations to improve understanding of the Earth's atmosphere, ocean, and climate. The group's models are higher in resolution than required by the Intergovernmental Panel on Climate Change Fourth Assessment Report (IPCC AR4).

Climate models from which data was gleaned for the IPCC AR4, released in 2007, were in the 100-kilometer range for both ocean and atmosphere. Accurate regional climate forecasting is invaluable to areas such as the American Southwest, which experiences persistent droughts that are rapidly depleting water reservoirs. If climate models can prove that anthropogenic causes are mostly to blame, stakeholders can amend policies and practices to make the most of scarce resources.

Causes of climate change

Natural variations in global climate arise from phenomena including changes in solar activity, periodic alterations in the Earth's orbit, and sulfate aerosols from volcanic eruptions. However, the AR4 claims, "It is extremely unlikely (less than 5 percent) that the global pattern of warming during the past half century can be explained without external forcing, and very unlikely (less than 10 percent) that it is due to known natural external causes alone."

Humanity's effect on climate is broad. Of most concern to climate scientists and environmentalists are increasing carbon dioxide levels in the Earth's atmosphere. According to the IPCC AR4, two-thirds of the anthropogenic carbon dioxide emissions come from fossil fuel combustion, while the remaining third comes from land-use changes such as irrigation, deforestation, and ranching. Carbon dioxide is naturally present in the atmosphere along with water vapor, nitrous oxide, methane, ozone, and other trace gases collectively referred to as greenhouse gases. As their name suggests, these gases warm the Earth by trapping energy from sunlight just as a greenhouse does. Ever since the Industrial Revolution, greenhouse-gas emissions have increased, warming the Earth and affecting regional weather patterns.

The terms "climate" and "weather" describe the atmosphere through different timeframes. Climate refers to average, or statistical, behavior of the global atmosphere over a long period of time—say, year to year—whereas weather is a measure of day-to-day local atmospheric conditions. Through cutting-edge models, Balaji's project, dubbed CHiMES (for Coupled High-Resolution Modeling of the Earth System), aims to explore the relationship between climate and weather.

CHiMES is a collaborative effort between DOE and NOAA. If climate can be predicted at all, the project endeavors to forecast change over time periods of interest to resource managers (e.g., decades) and to understand the responses of phenomena like tropical storms to a warming climate.

Anatomy of a climate model

Climate models are composed of mathematical sets of instructions called differential equations. These equations assess interactions among various components that influence the Earth's climate—atmosphere, oceans, land, and ice. In short, climate models calculate the solar energy our planet absorbs and the radiation it emits, some of which escapes into space while some remains trapped by Earth's atmosphere. Differential equations that govern the system are solved in accordance with different boundary conditions, such as specifying sea surface temperature distributions, carbon dioxide concentrations, and atmospheric aerosol distributions.

Balaji's model—GFDL Flexible Modeling System—is actually two models built upon a flexible framework that allows different components of the climate system to be modeled by multiple scientists and code developers and assembled in a variety of ways. The atmospheric model, Finite Volume Cubed-Sphere, runs at a resolution of 25 kilometers. The oceanic model MOM4 (short for Modular Ocean Model) is what climate scientists call a quarter-degree model, which translates into resolution of roughly 25 kilometers as well. The resolution of the models can be increased in specific areas, and

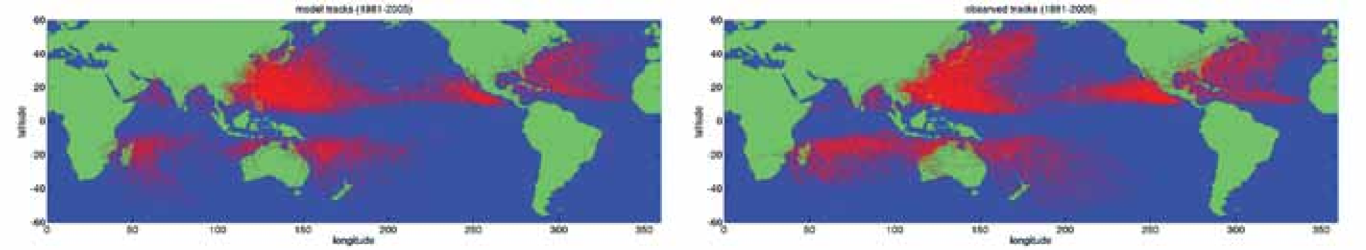

Simulation of global hurricane climatology and response to global warming in a new global High Resolution Atmospheric Model (HiRAM) at 50-kilometer resolution. Images courtesy: Ming Zhao, Isaac Held, Shian-Jian Lin, Gabe Vecchi, NOAA/Geophysical Fluid Dynamics Laboratory

Balaji and colleagues increased the granularity to 10-kilometer resolution in areas of particular interest.

The equations of any climate model are calculated at individual points on a grid covering the Earth. Traditionally, this grid has been dictated by lines of longitude and latitude in which each cell of the grid is defined as the intersection of a latitude line and a longitude line. Though easy to manage, a latitude/longitude-based grid creates problems for climate scientists. The distance between meridians (lines of longitude) decreases as they draw closer to the poles, requiring the dynamical algorithms to use exceedingly short time steps to keep the solutions stable in the polar regions. Although this stabilizes the calculation in the polar regions, it constrains the time step throughout the global domain, decreasing the overall efficiency of the calculation.

Balaji's team conquers this issue in its coupled model by using a cubed-sphere grid that projects a three-dimensional cubed grid onto the Earth, removing the clustered points around the poles and providing quasi-uniform resolution at each grid cell. Using a tripolar grid for the ocean model overcomes this "pole problem" by placing three "poles" over landmasses on the grid: one at the South Pole (as there is no ocean there), one over the Asian continent, and one over North America.

The atmospheric and oceanic models employ separate grids and solve different algorithms. The majority of the time, the models run independently but concurrently. A third grid is used to exchange data between the two models every 2 hours.

"Many other climate models only exchange data once a day or not at all, so anything that happens faster than a 24-hour timeframe will be missed," said Balaji. "Peaks in wind can change the ocean's circulation [in less than 24 hours], and we are able to resolve that."

Coupling climate models is computationally expensive, but the Jaguar XT5—the fastest supercomputer in the world at 2.33 quadrillion calculations per second and featuring nearly a quarter of a million processors—gave Balaji's team the power and speed it needed to frequently link its climate models.

Future forecasts

The most exciting results from the CHiMES project come from the team's study of tropical storm response to global climate change. In 2009 the team used 20 million processor hours on Jaguar to run approximately 500 years' worth of coupled-model simulations through an INCITE program allocation. Scaling its high-resolution models from 60,000 to 100,000 cores, Balaji's team was able to realistically duplicate year-to-year behavior of hurricanes, accurately simulating their seasonal peak in September.

The team is also resolving an issue that has plagued coupled climate models since their inception. The issue pertains to an area around the equator where winds originating in the northern and southern hemispheres meet, affecting the wet and dry seasons of many equatorial nations. The region is called the Intertropical Convergence Zone (ITCZ), and it appears in coupled climate models as two peaks in rainfall. In reality the ITCZ has only one peak in rainfall. As the team continues to increase the resolution of its models, it is slowly coming closer to eliminating the second "peak."

The team also hopes to use the ocean as a way to forecast climate on the scale of decades. The ocean influences the climate on a longer scale than do Earth's ice, land, and atmosphere. If models can capture the long-term behavior of the oceans, they can also capture the short-term behavior of Earth's other climate-influencing components.

While Balaji is unsure of whether his models are capable of forecasting climate on the decadal scale, the team has been awarded another 20 million hours on Jaguar through INCITE for 2010. Balaji's team plans to complete runs to determine if decadal predictability is possible, but it also hopes to bring even higher-resolution models online—ones capable of resolving fine-scale weather such as cloud convection and shifting winds.

Ultimately Balaji and his colleagues hope to demonstrate that global climate change is not a remote concept, but a tangible problem with significant consequences that the world's experts must solve together.

"Local changes will be far more intense; you might get warming here and cooling there," he explained. "The term global warming seems like a temperature problem, but precipitation patterns change—some areas will get wetter, some dryer. We have to make [climate change] more relevant to people by bringing it down to the regional scale. In order to help people adapt, we want to tell them what to be prepared for."—*by Caitlin Rockett*

JAGUAR EXPLORES Carbon–Water Union

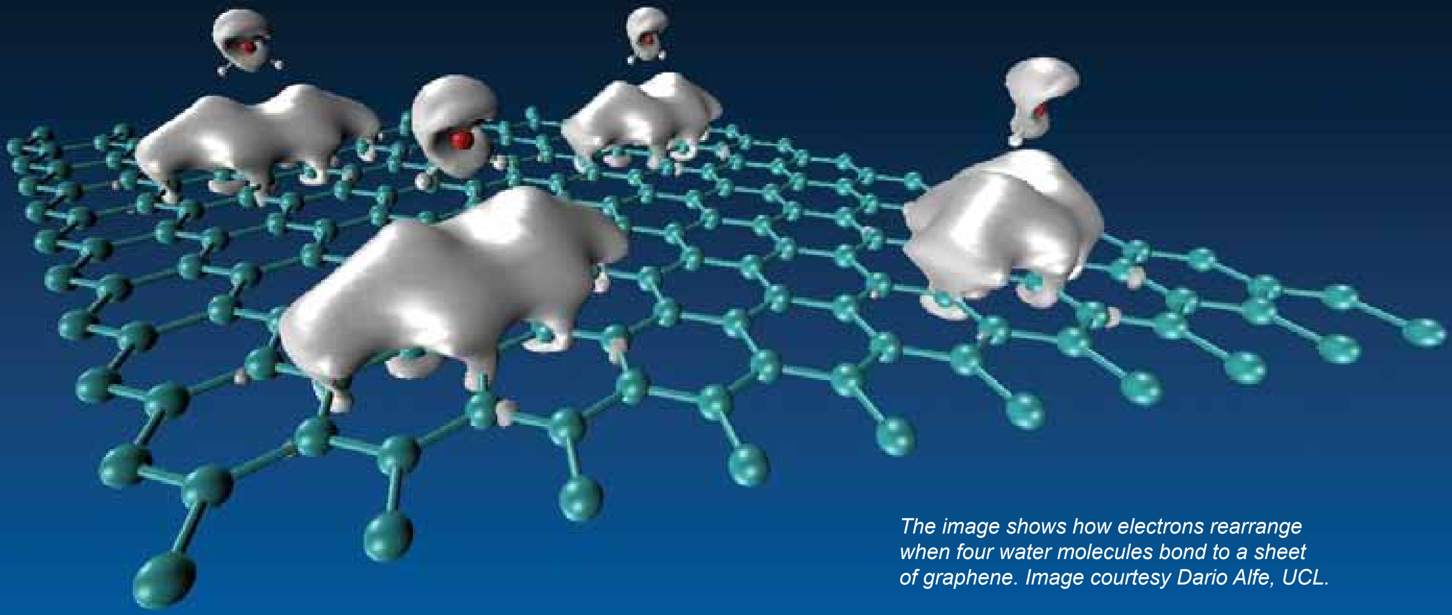

The image shows how electrons rearrange when four water molecules bond to a sheet of graphene. Image courtesy Dario Alfe, UCL.

Researchers hit jackpot with CASINO code

Some research has potential that's difficult to quantify. The ramifications of a better understanding of a particular phenomenon can sometimes reverberate throughout science, making everlasting waves in multiple arenas.

Take the seemingly trivial properties of a water molecule on graphene, or a single layer of graphite. Water is everywhere, as is carbon. Big deal, right? Actually, it is. Understanding the adsorption, or accumulation, of a water molecule on graphene could influence everything from hydrogen storage to astrophysics to corrosion to geology.

In 2009 a team made up of Dario Alfe, Angelos Michaelides, and Mike Gillan of University College London (UCL) used the OLCF's Early Petascale Science period to obtain a full-binding energy curve between water and graphene, research that will impact multiple domains and could only be performed on the world's fastest supercomputer. The project is a collaboration between ORNL and UCL and is supported by the Engineering and Physical Sciences Research Council of the United Kingdom.

Adsorption is the accumulation of atoms or molecules on the surface of a material. The process is possible because unlike the atoms in a material that bond with one another, surface atoms are not wholly surrounded by other atoms and can attract adsorbates, in this case a water molecule.

While the process may sound a bit strange, it turns out that this is the case with most natural surfaces—they are covered by a film of water. It is for precisely this reason that this research is so fundamental. "Water interacts with every surface that we see," said Alfe.

> "Jaguar is a fantastic machine, allowing us to do things we couldn't do without it."
> – Dario Alfe, Principal Investigator

Because of the large numbers of particles and multiple variables involved in the interaction, these systems are difficult to describe. However, science could greatly benefit from a more precise explanation of such systems simply because they occur so often in the natural world.

And as science benefits, so do we. For instance, a better understanding of this particular adsorption process could pave the way for more efficient hydrogen batteries, an integral part of a potential clean energy economy. Catalysis, the process of adding an agent to initiate a chemical change, is an important step in countless industrial processes—it could be streamlined to make products safer and cheaper. Furthermore, adsorption is key to understanding lubrication to reduce friction between moving parts, which plays a role in nearly every nook and cranny of our transportation sector.

And that's just energy. The research will likewise shed light on some of science's most fundamental questions, such as how the Earth forms materials at the extreme conditions deep in its interior, or how molecules form in the interstellar medium of outer space. And the methodological advances that result will lead to better predictions of the interaction of water with other materials such as clay particles and proteins, both of which are central to biological and environmental research programs.

Skinning Schrödinger's Cat

Alfe and his team received 2 million hours on Jaguar to calculate a water molecule's adsorption energy and geometry on a sheet of graphene. Using the quantum Monte-Carlo (QMC)–based CASINO code developed at Cambridge University, the team produced an unprecedented depiction of water-graphene adsorption.

By calculating the adsorption energy and geometry in 30 individual steps as the molecule moves closer to the surface, the team made history. "We've been able to obtain a full-binding energy curve between water and graphene," said Alfe. This breakthrough was only possible using CASINO's QMC-based methodology.

Using QMC methodology scientists attempt to solve the quantum mechanical Schrödinger equation, which describes the evolution of a quantum state over time, by simulating a mathematically equivalent random process, such as flipping a coin over and over to predict the chance of achieving heads. The more times one flips the coin, the more precise the prediction.

While there are several ways to solve the equation, QMC delivers the most accurate representation. It is computationally intensive, though, requiring the generation of random numbers across tens of thousands of compute nodes—a technique 1,000 to 10,000 times more expensive, but 10 times more accurate—than standard methods such as density functional theory (DFT).

These standard tools are "okay in giving you answers about the interaction energy between molecules and surfaces where the interactions are strong," said Alfe. If the interaction is weak one must look elsewhere for answers, hence the need for QMC. "QMC is more suited to Jaguar's architecture," said Alfe. Furthermore, he added, it is more befitting to leadership-class machines due to its complexity; and while it is currently more expensive, in the future it will be the more common method in conjunction with the advanced speed of modern supercomputers.

Fortunately, QMC and Jaguar make a good match. CASINO is very parallel in nature, much more so than its DFT counterparts. CASINO regularly runs on 10,000 to 20,000 cores with almost perfect scaling. In fact, the researchers recently scaled CASINO to 40,000 cores on Jaguar and envision additional improvements that will enable the code to scale to more than 100,000 cores, making it one of the premier codes in the world.

QMC was also relatively easy to port to Jaguar, said Alfe, primarily because the application was previously running on HECToR, a Cray XT4 located at the University of Edinburgh. Given the similar platforms, the researchers didn't even have to recompile their code. And from the looks of things, the team is pleased with Jaguar's performance.

"Jaguar is a fantastic machine," said Alfe, "allowing us to do things we couldn't do without it."

From our origins to our future, Alfe's work on Jaguar cuts a wide swath of discovery. Understanding nature's smallest, most complex systems requires the petascale power of mankind's most powerful computer, making Alfe's project and the OLCF's flagship system a perfect match.—*by Gregory Scott Jones*

Helping Superconductors Find the Path of Least Resistance

Researchers have made steady progress over three decades with the world's most advanced superconductors.

Composed of thin layers of copper and oxygen sandwiched between layers of insulating material, cuprates lend themselves to almost limitless manipulation and experimentation. Get the right combination and you've revolutionized broad swaths of modern life, from particle physics to transportation to medical diagnostics. Get the wrong combination and you've created a very expensive paperweight.

A team led by Thomas Schulthess of ORNL and the Swiss National Supercomputing Centre is applying the world's most powerful supercomputer to the nanoscale study of superconducting cuprates. Using ORNL's Cray XT5 Jaguar system, Schulthess and colleagues are working from the perspective of individual atoms and electrons to evaluate the effect of disorder among electrons in the system. The team has also verified the superconducting properties of one particular promising material composed of lanthanum, barium, copper, and oxygen (LaBaCuO for short), publishing its findings June 16, 2010, in the journal *Physical Review Letters*.

In the practical world of superconductors, "promising" means the material is superconducting in a slightly less frigid environment. The benefit of superconducting materials is that they are able to conduct electricity without loss, but they do so only when they are below a very, very low transition temperature. So far the highest transition temperature of a known superconducting material is only 135 degrees Kelvin, meaning the material can be no warmer than −216 degrees Fahrenheit.

The earliest discovered superconductors had to be kept below 20 degrees Kelvin (or −424 degrees Fahrenheit) and were cooled with liquid helium—a very expensive process. Later materials remained superconducting above liquid nitrogen's boiling point of 77 degrees Kelvin (−321 degrees Fahrenheit), making their use less expensive. The discovery of a room-temperature superconductor would revolutionize power transmission and make energy use ultraefficient.

Researchers develop new superconducting cuprates by varying their insulating layers, which in turn affects the conducting copper-oxygen layer both by removing electrons from it (with the locations where electrons used to be known as "holes") and by altering the local energies of remaining electrons. This manipulation potentially encourages the remaining electrons to join into the pairs, known as Cooper pairs, that carry a current without resistance.

"The material becomes superconducting if you pull out enough electrons from the planes," noted ORNL physicist Thomas Maier, a member of Schulthess's team. "You have to do so many substitutions, then the material becomes superconducting. It starts out superconducting at a very low temperature, but once you substitute more and more material in between there, your transition temperature goes up, reaches a maximum temperature, and then goes down again."

A material known as a dopant is added to the insulator until the cuprate reaches its highest possible transition temperature; this recipe is known as optimal doping. In most cuprates optimal doping pulls between 15 and 20 percent of the electrons from the conducting copper-oxygen layer.

Schulthess's group has been studying the arrangement of electron holes in a variety of cuprates. Generally, the holes are scattered more or less randomly throughout the conducting copper-oxygen layer. However, when barium is substituted for one-eighth of the lanthanum in a lanthanum-copper-oxygen cuprate (creating an optimal LaBaCuO cuprate), the holes line up in regular stripes. While researchers had known for some time about this striping, until very recently they believed that it would not help them meet their goal of higher-temperature superconductors.

"Early experiments showed striping at one-eighth doping and saw that transition temperature went way down," Maier explained, "and they concluded that striping was bad.

"That was the common wisdom until recently, when two independent experiments—one performed at Brookhaven National Laboratory, the other at Brookhaven, Lawrence Berkeley National Laboratory,

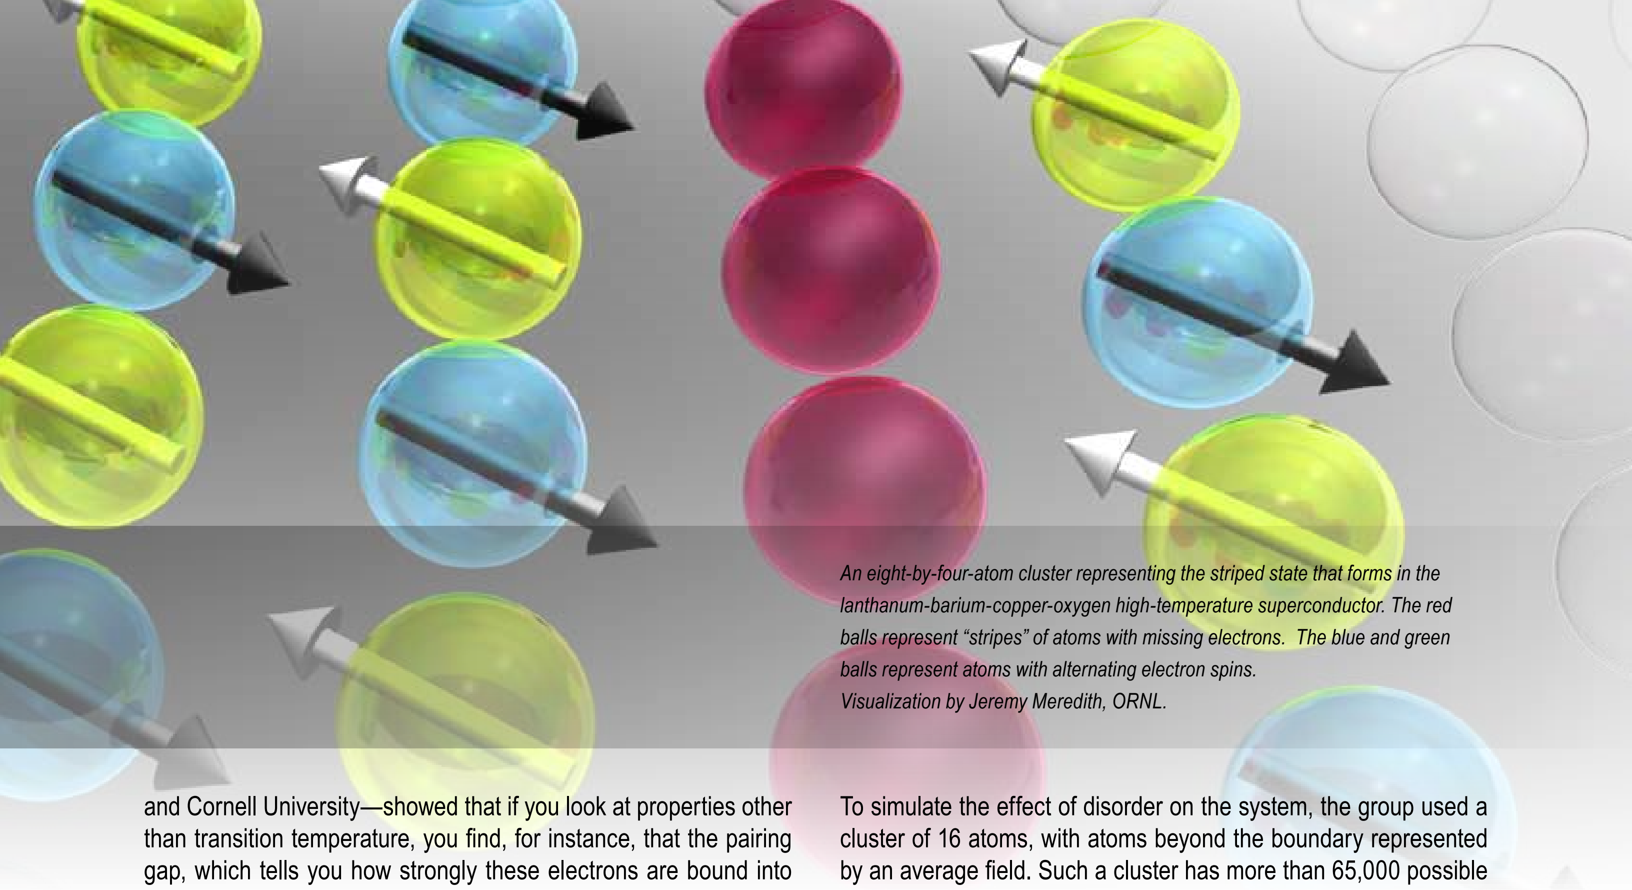

An eight-by-four-atom cluster representing the striped state that forms in the lanthanum-barium-copper-oxygen high-temperature superconductor. The red balls represent "stripes" of atoms with missing electrons. The blue and green balls represent atoms with alternating electron spins.
Visualization by Jeremy Meredith, ORNL.

and Cornell University—showed that if you look at properties other than transition temperature, you find, for instance, that the pairing gap, which tells you how strongly these electrons are bound into Cooper pairs, is highest at this doping."

Maier said more recent experiments have shown that resistance in this material drops at a much higher temperature (a good result) if you are able to look at single layers of atoms. It was this result that Schulthess's team was able to verify computationally using a simplified model known as the two-dimensional Hubbard model.

In other words the superconducting performance of LaBaCuO is great in two dimensions but rotten in three, noted Maier, possibly because the stripes of adjacent layers don't line up and the layers are electronically decoupled.

Up to this point, insight into the material's stellar performance in two dimensions might have been nothing more than a fascinating curiosity with no practical value. After all, the real world takes place in three dimensions. However, as Maier added, manufacturing technology has advanced to the point at which we are able to lay down layers so thin they are essentially two dimensional.

"It's important from an application point of view," he said, "because with recent advances in quantum engineering, one might be able to structure a system like this in the near future. It's not possible yet, but what can be done is to grow these systems layer by layer—you grow them layer by layer, and you can tune the properties layer by layer. You can have different layers with different doping."

Schulthess's team performed the simulations on ORNL's XT5 Jaguar system, as of November 2009 the fastest supercomputer in the world. The system sports 224,256 processing cores, features 300 terabytes of system memory, and operates at a peak performance of 2.3 thousand trillion calculations a second (or 2.3 petaflops). The team solves the two-dimensional Hubbard model with an application called DCA++, which earned Schulthess and colleagues the 2008 Gordon Bell Prize for the world's fastest scientific computing application.

To simulate the effect of disorder on the system, the group used a cluster of 16 atoms, with atoms beyond the boundary represented by an average field. Such a cluster has more than 65,000 possible disorder configurations (i.e., 2^{16}), but the team chose 128 arrangements as a representative sample.

Maier noted that the mean field used to represent the material outside the 16-atom cluster changes as the simulation continues, with results from each iteration used to adjust and improve the field. To simulate the striped state of LaBaCuO, he said, the team had to simulate a larger cluster of atoms—in this case 32—but was required to include only one configuration.

The team performed these simulations as part of the Jaguar system's Early Petascale Science program. Maier said DCA++ ran with nearly the same performance as it had during the Gordon Bell competition and managed about 85 percent efficiency on about 65,000 processing cores.

In the future, he said, the team will apply its calculations to iron-based high-temperature superconductors, discovered in 2008. Maier said these superconductors will be far more challenging to simulate than cuprate ones. While the superconducting properties of a cuprate superconductor rely on only a single electron orbit, this is not the case with the iron-based superconductors.

"To some extent we've been spoiled in theory with the cuprates," he said. "In the cuprates there is really only one orbital where the action is going on. That makes it easy for us, because in the models we look at we only have to take into account one orbital. In these new superconductors, the iron pnictides, that's not the case. There are essentially five orbitals of the iron that are all active. It makes it a lot harder."—*by Leo Williams*

COMPUTING A NUCLEAR RENAISSANCE

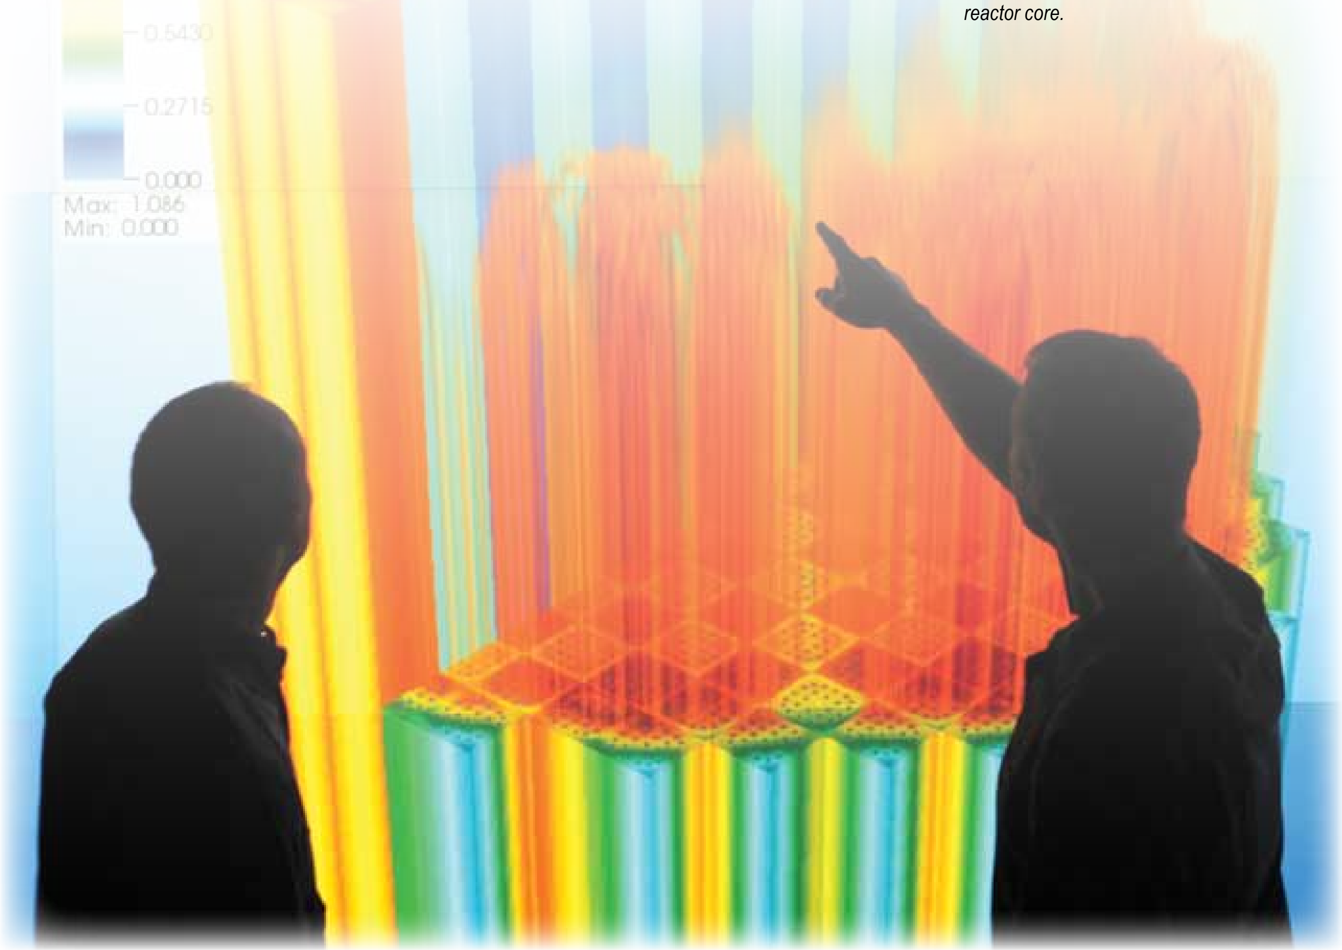

ORNL's John Turner and Tom Evans view Evans's visualization of power distribution in a Westinghouse PWR900 pressurized water reactor core.

The nuclear power industry appears to be at the brink of a resurgence, driven by uneasiness over America's reliance on foreign sources of energy and concern over the role of fossil fuels in climate change.

Funding for research is up, as are the number of job openings for nuclear scientists and engineers and enrollment in university programs. Even President Obama foresees a nuclear renaissance.

"But to create more of these clean energy jobs," he said in his 2010 State of the Union address, "we need more production, more efficiency, more incentives. And that means building a new generation of safe, clean nuclear power plants in this country."

It's good news for an industry that's been in a long stall. Nuclear power has been hampered for decades by skyrocketing costs and public uneasiness. The last construction permit for a commercial nuclear power plant in the United States was issued in 1978, and the accident at the Three Mile Island plant the following year seemed to drain whatever enthusiasm the country had for nuclear power.

As a result, America's nuclear power industry has been hibernating, and the power plants on which we now rely for nearly 20 percent of our electricity are old. If nuclear power is to play a major role in supplying the secure, Earth-friendly energy we need, existing plants must be made to operate more efficiently and new plant designs must be created that inspire public confidence and avoid unnecessary costs.

Fortunately, the tools we have for creating such designs are massively more powerful than the tools available in nuclear power's heyday, forty-plus years ago. The most important of these may be supercomputer simulation, which is now able to replicate the operation of nuclear reactors with a detail and accuracy that was almost unimaginable in those early days.

One of the most promising of these nuclear engineering applications was created by ORNL's Thomas Evans. Known as Denovo, this relatively young application promises to help bring nuclear power out of its long slump and into the forefront of 21st century computational science.

"That industry was in the vanguard of engineering in the 1960s and '70s," Evans explained, "and it stalled. A lot of the process that drives innovation kind of stopped in 1977. Since then it's been piecemeal—what can we do with what we've got—because nobody was giving us any big investments."

In technical terms, Denovo is a discrete ordinates linear radiation transport code. This means it simulates any radiation that does not interact strongly with the source of that radiation, whether it be photons, electrons, or the neutrons responsible for nuclear fission. The application solves the Boltzmann equation, which describes the statistical distribution of individual particles in a fluid.

In its current work Denovo is focusing on the neutrons that split atoms in a nuclear reactor. The application was among the first to get access to ORNL's Jaguar supercomputer when it was upgraded to the petascale. Taking advantage of the Early Petascale Science program, Denovo proved its ability to run on the world's most powerful scientific supercomputer by using as many as 60,000 processor cores at a time to run problems with as many as 280 trillion unknowns.

Evans used this period to simulate the workings of nuclear reactors and their shielding with unprecedented detail. He and colleagues have produced two major papers detailing their experience. "A Transport Acceleration Scheme for Multigroup Discrete Ordinates with Upscattering" is scheduled to appear in July 2010 in *Nuclear Science and Engineering.* "Denovo—A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE" is scheduled to appear in August 2010 in *Nuclear Technology*.

Denovo has been able to model nuclear reactor assemblies in quarter-inch-high slices with a radius of only one thousandth of an inch—a resolution several orders of magnitude finer than previous simulations. Evans believes such performance will help bring down the cost of new reactors by moving the heavy design work onto supercomputers and away from expensive prototypes.

"A lot of those early reactors were designed with 2-D calculations, and they use a lot of engineering knowledge and know-how to predict how it's going to extrapolate to the real system. Then they built prototypes to take measurements and validate the process, and so on.

"In engineering analysis, the goal isn't necessarily to remove prototyping and experiments. The idea is to reduce the cost."

Denovo was first conceived as a tool to evaluate the reliability of radiation shielding. For example, a reactor operator might want to know the risk to people nearby if a reactor's containment building were hit by a tornado or an airplane. Evans has also applied Denovo to shielding problems for the planned ITER fusion reactor, a multinational prototype located in France and scheduled to begin operation in 2018.

Evans said Denovo's shielding calculations, like its reactor simulations, are far more detailed than previous efforts.

"We're just starting to do transport problems in the nuclear reactor community of this size," he said, noting that the simulation with 280 trillion unknowns was a shielding analysis. "When we were first approached about doing these kinds of problems for a pressurized water reactor [the most common design for American nuclear power plants], they didn't think it could be done.

"What people have done in the past is run Monte Carlo calculations of problems of that size. It takes weeks of run time, and they get the answer at one point. So, for example, they set up the thing, and they say, 'OK, the guy's standing next to the turbine, we can get you an answer there.' But they don't have an answer anywhere else. With Denovo we have a true global solution."

Denovo has at least two attributes that set it apart from other leading supercomputing applications. First, it is very young, having been created only three years ago, in 2007; and second, it runs on a wide range of systems.

Evans said the application's youthfulness has its good points: It is free of the long-term evolution that makes portions of older applications seem opaque, and it incorporates the latest and most advanced computational tools.

"The first line of code was laid down in August 2007, so it's not an old code by any stretch. I would say this is very good from a science and engineering perspective. When these codes start up, they're all clean, and everybody understands each part. Over time, however, stuff builds in, and then you build up institutional archeology about how these codes are working.

"In an ideal world, you'd like to see these codes rewritten from scratch every five to ten years. Now obviously it's not an ideal world, and that's not realistic. But we had an opportunity to write a new code from scratch, and I think that was good."

The downside, he said, is that users don't have the long history and familiarity that they have with older applications.

"There's a lot of institutional trust in various tools out there," Evans noted, "People have been running a certain code for years."

He noted also that Denovo's user community spans the range from consumer laptops to Jaguar.

"With Denovo, the majority of our customers are running on laptops, PCs, and small clusters," Evans said. "The majority of runtime hours is on Jaguar, but we're the guys that run on Jaguar, not some guy sitting out in Nevada running shield cask analysis to do shielding for fuel casks. He's not running on Jaguar; he's running on his desktop."

The future also looks bright. Denovo was chosen for the 2010 Joule Metric, a year-long exercise in which Jaguar's capabilities are demonstrated by analyzing and improving codes in application areas important to the Department of Energy. Denovo is also one of the first applications being ported to run on future systems.

"I believe personally we've achieved a great deal in two years," Evans said. "Not many codes go from nothing to running problems of this size in that time. At the time of the Early Petascale Science program we were a year and a half into it. We were running billion-zone problems on Jaguar in a year and a half."—*by Leo Williams*

LIFE AND ITS HALF-LIFE

Physicists explore what makes carbon-14 tick

Carbon-14 allows us to put a timeline on our unwritten history.

When, in 1994, researchers discovered the Chauvet Cave in southern France filled with Stone Age artwork, carbon-14 told them the charcoal drawings on the cave walls were 31,000 years old, more or less. Earlier, in 1988, when scientists examined the Shroud of Turin, carbon-14 analysis indicated the relic—a linen cloth believed by many to have been placed over Jesus's body at the time of his burial—was actually created more than a millennium later, in the Middle Ages.

For the past half-century, carbon-14 has allowed scientists to date the flotsam of human history: skeletons, ruins, anything that was once part of a plant or body. By existing in all living things and decaying at a steady rate, carbon-14 gives researchers the ability to look at a long-abandoned community, tool, or other artifact and tell us how old it is. And because—for reasons not yet understood—carbon-14 decays far more slowly than most isotopes in its weight class, it allows researchers to confidently date items as far back as 60,000 years.

Now a team led by ORNL's David Dean is using the Jaguar supercomputer to examine the carbon-14 nucleus. The team, which includes Hai Ah Nam of ORNL, James Vary and Pieter Maris of Iowa State University, and Petr Navratil and Erich Ormand of Lawrence Livermore National Laboratory, hopes to both explain carbon-14's long half-life—about 5,700 years—and advance our understanding of what holds all nuclei together.

"Carbon-14 is interesting to us because the physics says it should decay quickly; however, the measured half-life is much longer than expected," explained Nam, who is a physicist with ORNL's National Center for Computational Sciences. "The theoretical models people have been using to describe light nuclei such as lithium, with six particles, or boron, which has ten, have been getting some pretty good results. But carbon-14, also a light nucleus, has been elusive, and the existing theoretical models don't do so well at coming up with the same value as what's measured experimentally. That means we're not capturing all of the physics."

Carbon-14 has three qualities that make it a boon to archeologists. First, new carbon-14 is constantly being produced. Cosmic rays bombard the atmosphere and set in motion a process that turns an occasional atom of nitrogen-14 into an atom of carbon-14. As a result about one of every trillion carbon atoms in the atmosphere is carbon-14. Second, plants use carbon dioxide in photosynthesis, exchanging carbon, including carbon-14, with the atmosphere throughout their lifetimes. Animals, including humans, participate in this process by eating plants. When a plant, animal, or human dies, however, this exchange ends. We are assumed to die with our bodies containing the same ratio of carbon-14 to total carbon atoms as the atmosphere, say one in a trillion. Third, carbon-14 decays at a constant rate. If at the time of your death, one carbon atom per trillion in your body is carbon-14, you can assume that the proportion will drop to one per two trillion after 5,730 years, one per four trillion after 11,460 years, and so on.

Carbon is carbon because its nucleus contains six protons. Carbon-14, with eight neutrons, is one of three naturally occurring carbon isotopes and the only one subject to radioactive decay, in this case through a process known as beta decay. Through beta decay carbon-14 emits a negatively charged electron and an antineutrino; at the same time an uncharged neutron becomes a positively charged proton and the atom itself goes back to being nitrogen-14.

Parsing the nuclear attraction

An isotope's half-life, then, is the time it takes half the atoms in a sample to decay. For most light isotopes the half-life is typically minutes or even seconds, so carbon-14, with a half-life pushing 6,000 years, is an anomaly. A simulation that can make us understand why the half-life of this isotope is so long has the potential to illuminate all half-lives, long and short, and help us better understand how we and the observable matter in the universe around us are put together.

The task is especially challenging because we don't quite know how an atom's nucleus is held together. We know that the nucleus is made up of protons and neutrons, known generically as nucleons. We know further that these nucleons are made up of even smaller particles known as quarks and gluons, which hold together through the "strong force," the strongest known to physics.

The tail of the strong force yields interactions among nucleons that hold the nucleus together. A longstanding goal of nuclear physics, then, is to theoretically describe the properties of all nuclei, stable and unstable, mundane and exotic, large and small.

Dean and his colleagues have an allocation of 30 million processor hours on Jaguar to dissect the secrets of carbon-14 with an application known as Many Fermion Dynamics, nuclear (MFDn), created by Vary at Iowa State. According to Nam, MFDn is an especially good code for this application because it scales very well. Dean's team used Jaguar's entire XT5 partition (nearly 150,000 processing cores at the time of the simulation).

The team is using an approach known as the nuclear shell model to describe the nucleus. Analogous to the atomic shell model that explains how many electrons can be found at any given orbit, the nuclear shell model describes the number of nucleons that can be found at a given energy level. Generally speaking, the nucleons gather at the lowest available energy level until the addition of any more would violate the Pauli exclusion principle, which states that no two particles can be in the same quantum state. At that point nucleons occupy the next higher energy level, and so on. The force between nucleons complicates this picture and creates an enormous computational problem to solve.

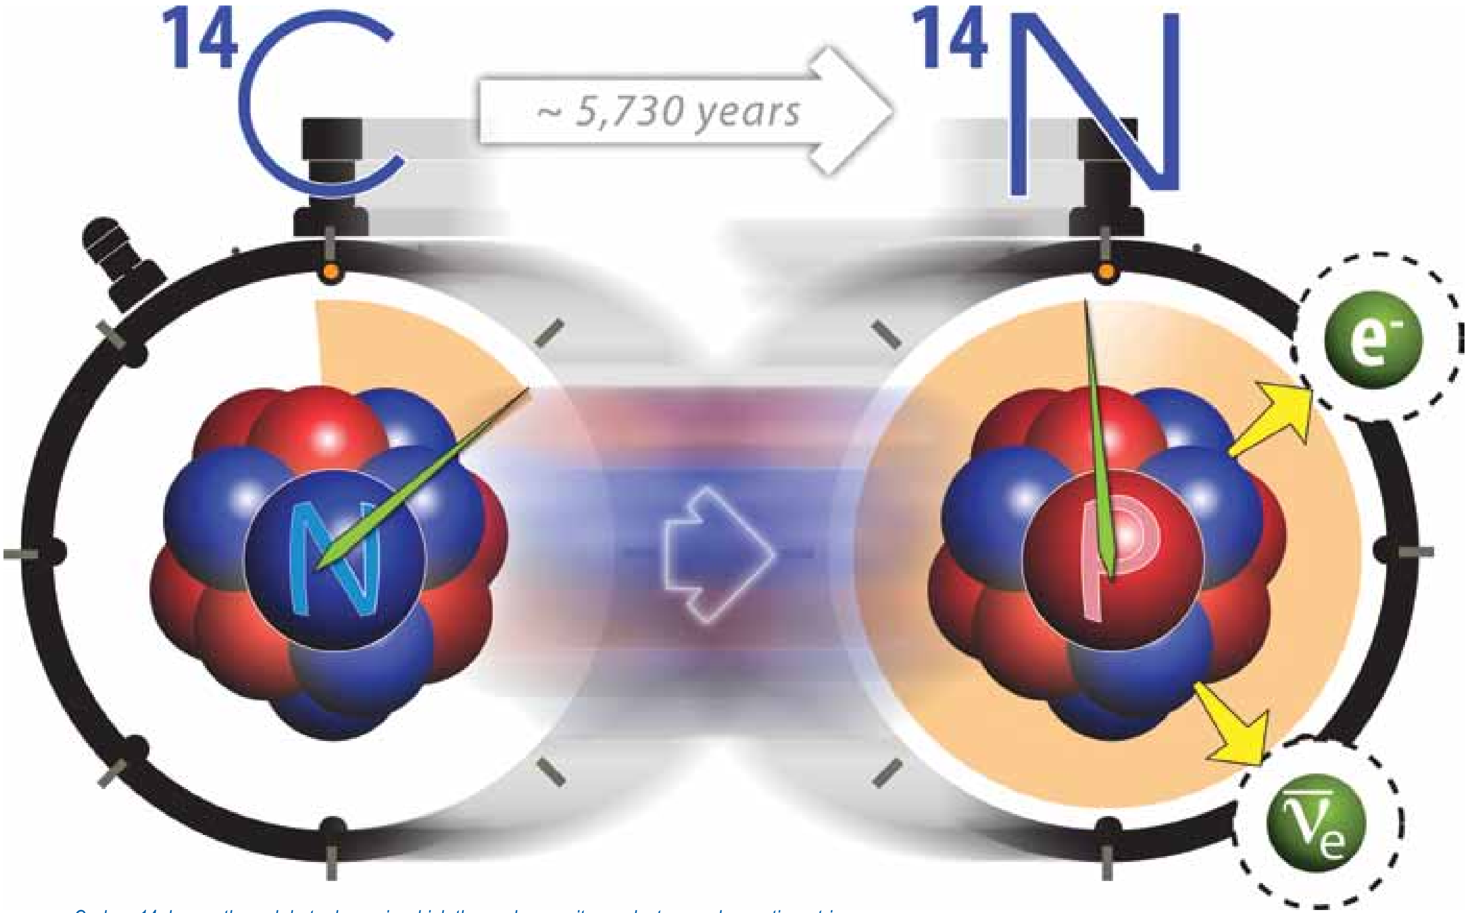

Carbon-14 decays through beta decay, in which the nucleus emits an electron and an anti-neutrino and becomes a nitrogen-14 nucleus. Visualization by Hai Ah Nam and Andrew Sproles, ORNL.

Using the power of the petascale

Jaguar allows the team to depart from other nuclear structure studies in a variety of respects. For one thing the project takes a microscopic look at the nucleus, working from its smallest known constituents. Nuclear models have been moving in this direction for seven decades, from the liquid drop model of Niels Bohr, which treated the nucleus as a single drop of nuclear fluid, to models that looked at the protons and neutrons separately. Now, the team is able to go even deeper, taking an ab initio (or first principles) approach, working from the strong-force interactions of the quarks and gluons within each nucleon. In addition, Dean's team takes a "no-core" approach that incorporates all 14 nucleons—without assuming an inert set of particles—and includes more energy levels in the model space. And lastly, the simulations go beyond two-body forces, which include the interaction of every nucleon with every other nucleon two at a time, to incorporate three-body forces.

"Previously we could only consider two-nucleon interactions because the number of combinations needed to describe all the different interactions is really big, even for only two particles at a time," Nam explained. "And while two-particle interactions are the dominant way that these particles interact, there are some nuclear phenomena, like the half-life of carbon-14, that can't be explained using a two-nucleon interaction only. Three-particle interactions or higher can also be at play.

"So this project is probing whether these two approaches—using the ab initio methods and the higher number of interactions—will better describe why carbon-14 has such a long half-life and in general explain how all nuclei are put together."

Jaguar makes these calculations possible not only because it is capable of up to 2.3 thousand trillion calculations a second, making it the world's fastest scientific supercomputer, but also because, at 362 terabytes, it has three times more memory than any other system on the planet. Before the system was installed in late 2008, such a simulation of the carbon-14 nucleus working from its smallest known constituents would have been unthinkable.

"These types of calculations for carbon-14 were previously not possible because it's a memory-intensive calculation," explained Nam. "Accounting for the three-nucleon force amounts to storing tens of trillions of elements … that's hundreds of terabytes of information."

By making use of Jaguar's power, the team hopes to push us a little closer to an understanding of the atom's nucleus. In doing so it will make carbon-14 an even bigger star than it already is.—*by Leo Williams*

Talented computational researchers will continue to accelerate the expansion of human knowledge as long as they are given sufficient tools. We at the OLCF have been honored with the task of providing these tools, including supercomputers of unprecedented power and file systems of unmatched data storage.

INSIDE THE OLCF

JAGUAR FACTS

- Cray XT computer system
- Top500 rank in November 2009: 1st for XT5 component and 16th for XT4 component
- 2.595 petaflop/s peak theoretical performance for the combined system (2.332 petaflop/s from XT5 and 0.263 petaflops from XT4)
- Superlative speed: 1.759 petaflop/s actual performance for the XT5 and 0.205 petaflop/s for the XT4 on HPL benchmark program
- 255,584 processing cores
- XT5: 37,376 AMD six-core Istanbul Opteron™ 2.6 gigahertz processors (224,256 compute cores)
- XT4: 7,832 AMD four-core Budapest Opteron™ 2.1 gigahertz processors (31,328 compute cores)
- InfiniBand network connects XT5 and XT4 components
- Cray SeaStar network interface and router
- System memory: 362 terabytes (almost three times that of the second largest system)
- Unmatched input/output bandwidth to read and write files: 240 gigabytes per second
- Sizable storage: Spider, a 10-petabyte Lustre-based shared file system
- Speedy Internet connections that enable users to access Jaguar from around the world
- High-end visualization to help users make sense of the data flood Jaguar generates

JAGUAR:

Delivering Science at the Petascale

Supercomputing has come a long way since the OLCF was founded in 2004 to strengthen American competitiveness in HPC. In 2002 Japan had deployed its Earth Simulator, at one time the fastest supercomputer in the world with a peak speed of 36 teraflops.

American supercomputers quickly regained the lead, and today the OLCF's Cray XT5, Jaguar, makes the United States the world's epicenter for computational science.

With its ascension to the petascale, the Cray XT5 is now able to give computational scientists unprecedented resolution for studying a whole range of natural phenomena, from climate change to fusion energy to nanotechnology.

The OLCF's petascale odyssey began when a series of upgrades took Jaguar from a 100-teraflop system to 250 teraflops and then past the petaflop milestone in 2008. Following that achievement, the OLCF replaced Jaguar's 37,376 AMD Opteron quad-core sockets with AMD Istanbul 6-core sockets, bringing the king of the HPC jungle to a blazing peak performance speed of 2.33 petaflops, making it the most powerful computer in the world.

With its newfound petascale power, Jaguar is revealing details of nature never before seen. Every day researchers in numerous fields have the opportunity for scientific breakthroughs that can only be accomplished on a system of Jaguar's strength, stability, and raw computing power.

GORDON BELL & HPC CHALLENGE *Awards*

Gordon Bell Prize winners Markus Eisenbach of ORNL, second from left, and Thomas Schulthess of ORNL and the Swiss National Computing Center, second from right, stand with Association for Computing Machinery Executive Director John White, far left, and Gordon Bell Prize Committee chair Mateo Valero of the Barcelona Supercomputing Center.

The OLCF's star player does more than just look good; Jaguar pushes computational science to its limits, accelerating discovery and picking up numerous honors along the way. Take SC09 in Portland, Oregon, for example. Not only was Jaguar officially recognized as the world's most powerful computer, but it also helped the OLCF rake in a number of awards at the world's premier supercomputing conference.

A team led by ORNL's Markus Eisenbach was named winner of the 2009 ACM Gordon Bell Prize, which honors the world's highest-performing scientific computing applications. Another team led by ORNL's Edo Aprà was also among nine finalists for the prize.

Eisenbach and colleagues achieved 1.84 thousand trillion calculations per second—or 1.84 petaflops—using an application that analyzes magnetic systems and, in particular, the effect of temperature on these systems. By accurately revealing the magnetic properties of specific materials—even materials that have not yet been produced—the project promises to boost the search for stronger, more stable magnets, thereby contributing to advances in such areas as magnetic storage and the development of lighter, stronger motors for electric vehicles.

INCITE

Jaguar's increased power and speed will mean larger allocations for researchers, primarily under the INCITE program, which is jointly managed by the DOE leadership computing facilities at Argonne and Oak Ridge national laboratories.

In 2009 the OLCF delivered nearly 470 million processor hours for INCITE projects. Several recent upgrades to the HPC system have allowed the OLCF to allocate in 2009 three times the 145 million processor hours allocated in 2008 and six times the 75 million processor hours allocated in 2007.

Thirty-eight separate projects advanced breakthrough research in critical areas such as climate studies, energy assurance, materials, and other areas of fundamental science. By year's end the program had allocated 889 million processor hours at national laboratories in Illinois, California, and Washington as well as at ORNL.

"From understanding the makeup of our universe to protecting the quality of life here on Earth, the computational science now possible using DOE's supercomputers touches all of our lives," said DOE Under Secretary for Science Raymond Orbach, who launched INCITE in 2003. "By dedicating time on these supercomputers to carefully selected projects, we are advancing scientific research in ways we could barely envision 10 years ago, improving our national competitiveness."

The program is continuing to grow. In 2010, for example, the OLCF will be responsible for more than 950 million hours, making INCITE the foremost HPC program in the world.

Networking and Storage

The minutely detailed simulations conducted on Jaguar require high-speed, reliable storage and retrieval of huge amounts of data. To keep abreast of the demands of the world's most powerful computer, the OLCF maintains one of the largest archival storage systems, known as the High Performance Storage System (HPSS).

The HPSS at the OLCF. As storage, network, and computing technologies continue to change, ORNL's storage system evolves to take advantage of new equipment that is both more capable and more cost-effective.

Data from simulations are first written to disks by high-speed data movers and then migrated to tapes. Building redundancy into HPSS is a current focus. OLCF staff are setting up the metadata servers as two mirrored systems so that all stored data will still be available if part of the system should fail. In 2009 the core processes that make up HPSS switched to Linux and a different brand of server for better performance and redundancy.

New disk and tape resources are being implemented to position HPSS to handle larger files and larger file sets that come in at faster rates. Also in 2009 the OLCF added an additional Storage Tek SL8500 storage silo (bringing the center's total to three such silos) to increase storage space for data archiving.

At the end of 2007, HPSS was storing about 2 petabytes of data. The amount stored has approximately doubled every year, and by

The application—known as WL-LSMS—achieved this performance by reaching nearly 80 percent of Jaguar's potential 2.33 petaflops.

WL-LSMS allows researchers to directly and accurately calculate the temperature above which a material loses its magnetism—known as the Curie temperature.

Aprà's team—the other finalist led by an ORNL researcher—achieved 1.39 petaflops on Jaguar in a first-principles, quantum mechanical exploration of the energy contained in clusters of water molecules. The team used a computational chemistry application known as NWChem developed at Pacific Northwest National Laboratory.

But the honors didn't stop with Jaguar's Gordon Bell contributions. The OLCF's flagship system also won first place in three categories at SC09's HPC Challenge Awards. The first gold was won for speed in solving a dense matrix of linear algebra equations by running the HPL software code at 1.53 petaflops.

The fastest cat in the HPC jungle also ranked first for sustainable memory bandwidth by running the STREAM code at 398 terabytes per second. STREAM measures how fast a node can fetch and store information.

Jaguar's third "gold" was for executing the fast Fourier transform, a common algorithm used in many scientific applications, at 11 teraflops.

Finally, Jaguar took third place for running the RandomAccess measure of the rate of integer updates to random locations in a large global memory array.

"It is very gratifying that Jaguar has been recognized as a very powerful machine," said OLCF Project Director Buddy Bland. "The HPC Challenge benchmarks are designed to give a better view of the entire system's performance. Jaguar was designed to be the most powerful system for scientific applications, and these results reflect that design and implementation."

While the OLCF's trophy case is filling up fast, it's the science that matters. All of these honors ultimately mean bigger, more detailed simulations capable of increasing our understanding of the natural world and helping researchers to find some of science's most sought after answers. The real rewards are just around the corner. *—by Gregory Scott Jones*

---

the end of 2009 the number had jumped to more than 8 petabytes. The OLCF has recently seen daily influxes of more than 50 terabytes.

As the OLCF continues to grow, so will the archival needs of the center, making HPSS an integral component of the world's leading computing center.

Another major issue for users of high-performance computers is the capability to move large data files from one system or location to another in a reasonable amount of time. For the OLCF with its many remote users, accurate, high-speed data transfer is essential—network capability must keep pace with computing capability. The center is working to put high-throughput networks in place among the different systems in the center and between the OLCF and other research institutions to ensure that moving data from the petascale computers doesn't become a roadblock to scientific discovery.

The OLCF has two 10-gigabit-per-second connections to ESNet, the primary internet provider for the DOE community. This connection gives ORNL high-speed access to other DOE research facilities and high-speed tiering with other research networks. The center is also a partner in the Advanced Networking Initiative, an American Recovery and Reinvestment Act–funded project to build a 100-gigabyte prototype network to connect DOE supercomputing centers, namely the OLCF, the Argonne Leadership Computing Facility, and the National Energy Research Scientific Computing Center. The project was begun last year and is expected to transition to an ESNet backbone network in 2012.

Deployment of a centerwide InfiniBand fabric began in 2007 to help meet bandwidth and scaling needs within the OLCF. InfiniBand, the new industry standard for high-performance networks, enables users to move large data sets from the simulation platforms to other OLCF platforms such as the Lustre file system, data storage, and analysis and visualization. The OLCF is one of the first sites to operate an InfiniBand network on a Cray XT system to support a move to petaflop computing.

Visualization, Data Exploration, and Analysis

The OLCF provides researchers with the Exploratory Visualization Environment for REsearch in Science and Technology (EVEREST) and its associated visualization cluster, Lens. A large-scale venue for data exploration and analysis, EVEREST measures 30 feet wide and 8 feet tall. Its main feature is a 27-projector PowerWall displaying 11,520 by 3,072 pixels, or a total of 35 million pixels, offering tremendous detail. The projectors, arranged in a 9-by-3 array, each provide 3,500 lumens. Lens is a 32-node Linux cluster dedicated to data analysis and high-end visualization of scientific data generated on Jaguar. Each node contains four quad-core 2.3 gigahertz AMD Opteron processors with 64 gigabytes of memory and two NVIDIA Tesla graphics processing units.

The EVEREST PowerWall provides researchers with premier data analysis and visualization capabilities.

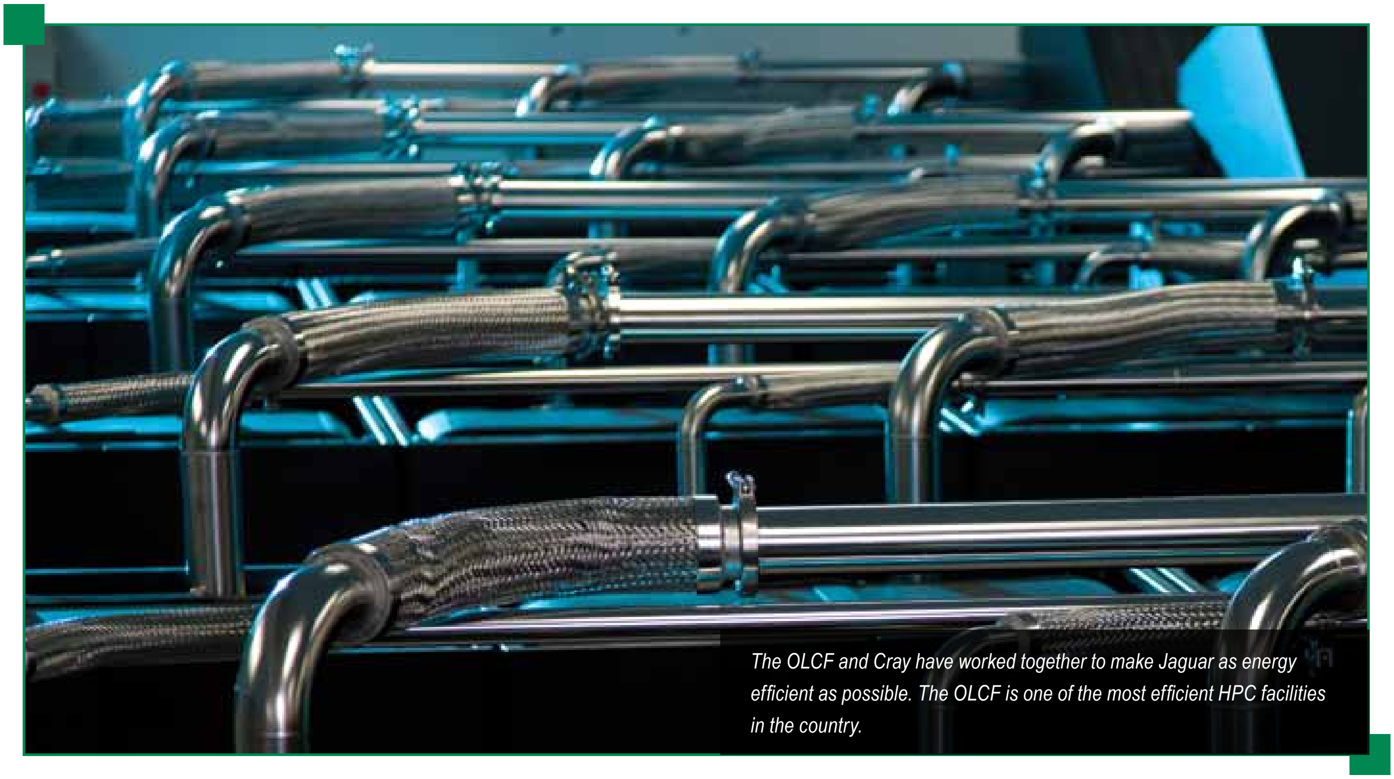

The OLCF and Cray have worked together to make Jaguar as energy efficient as possible. The OLCF is one of the most efficient HPC facilities in the country.

Energy Efficiency

Like other leading supercomputers, Jaguar requires robust cooling and support infrastructures that in turn demand tremendous power. Finding tomorrow's energy solutions requires substantial energy consumption today.

The OLCF has recently begun to reduce its resource footprint by harnessing energy savings wherever possible. As a testament to this philosophy, the OLCF is one of the most efficient petaflop-plus HPC centers in the country. In other words, the center gets the most computing bang for its power buck among the leading HPC centers, allowing it to tackle big science more quickly and efficiently.

This is a result of innovation across the entire HPC spectrum, from the building that houses Jaguar to the machine itself.

ALONG CAME A SPIDER

Spider, the world's biggest Lustre-based, centerwide file system, was fully tested to support Jaguar in 2009 and is now in full production.

An extremely high-performance file system, Spider has 10.7 petabytes of disk space and can move data at more than 240 gigabytes a second. "It is the largest-scale Lustre file system in existence," said Galen Shipman, Technology Integration Group leader at the OLCF. "What makes Spider different [from large file systems at other centers] is that it is the only file system for all our major simulation platforms, both capable of providing peak performance and globally accessible."

Ultimately, it will connect to all of ORNL's existing and future supercomputing platforms as well as off-site platforms across the country via GridFTP (a protocol that transports large data files), making data files accessible from any site in the system.

It all starts with the building. The OLCF's Computer Science Building (CSB) was among the first Leadership in Energy and Environmental Design (LEED)–certified computing facilities in the country, meaning that its design satisfies criteria used by the U.S. Green Building Council to measure the efficiency and sustainability of a building.

The facility has a power usage effectiveness (PUE) rating of 1.25— a recent study of 22 large-scale data centers by Lawrence Berkeley National Laboratory reflected an average PUE of 1.83 (the closer to 1 the more efficient). This metric means that for every 1.25 megawatts of energy consumed by the CSB, 1 megawatt is used to power the machines and accelerate science. The other 0.25 megawatt is used for lighting, the dispersion of heat generated by the machines, and various other support equipment. The CSB's PUE rating makes it almost 32 percent more efficient than its average counterpart, a direct result of the laboratory's energy conservation measures.

When it comes to the machines, it's all about keeping them cool. In 2009 the OLCF upgraded all of the fans in the computer room's air-conditioning units (CRUs), which are charged with helping to keep the computer room cool despite the heat generated by the machines.

The fan is essentially a 5-horsepower motor within the CRU that pushes air through the CRUs' coils to remove the heat, said Jim Rogers, director of operations for the OLCF. The new fans allow the center to run the CRUs at a more favorable speed, namely from 60 to 80 percent of their peaks. Instead of being on or off, said Rogers, the CRUs can now "operate more efficiently and in unison." This upgrade will pay for itself in a year and a half, after which it will save the center $150,000 every year throughout the life of the HPC center.

Combine that with Cray's latest ECOphlex cooling system and the savings start to add up. ECOphlex allows ORNL to reduce the amount of chilled water used to cool Jaguar by taking up to 95% of Jaguar's heat away in a direct heat-to-chilled-water transfer, versus a significantly less effective heat transfer method using only air. Considering that thousands of gallons of water per minute are necessary for cooling, a reduction in volume means a proportionate reduction in cooling costs. Simply put, these more efficient heat transfer methods can mean big energy savings for the OLCF and the taxpayer.

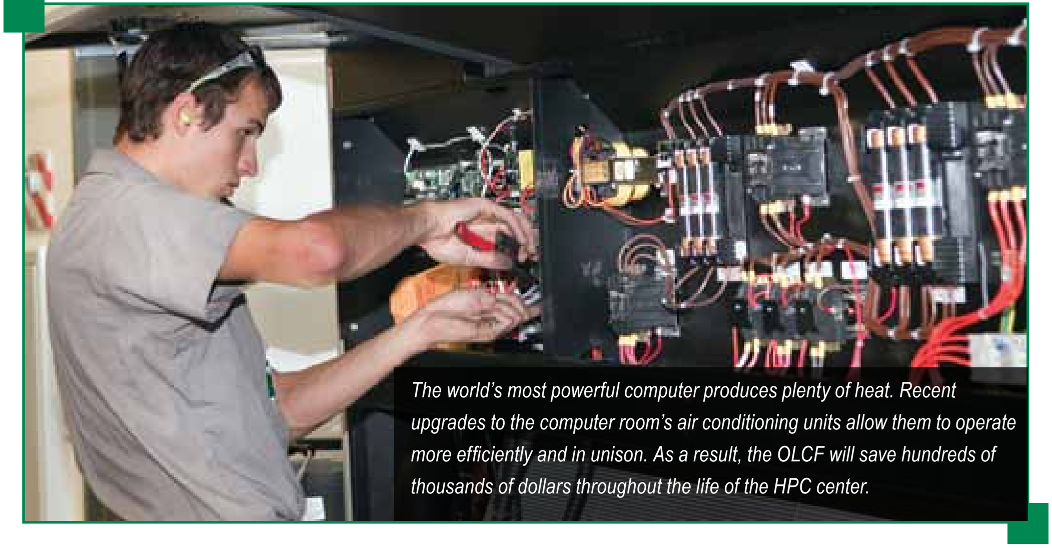

The world's most powerful computer produces plenty of heat. Recent upgrades to the computer room's air conditioning units allow them to operate more efficiently and in unison. As a result, the OLCF will save hundreds of thousands of dollars throughout the life of the HPC center.

Spider has demonstrated stability on the XT5 and XT4 partitions of Jaguar, on Smoky (the center's development cluster), and on Lens (the visualization and data analysis cluster). "We've had all these systems running on the file system concurrently, with over 26,000 compute nodes (clients) mounting the file system and performing I/O [input and output]. It's the largest demonstration of Lustre scalability in terms of client count ever achieved," said Shipman.

A phased approach was used to get Spider crawling. ORNL computer scientists and technicians (David Dillow, Jason Hill, Ross Miller, Sarp Oral, Feiyi Wang, and James Simmons) worked in close collaboration with partners Cray, Inc., Data Direct Networks (DDN), Sun Microsystems, and Dell to bring Spider online. Cray provided the expertise to make the file system available on both Jaguar XT4 and Jaguar XT5. DDN provided 48 DDN 9900 storage arrays, Sun provided the Lustre parallel file system software, and Dell provided 192 I/0 servers. The vendors' collaboration has produced a system which manages 13,000 disks and provides over 240 GB/s of throughput, a file system cluster that rivals the computational capability of many high-performance compute clusters.

The Spider parallel file system is similar to the disk in a conventional laptop— multiplied 13,000 times. A file system cluster sits in front of the storage arrays to manage the system and project a parallel file system to the computing platforms. A large-scale InfiniBand-based system area network connects Spider to each OLCF system, making data on Spider instantly available to them all.

"As new systems are deployed at the OLCF, we just plug them into our system area network; it is really about a backplane of services," Shipman said. "Once they are plugged into the backplane, they have access to Spider and to HPSS [the center's high-performance storage system] for data archival. Users can access this file system from anywhere in the center. It really decouples data access and storage from individual systems."

Spider will have both scratch space (short-term storage for files involved in simulations, data analysis, etc.) and long-term storage for each user. Shipman said the technology integration team is now working with Sun to prepare for future OLCF platforms with even more daunting requirements.—*by Agatha Bardoel*

Put together, the OLCF's energy reduction strategy is making an impact. By seeking energy savings in a variety of different arenas, the center is reducing computational science's carbon footprint while accelerating solutions to some of tomorrow's most daunting energy challenges.

Future Systems

Imagine a computer 500 times more powerful than Jaguar. While it may be hard to do, that is the vision currently being pursued by the OLCF. To further our knowledge in climate research, alternative energy development, and other critical areas, we will need exascale systems, which will be capable of more than a million trillion calculations each second, or 1,000 petaflops.

The OLCF is leading the way. OLCF staff are working with vendors to design future architectures and help identify the infrastructure needs of future systems such as storage, file systems, and networking. The center itself is hosting workshops to address the needs and challenges of achieving exascale computing with input from the scientific and application communities.

"It's becoming clear that getting to exascale will require paradigm shifts in how we do high-performance computing," said OLCF staff scientist Wayne Joubert. "We need to begin preparing now for the hardware and software changes required for exascale computing."

Supercomputers able to handle these problems will be the result of collaborations among many communities. Manufacturers will combine different, but complementary, types of computer chips to get the highest possible performance. Developers of programming languages and software tools will give application programmers the resources they need to take advantage of new, mixed architectures. And programmers themselves will need to adapt and, in some cases, rewrite codes that are already enormously complex. In many areas, these communities will need to develop creative new technologies and new ways of thinking to overcome looming roadblocks.

A thermoelectric generator in a Chevy Suburban would provide up to a 5 percent improvement in fuel economy. (Source: General Motors)

HPC INDUSTRIAL REVOLUTION

In today's complex and highly-interconnected global economy, modeling and simulation with HPC has become a critical ingredient in the recipe for industrial success. The OLCF HPC Industrial Partnerships Program helps companies supercharge their competitiveness through access to world-class computational resources and domain expertise not available elsewhere.

The program is a triple win, benefiting ORNL, industry, and the country. It links the lab to some of the best thinking within U.S. companies, and many of these firms are pursuing scientific challenges similar or complementary to those at the lab in their quest to develop innovative products and services. Companies benefit as they tackle complex, strategic problems that cannot be solved on their internal systems but that will help them "break out of the pack" if successfully addressed.

As the lab and companies advance together in their problem-solving abilities, they strengthen the nation's capacity to address national grand challenges, like substantially reducing energy consumption.

"We've reached the limit of getting gains by increasing the clock speed of processors," Joubert said. "Many people believe we will need some sort of accelerator technology to get to exascale."

At the OLCF the accelerator technology will come from graphics processing units (GPUs), originally created for video applications. Also known as graphics accelerators, GPUs are parallel processing on steroids, with each chip containing hundreds of small cores able to race through simple instructions. They will be combined with traditional processors—central processing units— to give the 2012 OLCF system a peak speed of 10 to 20 petaflops. The OLCF's exascale roadmap envisions a 100- to 250-petaflop system by 2015 and an exaflop machine sometime around 2018.

Tools have already been developed for this new computing environment, and they are continually being improved. Graphics processor manufacturer NVIDIA has created a programming environment called CUDA (for Compute Unified Device Architecture) that allows programmers to make the most of these mixed systems while still writing in traditional languages such as C and Fortran. A similar, open-source environment is known as OpenCL.

Even with these tools, however, it will be a massive undertaking for developers to adapt applications that typically contain tens of thousands of lines of code. Again, the OLCF is front and center in leading the application readiness process by identifying the obstacles in future software development.

Whatever the future holds, any eventual exascale system will ultimately require a number of innovations. The OLCF's pioneering past and current leadership status have it poised to help the HPC community surpass any roadblocks on the way to the exascale, thus enabling a new age in computational science.—*by Gregory Scott Jones*

There are three venues through which a company can apply for time on Jaguar:

- For smaller allocations (up to 2 million hours of compute time) companies can apply for the Laboratory "Director's Discretionary" Allocation.
- For larger problems (up to 20 million hours) that are directly related to either the DOE energy mission or national emergencies or that broaden the community of researchers capable of using leadership computing resources, companies can apply through the ASCR Leadership Computing Challenge.
- Companies posing the most complex scientific challenges (requiring more than 20 million hours and using 20% or more of Jaguar in their production runs) can apply for time through the INCITE program.

Regardless of the venue, companies should propose high-risk/high-return problems whose solutions will provide opportunities for breakthrough insights and competitive leaps forward.

Through the HPC Industrial Partnerships Program, BMI Corporation and General Motors are using Jaguar to develop more fuel-efficient transportation.

A research team lead by General Motors physicist Jihui Yang is researching a thermoelectric material that can trap the 70 percent of energy lost to waste heat in conventional car engines and convert it to useable energy. Such a step forward in gasoline efficiency could save hundreds of millions of gallons every year.

BMI Corporation, an engineering services firm in South Carolina, is using Jaguar to develop add-on parts that will substantially reduce the fuel consumption of Class 8 long haul trucks. Access to Jaguar enabled BMI to develop and use the most detailed and accurate numerical model of a Class 8 truck and trailer ever created. Trucks retrofitted with such parts can achieve an increase in fuel efficiency equivalent to a 10% to 17% increase in MPG. The result will be reduced fuel consumption by up to 3,700 gallons per truck per year, and a reduction in CO_2 by up to 41 tons (82,000 lb) per truck per year.

HPC is a game-changing technology and the HPC Industrial Partnerships Program provides companies a gateway to the best. General Motors and BMI Corporation are just two of the firms realizing transformational results from access to the leadership computing systems and expertise at OLCF. —*by Wes Wade*

James Hack, NCCS Director

FORGING the FUTURE:

Leadership computing begins move to exascale

In November 2009 the Oak Ridge Leadership Computing Facility's (OLCF's) Cray XT5 system, Jaguar, was ranked number one among the world's 500 fastest supercomputers, with a peak performance of more than 2 thousand trillion calculations a second, or 2 petaflops. This is a proud accomplishment for Oak Ridge National Laboratory, but more important than the ranking is the groundbreaking research now being conducted with this system. The world's leading computational scientists are using pioneering applications to advance our understanding of pressing scientific and technological problems, and they are doing so with unmatched speed and simulation fidelity.

The power of Jaguar is not theoretical. There was a time when most in the high-performance computing community regarded the synthetic benchmark application known as High-Performance Linpack (HPL), used to rank systems for the TOP500 list, to define the "speed limit" of a supercomputer. No longer. This past fall we saw two Jaguar applications focused on real-world problems eclipse computational performance exhibited by HPL. One is designed to explore the world's most advanced superconducting materials, while the other analyzes magnetic systems such as those found in computer hard drives. A third highly sophisticated computational chemistry application was close behind, running nearly as fast as the synthetic HPL benchmark.

It is no coincidence that researchers have seen remarkable computing performance with scientific applications. At the same time as we have striven to provide the world's most powerful, well-balanced systems for scientific computing, we have made the necessary intellectual investments with many of our computational science research teams to ensure that our users get the most out of these resources. The partnering of research teams with scientific computing liaisons, who aid in code development and optimization, streamline workflow, and troubleshoot problems both known and unforeseen, is now a proven recipe for facilitating breakthrough science, as demonstrated by the many Jaguar success stories.

Examples include the work conducted by a team led by Venkatramani Balaji from the National Oceanic and Atmospheric Administration's Geophysical Fluid Dynamics Laboratory. Petascale computing enabled this team to explore new formulations of their global climate model and scale it to exceptionally high resolution, bringing science a step closer to characterizing climate change on regional space scale, and possibly forecasting climate on decadal time scales. Another team, led by ORNL's Thomas Evans, conducted the largest simulation to date of a nuclear fission reactor, explicitly resolving what had been only approximations of individual neutron energies and directions of movement. And a third team, led by ORNL's Jeremy Smith, used a highly scalable molecular dynamics application to study the structure, dynamics, and degradation of lignocellulosic biomass—a potential biofuel.

Despite the many success stories, the scientific community continues to press for new high-performance computing capabilities to support their simulation needs. To address those needs we must look beyond petascale computing. The community is now examining what will be required to add another thousand-fold increase in computational performance before the end of the next decade. These systems, referred to as exascale systems, will be capable of a million trillion calculations a second.

There are many challenges to overcome on the way to achieving this ambitious goal. A first and necessary step is collaboration. As systems become more specialized, scientists must participate in the architectural design. We want to be able to choose the HPC architectures that are best suited for the applications, not the other way around. This task will be challenging because exascale systems will most certainly differ considerably from the parallel architectures that have evolved over the last decade. Future systems are likely to include hybrid processing nodes, incorporating technology from the worlds of gaming and video, using graphics processing units to provide low-power, floating point acceleration mechanisms. Such systems are expected to deploy at the OLCF in as little as two years. As we strive to take advantage of additional degrees of parallelism at the node level, successful exploitations of next-generation architectures will involve not only re-examination of underlying algorithms in scientific codes, but the creation of new computational environments including languages, compiler technologies, performance tools, and debuggers.

En route to delivering revolutionary science at the exascale, one goal may well be most important: We must ensure that the computational science community can use future systems as or more effectively than current petascale systems. Every gain in peak computing power must be accompanied by an equivalent gain in application performance. Only then will we realize the promise of these systems.

By combining the talents of our staff and user community, the OLCF achieved great things in 2009. Through a continued close partnership with our scientific application community, the best is still to come.

Buddy Bland, OLCF Director

GROWING at the Petascale

Leadership computing champions new era of science

The OLCF saw a number of critical breakthroughs in 2009. Early that year 20 research teams began testing Jaguar's initial petascale capabilities. Many of their applications used 100,000 or more of Jaguar's 150,000 cores, and the results are already advancing the frontiers of some of the world's most pressing problems in climate change, energy, and basic science.

Jeremy Smith and his team from ORNL and the University of Tennessee used Jaguar to develop a molecular-level understanding of structure, dynamics, and degradation pathways of cellulosic and lignincellulosic materials and their potential as sources of biofuels, which could help lessen the amount of oil we have to import. Max Suarez and a team of scientists from NASA's Goddard Space Flight Center and NOAA's Geophysical Fluid Dynamics Laboratory used Jaguar to determine the effects of individual weather events on global climate using high-resolution climate models. Jacqueline Chen and her team from Sandia National Laboratories, ORNL, and the University of Connecticut used it to study flames in diesel engines to help reduce soot level and increase fuel efficiency. Lawrence Berkeley National Laboratory's Lin-Wang Wang and his team used it to calculate the electronic properties of large nanocrystaline structures, and this new understanding may pave the way for dramatically more efficient solar cells.

Yet this was only the beginning. In late summer, Jaguar underwent an upgrade to increase its power by 70 percent. Committed to the scientific community ahead of all else, we kept major portions of Jaguar up and available to our users during this transition. Now we continue to deliver groundbreaking science with more than 224,000 processing cores, 300 terabytes of system memory, and a peak performance of 2.3 petaflops.

In November Jaguar placed number one on the TOP500 list of the world's most powerful supercomputers. A team led by ORNL researcher Markus Eisenbach won the Gordon Bell Prize for fastest scientific application with a magnetic-properties code that achieved 1.84 petaflops. A finalist for that same prize was a team lead by Eduardo Apra, another ORNL researcher, whose work on an extremely complex chemistry application proved to be an extraordinary achievement in scalability.

Due to Jaguar's enormous computing capability, managing the petabytes of data it can produce could pose serious problems. We've dealt with this challenge in two areas. First, we developed and deployed the world's largest and fastest Lustre-based shared file system, Spider, which provides centralized access to all major OLCF platforms, eliminating the need for large data transfers and making our scientists more productive. Second, we upgraded our High-Performance Storage System to support the increasing need for long-term data storage and double the speed with which we can store and retrieve the data.

As a result of these new capabilities, the OLCF is now the world's most powerful computational user facility, joining the ranks of other world-leading DOE Office of Science user facilities at ORNL such as the Spallation Neutron Source, which produces the world's most intense beam of pulsed neutrons, and the High Flux Isotope Reactor, one of the highest steady-state neutron fluxes among roughly 250 other research reactors in the world. Because we share fundamental research objectives with these facilities, progress made at one center is often transferred across campus to broaden the scientific knowledge base at large.

In the coming year we will remain focused on maintaining operational excellence and delivering the best scientific results possible. Modeling and simulation needs will only increase as we move forward in our battle against the most pressing problems in energy, health, the environment, and national security. The OLCF will continue to upgrade its computing facilities to keep pace with the scientific needs of the nation. In fact, the Office of Science has already signed a preliminary plan for a next-generation HPC system that we expect to be in the 10- to 20-petaflop range. Science isn't slowing down a bit, and neither are we.

Education and Collaboration

THE RESEARCH ALLIANCE IN MATH AND SCIENCE

The Research Alliance in Math and Science (RAMS) program, sponsored by the Office of Advanced Scientific Computing Research, hosted 28 undergraduate and graduate students during summer 2009. The program is designed to provide collaborative research experiences between ORNL researchers and underrepresented students majoring in computer science, the computational sciences, math, and related disciplines. The RAMS program seeks to improve U.S. competitiveness in research and to increase the number of U.S. citizens from underrepresented populations who hold advanced degrees in science, mathematics, engineering, and technology.

Each student is assigned an ORNL research scientist mentor and conducts a project of mutual interest, often an important aspect of a larger ongoing project, over the 12-week internship period. For instance, physics graduate Jessica Travierso worked with OLCF astrophysicist Bronson Messer to build a test system for a physics code used by researchers around the world. Students keep journals of their activities and experiences, attend seminars and workshops, summarize the results of their research projects in papers and oral presentations, and participate in poster presentations at ORNL and national conferences. The program is administered through the Computing and Computational Sciences Directorate office and has provided research and conference participation opportunities to students from both majority- and minority-serving institutions across the United States, including Puerto Rico.

Training the next generation of computational scientists: 2009 students from ORNL's Computing and Computational Sciences Directorate.

EDUCATION/OUTREACH

The OLCF's commitment to the scientific community includes a dedication to fostering the development of future researchers and computational scientists.

In spring 2009, Bobby Whitten of the OLCF User Assistance and Outreach Group taught the first in a series of planned online HPC courses for Morehouse College students. The course, which stemmed from the mentor–protégé partnership formed in summer 2007 between the college and ORNL, gave students a broad overview of current HPC topics. The topics included parallel programming models, visualization of data, and an introduction to astrophysics and molecular dynamics computer models. The course culminated in a hands-on session during which students built a cluster, compiled and executed a scientific application on a data set, and created visualizations of and analyzed their results.

The OLCF also witnessed the growth of its partnership with the Appalachian Regional Commission, a body of educators and business people whose goal is to support the educational development of the Appalachian region by offering college and career options to underprivileged and minority high school students. This partnership program, now in its second year, accepts 26 high school students and 26 college and high school instructors annually from a 205,000-acre stretch of land covering 15 states. The 2-week workshop, which ran from July 11 to 24, engaged the students and instructors with activities and projects in areas such as biology, astronomy, computational science, and chemistry. With the help of

The OLCF hosted over 400 tours of the facility in 2009, giving dignitaries, visiting scientists, and community groups the opportunity to see the machines that are enabling today's groundbreaking science.

interns Jesse Mays and Stephanie Poole and their mentors, six of these students learned the basics of hardware, software, and programming and how they are applied in an HPC environment. They also received a crash course in parallel processing.

The facility also offered several seminars in 2009 as part of its series, which typically hosts world-renowned researchers as guest speakers. The center tries to schedule at least one seminar per month with topics that include areas of current interest in HPC, focusing particularly on performance at the petascale. Ten seminars took place in 2009, including one with guest speaker Cecilia Bitz from the University of Washington, who spoke on the survival of Arctic summer sea ice in the twenty-first century, and another with Timothy C. Germann from Los Alamos National Laboratory, who gave a talk entitled “From atoms to agents: adventures in petascale computing.”

The OLCF sometimes acts as a liaison between its user community and the community at large. Through user meetings, public tours (more than 400 in 2009), and speeches to outside groups through the ORNL Speakers Bureau, the facility supports a strong connection with the public and the vital exchange of information such a connection enables.

MEETINGS/WORKSHOPS

Often working closely with vendors such as Cray, IBM, and Data Direct Networks, the OLCF hosts regular workshops, meetings, and education courses geared toward current and potential users. Topics are designed to improve and enhance the HPC experience and range from crash courses in supercomputing to scalability and hardware issues.

Lustre Scalability Workshops

February 10–11, May 19–20

As computing resources continue to expand from petascale to exascale capabilities, storage systems must scale to these same levels. The OLCF collaborated with Sun Microsystems and Cray to host a two-part workshop on the design of next-generation Lustre-based storage systems to meet these future needs.

The workshops brought together key members of the Lustre community to identify and address scalability issues in an effort to keep ORNL’s storage potential up to date with the latest technological advances. Attendees developed a feasible plan to deliver a system with multiple terabytes/second of bandwidth capabilities and the capability to manage exabytes of storage by 2015.

OLCF Users Meeting

April 16

A forum for current and prospective users of the facility’s leadership resources, the annual NCCS users meeting allows the center to establish and better facilitate the needs of the user and broader scientific communities at large. The 2009 meeting gave researchers a unique opportunity to engage one-on-one with center staff and fellow colleagues.

Principal Investigators and their collaborators from each DOE INCITE project were invited to deliver a 20 minute presentation on recent and upcoming developments. Users were also encouraged to attend the “Climbing to Petaflop on Cray XT” workshop that preceded the meeting, April 14–16.

Visualization with VisIt

June 4

Nearly 50 users attended the VisIt workshop to become more familiar with this scientific visualization software application. The course alternated between lecture discussions and hands-on sessions designed to give participants an introduction to the software in an interactive setting. Active sessions included working with ORNL’s visualization analysis cluster, LENS.

Crash Course in Supercomputing

June 16–17

About 40 students were introduced to the parallel programming basics necessary for designing applications to run on today’s leadership computers. Experienced staff from both the NCCS and NICS lectured and assisted students during hands-on programming activities with the facilities’ leadership systems. The course is designed for those with little or no programming background, as an environment for interested students to experience supercomputing first-hand.

Allinea DDT Training for Cray XT5

July 21–22

Donald Frederick from the NCCS User and Assistance Outreach Group helped introduce users to the first debugger to run on over 100,000 cores. Attendees were provided with soft copies of the debugger as well as user guides and additional documentation prior to this on-site training class.

Cray XT5 Hex-core Workshop

December 7–9

Staff from the NCCS, NICS, and Cray lectured and provided assistance to users during hands-on sessions with the latest XT5 architectures. To ensure users full access to the novel potential Jaguar and Kraken offer the scientific community, the features of these new systems must be fully understood. The 3-day workshop covered important issues in obtaining increased performance, including programming for the AMD six-core CPU and addressing bottlenecks and scalability issues.

—by Wes Wade

Breakthrough science is built on teamwork. At the OLCF, our researchers and staff are second to none.

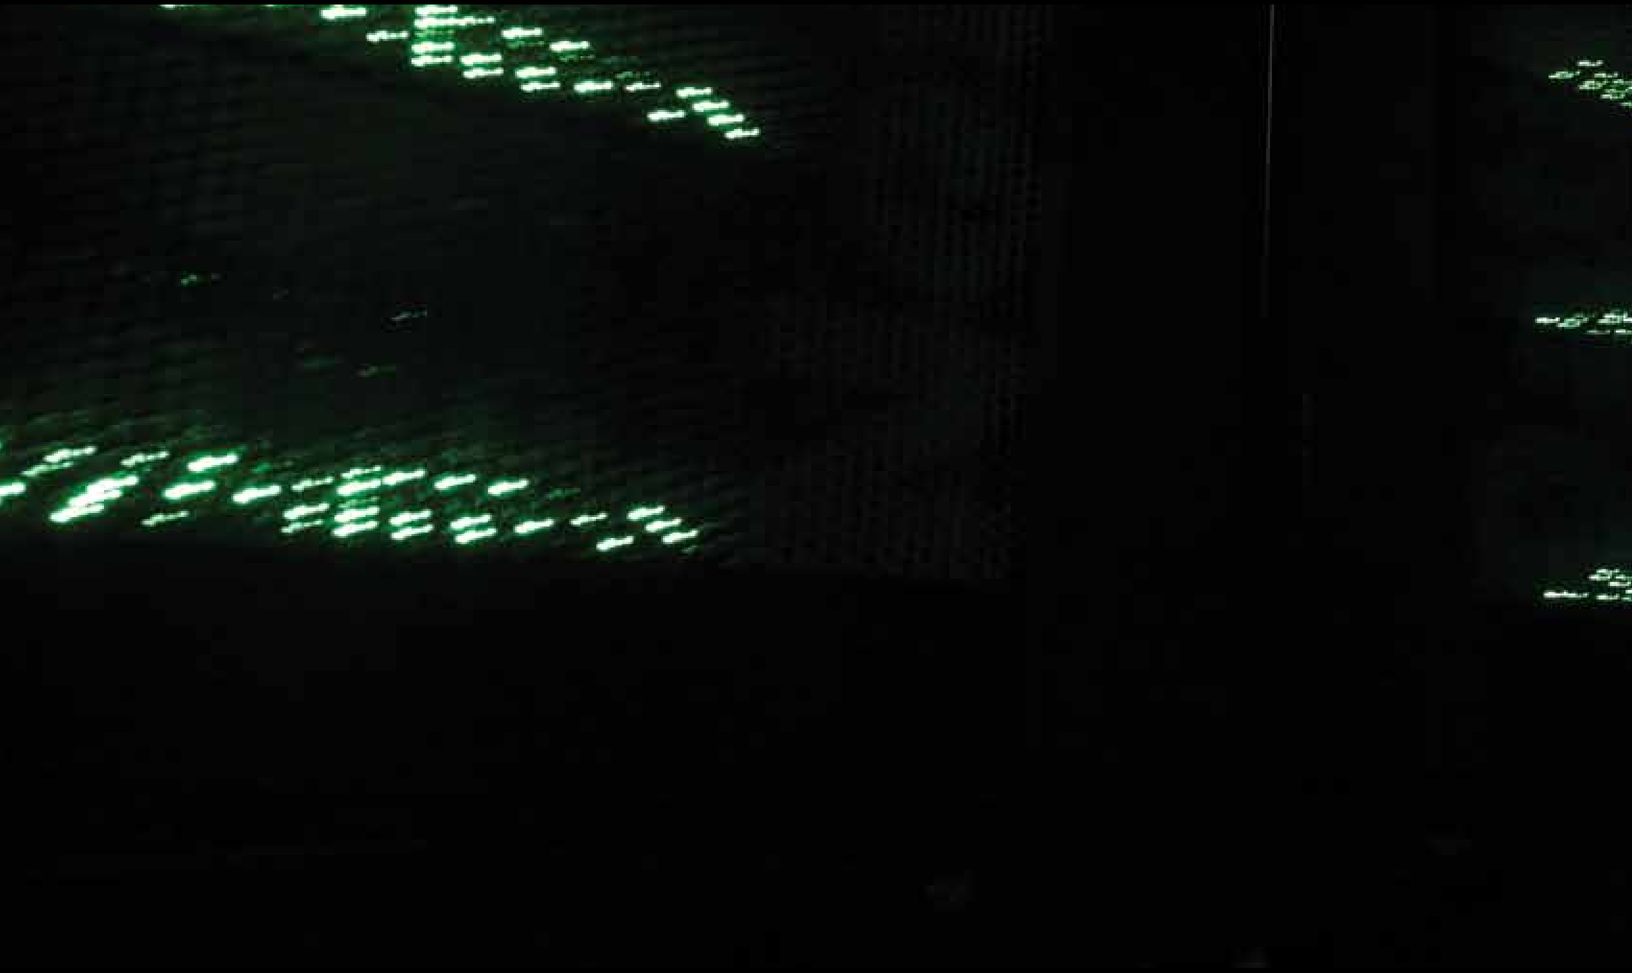

APPENDICES

ASTROPHYSICS

Multidimensional Simulations of Core Collapse Supernovae
Anthony Mezzacappa, Oak Ridge National Laboratory
Jaguar: 34,000,000 hours

Turbulent Heating of Astrophysical Plasmas
Gregory Howes, University of Iowa
Jaguar: 12,000,000 hours

Multidimensional Models of Type Ia Supernovae from Ignition to Observables
Stan Woosley, University of California–Santa Cruz
Jaguar: 5,000,000 hours

The Via Lactea Project: A Glimpse into the Invisible World of Dark Matter
Piero Madau, University of California–Santa Cruz
Jaguar: 5,000,000 hours

BIOLOGY

Sculpting Biological Membranes by Proteins
Klaus Schulten, University of Illinois at Urbana-Champaign
Jaguar: 25,000,000 hours

Cellulosic Ethanol: Physical Basis of Recalcitrance to Hydrolysis of Lignocellulosic Biomass
Jeremy Smith, Oak Ridge National Laboratory
Jaguar: 25,000,000 hours

Interplay of AAA+ Molecular Machines, DNA Repair Enzymes and Sliding Clamps at the Replication Fork: A Multiscale Approach in Modeling Replisome Assembly and Function
Ivaylo Ivanov, University of California–San Diego
Jaguar: 4,000,000 hours

Sequencing DNA Using MspA
Aleksei Aksimentiev, University of Illinois at Urbana-Champaign
Jaguar: 10,000,000 hours

CHEMISTRY

Molecular Simulation of Complex Chemical Systems
Christopher Mundy, Pacific Northwest National Laboratory
Jaguar: 10,000,000 hours

An Integrated Approach to the Rational Design of Chemical Catalysts
Robert Harrison, Oak Ridge National Laboratory
Jaguar: 75,000,000 hours

CLIMATE

Assessing Transient Global Climate Response of the NCAR-CCSM3: Climate Sensitivity and Abrupt Climate Change
Zhengyu Liu, University of Wisconsin–Madison
Jaguar: 5,000,000 processor hours

Simulation of Global Cloudiness
David Randall, Colorado State University
Jaguar: 3,000,000 hours

High Resolution Ensemble Simulations of Hurricanes
Robert Gall, NOAA
Jaguar: 20,000,000 hours

CHiMES: Coupled High-Resolution Modeling of the Earth System
Venkatramani Balaji, NOAA/GFDL, Princeton University
Jaguar: 20,000,000 hours

Climate-Science Computational End Station Development and Grand Challenge Team
Warren Washington, National Center for Atmospheric Research
Jaguar: 70,000,000 hours

COMPUTER SCIENCE

HPC Colony: Removing Scalability, Fault, and Performance Barriers in Leadership Class Systems Through Adaptive System Software
Terry Jones, Oak Ridge National Laboratory
Jaguar: 4,000,000 hours

Performance Evaluation and Analysis Consortium End Station
Patrick Worley, Oak Ridge National Laboratory
Jaguar: 20,000,000 hours

Scalable System Software Research for Extreme-Scale Computing
Ron Oldfield, Sandia National Laboratories
Jaguar: 5,000,000 hours

ENGINEERING

Interaction of Turbulence and Chemistry in Lean Premixed Laboratory Flames
John Bell, Lawrence Berkeley National Laboratory
Jaguar: 8,000,000 hours

High-Fidelity Simulations for Clean and Efficient Combustion of Alternative Fuels
Jacqueline Chen, Sandia National Laboratories
Jaguar: 65,000,000 hours

Clean and Efficient Coal Gasifier Designs Using Large-Scale Simulations
Madhava Syamlal, National Energy Technology Laboratory
Jaguar: 6,000,000 hours

Petascale Adaptive Computational Fluid Dynamics for Applications with High Anisotropy
Kenneth Jansen, Rensselaer Polytechnic Institute
Jaguar: 10,000,000 hours

Petascale Simulation of Nano-Electronic Devices
Gerhard Klimeck, Purdue University
Jaguar: 18,000,000

Next Generation Multi-Scale Quantum Simulation Software for Strongly Correlated Materials
Mark Jarrell, Louisiana State University
Jaguar: 17,000,000 hours

Understanding the Ultimate Battery Chemistry: Rechargeable Lithium/Air
Jack Wells, Oak Ridge National Laboratory
Jaguar: 12,000,000 hours

Computational Surface Science at High Accuracy with Quantum Monte Carlo
Dario Alfe, London Centre for Nanotechnology
Jaguar: 20,000,000 hours

Magnetic Structure and Thermodynamics of Low Dimensional Magnetic Structures
Markus Eisenbach, Oak Ridge National Laboratory
Jaguar: 21,000,000 hours

FUSION

Verification and Validation of Petascale Simulation of Turbulent Transport in Fusion Plasmas
Patrick Diamond, University of California–San Diego
Jaguar: 35,000,000 hours

Gyrokinetic Simulation of Energetic Particle Turbulence in ITER Burning Plasmas
Zhihong Lin, University of California–Irvine
Jaguar: 20,000,000 hours

High-Fidelity Tokamak Edge Simulation for Efficient Confinement of Fusion Plasma
C.S. Chang, New York University
Jaguar: 50,000,000 hours

Investigation of Multi-Scale Transport Physics of Fusion Experiments Using Global Gyrokinetic Turbulence Simulations
Weixing Wang, Princeton Plasma Physics Laboratory
Jaguar: 34,000,000 hours

Validation of Plasma Microturbulence Simulations for Finite-Beta Fusion Experiments
William Nevins and Greg Hammett,
Lawrence Livermore National Laboratory
Jaguar: 30,000,000 hours

GEOSCIENCES

Ultrascale Simulation of Basin-Scale CO_2 Sequestration in Deep Geologic Formations and Radionuclide Migration Using PFLOTRAN
Peter Lichtner, Los Alamos National Laboratory
Jaguar: 18,000,000 processor hours

Deterministic Simulations of Large Regional Earthquakes at Frequencies up to 2 Hz (2009)
Thomas Jordan, University of Southern California
Jaguar: 20,000,000 hours

MATERIALS

A Petascale Study of Turbulent Mixing in Non-Stratified and Stratifed Flows
Pui-Kuen Yeung, Georgia Tech
Jaguar: 20,000,000 hours

Predictive and Accurate Monte Carlo Based Simulations for Mott insulators, Cuprate Superconductors, and Nanoscale Systems
Thomas Schulthess, Oak Ridge National Laboratory
Jaguar: 70,000,000 hours

Electronic, Lattice, and Mechanical Properties of Novel Nano-Structured Bulk Materials
Jihui Yang, GM R&D Center
Jaguar : 14,000,000 hours

Development and Correlations of Large Scale Computational Tools for Transport Airplanes
Moeljo Hong, The Boeing Company
Jaguar: 6,000,000 hours

Electronic Structure Calculations for Nanostructures
Lin-Wang Wang, Lawrence Berkeley National Laboratory
Jaguar: 9,000,000 hours

PHYSICS

Computational Nuclear Structure
David Dean, Oak Ridge National Laboratory
Jaguar: 25,000,000 processor hours

Petascale Computing for Terascale Particle Accelerator: International Linear Collider Design and Modeling
Lie-Quan Lee, Stanford Linear Accelerator Center
Jaguar: 12,000,000 processor hours

Lattice QCD
Paul Mackenzie, Fermilab
Jaguar: 40,000,000 hours

Advanced Reactor Thermal Hydraulic Modeling
Paul Fischer, Argonne National Laboratory
Jaguar: 2,000,000 hours

Uncertainty Quantification for Three-Dimensional Reactor Assembly Simulations
Thomas Evans, Oak Ridge National Laboratory
Jaguar: 8,000,000 hours

Petascale Particle-In-Cell Simulations of Plasma Based Accelerators
Warren Mori, University of California–Los Angeles
Jaguar: 8,000,000 hours

U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

OLCF Staff

MANAGEMENT TEAM

The OLCF management team ensures that world-class researchers get the most from OLCF resources and that the facility effectively manages and supports testing, evaluation, and operation of HPC systems for DOE's Office of Science. Director Jim Hack guides the overall vision of the center. Project Director Buddy Bland supervises installation and upgrades of the OLCF supercomputers. Kathlyn Boudwin is deputy director of the OLCF project. Director of Science Doug Kothe guides the research teams using the computing systems. And Director of Operations Jim Rogers manages day-to-day operations and planning for future systems and infrastructure.

SCIENTIFIC COMPUTING

The Scientific Computing Group (SCG) works concurrently with the users of HPC systems to help them obtain optimal results from the OLCF's computational resources. The SCG is made up of research scientists, visualization specialists, and work-flow experts who are trained in chemistry, physics, astrophysics, mathematics, numerical analysis, or computer science. Each research team using an OLCF system is assigned an SCG liaison who is familiar with the field of research. SCG liaisons actively participate in the research, help design and optimize code for the users' applications, streamline the work flow, and solve any computer issues that arise. Visualization specialists capture the resulting data in images and help the users analyze it. Contact Ricky A. Kendall, group leader, at kendallra@ornl.gov.

USER ASSISTANCE AND OUTREACH

The User Assistance and Outreach (UAO) Group provides support to the OLCF users, mediates between the outside world and the OLCF, and acquaints the public with the work conducted at the OLCF. The group creates accounts for news users, provides Level I and Level II technical support to the research teams, and generates documentation on OLCF systems access, policies, and procedures. UAO staff members attend the meetings of the other OLCF groups, communicating suggestions from users and participating in decisions on the users' behalf. The group also creates science research highlights, writes *SciDAC Review* and *HPCwire* articles, produces podcasts for scientists and the public, and connects the OLCF with universities across the nation. Contact Ashley D. Barker, group leader, at ashley@ornl.gov.

TECHNOLOGY INTEGRATION

The Technology Integration (TechInt) Group is responsible for updating and integrating the networks, file systems, and archival storage infrastructure into the OLCF computing systems. The group researches and evaluates emerging technologies and provides system programming to seamlessly integrate new technologies and tools into the infrastructure as they are adopted. TechInt co-developed the OLCF's High-Performance Storage System (HPSS) and is constantly working to increase the speed of data transfer and implement cybersecurity measures for the OLCF's area-wide network. As the OLCF computing resources continue to scale up, the TechInt group works to develop tools such as compilers, debuggers, and performance-analysis tools that allow users to take full advantage of the leadership-class systems. Contact Galen Shipman, group leader, at gshipman@ornl.gov.

HIGH-PERFORMANCE COMPUTING OPERATIONS

The High Performance Computing Operations (HPCO) Group keeps the OLCF leadership supercomputing systems running. Members of the group monitor all systems 24 hours a day, seven days a week, 365 days a year, and are responsible for administration, configuration management, and cybersecurity. The staff works with the infrastructure systems as well as with the Cray XT4/XT5 Jaguar and other OLCF supercomputers. The HPCO tests the systems when they are installed and upgraded, and uses diagnostic tools to continually monitor them. They anticipate problems before they arise and identify components that are near failure. The group also ensures that all systems conform to ORNL cybersecurity policy. Contact Ann Baker, group leader, at bakerae@ornl.gov.

APPLICATION PERFORMANCE TOOLS

The Application Performance Tools Group researches, tracks, and purchases a wide range of software tools that help science researchers access and improve the performance of their applications on current and emerging OLCF computing systems. The group also manages the contacts with the vendors for the purchase of new modeling tools, languages, middleware, and performance-characterization tools. The group primarily focuses on issues that arise for research applications when they are run on very large-scale systems, such as the OLCF's 2.33 petaflops Cray XT4/XT5 Jaguar supercomputer. Contact Rich Graham, group leader, at rlgraham@ornl.gov.

OLCF Resources

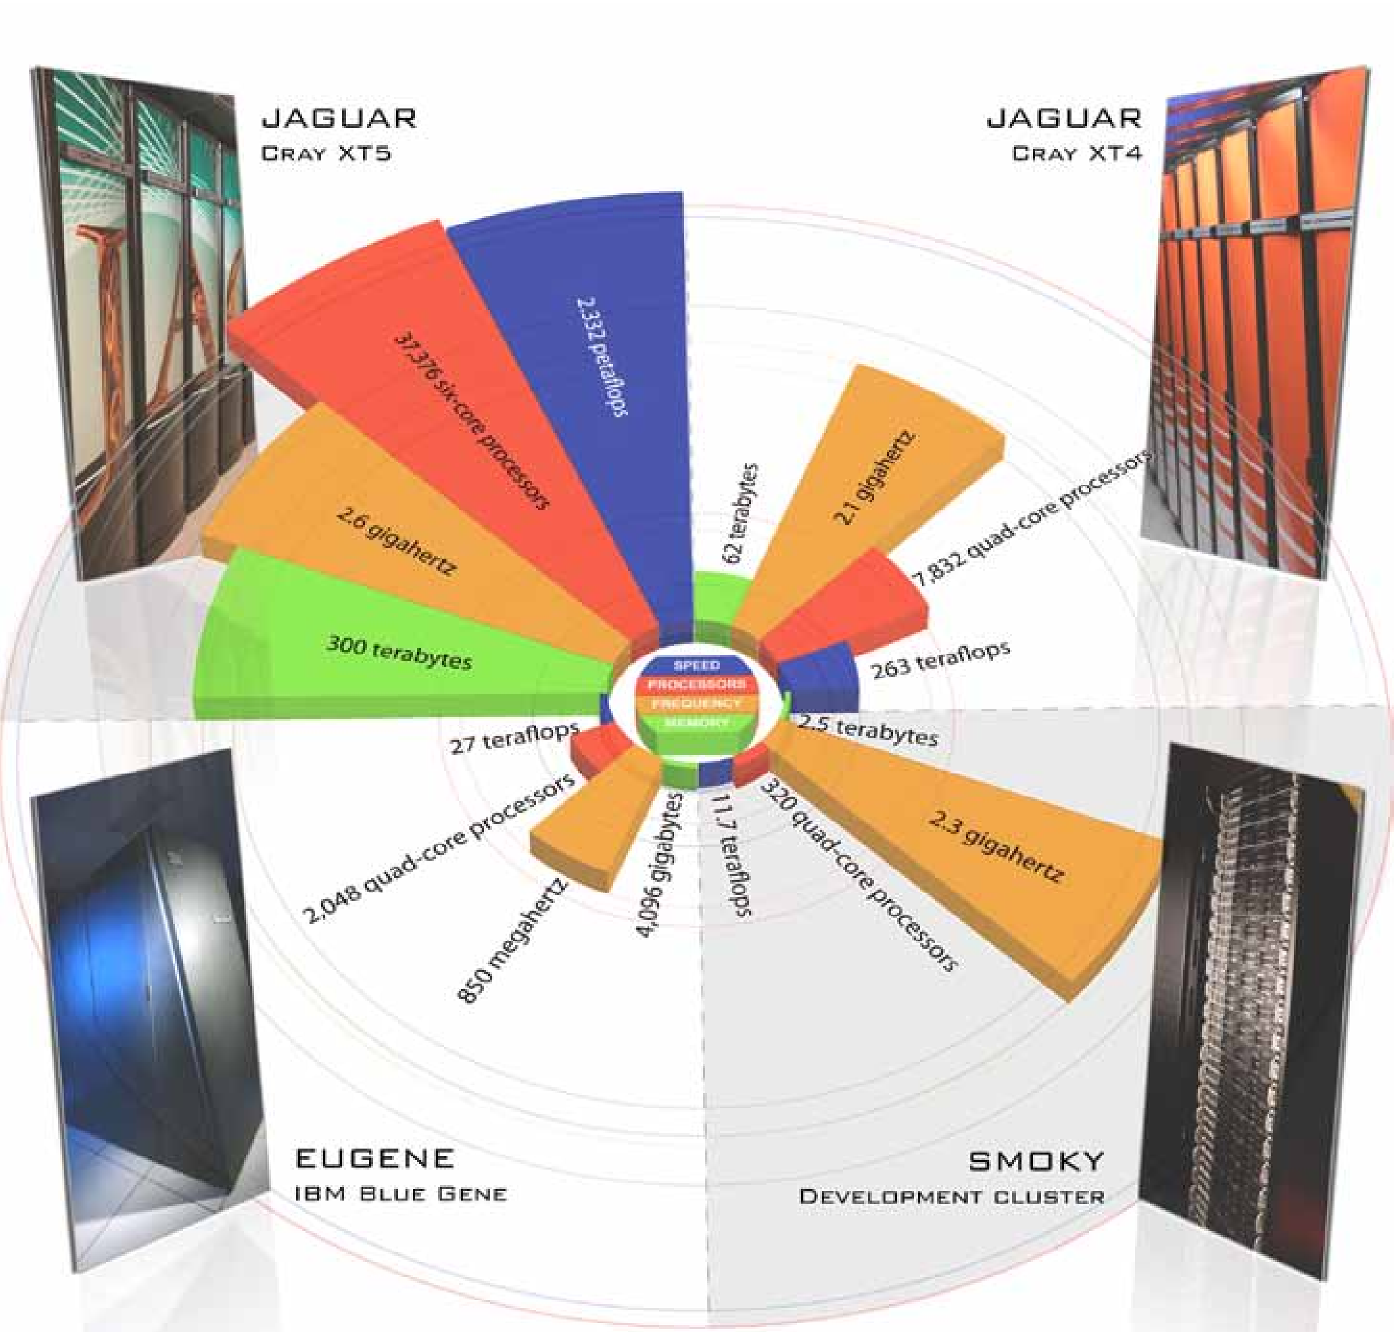

EVEREST
Scientific Visualization Lab

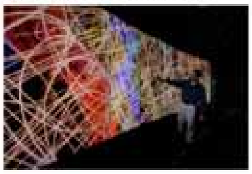

- 30 x 8 feet long
- 27-projector PowerWall
- 35 million pixels

HPSS
Back-up Storage

- Many storage devices supported
- Over 30 petabytes (PB) of capacity

LENS
Scientific Visualization Cluster

- 128 quad-core processors
- 2.3 gigahertz
- 2 terabytes of memory
- 2 NVIDIA 8800 GTX GPUs

COMPUTATIONAL SCIENCES BUILDING

- 40,000 square feet of raised floor
- 25 megawatts of power
- 6,600 tons of cooling
- 480v power
- 280 megawatt substation
- LEED certified by U.S Green Building

2009 INCITE Allocations by Discipline

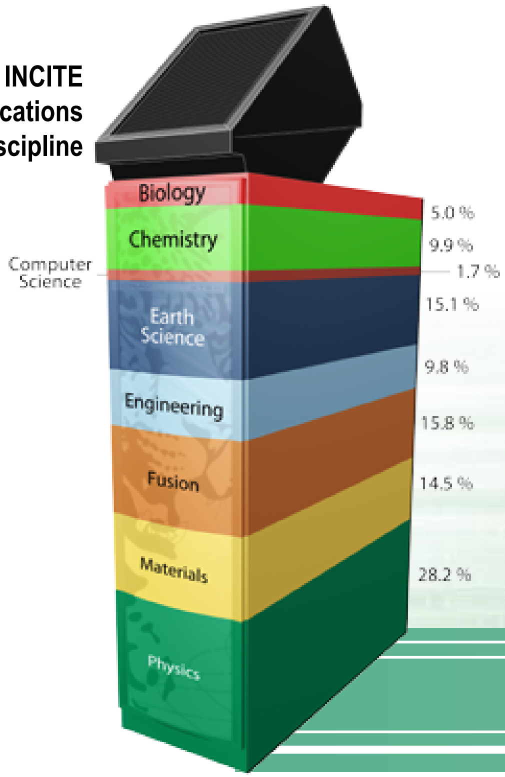

Active Users by Discipline

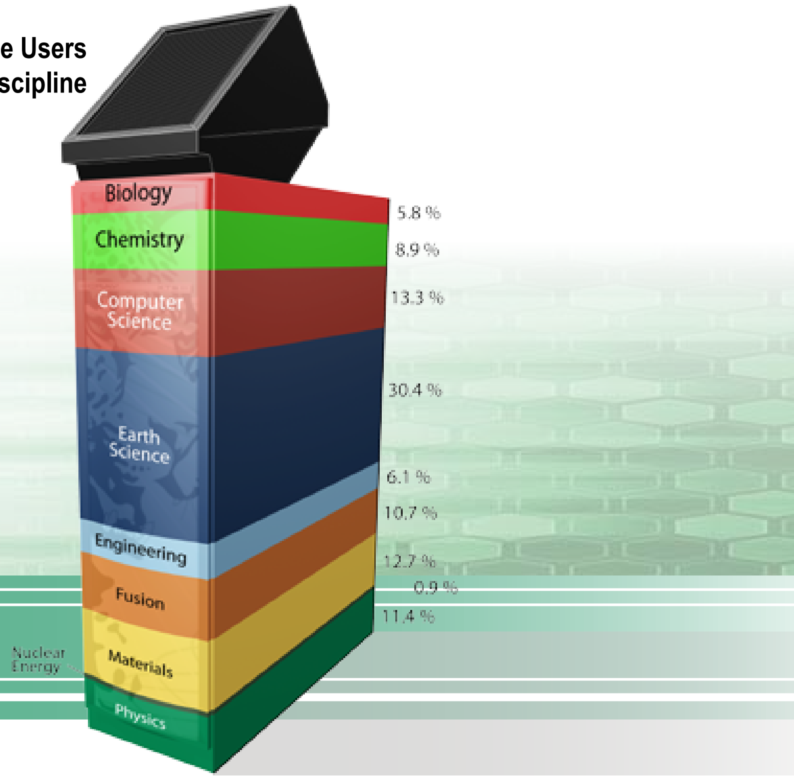

Active Users by Sponsor

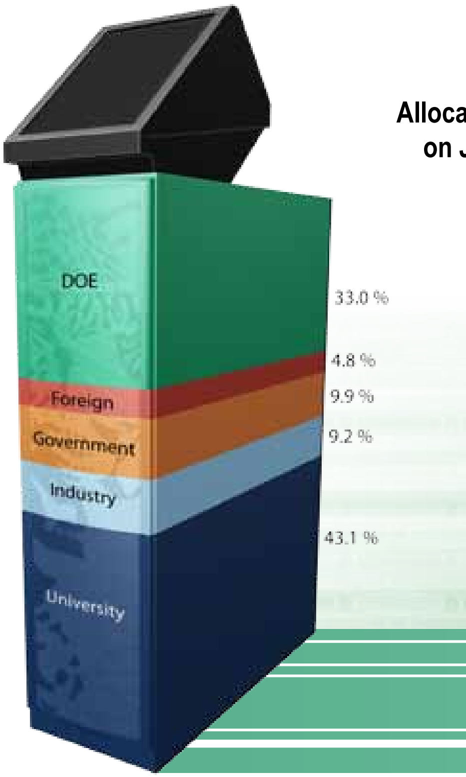

Allocation Hours on Jaguar XT5

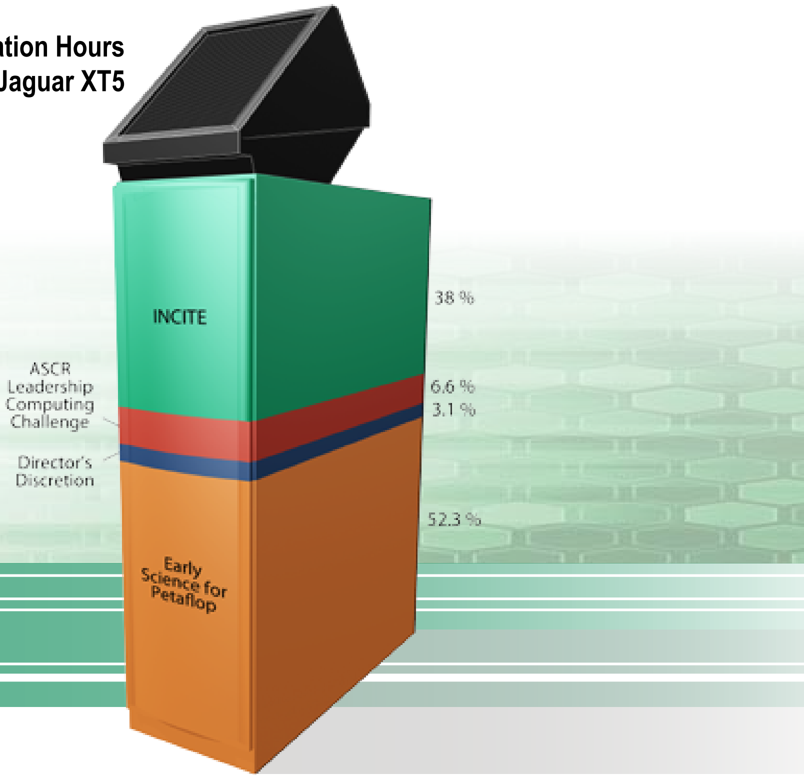

OLCF Publications

Researchers have produced numerous scientific breakthroughs using the resources of the OLCF, since the inception of the center. Listed below is a small sampling of hundreds of publications, grouped by related discipline, that highlights a portion of the work being achieved through the combination of talented researchers, leadership-class systems, and the dedicated staff of the OLCF.

Astrophysics

Hannam, M., et al. 2009. “Samurai project: Verifying the consistency of black-hole-binary waveforms for gravitational-wave detection.” Physical Review D **79**(8): 084025.

Liebendorfer, M., et al. 2009. “Supernovae as nuclear and particle physics laboratories.” Nuclear Physics A **827**(1–4): 573C–578C.

Zingale, M., et al. 2009. “Low mach number modeling of type Ia supernovae. IV. White dwarf convection.” Astrophysical Journal **704**(1): 196–210.

Chemistry

Cheng, X. L., et al. 2009. “Molecular-dynamics simulations of ELIC-a prokaryotic homologue of the nicotinic acetylcholine receptor.” Biophysical Journal **96**(11): 4502–4513.

Petridis, L., and J. C. Smith. 2009. “A molecular mechanics force field for lignin.” Journal of Computational Chemistry **30**(3): 457–467.

Climate

Liu, Z., et al. 2009. “Transient simulation of last deglaciation with a new mechanism for Bolling-Allerod warming.” Science **325**(5938): 310–314.

Washington, W. M., et al. 2009. “How much climate change can be avoided by mitigation?” Geophysical Research Letters **36**: L08703.

Fusion

Candy, J., et al. 2009. “Tokamak profile prediction using direct gyrokinetic and neoclassical simulation.” Physics of Plasmas **16**(6): 060704.

Zhang, W., L. Zhihong, and L. Chen. 2008. “Transport of energetic particles by microturbulence in magnetized plasmas.” Physical Review Letters **101**(9): 095001.

Materials

Mikelsons, K., et al. 2009. “Thermodynamics of the quantum critical point at finite doping in the two-dimensional Hubbard model studied via the dynamical cluster approximation.” Physical Review B **80**(14): 140505.

Nukala, P., et al. 2009. “Fast update algorithm for the quantum Monte Carlo simulation of the Hubbard model.” Physical Review B **80**(19): 195111.

Zhang, Y., et al. 2009. “Thermodynamic properties of PbTe, PbSe, and PbS: First-principles study.” Physical Review B **80**(2): 024304.

Nanosciences

Luisier, M., and G. Klimeck. 2009. “Atomistic full-band simulations of silicon nanowire transistors: Effects of electron-phonon scattering.” Physical Review B **80**(15): 155430.

Wang, L. W., Z. Zhao, and J. Meza. 2008. “Linear-scaling three-dimensional fragment method for large-scale electronic structure calculations.” Physical Review B **77**(16): 165113.

Physics

Hagen, G., T. Papenbrock, and D. J. Dean. 2009. “Solution of the center-of-mass problem in nuclear structure calculations.” Physical Review Letters **103**(6): 062503.

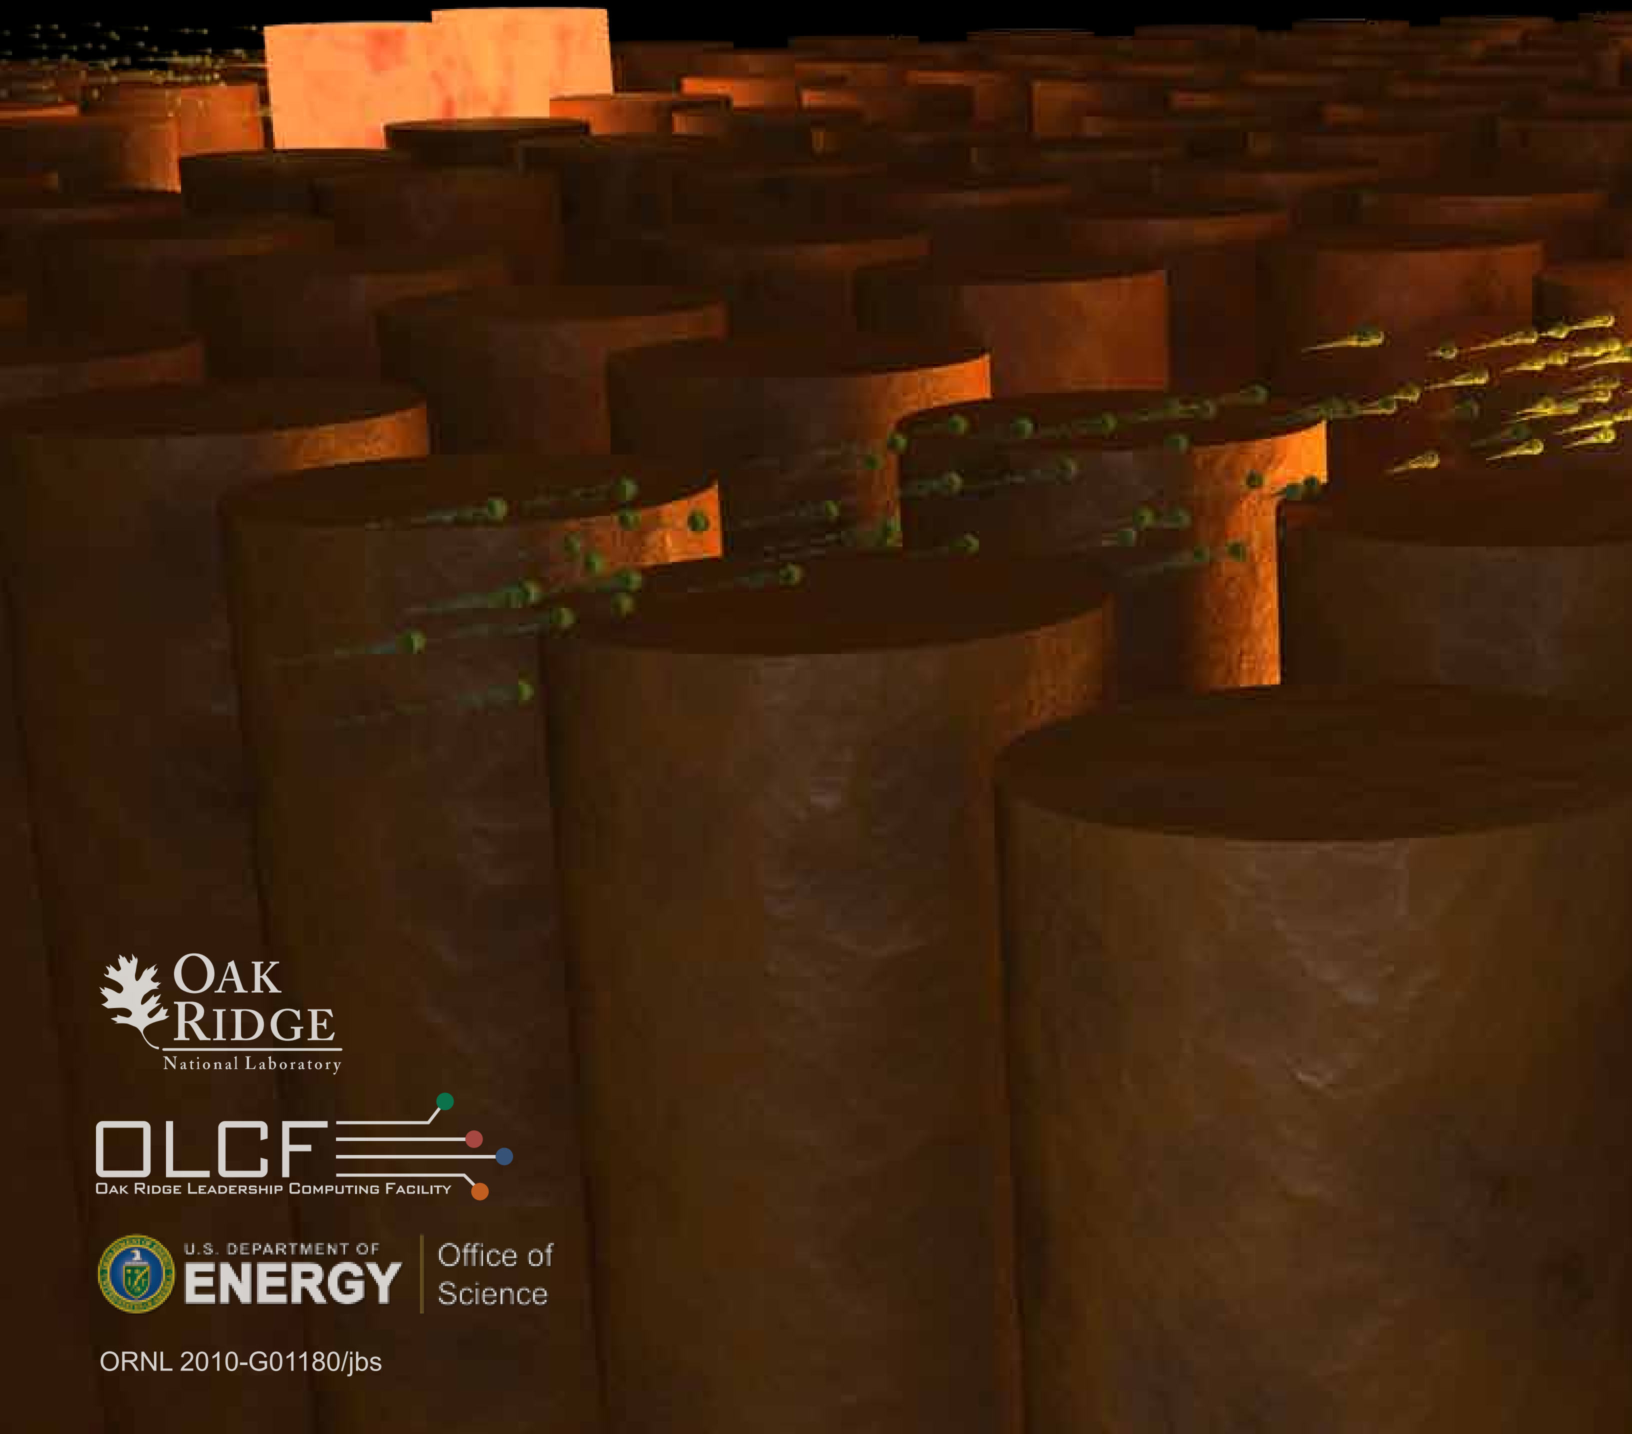

Oak Ridge National Laboratory

OLCF
Oak Ridge Leadership Computing Facility

U.S. Department of Energy
Office of Science

ORNL 2010-G01180/jbs