NCCS ANNUAL REPORT 2008

LEADERSHIP COMPUTING FOR SCIENCE

DOE DELIVERS PETASCALE SYSTEM IN 2008

NCCS
NATIONAL CENTER FOR COMPUTATIONAL SCIENCES

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

Contents

James J. Hack, Director
National Center for Computational Sciences

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.nccs.gov

The research and activities described in this report were performed using resources of the National Center for Computational Sciences 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.

NCCS's Jaguar, a 1.64-petaflop Cray XT5.

Beyond Linpack: ORNL Supercomputer Delivers Science as Well as Speed

'Jaguar' runs swiftest scientific simulation ever, gives insight into superconductors

Powerful. Versatile. Fast. Smart. For unclassified research the fastest cat in the high-performance computing (HPC) jungle is "Jaguar," a Cray XT system at the National Center for Computational Sciences (NCCS) located on the U.S. Department of Energy's (DOE's) Oak Ridge National Laboratory (ORNL) campus. Science's quickest supercomputer has a theoretical peak calculating speed of 1.64 petaflops (quadrillion calculations per second) and more memory than any other HPC system. With approximately 182,000 AMD Opteron processing cores working together, Jaguar can do in 1 day what it would take every man, woman, and child on Earth 650 years to do—assuming each could complete one calculation per second.

In November 2008 the Jaguar system was upgraded, and its XT5 and XT4 components were ranked second and eighth, respectively, on the TOP500 list of the world's swiftest supercomputers. DOE's "Roadrunner" IBM supercomputer at

Cray personnel installing new XT5 cabinets.

Los Alamos National Laboratory, specialized to run classified applications, took the top spot.

The TOP500 list ranks HPC systems based on speed alone. To earn rankings, supercomputers run High-Performance Linpack, a benchmark code that solves linear algebra problems. But to produce useful science results, it is not enough for a machine to be fast. It must also be balanced. In Austin at SC08, the premier international HPC conference, Jaguar proved it could also easily run a breadth of application codes tackling a variety of challenges in science and engineering. At the November 2008 conference, the Jaguar XT5 earned the Gordon Bell Prize for the world's fastest scientific computing application. The winning researchers used virtually all of the machine's approximately 150,000 processing cores running at a sustained speed of 1.352 petaflops to simulate superconducting materials. Jaguar XT5 also placed first in three of the six HPC Challenge benchmarks for excellence in handling computing workloads.

"The HPC community has finally evolved beyond Linpack," said ORNL's Thomas Zacharia, who during Jaguar's triumphs was associate laboratory director for computing and computational sciences before a 2009 promotion to coordinate ORNL's major research and development programs. "Jaguar's performance in the Gordon Bell and HPC Challenge competitions demonstrates what makes it the world's best machine for open science—an unprecedented balance among computing speed, memory size, and communication bandwidth."

Jaguar's architecture has remained consistent with its upgrades, so scientists have not had to dramatically modify their software codes as the number of processing cores in the system has increased. They can focus on their science instead of the mechanics of petascale computing. Through DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, researchers at universities, companies, and national labs will gain access to their first petascale supercomputer. In 2009 INCITE will make nearly 470 million processor hours on Jaguar available to 38 projects. That's more than half the 889 million processor hours INCITE allocates nationwide. Moreover, to speed product innovations and reduce production costs, ORNL has established an industrial partnership program to help American companies exploit petascale computing to gain or retain a competitive advantage.

Past scientific discoveries were built on the twin pillars of theory and experiment. Tomorrow's discoveries may also rest on a third pillar—simulation. The unprecedented speed of petascale systems means simulations can arrive at answers faster than ever and explore complex, dynamic systems in ways never before possible. Researchers can simulate complex phenomena on the smallest and largest conceivable scales of time and space. Their investigations can mirror the one-hundred billionths of a second of a supernova's core collapse or the billions of years of the universe's evolution. They can computationally probe systems that are minuscule (for example, correlated electron spin in exotic materials) or majestic (such as the earth's changing climate).

The Path to Petascale

In 2002 Japan installed the Earth Simulator, an NEC supercomputer that could model planetary processes with a then-unprecedented speed of 36 teraflops (trillion calculations per second). In response the United States built the Jaguar XT, born in 2005 as a 26-teraflop system at the NCCS, which selects, procures, deploys, and operates ORNL's leadership computing systems for DOE.

Over the years Jaguar was expanded through a series of upgrades, all the while keeping the architecture consistent to facilitate users' software codes. By May 2008 its XT4 component had grown to 263 teraflops. The XT4 was then "accepted," meaning it had passed stringent performance tests required to authorize payment for a commissioned system.

THE UNPRECEDENTED SPEED OF PETASCALE SYSTEMS MEANS SIMULATIONS CAN ARRIVE AT ANSWERS FASTER THAN EVER AND EXPLORE COMPLEX, DYNAMIC SYSTEMS IN WAYS NEVER BEFORE POSSIBLE.

The XT4 subsequently entered a "transition-to-operations" phase, dubbed T2O, during which six computationally demanding applications each ran on as many as 31,000 processing cores for up to 4.5 million processor hours to generate data about fuel combustion, core-collapse supernovas, fusion plasmas, chemical catalysts, ocean turbulence, and materials science. The activities of the T2O pioneers smoothed the way for Jaguar to become available to general users with supercomputing allocations through the INCITE program.

Another upgrade, in November 2008, added a 1.382-petaflop XT5 component to Jaguar, taking the system to 1.64 petaflops and making it the first petascale system dedicated to use by the open scientific community. Jaguar has more memory than any other computer (362,000 gigabytes), a file system with enormous data-storage capacity (10 petabytes), and a speedy means of moving those data (578 terabytes per second of memory bandwidth). Its unprecedented input/output (I/O) bandwidth of 284 gigabytes per second quickly moves data into and out of processors, making it nimble at complex simulations in energy, climate, chemistry, and other key research domains.

Preparations to run acceptance codes on the XT5 set world records for two important scientific applications. The DCA++ code illuminated the behavior of high-temperature cuprate superconductors, and the LSMS Wang-Landau code elucidated a complex magnetic system. Both codes, employed by an ORNL research group led by computational materials scientist Thomas Schulthess, ran at speeds exceeding a petaflop.

On December 29, 2008, Jaguar's XT5 component was officially accepted. Achieving that milestone was the culmination of a Herculean effort by Cray and ORNL professionals working around the clock to deliver cabinets; connect power and cooling water; and test software, hardware, operating systems, I/O, and network function. Experts at the Cray Supercomputing Center of Excellence and ORNL staff optimized code performance, trained users, and created libraries of routines.

As with the XT4, the XT5 T2O then allowed eight select application codes to run on the petascale system to ready it for general use. These applications included a new, strategic user community (weather), XT4 T2O projects (materials science, climate, combustion, and fusion), and Gordon Bell winners and finalists (materials science and seismology). Many codes carried out the largest calculations ever performed for their scientific fields. Among the Gordon Bell winners given early access to Jaguar were researchers from DOE's Lawrence Berkeley National Laboratory who are creating an algorithm to illuminate the energy-harnessing potential of nanostructures. To predict the efficiency of a material for solar cells, over 2 days they ran an optimized application on the XT5 that achieved a speed of 442 teraflops using the majority (147,146) of Jaguar's processing cores. Their work demonstrates the ease of porting a scientific application to Jaguar and running it at sustained petaflops to produce insight-generating data.

Thought Leadership

If mathematics is the language of science, computation is its workhorse—and will be for the foreseeable future. As the number one petascale computing complex for open science, ORNL hosts world-class resources for accelerating discoveries. Besides the second- and eighth-ranked Jaguar components, in 2008 ORNL also became home to the National Science Foundation's (NSF's) Cray XT4, known as Kraken and ranked

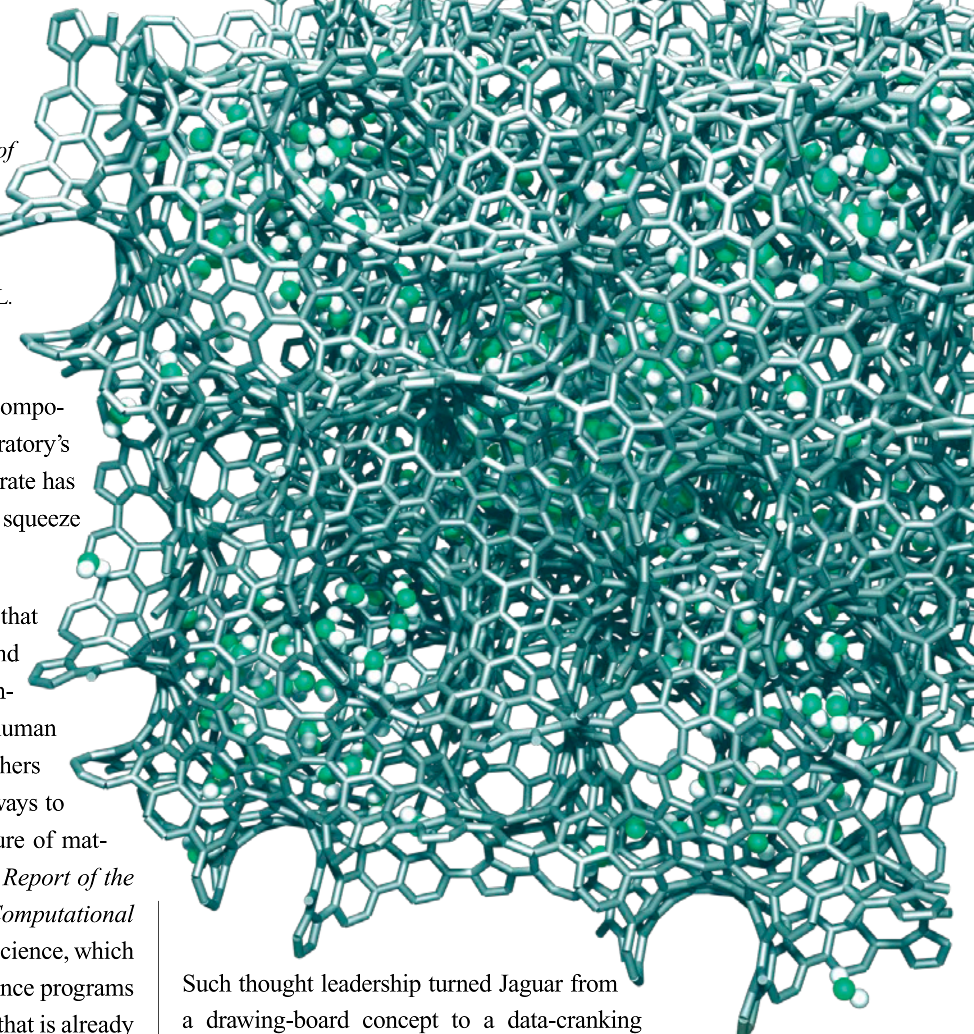

Molecular dynamics simulation of confinement and dispersion of small molecules within carbon nanostructures, mimicking the dynamics of electrolytes in porous carbon materials. The visualization was generated by the SciDAC code VisIt. Simulation by Dr. Vincent Meunier, ORNL; and visualization by Jeremy Meredith and Sean Ahern, ORNL.

fifteenth in the TOP500 list. In 2009 an XT5 component was added to Kraken. Additionally, the laboratory's Computing and Computational Sciences Directorate has more than 600 experts on-site to help researchers squeeze science from the supercomputers.

Jaguar's historic achievements in 2008 proved that petascale machines can produce data that lend insight into grand challenges in science and engineering. Those insights are likely to take human knowledge a giant leap forward as researchers reveal the future of regional climates, develop ways to tap new energy sources, and delve into the nature of matter and the origins of life. *Breakthroughs 2008: Report of the Panel on Recent Significant Advancements in Computational Sciences*, a 2008 report from the DOE Office of Science, which is America's largest funder of basic physical science programs at universities and national laboratories, showed that is already happening. Five of ten top advances in computational science, the report said, used ORNL supercomputers to gain insight into supernovas, combustion, fusion, superconductivity, dark matter, and mathematics.

Computational scientists and other experts at ORNL have deeply engaged with scientists worldwide in research areas relevant to meeting global energy demands. They have systematically studied the science drivers for computing at the petascale and exascale (a thousand times faster than petascale) to accelerate progress in strategic areas such as biofuels, fusion energy, battery technology, and climate science. They have freely shared their findings with the HPC community through publications, speeches, and service on steering committees.

Such thought leadership turned Jaguar from a drawing-board concept to a data-cranking dynamo poised to deliver transformational solutions in energy, climate, and health. Delivery of a petascale machine to America's top scientists in barely 4 years required experience, insight, strategy, and diligence. More of the same will be needed to guide the nation as it responds to scientists' demands for more computing power. DOE's leadership computing facilities are on track to deliver a projected capacity of hundreds of petaflops by 2015 and exaflops by 2018. And scientists say they will be able to use every bit of it.—*by Dawn Levy*

Preparing for Science at the *Petascale*

Six pioneering applications were chosen to perform science-at-scale simulations on the NCCS's Jaguar following its upgrade to 263 teraflops. The running of these applications is part of a "transition-to-operations" activity, dubbed T2O, which gave each of these pioneering applications up to 4.5 million hours for the first data-production runs to applications that could concurrently use the majority of Jaguar's 31,000 processing cores. This T2O period not only tested Jaguar's mettle, but also accomplished breakthrough science in a number of arenas.

Burn, Baby, Burn

Jaguar lights the way to energy independence

America's dependence on foreign oil is becoming more than a thorn in the country's side—it's now widely viewed as an unsustainable, and potentially dangerous, liability.

One of the more promising remedies for our import addiction is the deployment of engines, furnaces, and power-generation devices that burn alternative fuels and use advanced technologies. These next-generation combustion technologies will not only reduce our dependence on foreign oil but also likely help reduce the amount of pollution generated by traditional internal combustion engines in automobiles.

Considering that two-thirds of the petroleum Americans use goes for transportation, while the remaining one-third heats buildings and generates electricity in steam turbines, the necessity for advanced combustion devices is a no-brainer. Their deployment directly addresses two key DOE mission objectives: energy security and environmental responsibility.

To bring these technologies to market, a team led by Jacqueline Chen of Sandia National Laboratories (SNL) is simulating the combustion of various fuels on Jaguar, research that is creating a library of data that captures complex aero-thermo-chemical interactions and provides insight into how flames stabilize, extinguish, and reignite. Chen's data libraries will assist engineers in the development of models that will be used to design the combustion devices of tomorrow.

"If low-temperature compression ignition concepts employing dilute fuel mixtures at high pressure are widely adopted in next-generation autos, fuel efficiency could increase by as much as 25 to 50 percent," Chen said. With mechanical engineer Chun Sang Yoo of SNL and computational scientist Ramanan Sankaran of the NCCS, Chen used Jaguar to simulate combustion of ethylene, a hydrocarbon fuel. The simulation generated more than 120 terabytes (120 trillion bytes) of data.

Chen's simulations were a major contributor to the NCCS's T2O period. In an effort to revolutionize combustion technology, the NCCS granted Chen and her team time on Jaguar following the system's upgrade to 263 teraflops. Running computationally demanding applications after a major machine upgrade is part of a T2O activity that begins when a commissioned NCCS system passes a formal acceptance test and its performance is monitored and assessed. Essentially, researchers with compelling science problems and codes that could use a large fraction of the machine were selected to run on one of the most powerful supercomputers in the world. Chen's direct numerical simulation (DNS) code, known as S3D, was a perfect fit.

Exploring Lifted Flames

Advanced combustion technology depends on lifted flames, which result when cold fuel and hot air mix and ignite in a high-speed jet. If the speed increases too much, lifted

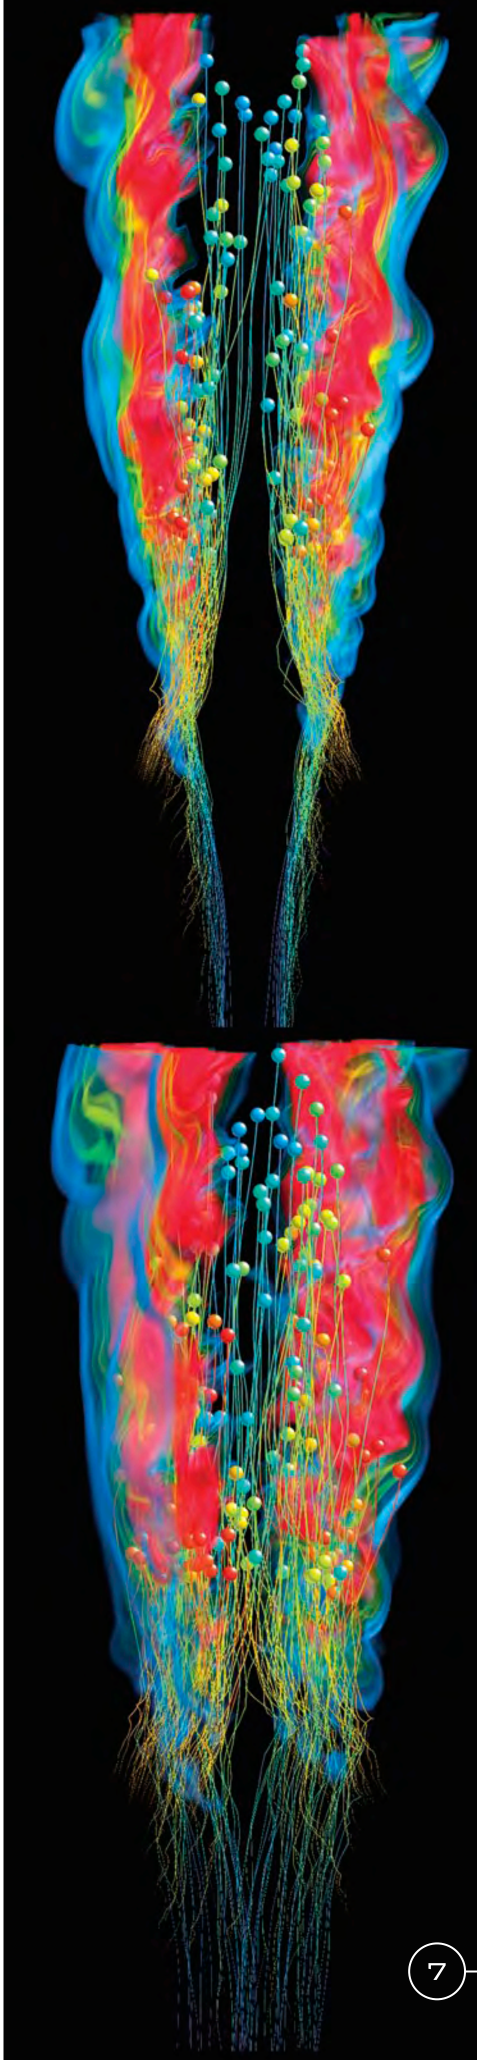

Simultaneous volume rendering of a lifted ethylene/air slot jet flame, where the lifted flame is represented by hydroxyl radical showing the flame stabilization point and representative particle tracks originating from the stabilization point and tracked forward and backward in time. The particles are colored by temperature: cold (blue), hot (red). Image courtesy Jackie Chen, SNL, and Kwan-Liu Ma, UC–Davis.

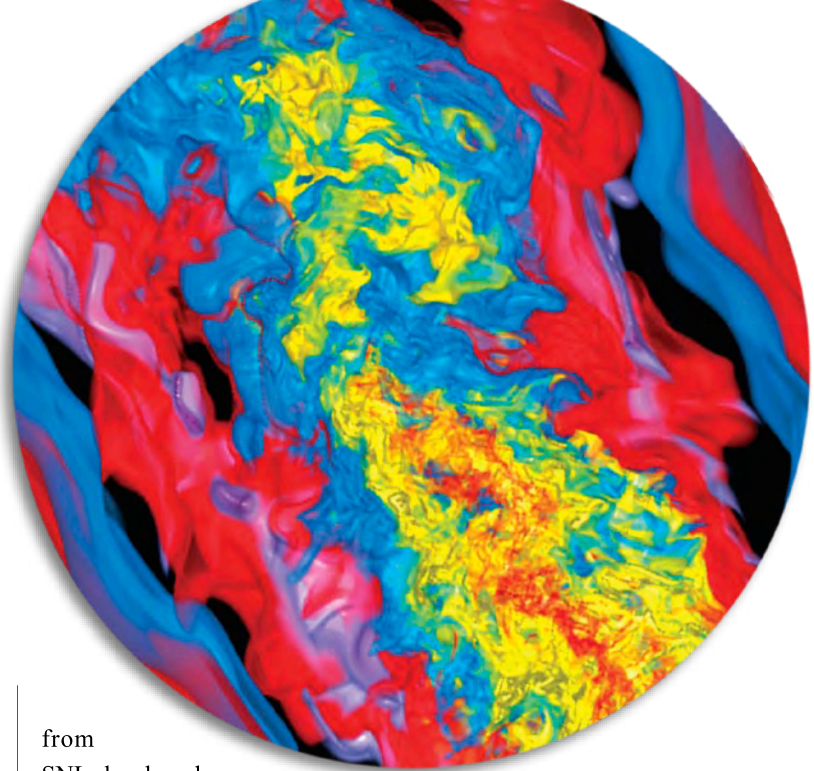

flames can blow out. For flames to stabilize, or continue to burn downstream from the burner, turbulence, which mixes fuel with air to enable burning, must exist in balance with key ignition reactions that occur upstream of where the flames appear.

Proper positioning of lifted flames in advanced engines could burn fuel so cleanly that emissions of nitrogen oxide, a major contributor to smog, would be nearly undetectable. To explore processes underlying ethylene combustion, the group uses DNS, a high-fidelity numerical approach that fully resolves all of the temporal and spatial scales associated with both turbulence and flames. The technique uses a computational mesh of more than a billion points to capture both fast ignition events and the slower motion of turbulent eddies.

"Direct numerical simulation is our numerical probe to measure, understand, or see things in great detail at the finest scales where chemical reactions occur," Chen said. "That's particularly important for combustion because reactions occurring at the finest molecular scales impact global properties like burning rates and emissions."

S3D runs on multiple processing cores to model compressible, reacting flows with detailed chemistry. The 2008 T2O simulations used up to 30,000 of Jaguar's 31,000 processing cores and 4.5 million processor hours. Furthermore, Chen and her colleagues created the world's first fully resolved simulation of small lifted autoigniting hydrocarbon jet flames, designed to represent some of the fine-grained physics relevant to stabilizing in a direct-injection diesel engine.

Chen's simulations also showcased the relationship between the NCCS liaisons and principal investigators. For the T2O allocation, the NCCS's Ramanan Sankaran developed a new feature for the S3D code that allowed the team to track tracer particles throughout the simulated flames. The tracer, which doesn't alter the data and can be programmed to track different phenomena in the flame, is analogous to dropping a bottle with a tracking device in the ocean, said Sankaran. As the ocean currents move the bottle across the sea, the tracking device records all sorts of information, not only telling the recipient where the bottle has been, but also what it has been through. "It gives us a better look at what's going on inside the flame," said Sankaran, adding that "it also enables a different view of the same data."

These next-generation combustion technologies will...reduce our dependence on foreign oil...

Because of the large number of particles transported, mechanical engineer Ray Grout and computer scientist Hongfeng Yu from SNL developed a graphics processing unit–enhanced framework for interactive Lagrangian particle query and analysis on the analytics machine at ORNL known as Lens.

The sheer scale of Chen's simulations has enabled an unprecedented level of quantitative detail in the description of both turbulence and chemistry and their interactions. And while ethylene is not a common transportation fuel, it is an important reference fuel in research. Ethylene's relatively simple chemistry makes it a good candidate for DNS. Moreover, it is easy to work with in the lab, and its chemical properties are well known, so it is feasible to test results from experiments against predictions from numerical simulations.

"DNS of turbulence-chemistry interactions requires the horsepower of machines like Jaguar and beyond to simulate the coupling between complex hydrocarbon ignition kinetics and turbulent mixing," said Chen. "It is particularly important to understand how differences in chemical properties of various fuels—biofuels and petroleum—impact combustion phenomena such as ignition and extinction in a turbulent environment."

Chen's simulations are helping engineers develop unprecedented models of combustion in a variety of environments. These models will eventually be used in engineering computational fluid dynamics simulations to optimize the design of engines and power-generation devices, a crucial ingredient in America's quest for both cleaner energy and energy independence.—*by Scott Jones*

Big Stars, Small Wonders

Researchers model role of neutrinos in core-collapse supernovas

Life as we know it is the product of core-collapse supernovas (CCSNs), or stars ten to twenty-five times more massive than our sun that implode and scatter the universe with the elements necessary for life. These astronomical heavyweights are the most important sources of elements between oxygen and iron and are responsible for half of the elements heavier than iron. We are, quite literally, the stuff of stars.

Now that they know where we come from, scientists want to know the details, such as how and why these stars explode. This knowledge will ultimately provide valuable clues to our origins and the nature of the universe. Researchers led by ORNL's Tony Mezzacappa are gaining ground; they used the 2008 T2O period to advance the state of the art in understanding these cosmic cataclysms.

> It's the only code in the world that joins multidimensional hydrodynamics with multienergy neutrino transport and detailed nuclear burning, leading to simulations with significant realism.

Specifically, the team used its T2O allocation to begin the world's first three-dimensional CCSN simulation with sophisticated neutrino transport. The importance of neutrinos is hard to overstate because CCSNs are neutrino-driven events. In fact, the necessary binding energy for the resulting neutron star following a CCSN is 1,053 ergs, almost the equivalent of a trillion trillion billion Fat Man atomic bombs.

According to observations, the kinetic energy and light combined account for roughly 1 percent of the neutron star's binding energy. Hence the neutrinos, these tiny, virtually undetectable particles, seem to account for 99 percent of the binding energy necessary for the postexplosion neutron star. When it comes to CCSNs, big things really do come in small packages.

Mezzacappa's team earlier discovered the standing accretion shock instability (SASI), a destabilizing phenomenon that had never been theorized and was revealed only through simulations conducted on the NCCS's Phoenix system, a retired Cray X1E, work that netted the team an article in *Nature*. However, neutrinos were not included in the first simulations because of the computational expense. With its T2O allocations, however, the team was able to gauge the effect of neutrinos on the SASI.

The team ultimately burned through 4.6 million processor hours and generated roughly 3 terabytes of data simulating both the transport of neutrinos through the star and the SASI. Essentially, said team member Bronson Messer, the new simulations showed that the previous SASI simulations were reasonable approximations because adding neutrinos and nuclear burning, which weren't included in the previous SASI runs, didn't seem to strongly suppress the phenomenon.

A Stellar Biography

The death of a massive star is built on the details of its evolution. Over millions of years the star burns through its nuclear fuel, fusing nuclei to create increasingly heavy elements. These elements form into layers, from hydrogen at the surface through helium, carbon, oxygen, and silicon to iron at the core. The process hits a roadblock at iron, which does not create energy when the nuclei fuse. Eventually the iron core becomes so heavy it collapses under its own weight. When the core reaches a density exceeding that of an atomic nucleus (10^{14} grams per cubic centimeter, or roughly the weight of the whole human race in something the size of a sugar cube), it bounces, creating a shock that eventually blows the star apart. Recreating this pro-

Volume-rendered sequence of the development of the standing accretion shock instability in a 3D simulation. Visualization by David Pugmire, ORNL.

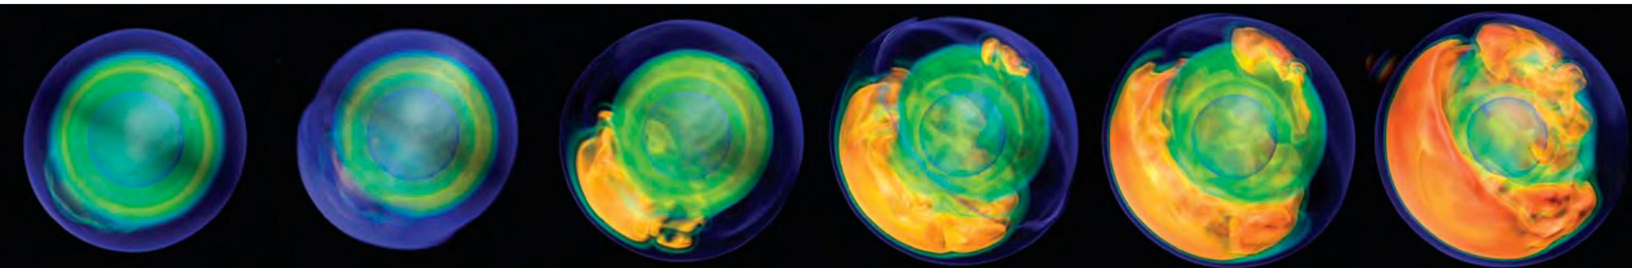

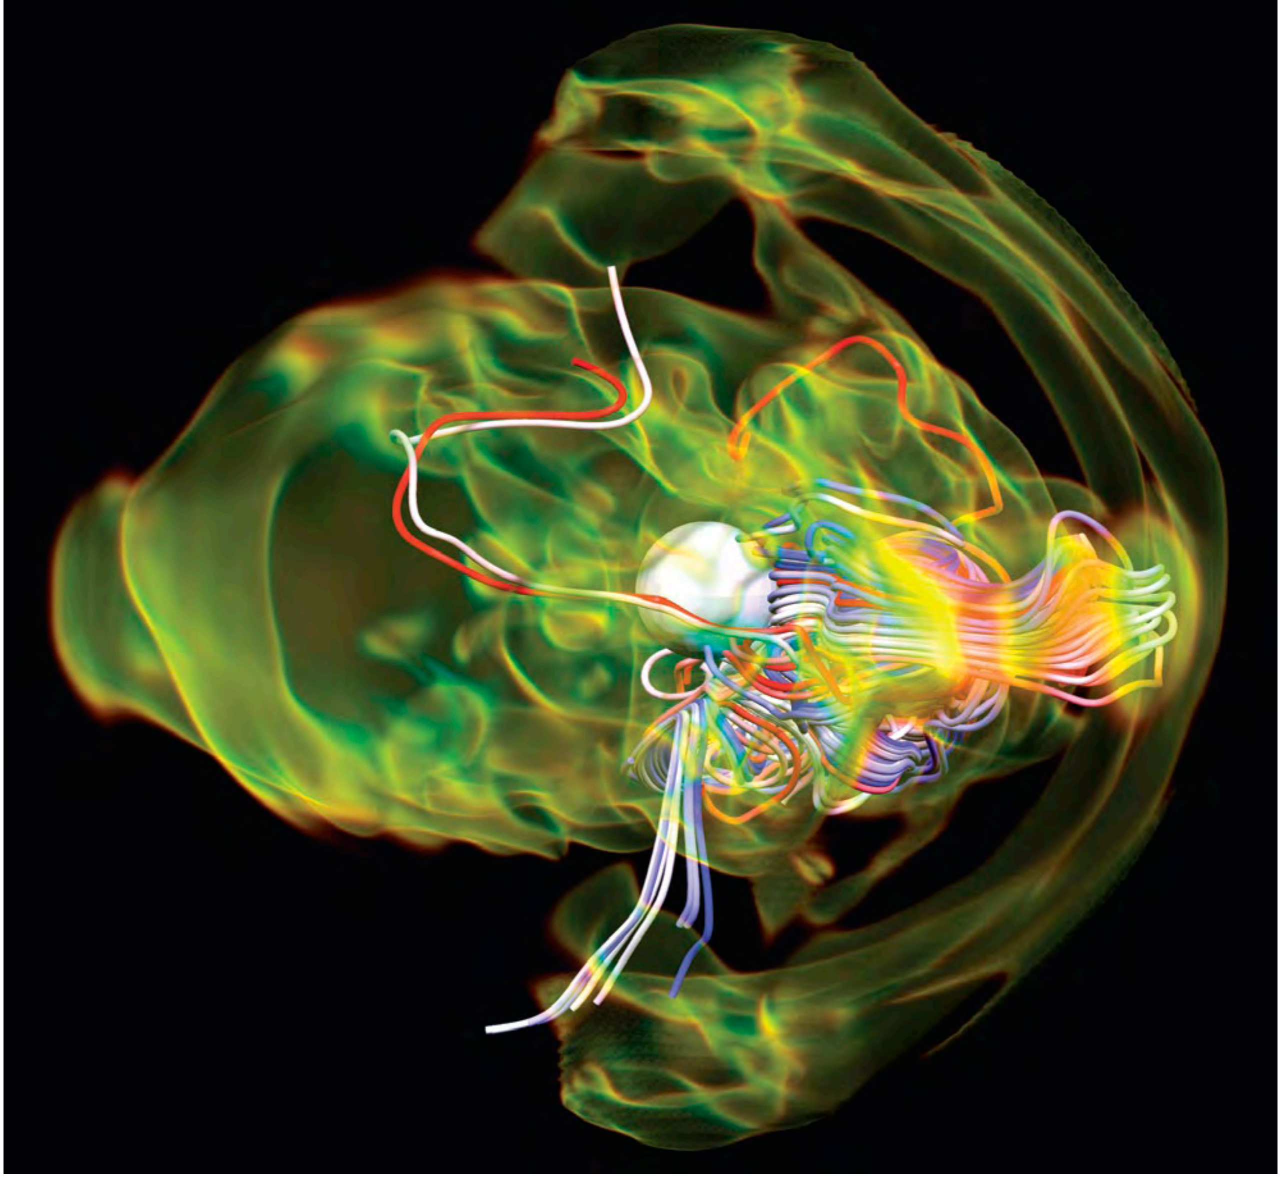

Fluid velocity streamlines during a type II supernova collapse. Visualization by D. Pugmire, ORNL, and simulation by E. Endeve, C. Cardall, R. Budiardja, ORNL and UT–Knoxville, and A. Mezzacappa, ORNL.

cess computationally is like reverse-engineering your favorite entrée; you know the dish but have no recipe.

"To use a cooking analogy," Mezzacappa explained, "you're adding key ingredients in a recipe. If a recipe works, you have a soufflé. But if you go in and remove a key ingredient, either you have no soufflé or you have something unfit to eat.

"What we're doing here is finding the recipe for a core-collapse supernova. You add the ingredients—you add the physics—and you then carry out the simulations and see what the outcomes are. At the end of the day, if your model makes predictions that match observables, then your model is good."

The cookbook currently used by Mezzacappa's team is a sophisticated software code known as CHIMERA, named after the three-headed beast in Greek mythology. Essentially, CHIMERA incorporates three important ingredients in a supernova explosion: hydrodynamics, neutrino transport, and the nuclear burning inside the star.

This beast's first head, known as MVH3, calculates the hydrodynamics and is tweaked for superior scaling to take advantage of Jaguar's processing power. The neutrino transport head, known as MGFLD-TRANS, monitors neutrino transport along radial rays extending from the center of the star. And the final head, the burning code XNET, tracks how the elements in the core of the star are fusing or breaking apart, giving researchers an invaluable look at the evolution of the star's nuclear composition. Taken together these codes make one formidable, extremely scalable beast.

"It's the only code in the world that joins multidimensional hydrodynamics with multienergy neutrino transport and detailed nuclear burning, leading to simulations with significant realism," said Messer.

The team is continuing to refine its simulations on the new petascale Jaguar, building on the theoretical confirmation obtained using its 2008 T2O allocation and proving once again that this preliminary period does much more than test the machine; it also achieves scientific results.—*by Scott Jones*

Simulating the Sun

T2O sheds light on mystery of electron transport

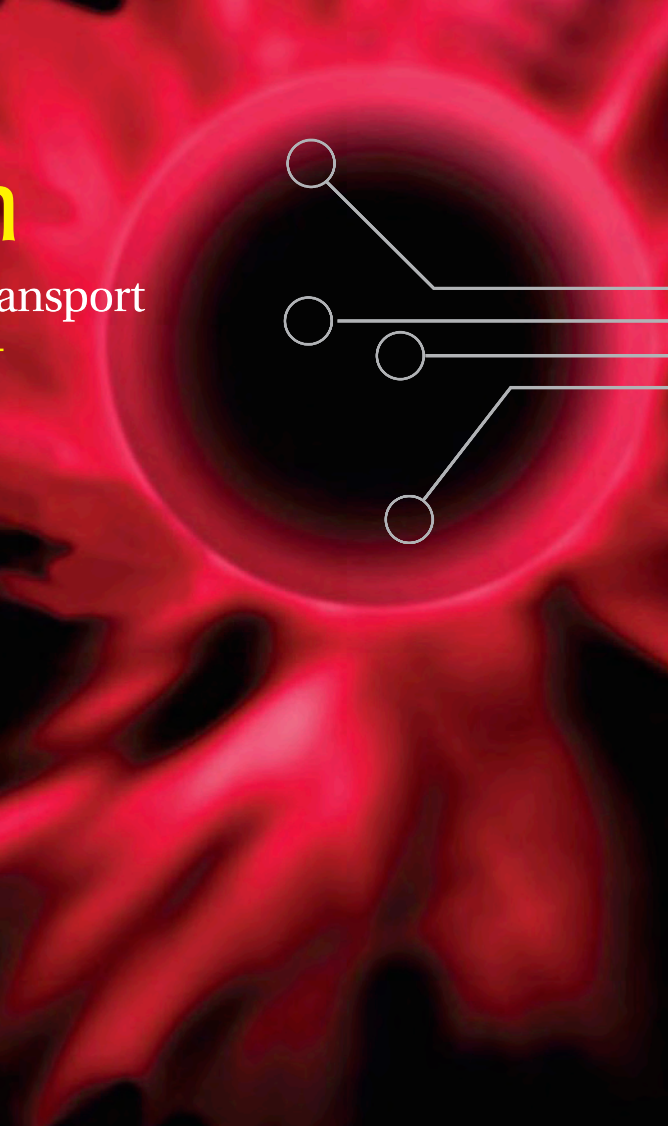

Few areas of science benefitted more from T2O or pushed Jaguar harder than fusion, an arena that looks to provide the world with a cleaner, virtually unlimited energy source in the form of ionized hydrogen.

The first breakthrough came when a team of researchers from the University of California–Irvine (UCI), in conjunction with NCCS staff, set the pace with a fusion simulation that was at the time the largest in history. Led by Yong Xiao and Zhihong Lin of UCI, the team used the classic Gyrokinetic Toroidal Code (GTC) to give researchers an unprecedented understanding of fusion plasmas.

Successful fusion energy depends on the ability to heat and electromagnetically confine the plasma inside the tokamak, or donut-shaped reactor, for a sufficient period of time. Unfortunately, previous research has shown that heat transport for both the ions and electrons in the plasma is greater than theory predicts (on the order of two magnitudes greater for electron heat transport). If current tokamak designs aren't modified to accommodate this larger-than-expected heat transport phenomenon, the plasma confinement will be insufficient and fusion energy will remain on the drawing board.

To successfully produce a fusion reaction, an extremely hot ionized gas known as a plasma must be magnetically confined for a sufficient period of time. Previous research has shown that heat transport, both within and out of the plasma and for both the ions and the electrons in the plasma, is far greater than theory predicts. This larger-than-expected heat transport out of the plasma could lead to significant heat and particle loss, and thus confinement failure, within 1 second, quickly rendering energy production impossible if reactor designs are not modified to accommodate it.

"The success of fusion research depends on good confinement of the burning plasma," said Xiao. "This simulation size is the one closest to ITER in terms of practical parameters and proper electron physics."

The researchers discovered, among other things, that for a device the size of ITER, the containment vessel will demonstrate GyroBohm scaling, meaning that the heat transport level—or amount of heat loss—will be inversely proportional to the device size. In other words the simulation supports the ITER design: a larger device will lead to more efficient confinement.

Because fusion simulation is such a complex process, today's leadership-class supercomputers such as Jaguar, with their fast CPU speeds and large memory, are an absolute necessity. However, the huge amounts of data produced by fusion simulations can create I/O nightmares: in one GTC run, the team can produce terabytes of data (in this case, 60 terabytes). To address this potential bottleneck, the team used ADIOS, a set of library files that allows for easy and fast file input and output, developed by a team led by the NCCS's Scott Klasky and Chen Jin and Georgia Tech's Jay Lofstead and Karsten Schwan.

"This huge amount of data needs fast and smooth file writing and reading," said Xiao. "With poor I/O the file writing takes up precious computer time, and the parallel file system on machines such as Jaguar can choke. With ADIOS the I/O was vastly improved, consuming less than 3 percent of run time and allowing the researchers to write tens of terabytes of data smoothly without file system failure."

ADIOS (see page 20) is a good example of the symbiotic relationship between researchers and NCCS staff. "My experience with Cray and the NCCS has been very good," said Xiao. "The account and operation staffs are very accessible and responsible, which enabled us to run the GTC code with a total of 28,000 cores smoothly for 2 days. [The NCCS's] Scott Klasky provided an effective channel for technical communication between the science application team and the computational support team."

A 3D simulation of plasma microturbulence in a fusion reactor, simulated on Jaguar. Visualization by Chris Ho, Chad Jones, and Kwan-Liu Ma of UC–Davis and the SciDAC Institute for Ultrascale Visualization.

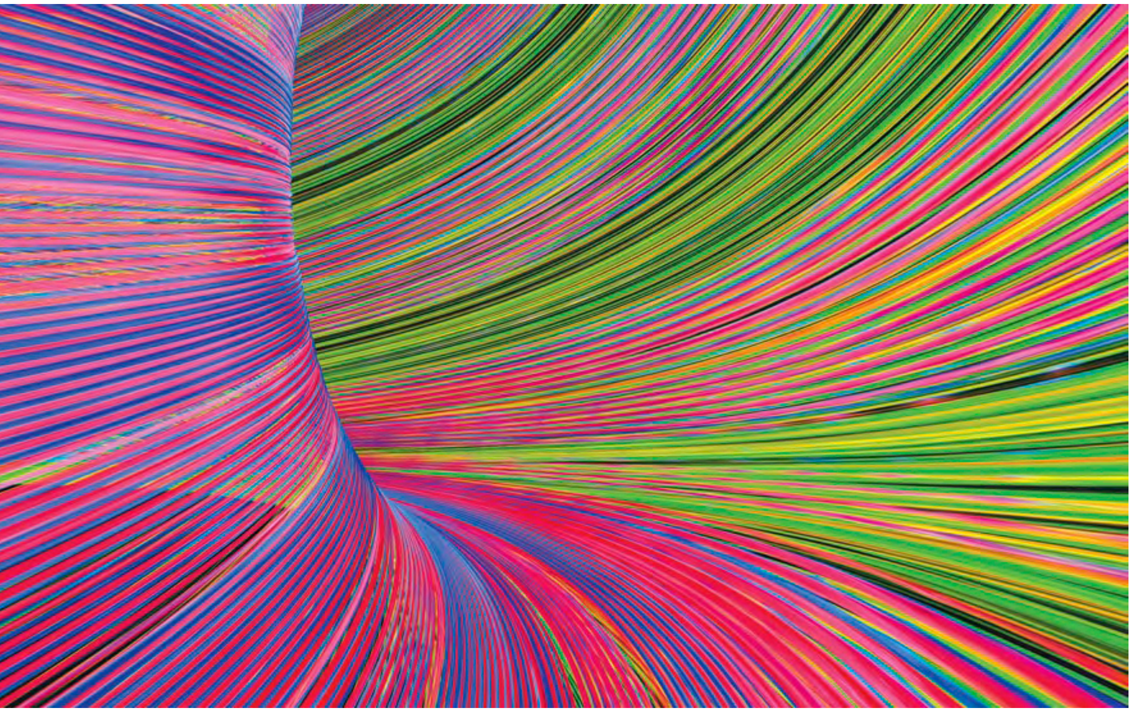

Despite the success of the simulation, there is still much work to be done. "Plasma turbulent transport is still an ongoing research area," said Xiao.

Records Were Made to Be Broken

The GTC record wouldn't stand for long, however. Within weeks a different team of researchers using a different code would set yet another record on Jaguar, bringing the world one step closer to commercially viable fusion power.

That team, led by physicist Weixing Wang and computational scientist Stephane Ethier of the Princeton Plasma Physics Laboratory (PPPL), used GTS, a fusion code designed to track energy and particle loss, or transport, in an experimental magnetic fusion device. One T2O simulation used 99.7 percent of the entire machine, running for 72 hours on 31,232 of Jaguar's cores, netting 2.25 million CPU hours and more than 100 terabytes of data.

The goal of these simulations, said Ethier, was to verify data gleaned in recent experiments. This is the advantage of GTS, which was designed to use experimental data and take into account the optimized "bean-shaped" cross-section of tokamak reactors, such as the upcoming ITER fusion reactor in France.

Specifically, the team simulated a plasma in the National Spherical Torus Experiment (NSTX) device, a magnetic fusion device located at PPPL. The NSTX was specifically designed to minimize transport, particularly in the ion channel, and is extremely efficient, said Ethier. However, fusion scientists want to know how energy is lost through electrons—perhaps fusion energy's greatest mystery and one that has puzzled scientists for decades. With its unique capability, the NSTX has focused on a particular instability that may be the cause of the high-level electron energy transport observed in the experiments.

By incorporating realistic parameters (plasma profiles, magnetic field, etc.) from the experimental data into the GTS code, the team was able to run the first nonlinear, global, kinetic simulation of electron-scale turbulence to directly compare apples to apples with NSTX experiments, thus making a true impact on our understanding of what causes electron transport.

"We confirmed that the turbulence is where the experiments said it would be," said Ethier. "However, a clear connection between the existing turbulence and the observed high level of electron energy transport has not been obtained yet and is the focus of ongoing experimental and computational research."

In essence, this breakthrough simulation did four things simultaneously: tested the mettle of Jaguar, validated the GTS code, confirmed the suspected instability, and eliminated it as the possible cause of the significant electron transport observed in experiments on the NSTX. The next step, said Ethier, is to continue to refine the code on the petaflop Jaguar and prepare for ITER-scale simulations. Theoretically, ITER will be the first fusion reactor to produce more power than is needed to heat the reaction and is considered an integral step in eventually producing large amounts of fusion power for commercial use.

"We need to add more physics effects to the code [such as the forces encountered by electrons and ions in the core] to bring our computer simulation experiments even closer to true experiments in the laboratory," said Ethier. While the code isn't yet perfect, Ethier believes it can play a large role in ITER-scale simulations.

Whatever the future holds, 2008 was a breakthrough year for fusion research, and Jaguar emerged as the leader in fusion simulation, providing researchers with support for ITER's size and design and bringing them one step closer to unraveling the mystery of electron transport and enabling commercial fusion energy production. Its highly parallel architecture enables unprecedented scaling for leading fusion codes, a relationship that might one day help to make fusion energy a reality.—*by Scott Jones*

Climate Moves to the Petascale

Unprecedented high-performance computing resources spur partnerships and progress in 2008

Every paramecium, petunia, and person has a stake in global warming. The actions of human stakeholders may make it difficult for the rest to thrive or even survive, indicates the Fourth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC), whose authors shared the 2007 Nobel Peace Prize. The report cites about 300 scientific papers that used data from simulations run on supercomputers in 2004 and 2005. It found a greater than 90 percent likelihood that increased greenhouse gases from human activities caused most of the rise in globally averaged temperatures during the 20th century.

In the 21st century the challenge will be to continue to give stakeholders information based on sound science to guide decisions about how to adapt to climate change or lessen its effects. Scientists predict the planet's average temperature will increase 5 to 12 degrees Fahrenheit this century, melting polar ice caps, raising sea levels, fueling extreme weather, displacing populations, shrinking food and water supplies, and driving more than a million species to extinction.

Stakeholders need a crystal ball of sorts to explore future climate scenarios. Among science's best prognosticators are supercomputers, such as those that produced the IPCC data at Oak Ridge and Lawrence Berkeley national laboratories, the National Center for Atmospheric Research (NCAR), and the Earth Simulator in Japan. Supercomputers help stakeholders ask "what if?" What if hurricanes and heat waves become more intense? What if, to displace fossil fuels, we convert forest and food croplands to biofuel crops? What if melting permafrost releases methane, which is 20 times more potent at heating the atmosphere than carbon dioxide?

Supercomputers help stakeholders get answers quickly and plan accordingly. DOE, NSF, the National Oceanic and Atmospheric Administration (NOAA), and the National Aeronautics and Space Administration (NASA) are among the U.S. government agencies tasked with assessing climate change. DOE's strength in building and running the world's fastest supercomputers has made its ORNL an interagency crossroads for climate modeling.

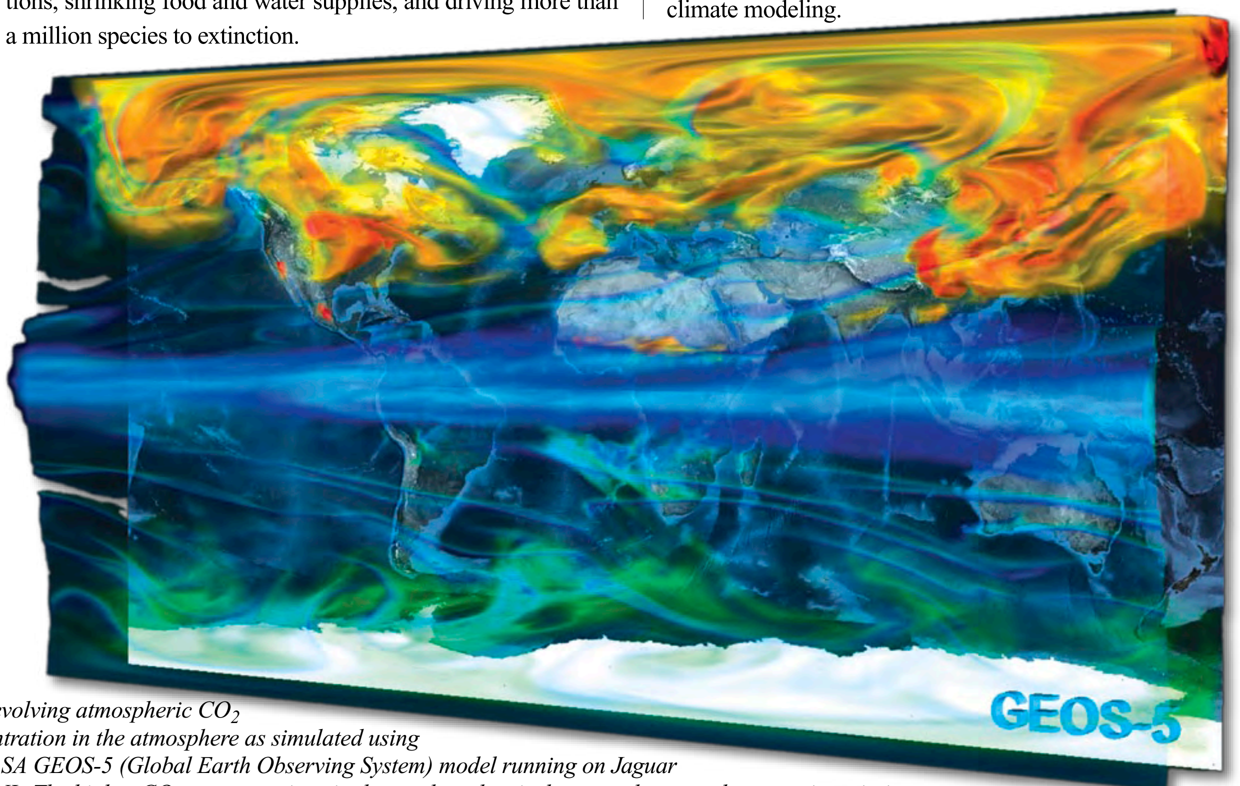

Time-evolving atmospheric CO_2 concentration in the atmosphere as simulated using the NASA GEOS-5 (Global Earth Observing System) model running on Jaguar at ORNL. The higher CO_2 concentrations in the northern hemisphere are due to anthropogenic emissions. This is a joint climate science ORNL-NASA project between David Erickson, ORNL, and Steven Pawson, NASA/Global Modeling and Assimilation Office/Goddard Space Flight Center. Visualization was completed by Jamison Daniel, ORNL.

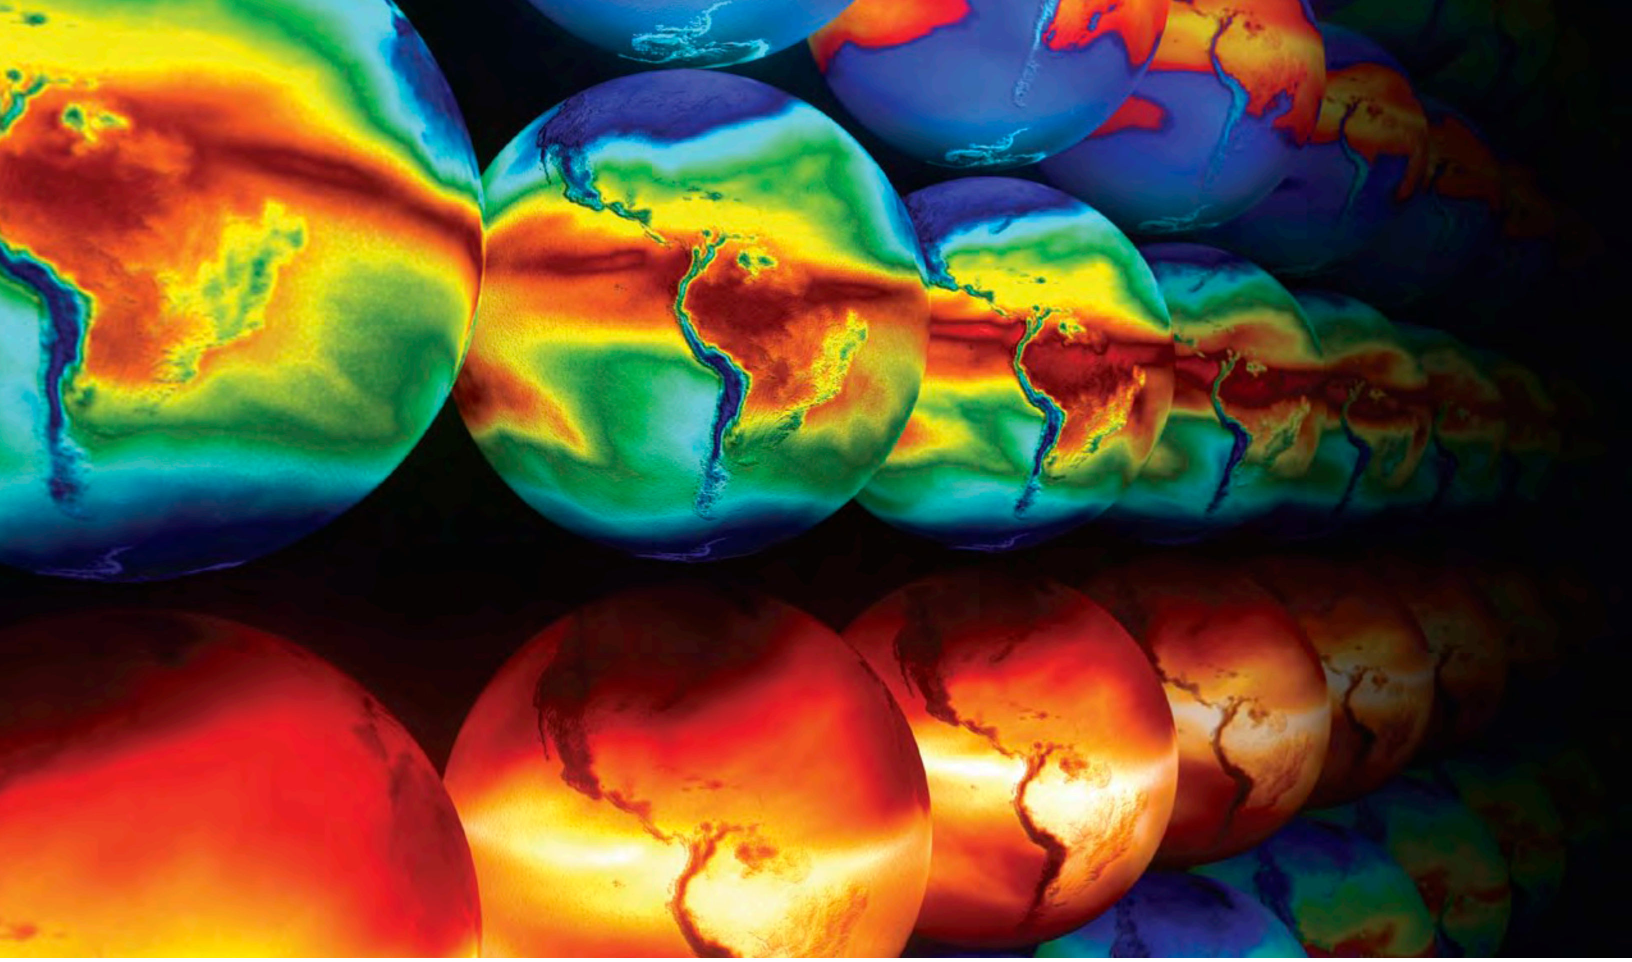

The proper simulation of the distribution of water vapor in the climate system is essential to the accurate treatment of the hydrological cycle and the planetary radiation budget. These images show the simulated monthly-averaged distribution of the total column water vapor from a high-resolution configuration of the CCSM Community Atmospheric Model. Visualization by Jamison Daniel, ORNL.

"We have started to get the required competencies here at ORNL that build a much stronger program and a capability to respond to national opportunities and national needs in climate," said James J. Hack, director of the NCCS.

"Jaguar provides an opportunity for modeling efforts from different organizations to come in contact with each other, leveraging their experience to accelerate progress in climate science," Hack said. In addition, the researchers get help from NCCS experts—scientific computing liaisons, modeling experts, software engineers, and computer scientists—to scale their algorithms to best use Jaguar's 182,000 processors.

"Working with domain scientists—to get their applications to scale to more efficiently exploit the computational platforms that we have and that we know are coming—is going to be a critical part of advancing climate science," Hack said. "Models are not going to get faster any more by waiting for CPU clock frequencies to get faster. It's going to happen by learning how to exploit parallelism in a much more efficient way."

Leveraged Investment

Partnerships between government agencies leverage the public investment in research. A case in point is the Climate Science Computational End Station, an alliance of DOE's ORNL, NSF's NCAR, and NASA that organizes efforts to produce simulation data for the fifth IPCC assessment, expected in 2014. ORNL's John Drake is the end station's chief computational scientist, and NCAR's Warren Washington is its chief climate scientist.

Another important interagency partnership in climate-change research, this one between NOAA and DOE, was established in 2008 to leverage complementary expertise in modeling atmospheric processes. According to a Memorandum of Understanding signed in September, the collaboration aims "to improve the quality of and quantify the uncertainty of climate and weather prediction, including improving the prediction of high-impact weather events, to provide the best science-based climate and weather information for management and policy decisions." NOAA's time on Jaguar rose from 18,171 processor hours used in 2008 to 25.1 million awarded for 2009.

The Climate Change Initiative at ORNL, led by Hack, also awarded its first Laboratory Directed Research and Development grants in 2008 to seed high-impact studies. Hack, who helped lead development of the Community Climate System Model (CCSM) much of America's climate science community uses to simulate Earth's atmosphere, spearheaded the initiative to accelerate discoveries about the planet's climate systems through labwide engagement of researchers from diverse directorates encompassing energy, environment, computing, and national security.

Among the new grant recipients were Auroop Ganguly (reducing uncertainty in climate-change models), Kate Evans (decadal prediction of Earth's climate after major volcanic eruptions), and Peter Thornton (biogeochemical modeling of the global carbon cycle when land use changes).

"We are technically and scientifically positioning ourselves to explore the skill of our modeling tools to predict climate change on regional spatial scales and decadal timescales," Hack said. "We never had the resources to do so in the past. Petascale computing provides a first opportunity to see what the modeling tools are capable of in terms of the current state of the science and related computational algorithms."

Accelerating Advancements

In 2008 several climate projects showed they had what it takes to scale up to use greater numbers of processors, providing the higher resolution needed to examine features on finer scales. Simulations on Jaguar's XT4 component were led by Paola Cessi at the University of California–San Diego (atmospheric effects on ocean circulation), Zhengyu Liu at the University of Wisconsin–Madison (abrupt climate change), Synte Peacock at the University of Chicago (first century-scale simulation to track chemicals moving throughout the world's oceans), and Warren Washington at NCAR (development of the Climate Science Computational End Station). In addition, Peacock, who moved to NCAR in 2008, gained early access to the XT4 when it was upgraded from 119 to 263 teraflops in June. Her team's simulation improved understanding of timescales and pathways of oceanic uptake of atmospheric greenhouse gases.

The 2008 accomplishments set the stage for renewed INCITE climate projects in 2009 on Jaguar by Liu (4 million processor hours) and Washington (30 million processor hours) as well as new projects by David Randall of Colorado State University (2 million processor hours to simulate global cloudiness), Venkatramani Balaji of NOAA's Geophysical Fluid Dynamics Laboratory at Princeton University (24 million processor hours to model the Earth system), and Gilbert Compo of the University of Colorado (1.1 million processor hours to study extreme weather).

In November Jaguar was upgraded again with the addition of a 1.4-petaflop Cray XT5 component. The combined XT4/XT5 system features an InfiniBand network to move large data sets from the supercomputer to other computing platforms and a Spider file system that can hold 1,000 times as much data as contained in the printed collection of the Library of Congress, assuming one word is 10 bytes. The massive effort to stand up the DOE Office of Science petascale center began in May 2004 when the federal government awarded ORNL the Leadership Computing Facility project. By November 2008 ORNL had broken the petascale barrier with an upgrade that produced a 1.64-petaflop XT combined system. Scientists immediately ran virtually all of the XT5's processing cores at a sustained speed of 1.352 petaflops to simulate superconducting materials, winning the Gordon Bell Prize for world's fastest scientific computing application.

Thought Leadership

Giving the scientific community the world's fastest HPC system in less than 4 1/2 years required great thought leadership and a sustained effort in delivering on that vision. That said, scientists are already demanding exascale computing systems—1,000 times faster than petascale machines—that promise even greater understanding of complex biogeochemical cycles that underpin global ecosystems and make life on Earth sustainable. One example of the type of thought leadership that will enable continued progress in ultrascale computing is expert testimony Hack delivered in May 2008 to the U.S. Senate Committee on Commerce, Science, and Transportation. "Computational capability remains a significant bottleneck and should remain a high-priority investment," he said, emphasizing the need for balanced expenditures in computational infrastructure, climate science and modeling, climate observations, computer science, and applied mathematics. Congress responded by raising DOE's climate-change-modeling budget from $31 million in 2008 to $45 million in 2009.

In the 21st century the challenge will be to continue to give stakeholders information based on sound science to guide decisions about how to adapt to climate change or lessen its effects.

UT-Battelle, LLC, which manages ORNL for DOE, has long displayed such visionary thinking, investing in climate-research facilities years before federal research grants became available. UT-Battelle built HPC systems in a town famous for putting the best minds of a generation to work on a global problem of importance to future generations. As a result of this thought leadership, well-positioned Oak Ridge is one of the few places thriving during a worldwide economic downturn. Thanks in part to the watershed achievements of 2008, ORNL announced in 2009 it would hire 1,000 people over 2 years, many of whom will develop new energy solutions and further explore the consequences of our energy use.—*by Dawn Levy*

An atomistic model of cellulose (blue) surrounded by lignin molecules (green) comprising a total of 3.3 million atoms. Water molecules that are present in the model are not shown. Image courtesy Jeremy Smith and Jamison Daniel, ORNL.

Fueling the Future

Simulations unlock the secrets of cellulosic ethanol

A consensus has recently emerged that America's dependence on petroleum-based fuels is damaging to both our economy and national security. DOE is exploring a number of options in an effort to reduce this dependence and achieve the department's goal of "twenty in ten": the reduction of American consumption of gasoline by 20 percent in 10 years through investment in alternative fuel sources.

One of the more promising solutions is the harvesting of ethanol from cellulose, the complex carbohydrate that forms the cell walls of plants and gives leaves, stalks, stems, and trunks their rigidity. Imagine filling your gas tank with the stuff of trees and weeds; there seems to be plenty of it. In fact, cellulose is the most abundant organic material on earth. With the help of Jaguar's computing muscle, ORNL's Jeremy Smith aims to make affordable cellulosic ethanol a reality.

"The simulations we are performing are designed to provide a picture of biomass that will help experimentalists design plants with new, less resistant cell walls and enzymes that break the cellulose down more efficiently," said Smith. "This is basic research designed to help underpin the current major, worldwide effort in renewable energy research."

And cellulosic ethanol, besides being plentiful, also carries with it an environmental advantage. Although ethanol combustion still emits carbon dioxide, studies show some crops planted to make ethanol convert nearly as much back into glucose through photosynthesis, making the fuel almost a net-zero emitter of carbon dioxide. Fossil fuels simply pollute.

Unfortunately, to get to the stored energy in cellulose, it must be broken down into its simpler sugar subunits—glucose, the unit fermented to make ethanol, included. However, cellulose strongly resists this breakdown; researchers call this natural resistance to decomposition biomass recalcitrance. Hydrolysis, a reaction that uses water to break down chemical bonds, must take place for cellulose to ultimately become glucose and affordably find its way into your gas tank.

In lignocellulosic biomass—or biomass composed of cellulose, hemicelluloses, and lignin—the noncellulosic components slow the hydrolysis of cellulose and must be removed before it can effectively occur.

Smith plans to ultimately create models of lignocellulosic biomass and cellulose that will show the structure, motion, and mechanics of the materials on a level never seen before.

The current processes for this breakdown are difficult, expensive, and time consuming, all of which are obstacles to realizing the potential of an ethanol-based transportation economy.

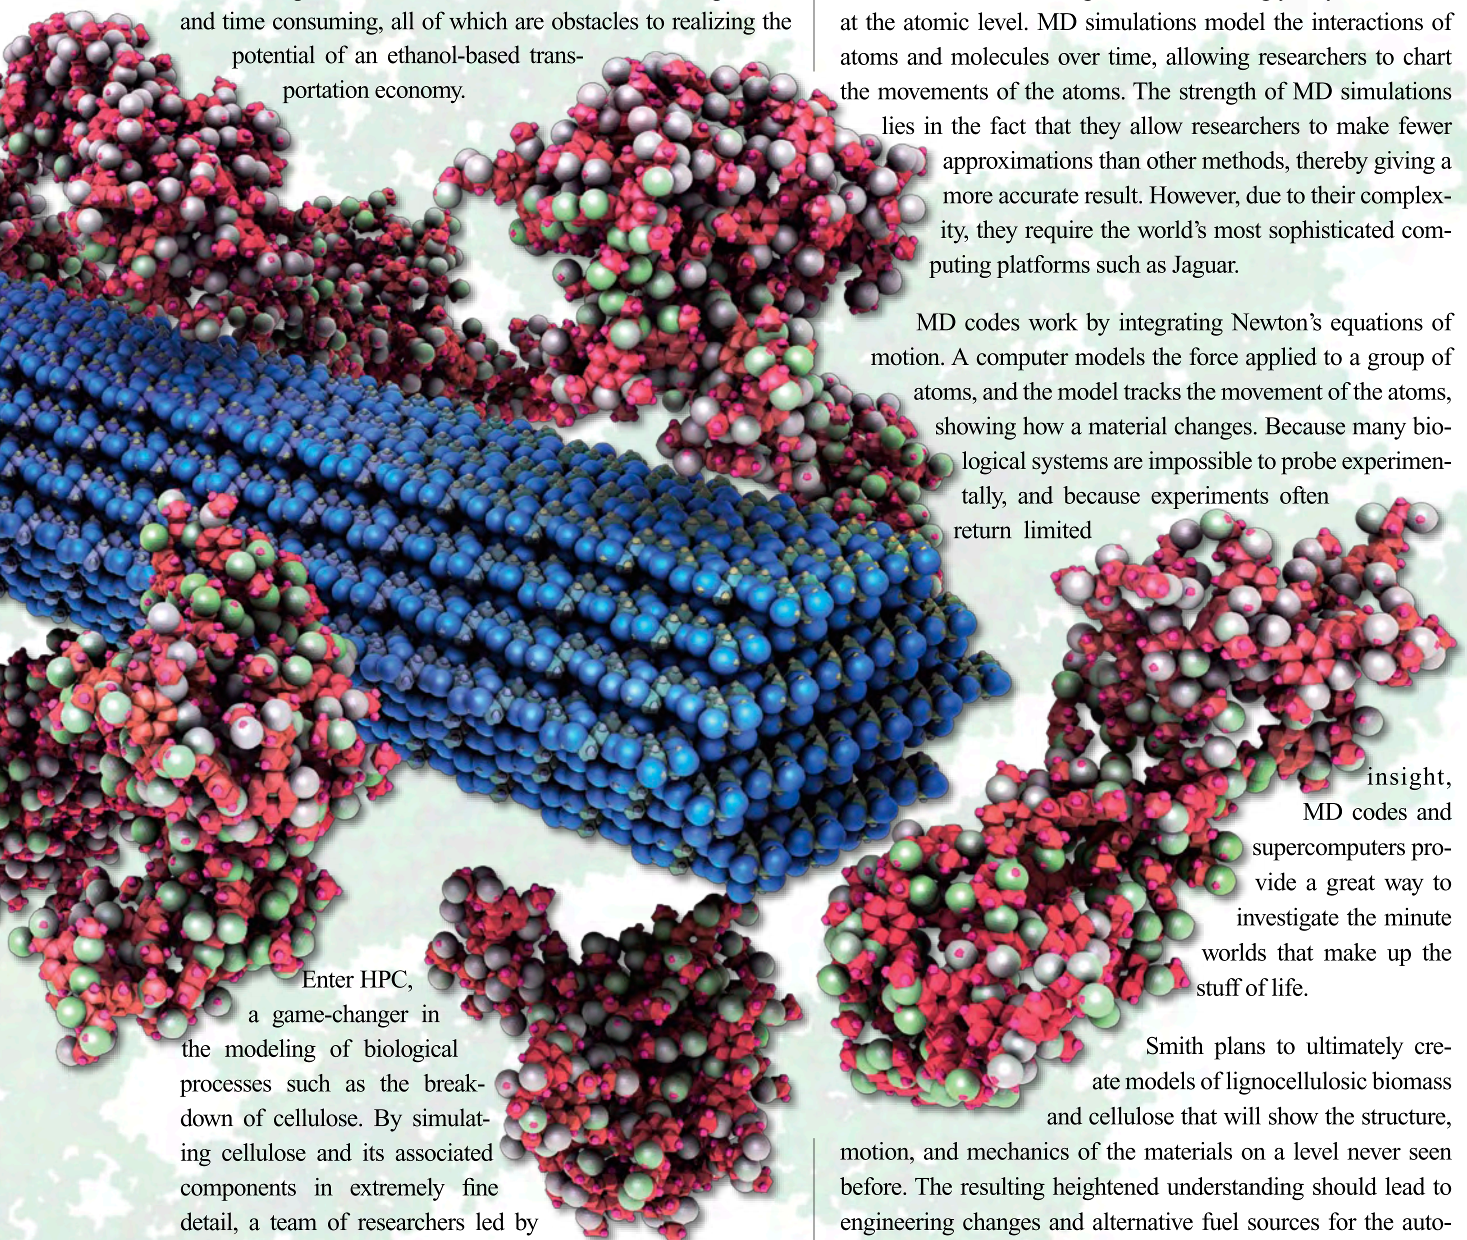

Enter HPC, a game-changer in the modeling of biological processes such as the breakdown of cellulose. By simulating cellulose and its associated components in extremely fine detail, a team of researchers led by Smith is paving the way for America's next-generation fuel source. Using Jaguar, Smith and his team have conducted groundbreaking simulations of the structure and mechanics of cellulosic systems.

Molecular Dynamics to the Rescue

In work funded by the ORNL Bioenergy Science Center and the DOE Office of Science, Smith and his colleagues used their 2008 allocation on the supercomputer to pair a fast, massively parallel code known as LAMMPS with Jaguar's unprecedented computing power to visualize the breaking down of cellulose and bring America one step closer to cleaner air and energy independence.

Specifically, Smith and his team used a molecular dynamics (MD) simulation method to calculate how the parts of lignocellulosic biomass are arranged and how strongly they are bonded at the atomic level. MD simulations model the interactions of atoms and molecules over time, allowing researchers to chart the movements of the atoms. The strength of MD simulations lies in the fact that they allow researchers to make fewer approximations than other methods, thereby giving a more accurate result. However, due to their complexity, they require the world's most sophisticated computing platforms such as Jaguar.

MD codes work by integrating Newton's equations of motion. A computer models the force applied to a group of atoms, and the model tracks the movement of the atoms, showing how a material changes. Because many biological systems are impossible to probe experimentally, and because experiments often return limited insight, MD codes and supercomputers provide a great way to investigate the minute worlds that make up the stuff of life.

Smith plans to ultimately create models of lignocellulosic biomass and cellulose that will show the structure, motion, and mechanics of the materials on a level never seen before. The resulting heightened understanding should lead to engineering changes and alternative fuel sources for the automobiles on which Americans so heavily depend.

"Putting designed plants into fields could mean cheaper ethanol production, getting more per acre, and more efficient microbes would also extract more ethanol per plant," said Smith. "It's a big effort—worldwide—to understand these ethanol production processes, and we will likely see some research findings take practical effect in the next 5 years or so."

Its potential availability and its net cleanliness make cellulosic ethanol a top prospect as an alternative fuel source. However, understanding the complex chemical processes necessary to reduce cellulose to glucose and ferment it, thus creating ethanol, is a monumental task. Big problems require big solutions, few of which are bigger than Jaguar.—*by Scott Jones*

Bridges to Discovery:

Experts Help Projects Scale Up to Close Knowledge Gaps

NCCS team helps research pioneers cross the frontiers of petascale computing

Computation has become such an important workhorse in research that it now trots alongside science's other trusty steeds, theory and experiment.

Jaguar, the NCCS's flagship system, is available to researchers in industry, academia, and the government through the DOE Office of Science's INCITE program. This program allocates processor hours to pioneering projects ranging from developing sustainable energy and understanding climate change to illuminating the underpinnings of disease and exploring the cosmos. Five of DOE's top ten computational science advances in 2008 used Jaguar to gain insight into supernovas, combustion, fusion, dark matter, superconductivity, and mathematics.

The interactions with the principal investigators are pretty seamless because the NCCS liaisons are essentially members of that community and have an idea of what the needs are.

Today's grand challenges in science and engineering are too complex for soloists to solve. Further breakthroughs may require simulations of scenarios impossible to explore without supercomputers. At the NCCS, experts in HPC operations, technology integration, scientific computing, and computing tools help users of the center's machines scale up simulations that once ran on 1,000 to 10,000 processors to now fully utilize the 150,000 processors of Jaguar's XT5 component.

"Using machines at this scale is unprecedented for most of the science teams involved," said Ricky Kendall, who leads NCCS's Scientific Computing Group. "Things that didn't get in your way in the past are now bottlenecks in the algorithms. It's my team's job to work at the level of algorithms, tools, and libraries to ensure codes become scalable to make effective use of such a huge resource."

Bridges to Research Success

At the forefront of scaling efforts, Kendall's group includes computational scientists, visualization specialists, workflow professionals, and performance modelers. "My team is a very sophisticated set of scientific users who are the bridge between the science community as a whole and the center," Kendall said. "We often find issues with the supercomputing resources before users do and then do whatever it takes to get them resolved. The team has quite a bit of experience in badgering the code and the machine until it tells you what it's doing wrong."

The team includes research scientists in chemistry, physics, astrophysics, mathematics, and computer science who are also experts in developing and optimizing code for the NCCS systems.

"The interactions with the principal investigators are seamless because the NCCS liaisons are essentially members of that community and have an idea of what the needs are," Kendall said. NCCS computational scientist Ramanan Sankaran, for example, coauthored the S3D code used in the combustion simulations of his principal investigator, Jacqueline Chen at Sandia National Laboratories. NCCS astrophysicist Bronson Messer publishes papers about core-collapse supernovas with supercomputer user Tony Mezzacappa of ORNL. And NCCS computational scientist Mark Fahey helps physicists Ron Waltz and Jeff Candy of General Atomics optimize their fusion simulations.

"There's this notion of a subset of algorithms that apply across all domains, and that's where our group leverages one another's expertise to try and make sure that we're doing things in a scalable and proper way to use a huge resource," Kendall said. "Particle-tracking techniques that have been used in combustion by Jackie [Chen]'s team and things that Ramanan [Sankaran] put into the code are applicable to, say, the fusion or astrophysics efforts in the same sort of way they are in combustion."

"We're a motley crew of people who don't really belong together other than we like to live close to the machine," said Messer, a liaison for projects that explore the birth and death of stars and the clumpy nature of dark matter. "Most of what

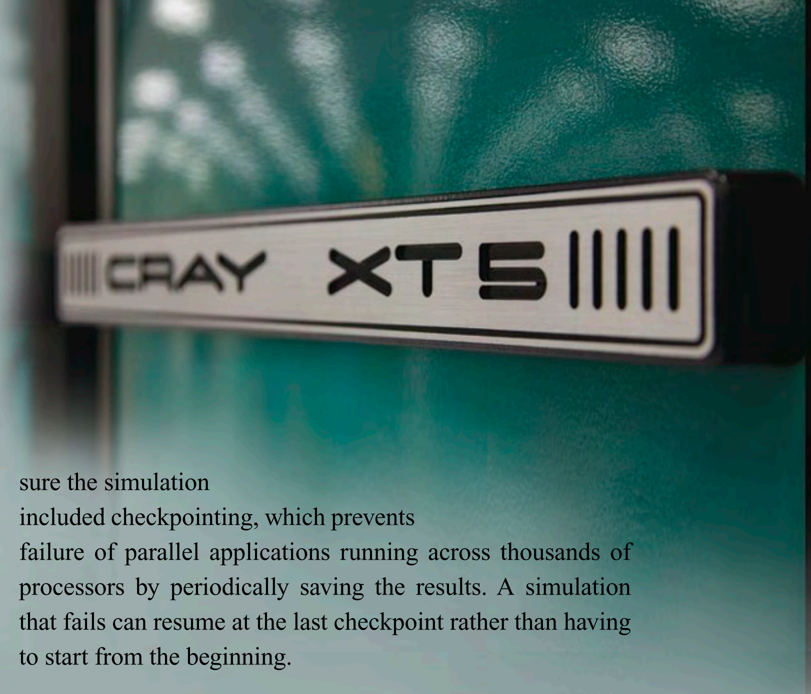

we trade with one another are stories about what we find that is particular to the machine—pathologies, weird things happening, stuff you should look out for."

Liaisons often bridge communication gaps between domain scientists and computational experts. "I try to translate back and forth between the computational folks and the astrophysicists," Messer said. "I try to make it clear to the physicists what shortcomings supercomputers have and what we should be able to do and help with technical problems, but I also try to teach people. That's giving people fish. I also try to teach them how to fish. I try my best to teach the astrophysicists methods to find out things for themselves on the computational side because that's how I learned."

Undeniably, liaison expertise is critical to wresting productive results from the powerful supercomputers. "Everybody agrees that computers should be easier to use as time goes by," Messer said. "However, if you expect this huge, one-of-a-kind machine will work like your desktop machine, that is an unrealistic expectation."

Liaison success is notable. Climate computational scientist Jim Rosinski, at the NCCS since July 2008, noted a case in which a bottleneck forced all calculations through a single processor. Rosinski modified the process to instead use multiple processors and wrote a wrapping layer of software to bypass a part of the code that didn't work. The fix made it possible to simulate fluid motion in the world's oceans in extremely fine detail. "What I've done helps them get results quicker at that resolution," Rosinski explained. That means results will be more quickly deposited into the central repository of data scientists use to produce the papers that will be cited in the next assessment report of the IPCC.

In 2008 computational mathematician Rebecca Hartman-Baker worked with James Vary of Iowa State University, a collaborator on David Dean's project at ORNL to explore atomic nuclei. In a simulation of physical forces on 30,000 processors of Jaguar's XT4 component, the job would fail half the time. "If a 4-hour job ran 3 hours 59 minutes and bombed, it'd have to completely start over," Hartman-Baker said. She made sure the simulation included checkpointing, which prevents failure of parallel applications running across thousands of processors by periodically saving the results. A simulation that fails can resume at the last checkpoint rather than having to start from the beginning.

Hai Ah Nam, a computational nuclear physicist, also helps Dean scale up his codes. A recent project used a groundbreaking computation to investigate why carbon-14, an isotope used in carbon dating, has an anomalously long half-life. Her hybrid talents in high-level computation and theoretical physics accelerated collaboration. When a computer scientist in the group talked with a theoretical physicist, conversation would die in just 10 minutes. Nam's ability to fill in the gaps allowed for a meaningful interaction between the two communities. Nam, who came to the NCCS in September 2008, chose the center because she could continue her own theoretical physics research here. "Working at the NCCS puts you at the nexus of collaborative research. It's a perfect place to learn and make connections, especially if you want to become a principal investigator for your own project," she said. "I'm definitely a bridge person, but I have my own island."

In 2008 NCCS's Markus Eisenbach and Don Maxwell worked with Thomas Schulthess and colleagues to speed the DCA++ application, allowing scientists to scrutinize more superconducting materials. At the SC08 supercomputing conference, Jaguar's XT5 component earned the Gordon Bell Prize for world's fastest scientific computing application by running this enhanced version at a sustained speed of 1.352 petaflops.

Another task of the Scientific Computing Group is performance modeling, led by Wayne Joubert. "When we do a procurement and acceptance, when we're planning for an upgrade, we need to understand how to tune the possible options of that upgrade to make the most effective use of the machine for the science," Kendall said.

The group also tries to manage the data flood

The NCCS's Scientific Computing Group.

Jaguar generates. [See sidebar below, "ADIOS Technology Lets Processing Cores Say Goodbye to Poor Performance."] Scott Klasky leads a team to streamline and automate the workflow and visualize and monitor ongoing simulations.

Once users have completed their simulation runs, NCCS scientific visualization specialists led by Sean Ahern help them make sense of the data. In 2008 the visualization team stood up Lens, a new cluster with 512 Opteron processing cores, 2 terabytes of aggregate memory, and 64 graphics cards for parallel analysis and remote visualization. With staff from Lawrence Berkeley National Laboratory, team members developed an accelerated data access technique for particle identification and tracking. And working with scientists of the Climate Science Computational End Station project, they helped analyze and explore global biogeochemical cycles of carbon, sulfur, and nitrogen.

Gargantuan Tasks

In 2008 under Oak Ridge Leadership Computing Facility Project Director Buddy Bland, Jaguar underwent two major upgrades, accelerating its XT4 component from 119 to 263 teraflops in the spring and adding a 1.4-petaflop XT5 component in the fall. The upgrades were gargantuan tasks involving dozens of workers at the NCCS and Cray. After an upgrade the machine must be accepted, meaning it must pass tests to ensure it is ready for the scientists.

The hardware acceptance process confirms proper installation and functioning of disks, interconnects, memory, and network—required for system administration, message passing, storage, and more. The NCCS's Ann Baker and Don Maxwell of HPC Operations and Galen Shipman, David Dillow, and Sarp Oral of Technology Integration were key to hardware acceptance.

In contrast to acceptance of the hardware, software acceptance verifies that DOE application codes run properly as the machine scales to hundreds of thousands of processing cores. For the petaflop upgrade, for example, could the machine meet benchmarks, such as running the High-Performance Linpack code at greater than 1 petaflop or using the POP transport code to simulate 15 years of ocean behavior in 1 day?

ADIOS Technology Lets Processing Cores Say Goodbye to Poor Performance

Meanwhile, 'dashboard' workflow monitor provides window on simulations in progress

Jaguar's computer cores work in concert at lightning speed to simulate diverse phenomena from fuel combustion to supernova collapse. The simulations generate a tsunami of data. Managing it—and understanding it—is a monumental challenge that encompasses data generation, collection, movement, storage, processing, and analysis.

"As data sets get larger, scientific data management gets more important," said computational scientist Scott Klasky of the NCCS. With collaborators at Georgia Tech, Klasky developed ADIOS, or the Adaptable I/O System. This middleware says goodbye to inefficient processing as data moves in and out of the computer to the I/O system.

Busy processors must write data to disks for storage, analysis, and visualization. They must also execute calculations. Time spent saving data steals from time spent computing. Less computing means less science can get done. ADIOS lets code writers use application programming interfaces, or subroutines, to simplify the I/O process.

Staff members from the Scientific Computing Group work with domain scientists to enable research.

Computational scientist Arnold Tharrington from Kendall's group led software acceptance for both 2008 upgrades, working with Sankaran, Nam, and Joubert. He wrote harness scripts to orchestrate simultaneous running of numerous codes to stress the machine. To pass the test, Jaguar's XT5 component completed more than 95 percent of its jobs while running a series of science problems.

In 2008 the NCCS team started acceptance tests in November and ended in December. But preparations for the acceptance tests began much earlier—in 2007. Jeff Larkin, John Levesque, Jeff Becklehimer, Cathy Willis, and more than a dozen other Cray experts were available to fix the problems that inevitably came up while running codes on so many processors for the first time.

Kendall's team is also charged with gathering scientific and computational requirements for future endeavors. Members generate a periodic report of national research priorities over the next few years and a roadmap to create HPC resources to support them.

"With these machines, speed always matters," said Cray applications engineer Nathan Wichmann. "We're always looking at the codes and asking, what about the next machine? If a code has a communications bottleneck and the tools do not exist to fix it, what tools do we have to build? Or do we adjust the hardware platforms?"

At the Cray Supercomputing Center of Excellence on the ORNL campus, Levesque and Larkin work with Wichmann, based in Minnesota, to optimize Jaguar's system software. Working with scientific users, they identify performance problems and brainstorm solutions to make applications run faster. Their eyes are always on the next-generation machine, processor, or code. This proactive approach is a major reason Jaguar's petaflop acceptance went so smoothly, Wichmann said.

Such forward-looking strategy and old-fashioned elbow grease on the part of hundreds of professionals at Oak Ridge and beyond made Jaguar today's fastest path to scientific solutions. "The return on investment is more science per dollar than could be accomplished without the effort," Kendall said.—*by Dawn Levy*

In traditional I/O, scientists typically write with one set I/O method, such as MPI-IO. If scientists want to change from one method to another for performance or historical reasons, they must program all of their methods in their code and, worse yet, maintain this programming for the lifetime of that code. With ADIOS scientists describe in an XML file the desired appearance of the output and change one word to the I/O implementation they will use.

The I/O has sped up on all codes that have used ADIOS. The improvements mean more productivity, more science. At the University of California–Irvine, scientists simulated fusion using the Gyrokinetic Toroidal Code with I/O managed by ADIOS. Running on the 31,000 processing cores of Jaguar's XT4 component, the code generated 60 terabytes of data a day.

Moreover, ADIOS gives scientists access to metadata, or data about data, which helps researchers keep track of what their data represent. It also facilitates multiplexing of data, allowing visualization of data while I/O is in process in a simulation.

In 2007 Klasky, with ORNL computational scientists Roselyne Barreto and Norbert Podhorszki and colleagues from DOE's Scientific Data Management Center, started developing "dashboards" to provide windows into running simulations. A dashboard shows a simulation on the fly, displaying results just a minute or two behind the simulation.

Scientists like seeing their results in near-real time. "Scientists can do a comprehensive analysis of the data they're seeing on the dashboard," Podhorszki said. With easy clicking and dragging, they can perform analyses and then get publication-quality images on the dashboard without worrying about the location of the data used for making the images.

The team is integrating ADIOS and dashboards into an automated workflows management system called Kepler. Researchers at the Center for Plasma Edge Simulation depend on Kepler to monitor simulations on Jaguar and Ewok, a satellite computer cluster used for analysis. Physicist Seung-Hoe Ku of New York University, who used 4 million processor hours on Jaguar's XT4 component in 2008, uses the ORNL tools to monitor simulations of hot, ionized gas at the edge of a nuclear fusion reactor. "The productivity has doubled with ORNL's workflow tools," Ku said. "If they were not operating, we could not accomplish the science mission on time."—*by Dawn Levy*

New System Ready to Rock 'N' Roll

November awards cement Jaguar's status

The newly upgraded Jaguar system demonstrated its primacy in November by dominating two major competitions associated with SC08, the supercomputing conference held in Austin, Texas. In the HPC Challenge, Jaguar went head to head with the world's most powerful supercomputers, taking first place in two of the four categories and third place in another. In addition, a team led by Thomas Schulthess, who now serves dual roles as an ORNL research scientist and director of the Swiss National Supercomputing Centre, came away with the 2008 Gordon Bell Prize for the world's fastest scientific application.

The competitions not only established Jaguar's place in the pantheon of supercomputers—both as the world's most powerful scientific system and as the first to run scientific problems at the petaflop (topping 1 quadrillion calculations a second)—but also demonstrated the ease of use that will allow researchers to achieve breakthrough science, even before they make major changes to their codes.

The competitions not only established Jaguar's place in the pantheon of supercomputers...but also demonstrated the ease of use that will allow researchers to achieve breakthrough science even before they make major changes to their codes.

"Researchers using this system are telling us that Jaguar is both easy to use and delivering terrific results," said Buddy Bland, project director of the Oak Ridge Leadership Computing Facility. "We love getting this kind of feedback,

but it's really no surprise. This system was designed from the ground up with the needs and requirements of scientific applications in mind."

Jaguar's performance in the HPC Challenge was especially impressive because the benchmark tests were not optimized for the ORNL system. Contestants in the challenge can modify the codes, thereby running them in "optimized" mode, or they can leave them unmodified and run them in "baseline" mode. The codes for Jaguar went unmodified, yet it was best in the world when running two of the four benchmarks: one that evaluated speed in solving a dense matrix of linear algebra equations and another that judged sustainable memory bandwidth. In addition, Jaguar placed third in executing the global fast Fourier transform, a common algorithm used in many scientific applications.

The Gordon Bell Prize, administered by the Association for Computing Machinery, recognizes leadership in computational science and engineering. Schulthess and colleagues Thomas Maier, Michael Summers, and Gonzalo Alvarez, all of ORNL, took the main prize for their work with DCA++, an application that simulates the behavior of high-temperature superconductors. By making efficient use of 150,000 of Jaguar's 180,000-plus processing cores, the team was able to reach a sustained speed of 1.35 petaflops, an achievement that was not only the world's fastest scientific simulation ever, but also the fastest computer application of any type.

Two other finalists in the Gordon Bell competition also conducted simulations on Jaguar. A team led by Lin-Wang Wang of Lawrence Berkeley National Laboratory won a special Gordon Bell Prize for algorithm innovation with an application that conducts first-principles calculations of electronic structure. The model reached 442 teraflops on nearly 150,000 cores during the run. Another team, led by Laura Carrington of the San Diego Supercomputing Center, used the system to shatter its own record in modeling seismic waves traveling through the earth using a highly realistic model that includes the planet's rotation and gravity and the complexity of forces deep underground. Carrington's goal is to simulate waves with a period as low as 1 second, the fastest seismic waves likely to make it through the entire planet. With Jaguar her team cut its best time by more than two-thirds, reaching a period of only 1.15 seconds. The experience of these leading computational scientists bodes well for groundbreaking science in the years to come.

"It's amazing that a code this complex could be ported to such a large system with so little effort," noted Carrington. "This was a landmark calculation on the ORNL petaflop system that enables a powerful new tool for seismic-wave simulation."—*by Leo Williams*

Mike Summers and Thomas Schulthess receiving the Gordon Bell Prize at SC08.

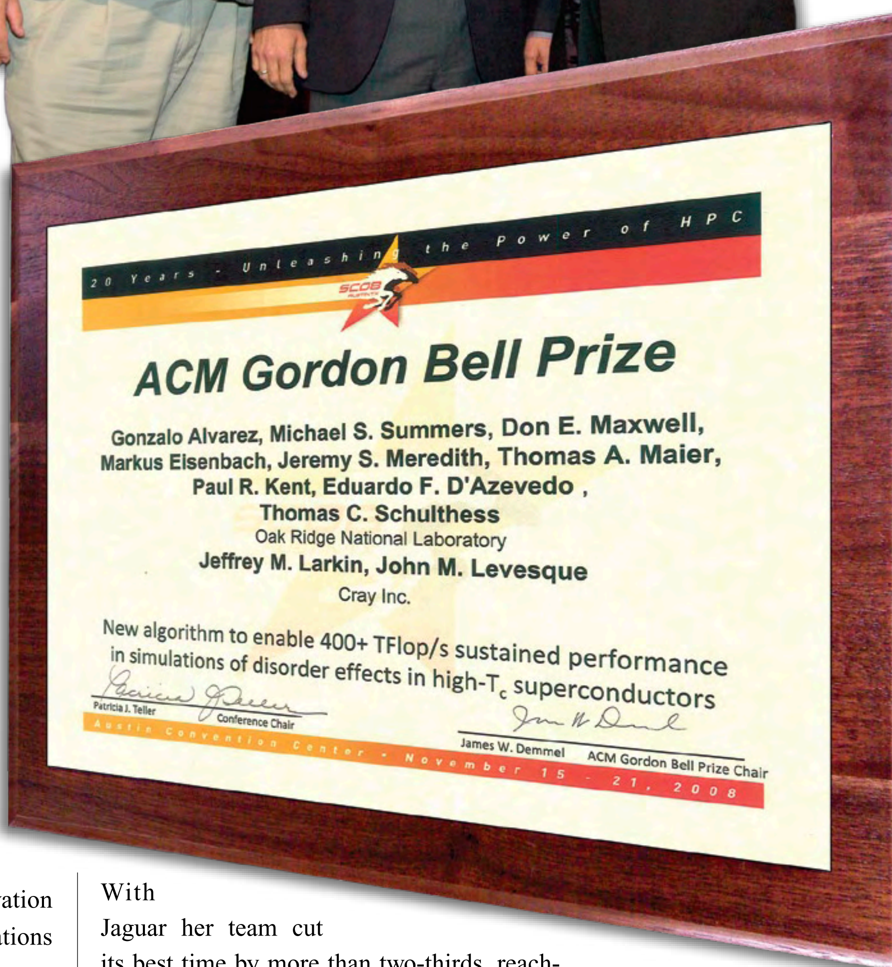

Other members of the Gordon Bell Prize application team included (back row) Thomas Maier, Jeremy Meredith, Don Maxwell, John Levesque, Paul Kent and (front row) Ed D'Azevedo, Gonzalo Alvarez, Jeff Larkin, Markus Eisenbach.

Made for Science

Jaguar upgrades show what a scientific supercomputer can be

The NCCS increased the power of its Jaguar supercomputer more than tenfold in 2008, standing up the most powerful scientific computer the world has ever seen.

In doing so we at the center also stood up a system that was ready to deliver groundbreaking science on day one. Weeks before Jaguar completed acceptance testing, leading computational scientists were telling us their experience on the system was surprisingly uncomplicated for a new supercomputer on the cutting edge of scientific research.

During these early days Jaguar ran the first two scientific simulations to break the petaflop barrier—materials applications that topped 1 quadrillion calculations a second (that's a 1 followed by 15 zeroes). One took the prestigious Gordon Bell Prize for fastest application in the world. In fact, of the six Gordon Bell finalists, three had their best runs on Jaguar.

During these early days Jaguar ran the first two scientific simulations to break the petaflop barrier...

These achievements by some of the world's most talented computational scientists demonstrate Jaguar's real-world potential. They focus on areas such as superconducting materials (which promise to revolutionize fields as diverse as power transmission, rail transportation, and medical imaging), electronic structures (whose potential applications include quantum dots and future solar cells), and the propagation of seismic waves through the entire planet (the closest thing we're likely to get to a three-dimensional picture of the earth's interior).

Nor are these the only success stories the NCCS can point to. Jaguar has already shown its prowess in an eclectic collection of other fields as well, including weather research, combustion science, and fusion research. Each of these areas has seen applications scale to more than 100,000 processing cores.

All told, Jaguar went through two upgrades in 2008. The system was converted from dual-core processors to quad-core ones in the spring, bringing it to more than 30,000 processing cores and a peak performance of more than 250 teraflops. It was upgraded later in the year, with more than 180,000 processors delivering a peak performance of 1.64 petaflops. The system now has 362 terabytes of memory, a 10-petabyte file system, 578 terabytes per second of memory bandwidth, and an unprecedented I/O bandwidth of 284 gigabytes per second.

This is a machine designed from the ground up with the needs and requirements of scientific applications in mind. It has powerful processors, allowing scientists to get more science out of fewer of them. It has three times more memory than any other system on the planet, preparing it to solve very, very large problems. And it has massive file-system bandwidth, ensuring there are no performance bottlenecks.

We are very proud of the resource we have been able to bring to the scientific computing community. We fully expect to see great things from these researchers and others.—*by Buddy Bland*

Buddy Bland, OLCF project director.

HPC Enters Petascale Era

Select applications will put Jaguar to the test

With Jaguar's recent upgrade, HPC has officially entered the petascale. Boasting a peak performance of more than 1.6 petaflops, Jaguar is now the most powerful computer in the world for open science and the first to enter this new era—one that will certainly redefine the potential of HPC research.

To facilitate the most able codes and test the mettle of the new system, the NCCS has granted early access to a number of projects that can use a majority of the machine and take it, and science, to its limits.

We have three principal goals during this petascale early science phase: (1) deliver important, high-impact science results and advancements; (2) harden the system for production; and (3) embrace a broad user community capable of and prepared for using the system.

This priority phase will consist of more than 20 projects and run for approximately 6 months. The first projects cover a broad range of science and include two climate projects and one project each in the domains of chemistry, biology, combustion, and materials science.

Both climate projects will incorporate atmospheric and oceanic elements and produce results at higher resolutions, bringing researchers one step closer to understanding the dynamics of the earth's complex climate system. Ultimately the data gleaned from these allocations will provide policymakers with better information to help them plan for predicted future climate fluctuations, such as regional climate change on decadal timescales.

The featured chemistry project will focus on enhanced energy storage in nanostructured systems, possibly revolutionizing battery and related technologies; the biology one will explore a more efficient means of converting cellulose to ethanol, a process that could one day make economically feasible biofuels a reality; combustion research will dissect the properties of ignition and flame dynamics, paving the way for more efficient future engines; and the materials science research will use the classic Hubbard model for the design of high-temperature superconductors, a technology that would increase the efficiency of energy transport several times over.

Fusion, nuclear energy, astrophysics, and geosciences will also be explored in the coming months. All of these projects will require more than 20 percent of the entire petascale resource to achieve their stated goals and are represented by more than 100 of the world's premiere computational scientists. In all, more than 500 million CPU hours will be consumed, twice the amount of last year's entire INCITE program allocation. This accomplishment is a direct tribute to Jaguar's unprecedented speed, which makes possible nearly 4 million hours of compute time daily.

Each of the petascale early science projects represents high-impact research of national importance, and we expect each one to deliver substantial results. Because they will be led by the community's most sophisticated users and prominent scientists, early simulations on the machine will also help us harden the system for a broader collection of projects later in the year.

These early pioneering research applications will push computational science to new limits, giving researchers a peek at the possibilities of petascale simulation and prepping Jaguar for the nearly 40 INCITE projects that will explore domains from biology to fusion energy in unprecedented detail. Now that the petascale era has finally arrived, the work really begins.—*by Doug Kothe*

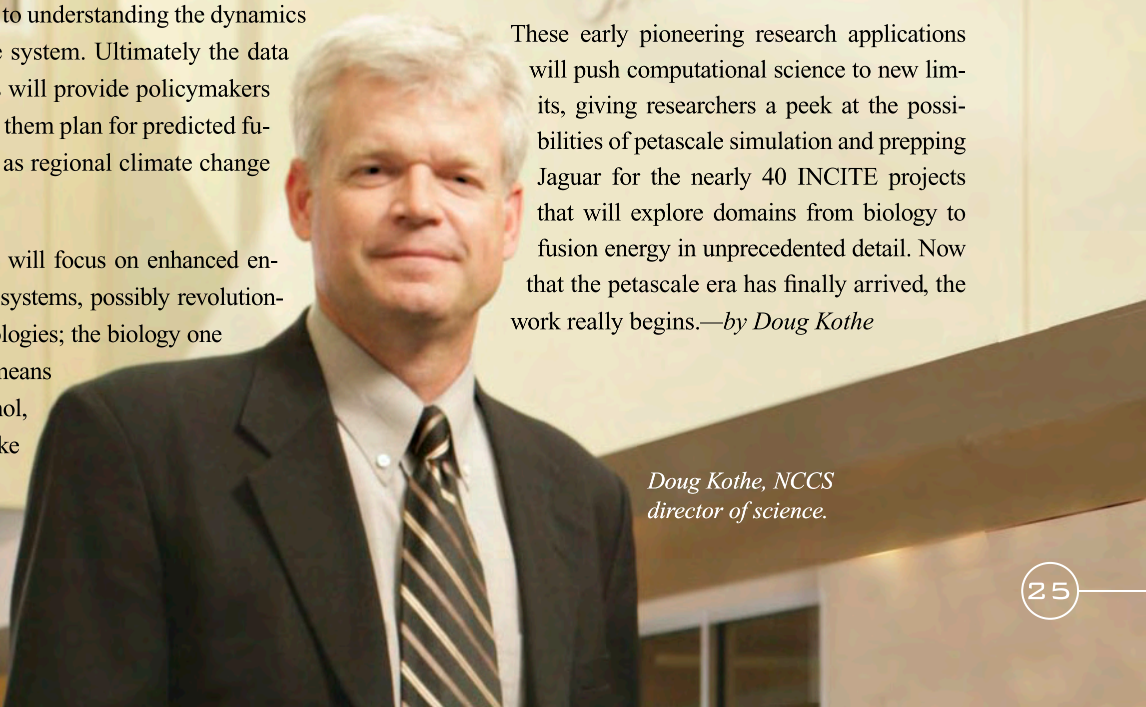

Doug Kothe, NCCS director of science.

Looking Forward

Preparing to make exascale computing a reality

It was only a few years ago that petascale computing was a dot on the horizon. Today we have arrived; computation at the petascale has become daily business at the NCCS. Jaguar, the world's most powerful computer for open science, is allowing researchers to study a wide range of systems in unprecedented detail.

We must not, however, be satisfied with this accomplishment. While Jaguar is currently revealing more about our world than we ever thought possible, it is likewise demonstrating the limitations of our most sophisticated computing platforms. Even at the petascale, large portions of the universe remain out of the reach of exploration.

In essence a supercomputer is a virtual laboratory. As machines get more and more powerful, they will allow scientists to incorporate more realism into the virtual environment. The most difficult problems we face are multiple-length-scale, multiple-physics ones that require complex simulation capabilities. They span a range of disciplines and require enormous amounts of computing muscle. Take climate for instance. An ideal climate model requires length scales ranging from 10^7 to 10^{-7} meters, or 14 orders of magnitude. Even at the petascale, attaining this level of detail is unthinkable.

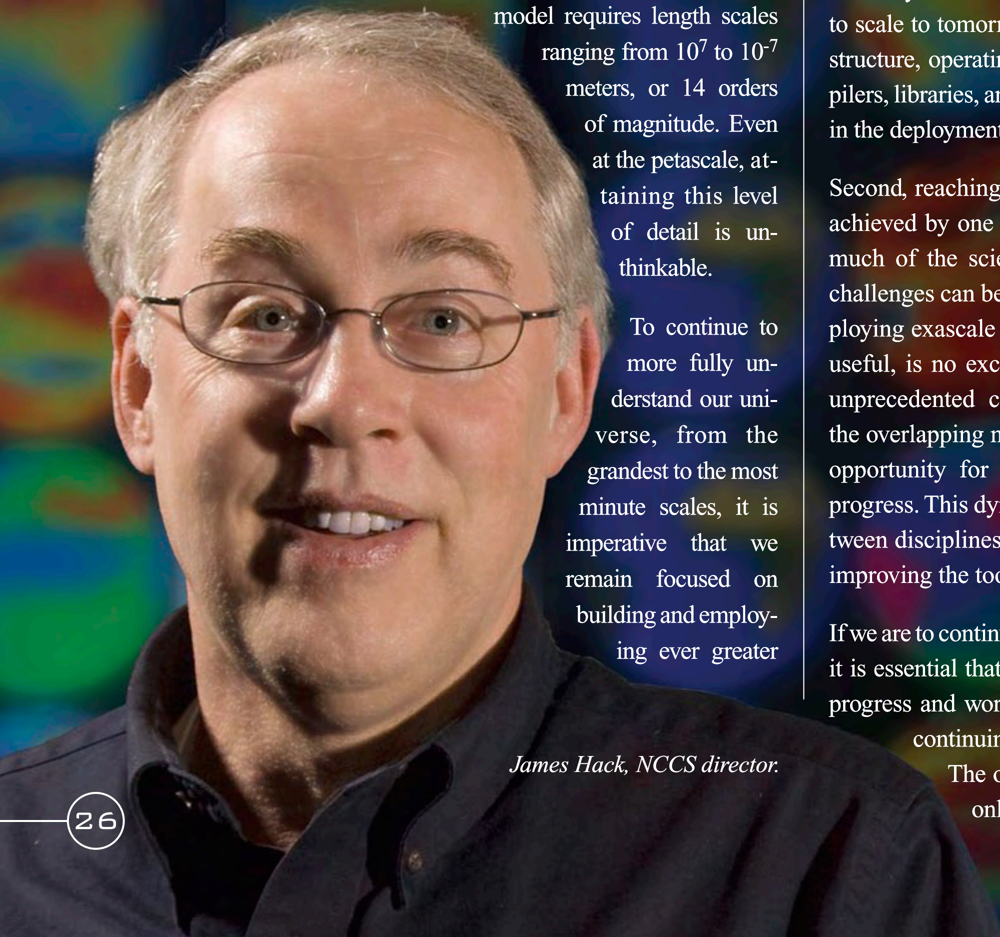

To continue to more fully understand our universe, from the grandest to the most minute scales, it is imperative that we remain focused on building and employing ever greater computing resources for scientific progress. Whereas today's petascale systems are capable of accommodating tens of thousands of threads, or independent processes, in a parallel code, next-generation exascale systems could reach the millions. That sort of detail is today's wishful thinking, much like the petascale of 5 years ago. Yet here we are, and there we will be.

James Hack, NCCS director.

The NCCS is currently preparing to make exascale computing a reality. The near future will see machines increase to approximately 20 petaflops, and we have every expectation that in 10 years exascale systems will begin tackling the world's most complex scientific challenges. These machines will be an entire order of magnitude more powerful than current petascale systems and will truly change the face of scientific endeavor.

However, two things are needed to realize this vision and continue on the path toward a more perfect virtual laboratory. First and foremost, a proportional investment is needed in the various intellectual efforts necessary to move the whole field forward, such as applied mathematics, computer science, and the respective domain sciences. The exascale will never be reached simply through a brute-force investment in hardware, and it is not currently clear whether today's software packages will be able to scale to tomorrow's systems. Investments in software infrastructure, operating systems, and software tools such as compilers, libraries, and debuggers will likewise be a key ingredient in the deployment of next-generation supercomputers.

Today we have arrived; computation at the petascale has become daily business at the NCCS.

Second, reaching the exascale and beyond is not going to be achieved by one organization. There is now a recognition in much of the scientific community that the biggest science challenges can be conquered only through collaboration. Employing exascale systems that are not only powerful, but also useful, is no exception. Such a daunting challenge requires unprecedented cooperation across government agencies—the overlapping needs of these organizations offer the perfect opportunity for interaction, collaboration, and accelerated progress. This dynamic works both ways. The cooperation between disciplines and agencies improves the science, thereby improving the tools on which that science is conducted.

If we are to continue to push the limits of computational science, it is essential that we invest in the people necessary to realize progress and work together to tackle the problem, while also continuing to make the most of our current resources. The obstacles ahead are indeed large and dwarfed only by the possibilities.—*by James Hack*

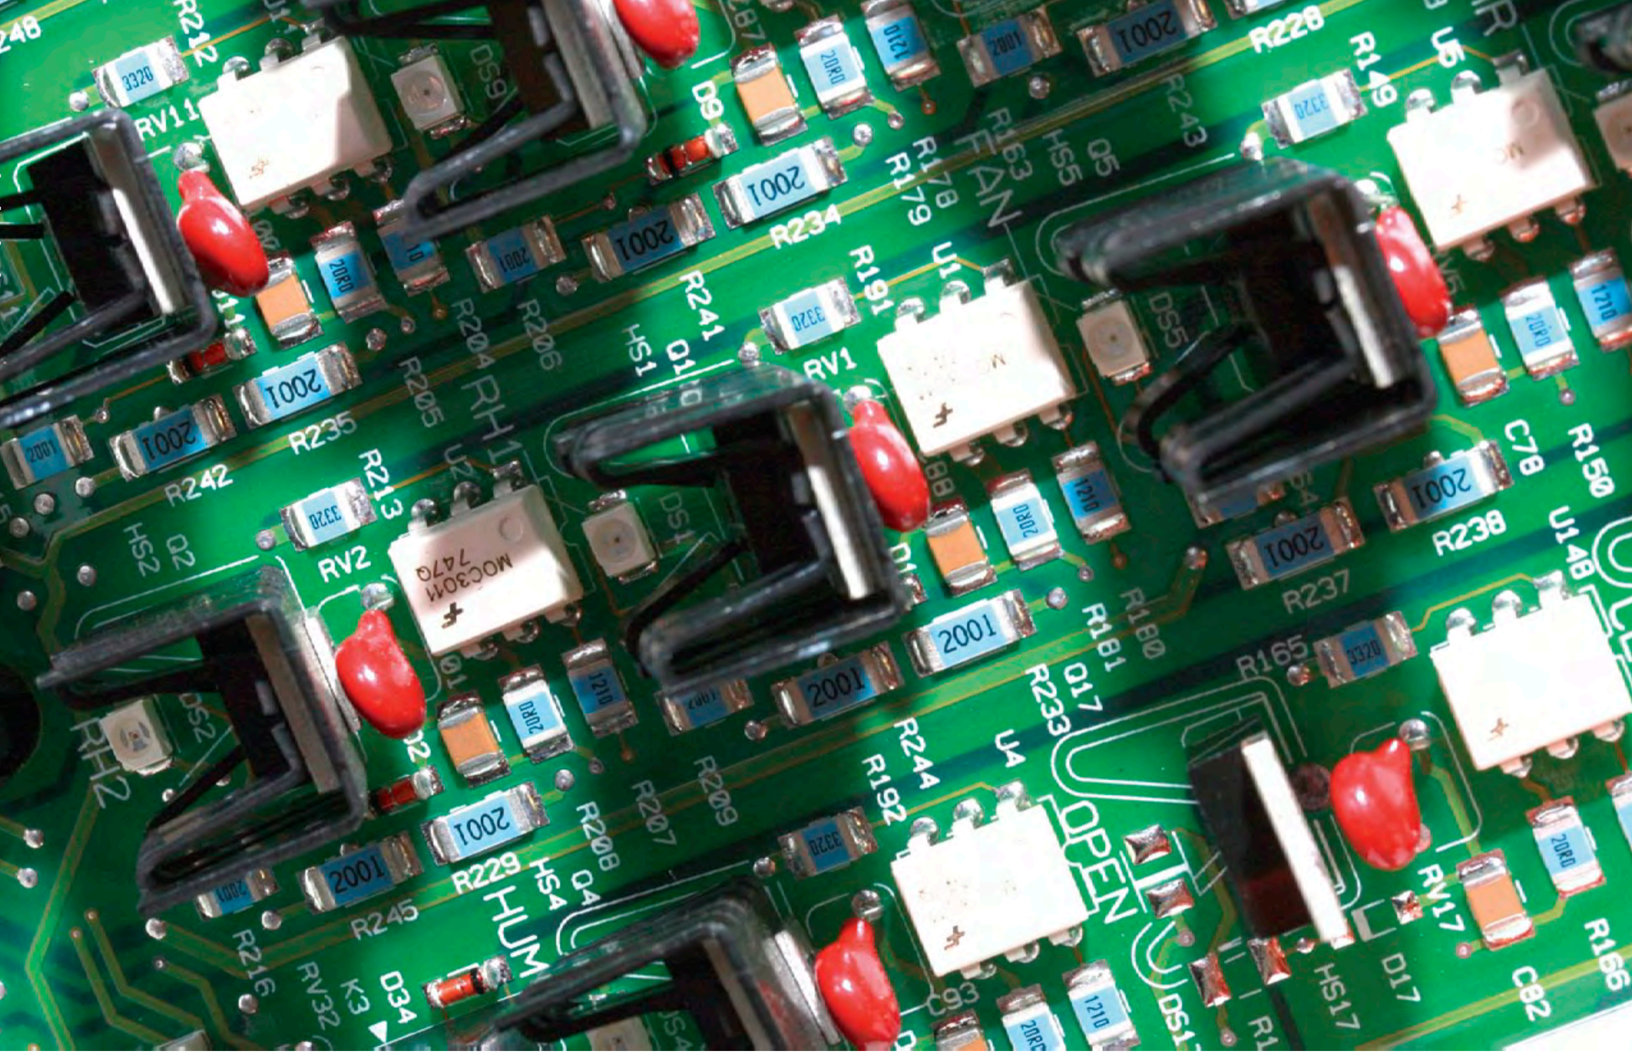

Eyeing the Leap to Exascale

Computing experts work on getting from here to there

With Jaguar's petascale upgrade, the NCCS has accomplished the seemingly impossible: it has stood up the world's most powerful scientific computer without forcing developers to go back to the drawing board. Weeks before it was subjected to acceptance testing, Jaguar ran record-breaking simulations in a diverse range of disciplines, including the first sustained petascale application to ever run on any system.

Impressive as that achievement was, the scientific computing community faces a far higher hurdle as it contemplates moving up three orders of computing power to the next major milestone—exascale computers capable of a quintillion (or a million trillion) calculations every second.

The potential payoff of exascale computing is exhilarating—batteries that let you drive 500 miles on a single charge; designer materials created with specific properties in mind; and accurate, decade-long regional climate predictions, to name a few. But the challenges are daunting. Practices now considered wise (but time-consuming) will be indispensable, and the NCCS and its counterparts throughout the HPC community will have to overcome imposing obstacles in the design of both systems and algorithms. Developers will need new tools as systems move from hundreds of thousands of processors to millions and tens of millions. Otherwise applications will get hopelessly bogged down on systems that should be delivering revolutionary science.

"What we as a center can do is identify these high hurdles, communicate the sense of urgency to the funding agencies, and put our best people on it," explained NCCS Director of Science Douglas Kothe. "None of these challenges look insurmountable. There are a few that look pretty daunting, but I think we ought to be able to tackle them with the right people and resources."

Fortunately developers are getting help from efforts such as the Institute for Architectures and Algorithms (IAA), a collaboration among HPC experts at ORNL and Sandia National Laboratories. According to ORNL's Al Geist, the goal of the institute is to fundamentally change the way supercomputers and algorithms are designed in the future.

"Supercomputers have been designed to be more powerful at all costs," Geist said, "not necessarily to make them better for science. And the scientists often design their algorithms without regard to how efficiently they will work on a million-processor computer. IAA is trying to foster the codesign of algorithms and computer architectures—each making changes and compromises to better scale the exascale mountain."

The Computational Sciences Building at ORNL. Upgrades to infrastructure allowed the NCCS to field the Jaguar supercomputer.

The challenges arise because supercomputers are becoming ever more complex as they become ever more powerful. In its current incarnation, Jaguar's 180,000-plus processing cores churn through calculations at an unimaginable pace, passing information back and forth and storing it in various levels of memory. Each time information is passed, time is lost; the farther the pass, the greater the loss.

Imagine a badly organized automobile assembly plant. A car's first part might be attached at one end of the plant, its second at the opposite end, its third somewhere in between, and so on. Each worker would attach several parts, but the car would be hauled away and brought back separately for each. And every time a worker called in sick, everyone would take the day off.

It Would Be Henry Ford's Nightmare

Yet we face an analogous fate in HPC. Applications regularly pass information between processors without knowing whether they are close together or far apart. Processors are repeatedly told to pull the same line of data from memory, reading only one number at a time. And when a single processor goes down, it takes the whole simulation with it. (In fact, supercomputing applications must regularly stop to take a time- and disc-intensive snapshot of the application, creating what is known as a restart file, in case one component of the system fails and crashes the whole simulation.)

In the IAA algorithms project, Geist and about two dozen colleagues are working to improve application performance, both in the near term and over the long haul. In the near term the team is creating architecture-aware algorithms for climate, chemistry, and nanotechnology applications. These new algorithms are designed to run efficiently on multicore computers, such as Jaguar, and will be made available to all applications through publicly available numerical libraries. In the longer term the project is developing simulators to study the effects of different hardware and software design choices in hopes of influencing the future design of both algorithms and supercomputers to ensure that science in the next generation achieves its potential.—*by Leo Williams*

Engaging New Users

Bringing new members into the community

The Research Alliance in Math and Science (RAMS) once again brought undergraduates, graduate students, and faculty to ORNL in summer 2008 for assignments ranging from 10 to 12 weeks. The program, which allows participants to gain experience working with individual ORNL scientist mentors, aims to increase the number of minority individuals taking advanced degrees in science, mathematics, engineering, and technology.

ORNL's Computing and Computational Sciences Directorate has worked for several years to establish and expand partnerships with historically black colleges and other minority educational institutions. "The goal is to bring these students back year after year to advance their hands-on experience and add to their academic experience," said ORNL's Debbie McCoy, RAMS program manager, "with the end goal of bringing them on payroll after they have received their advanced degrees. A number of former RAMS students are now in PhD programs, which is an indicator of our success."

Each student is assigned a research mentor and works on a project of mutual interest, often an integral task of a larger ongoing project. The students keep journals, attend seminars and workshops, write papers, and make oral and poster presentations at ORNL and national venues. Academic credit may be arranged for their summer intership.

Two of last year's students, Jessica Travierso of Austin Peay State University and Tara McQueen of Delaware State University, attended the SC08 supercomputing conference in Austin, Texas, in November. They participated in the recruiting fair as spokespersons for the RAMS program, as well as attending sessions in their areas of interest. Travierso, a veteran of several years in the RAMS program, won the undergraduate poster contest at the TeraGrid '08 annual meeting with her 2007 RAMS poster. She is now beginning a master's degree in astrophysics and will return to ORNL to work with NCCS computational astrophysicist Bronson Messer. Andrea Rocha, a RAMS student from the University of South Florida and a committee member for the Society for the Advancement of Chicanos and Native Americans in Science, attended the society's annual meeting and recruited several students to RAMS.

McCoy said the program passed a DOE peer review in December with flying colors. Five former RAMS students made

Students from the RAMS program visit the Spallation Neutron Source at ORNL.

presentations, and the students basically sold the program, she said. RAMS is sponsored by DOE's Office of Advanced Scientific Computing Research (ASCR).

Keeping Skills Up to Date

Cray XT Workshop
April 14–16

About two dozen computational scientists from universities, industry, and national laboratories came to ORNL to attend a workshop jointly sponsored by the NCCS and National Institute for Computational Sciences (NICS). The goal of the workshop was to prepare NCCS and NICS users for the upcoming Cray XT5 Jaguar and Kraken supercomputers.

NCCS Users' Meeting
April 17–18

NCCS users met with experts from ORNL and Cray Inc. to discuss approaches and challenges related to the NCCS systems and computational methods likely to be successful for the upcoming petascale system.

BlueGene Workshop
July 29–31

About 40 users and developers met at ORNL to discuss the features and architecture of ORNL's IBM BlueGene Eugene system.

Cray XT Quad-Core Workshop
October 15–17

About three dozen NCCS and NICS users met to discuss issues related to the use of quad-core processors on Cray XT4 and XT5 systems.

Supercomputing Crash Course 2008
June 16

Arnold Tharrington and Rebecca Hartman-Baker of the NCCS Scientific Computing Group gave more than 50 students an introduction to writing programs for supercomputers. The day was broken into a beginning course and an advanced one.

NCCS Lustre Workshop
April 16

The NCCS and Sun Microsytems hosted a 1-day workshop to help application scientists get the most from the Lustre File System. The workshop, focused on improving application input and output, included several case studies and an open discussion. *—by Agatha Bardoel and Leo Williams*

Kraken, a Cray XT5, is the NICS flagship system.

HPC researchers gather for a recent conference at the NCCS.

ORNL's booth at SC08.

HIGH-PERFORMANCE COMPUTING PROJECTS

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

Astrophysics

First Principles Models of Type Ia Supernovae
Stan Woosley, University of California–Santa Cruz
Jaguar: 3,000,000 hours

Multidimensional Simulations of Core Collapse Supernovae
Anthony Mezzacappa, Oak Ridge National Laboratory
Jaguar: 75,000,000 hours

Numerical Relativity Simulations of Binary Black Holes and Gravitational Radiation
Joan Centrella and Jim VanMeter, National Aeronautics and Space Administration
Jaguar: 500,000 hours

Intermittency and Star Formation in Turbulent Molecular Clouds
Alexei Kritsuk, University of California–San Diego
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

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: 2,600,000 hours

Cellulosic Ethanol: Physical Basis of Recalcitrance to Hydrolysis of Lignocellulosic Biomass
Jeremy Smith, Oak Ridge National Laboratory
Jaguar: 6,000,000 hours

Gating Mechanisms of Membrane Proteins
Benoit Roux, University of Chicago
Jaguar: 15,000,000 hours

Chemistry

Molecular Simulation of Complex Chemical Systems
Christopher Mundy, Pacific Northwest National Laboratory
Jaguar: 2,000,000 hours

An Integrated Approach to the Rational Design of Chemical Catalysts
Robert Harrison, Oak Ridge National Laboratory
Jaguar: 30,000,000 hours

Dynamically Tunable Ferroelectric Surface Catalysts
Andrew Rappe, University of Pennsylvania
Jaguar: 2,283,200 hours

Structural and Dynamical Studies of Hydronium and Hydroxide Ions in Bulk Water and at the Water/Air Interface Using Ab Initio Path Integral Simulations
Thomas Miller, California Institute of Technology
Jaguar: 12,000,000 hours

Physics

Computational Nuclear Structure
David Dean, Oak Ridge National Laboratory
Jaguar: 15,000,000 hours

Petascale Computing for Terascale Particle Accelerator: International Linear Collider Design and Modeling
Lie-Quan Lee, Stanford Linear Accelerator Center
Jaguar: 8,000,000 hours

Lattice QCD
Robert Sugar, University of California–Santa Barbara
Jaguar: 20,000,000 hours

High-Fidelity Computations for Complex Biological Membranes
Michael Heroux, Sandia National Laboratories
Jaguar: 1,000,000 hours

Materials Sciences

Development and Correlations of Large Scale Computational Tools for Transport Airplanes
Moeljo Hong, The Boeing Company
Jaguar: 1,000,000 hours

Bose-Einstein Condensation vs. Quantum Localization in Quantum Magnets
Tommaso Roscilde, Max-Planck Gesellschaft
Jaguar: 1,200,000 hours

Predictive and Accurate Monte Carlo-Based Simulations for Mott Insulators, Cuprate Superconductors, and Nanoscale Systems
Thomas Schulthess, Oak Ridge National Laboratory
Jaguar: 45,000,000 hours

Electronic, Lattice, and Mechanical Properties of Novel Nano-Structured Bulk Materials
Jihui Yang, GM R&D Center
Jaguar: 15,000,000 hours

Linear Scale Electronic Structure Calculations for Nanostructures
Lin-Wang Wang, Lawrence Berkeley National Laboratory
Jaguar: 2,000,000 hours

Geosciences

Modeling Reactive Flows in Porous Media
Peter Lichtner, Los Alamos National Laboratory
Jaguar: 10,000,000 hours

Computer Sciences

Performance Evaluation and Analysis Consortium End Station
Patrick Worley, Oak Ridge National Laboratory
Jaguar: 8,000,000 hours

Climate

Simulation of Global Cloudiness
David Randall, Colorado State University
Jaguar: 2,000,000 hours

CHiMES: Coupled High-Resolution Modeling of the Earth System
Venkatramani Balaji, NOAA/GFDL, Princeton University
Jaguar: 24,000,000 hours

Assessing Transient Global Climate Response of the NCAR-CCSM3: Climate Sensitivity and Abrupt Climate Change
Zhengyu Liu, University of Wisconsin–Madison
Jaguar: 4,000,000 hours

Climate-Science Computational End Station Development and Grand Challenge Team
Warren Washington, National Center for Atmospheric Research
Jaguar: 30,000,000 hours

Surface Input Reanalysis for Climate Applications (SIRCA) 1850–2011
Gilbert Compo, University of Colorado
Jaguar: 1,100,000 hours

Engineering

High-Fidelity Simulations for Clean and Efficient Combustion of Alternative Fuels
Jacqueline Chen, Sandia National Laboratories
Jaguar: 30,000,000 hours

Petascale Simulation of Nano-Electronic Devices
Gerhard Klimeck, Purdue University
Jaguar: 5,000,000 hours

Clean and Efficient Coal Gasifier Designs Using Large-Scale Simulations
Madhava Syamlal, National Energy Technology Laboratory
Jaguar: 13,000,000 hours

Landmark Direct Numerical Simulations of Separation and Transition for Aerospace-Relevant Wall-Bounded Shear Flows
Hermann Fasel, University of Arizona
Jaguar: 500,000 hours

Propulsor Analyses for a Greener, High Bypass Ratio, Aircraft Gas Turbine Engine
Robert Maleki, Pratt & Whitney
Jaguar: 1,500,000 hours

Fusion

Gyrokinetic Steady State Transport Simulations
Jeff Candy, General Atomics
Jaguar: 2,000,000 hours

High Power Electromagnetic Wave Heating in the ITER Burning Plasma
E. Fred Jaeger, Oak Ridge National Laboratory
Jaguar: 2,000,000 hours

Verification and Validation of Petascale Simulation of Turbulent Transport in Fusion Plasmas
Patrick Diamond, University of California–San Diego
Jaguar: 30,000,000 hours

High-Fidelity Tokamak Edge Simulation for Efficient Confinement of Fusion Plasma
C.S. Chang, New York University
Jaguar: 20,000,000 hours

Validation of Plasma Microturbulence Simulations for Finite-Beta Fusion Experiments
William Nevins and Greg Hammett, Lawrence Livermore National Laboratory
Jaguar: 20,000,000 hours

NCCS Staff

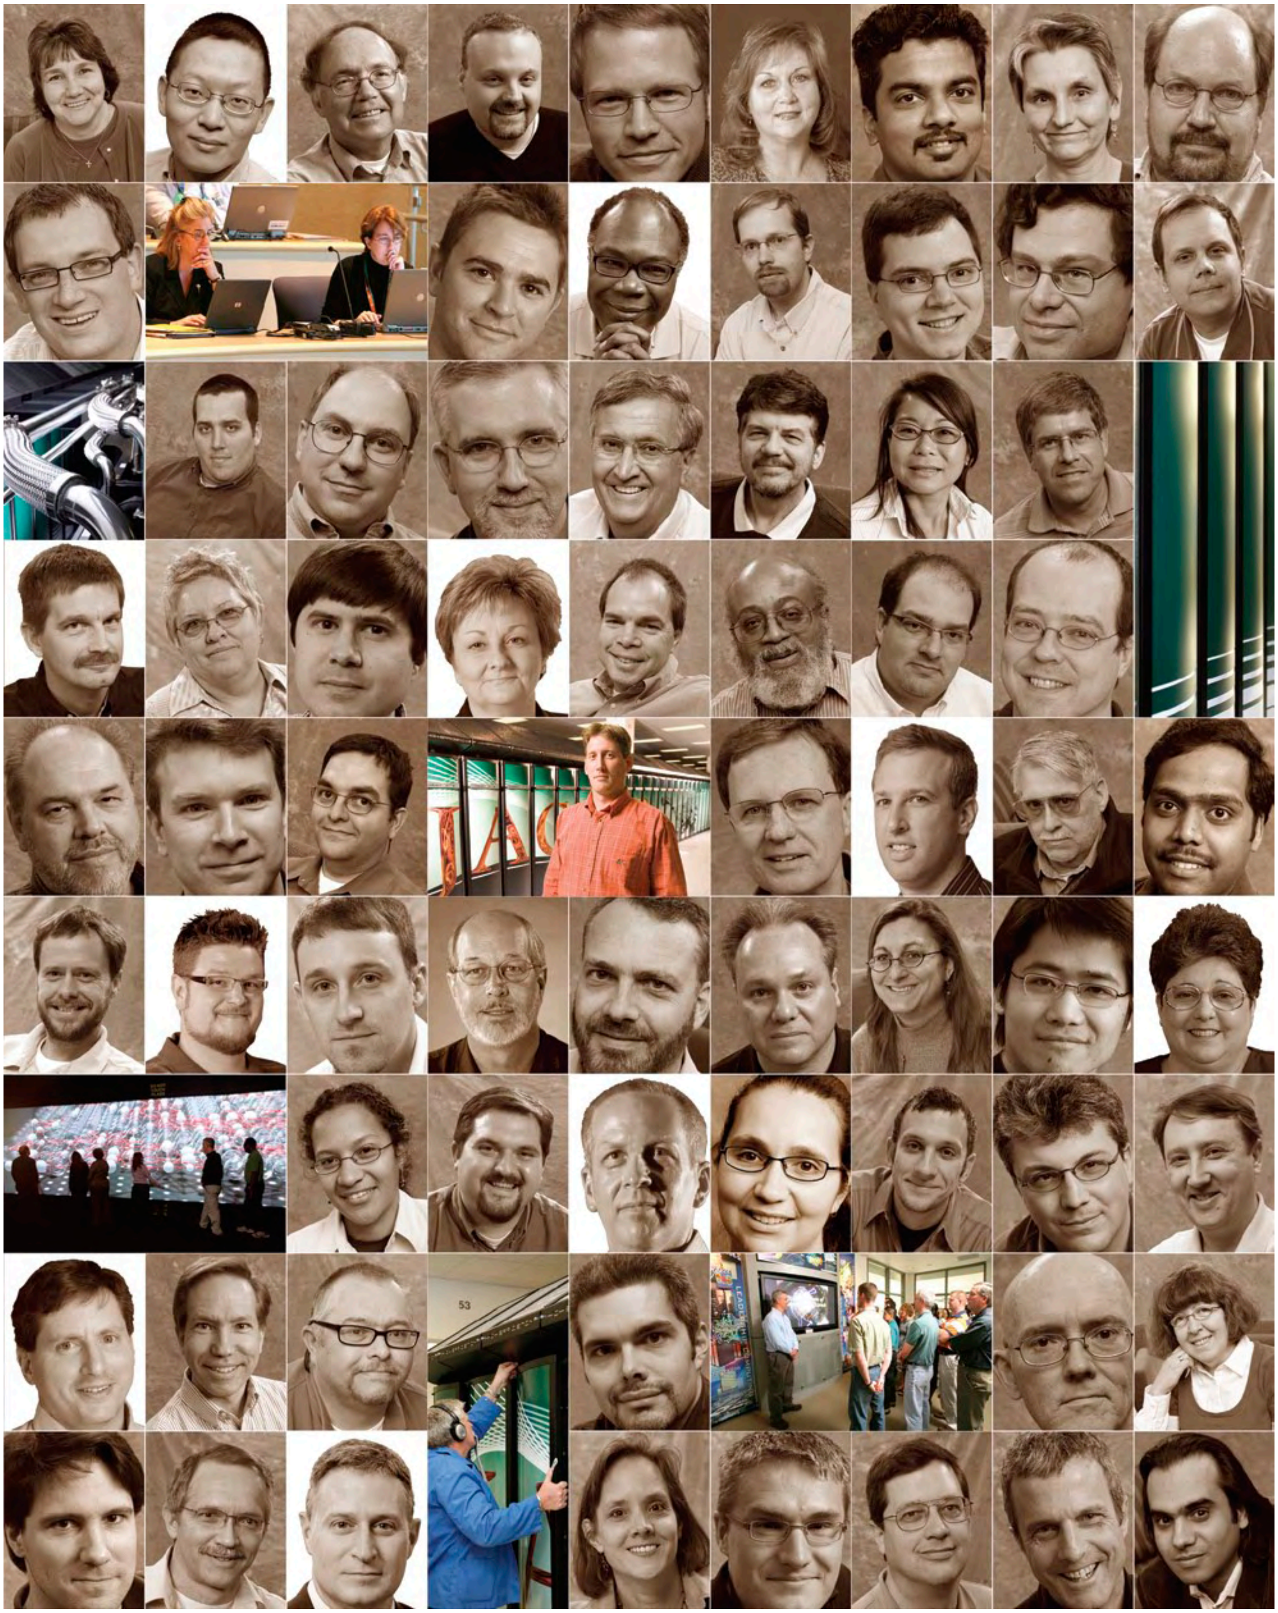

NCCS Organization

NCCS Management Team

The NCCS manages and supports major activities for DOE's Office of Science, such as testing and evaluation of leadership computing systems and operation of these machines for scientific research.

Helping the nation's premier researchers make the most of NCCS resources is an impressive management team. NCCS Director Jim Hack guides the overall vision of the center. Oak Ridge Leadership Computing Facility (OLCF) Project Director Buddy Bland oversees installation and upgrades of the NCCS 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 undertakings and planning.

Scientific Computing Group

The Scientific Computing Group (SCG) works with the users of high-performance computers to help them get the best results from the NCCS's computational resources. The SCG includes research scientists, visualization specialists, and work-flow experts who are themselves trained in chemistry, physics, astrophysics, mathematics, numerical analysis, or computer science. Each research group accepted to work on NCCS systems is assigned an SCG liaison who is familiar with its field, actively participates in the research, and solves the group's computer problems. The liaison designs code, optimizes it for the users' applications, and streamlines the work flow. 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 Group

The User Assistance & Outreach Group facilitates the scientific users' access to the NCCS and explains both the research science being done and the computational tools used to deliver it. It provides Level I and Level II technical support to researchers; writes technical documents explaining the systems, its access policies, and procedures; and answers the public's questions. Its staff attend the meetings of the other NCCS 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 NCCS with the nation's universities. Contact James B. White III, acting group leader, at *trey@ornl.gov*.

High Performance Computing Operations Group

The High Performance Computing Operations Group keeps the NCCS leadership supercomputing systems running. The group supports around-the-clock computer operations and is 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 supercomputers on-site. Its members test the systems when they are installed and upgraded and use diagnostic tools to continually monitor them. They anticipate problems before they arise and identify components that are near failure. The group also oversees cybersecurity for the NCCS, ensuring that systems conform to ORNL cybersecurity policy. Contact Ann Baker, group leader, at *bakerae@ornl.gov*.

Technology Integration Group

The Technology Integration (TechInt) Group works to update and integrate the networks, file systems, and archival storage infrastructure into the NCCS computer systems. Its members research and evaluate emerging technologies and write the programs that make new technologies—such as the new Lustre file system—work seamlessly with current systems. TechInt developed the NCCS's massive High-Performance Storage System (HPSS) and is constantly working to increase the speed of data transfer and implement cybersecurity measures for the NCCS's area-wide network. As NCCS systems continue to scale up, the TechInt team is developing tools (e.g., compilers, debuggers, and performance-analysis tools) that enable users to take full advantage of the leadership-class systems. Contact Galen Shipman, group leader, at *gshipman@ornl.gov*.

Application Performance Tools Group

The Application Performance Tools Group tracks and purchases a wide range of software tools that help science researchers assess and improve the performance of their applications on current and emerging NCCS HPC systems. The group researches the software and 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 NCCS's new 1.64 petaflops Cray XT4/XT5 Jaguar supercomputer. Contact Rich Graham, group leader, at *rlgraham@ornl.gov*.

RESOURCES

CURRENT SYSTEMS

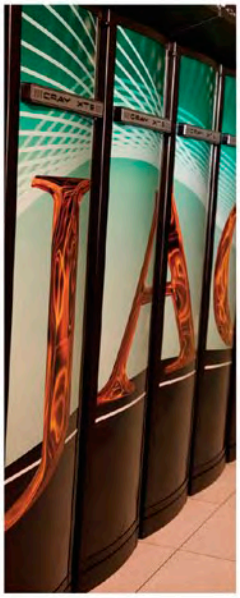

JAGUAR
Cray XT5

JAGUAR
Cray XT4

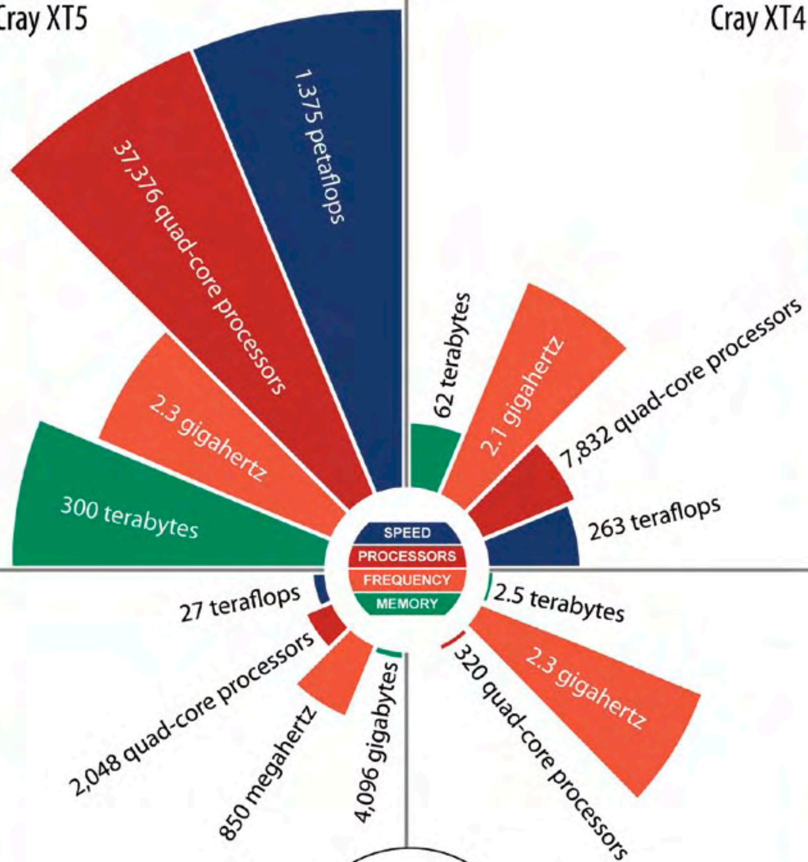

EUGENE
IBM BlueGene

"The World's Most Powerful Open Scientific Computing Complex"
Dr. Arden Bement, Director, National Science Foundation

SMOKY
Development cluster

SUPPORT SYSTEMS

EVEREST
Scientific Visualization Lab

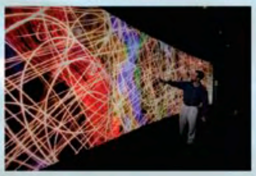

- 30 x 8 feet long
- 27-projector PowerWall
- 35 million pixels

HPSS
Backup Storage

- Many storage devices supported
- More than 30 petabytes of capacity

LENS
Scientific Visualization Cluster

- 128 quad-core processors
- 2.3 gigahertz
- 64 gigabytes 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
- 480-volt power
- 280-megawatt substation
- LEED certified by U.S Green Building Council

STATISTICS

The charts on this page present a statistical overview of the NCCS in 2008: the research conducted here, the scientists conducting the research, and the organizations supporting it. These charts clearly show that leadership computing at the NCCS spans a large range of scientific disciplines and research organizations.

INCITE Allocation Hours on Jaguar XT4 by Discipline

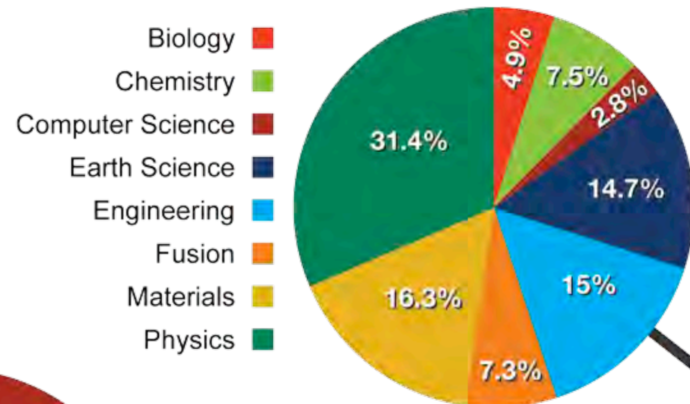

Active Users by Sponsor

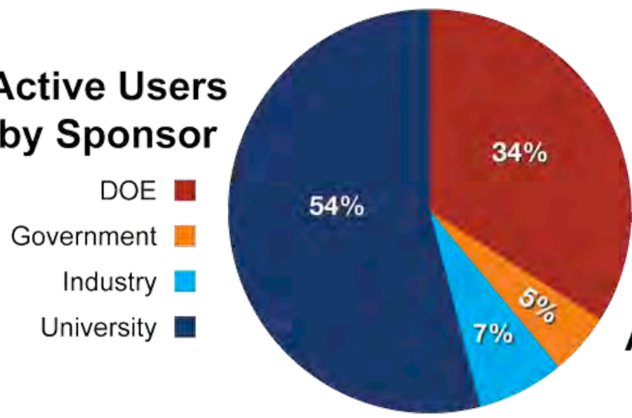

Allocation Hours on Jaguar XT4

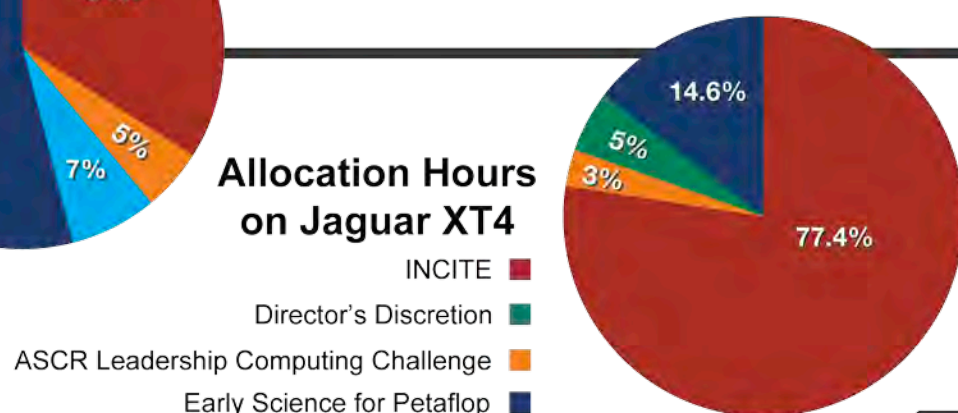

Active Users by Discipline

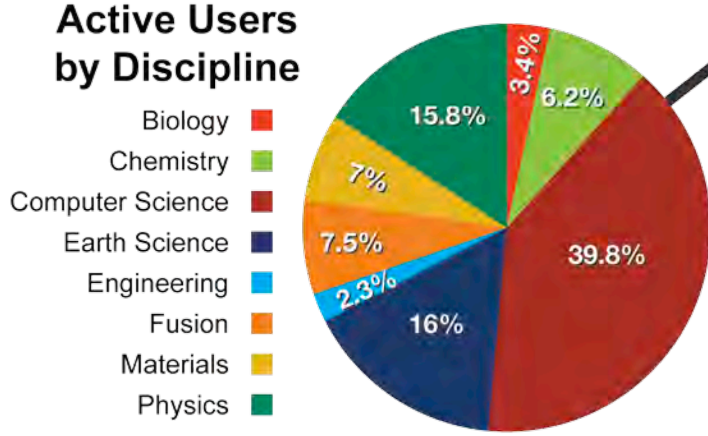

Selected Publicatons

Using the resources of the NCCS, researchers have produced numerous scientific breakthroughs 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 NCCS.

Astrophysics

Blondin, J.M., and A. Mezzacappa. 2007. "Pulsar spins from an instability in the accretion shock of supernovae." *Nature* **445**(7123), 58–60 (Feb.).

Diemand, J., et al. 2008. "Clumps and streams in the local dark matter distribution." *Nature* **454**(7205), 735–738 (May).

Ropke, F.K., S.E. Woosley, and W. Hillebrandt. 2007. "Off-center ignition in Type Ia supernovae. I. Initial evolution and implications for delayed detonation." *Astrophysical Journal* **660**(2), 1344–1356 (May).

Chemistry

Cui, S.T., et al. 2007. "A molecular dynamics study of a Nafion polyelectrolyte membrane and the aqueous phase structure for proton transport." *Journal of Physical Chemistry B* **111**(9), 2208–2218 (Feb.).

Romo-Herrera, J.M., et al. 2008. "An atomistic branching mechanism for carbon nanotubes: Sulfur as the triggering agent." *Angewandte Chemie-International Edition* **47**(16), 2948–2953 (April).

Climate

Qian, Y., et al. 2009. "Effects of soot-induced snow albedo change on snowpack and hydrological cycle in western United States based on weather research and forecasting chemistry and regional climate simulations." *Journal of Geophysical Research-Atmospheres* **114** (Feb.).

Computer Science

Alam, S.R., et al. 2008. "An evaluation of the Oak Ridge National Laboratory Cray XT3." *International Journal of High Performance Computing Applications* **22**(1), 52–80 (Feb.).

Engineering

Hawkes, E.R., et al. 2007. "Scalar mixing in direct numerical simulations of temporally evolving plane jet flames with skeletal CO/H-2 kinetics." *Proceedings of the Combustion Institute* **31**, 1633–1640.

Fusion

Waltz, R.E., J. Candy, and M. Fahey. 2007. "Coupled ion temperature gradient and trapped electron mode to electron temperature gradient mode gyrokinetic simulations." *Physics of Plasmas* **14**(5), 056116 (May).

Materials

Kent, P.R.C., et al. 2008. "Combined density functional and dynamical cluster quantum Monte Carlo calculations of the three-band Hubbard model for hole-doped cuprate superconductors." *Physical Review B* **78**(3), 035132 (July).

Nanosciences

Jiang, D.E., B.G. Sumpter, and S. Dai. 2007. "First principles study of magnetism in nanographenes." *Journal of Chemical Physics* **127**(12), 124703 (Sept.).

Jiang, D.E., B.G. Sumpter, and S. Dai. 2007. "Unique chemical reactivity of a graphene nanoribbon's zigzag edge." *Journal of Chemical Physics* **126**(13), 134701 (April).

Physics

Hagen, G., et al. 2007. "Coupled-cluster theory for three-body Hamiltonians." *Physical Review C* **76**(3), 034302 (Sept.).

Sen, C., G. Alvarez, and E. Dagotto. 2007. "Competing ferromagnetic and charge-ordered states in models for manganites: The origin of the colossal magnetoresistance effect." *Physical Review Letters* **98**(12), 127202 (March).

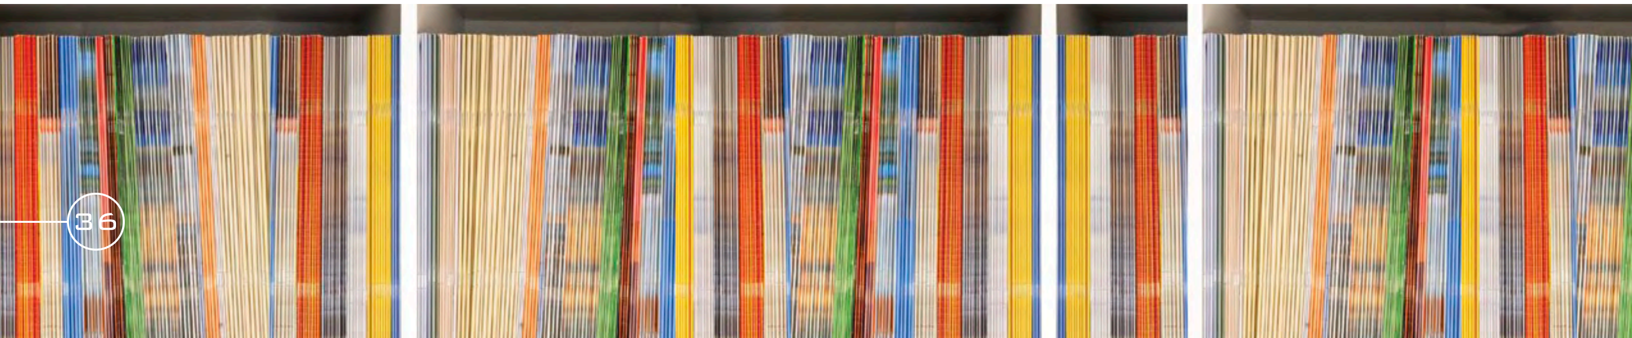

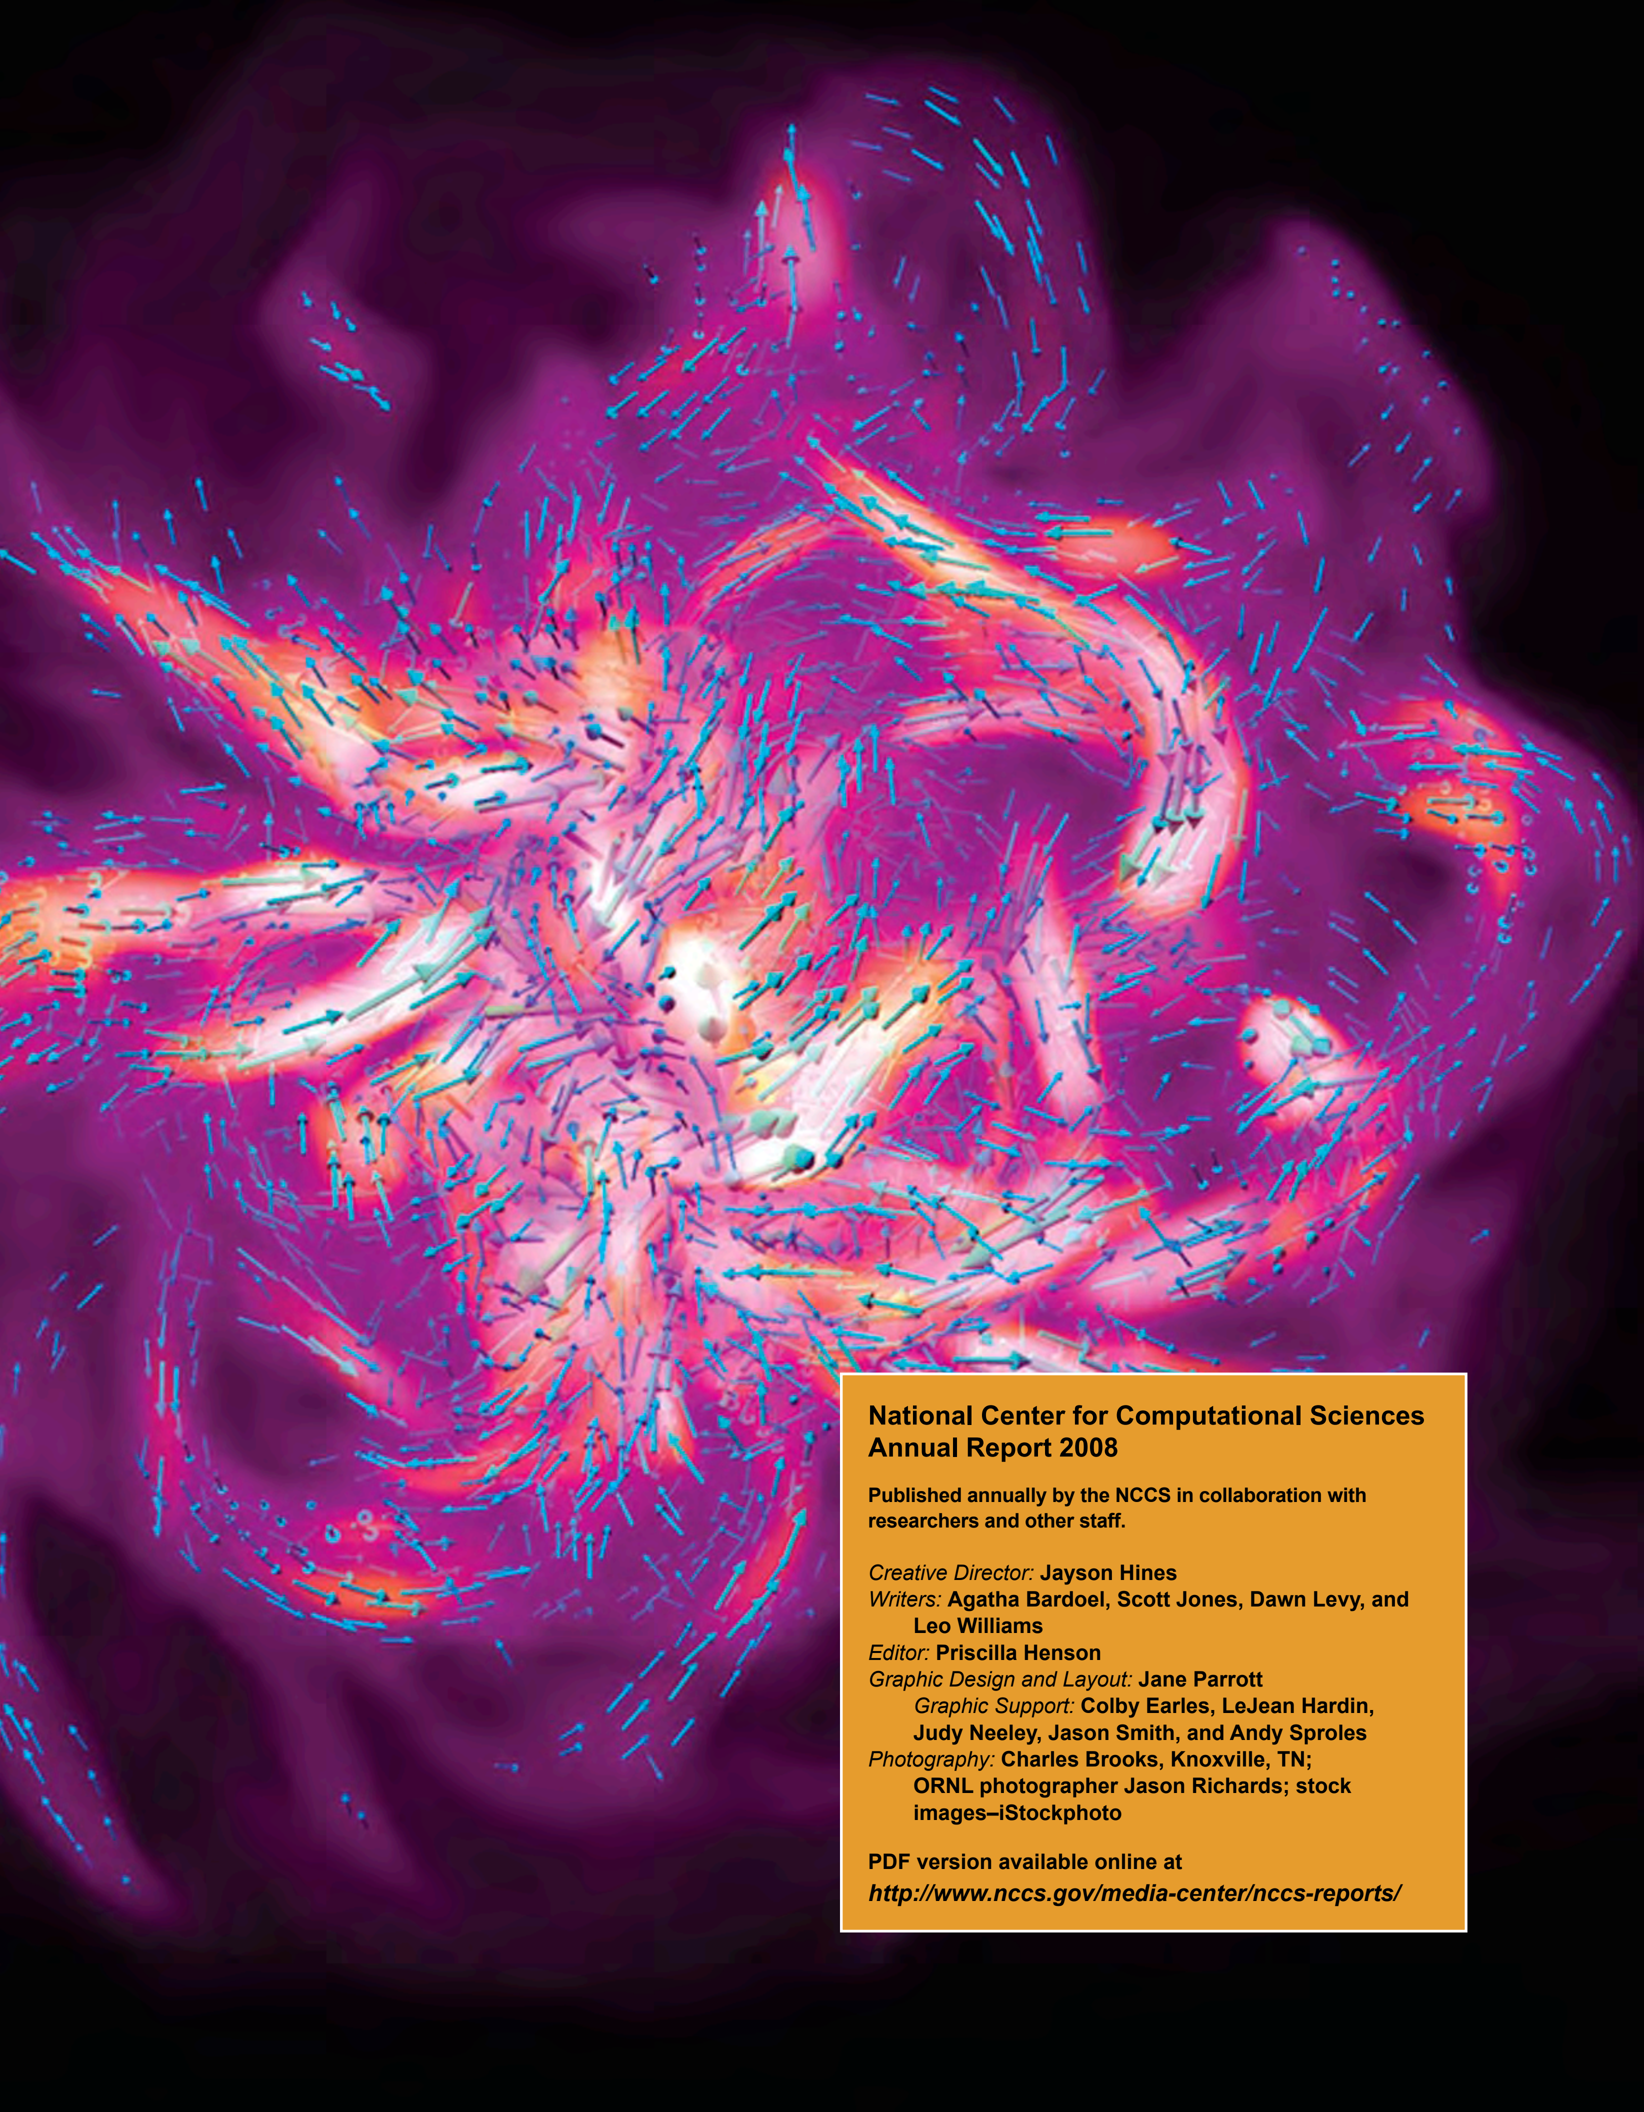

National Center for Computational Sciences Annual Report 2008

Published annually by the NCCS in collaboration with researchers and other staff.

Creative Director: **Jayson Hines**
Writers: **Agatha Bardoel, Scott Jones, Dawn Levy, and Leo Williams**
Editor: **Priscilla Henson**
Graphic Design and Layout: **Jane Parrott**
Graphic Support: **Colby Earles, LeJean Hardin, Judy Neeley, Jason Smith, and Andy Sproles**
Photography: **Charles Brooks, Knoxville, TN; ORNL photographer Jason Richards; stock images–iStockphoto**

PDF version available online at
http://www.nccs.gov/media-center/nccs-reports/

ORNL 2009-G00827/jcp