2007 ANNUAL REPORT

NATIONAL CENTER FOR COMPUTATIONAL SCIENCES

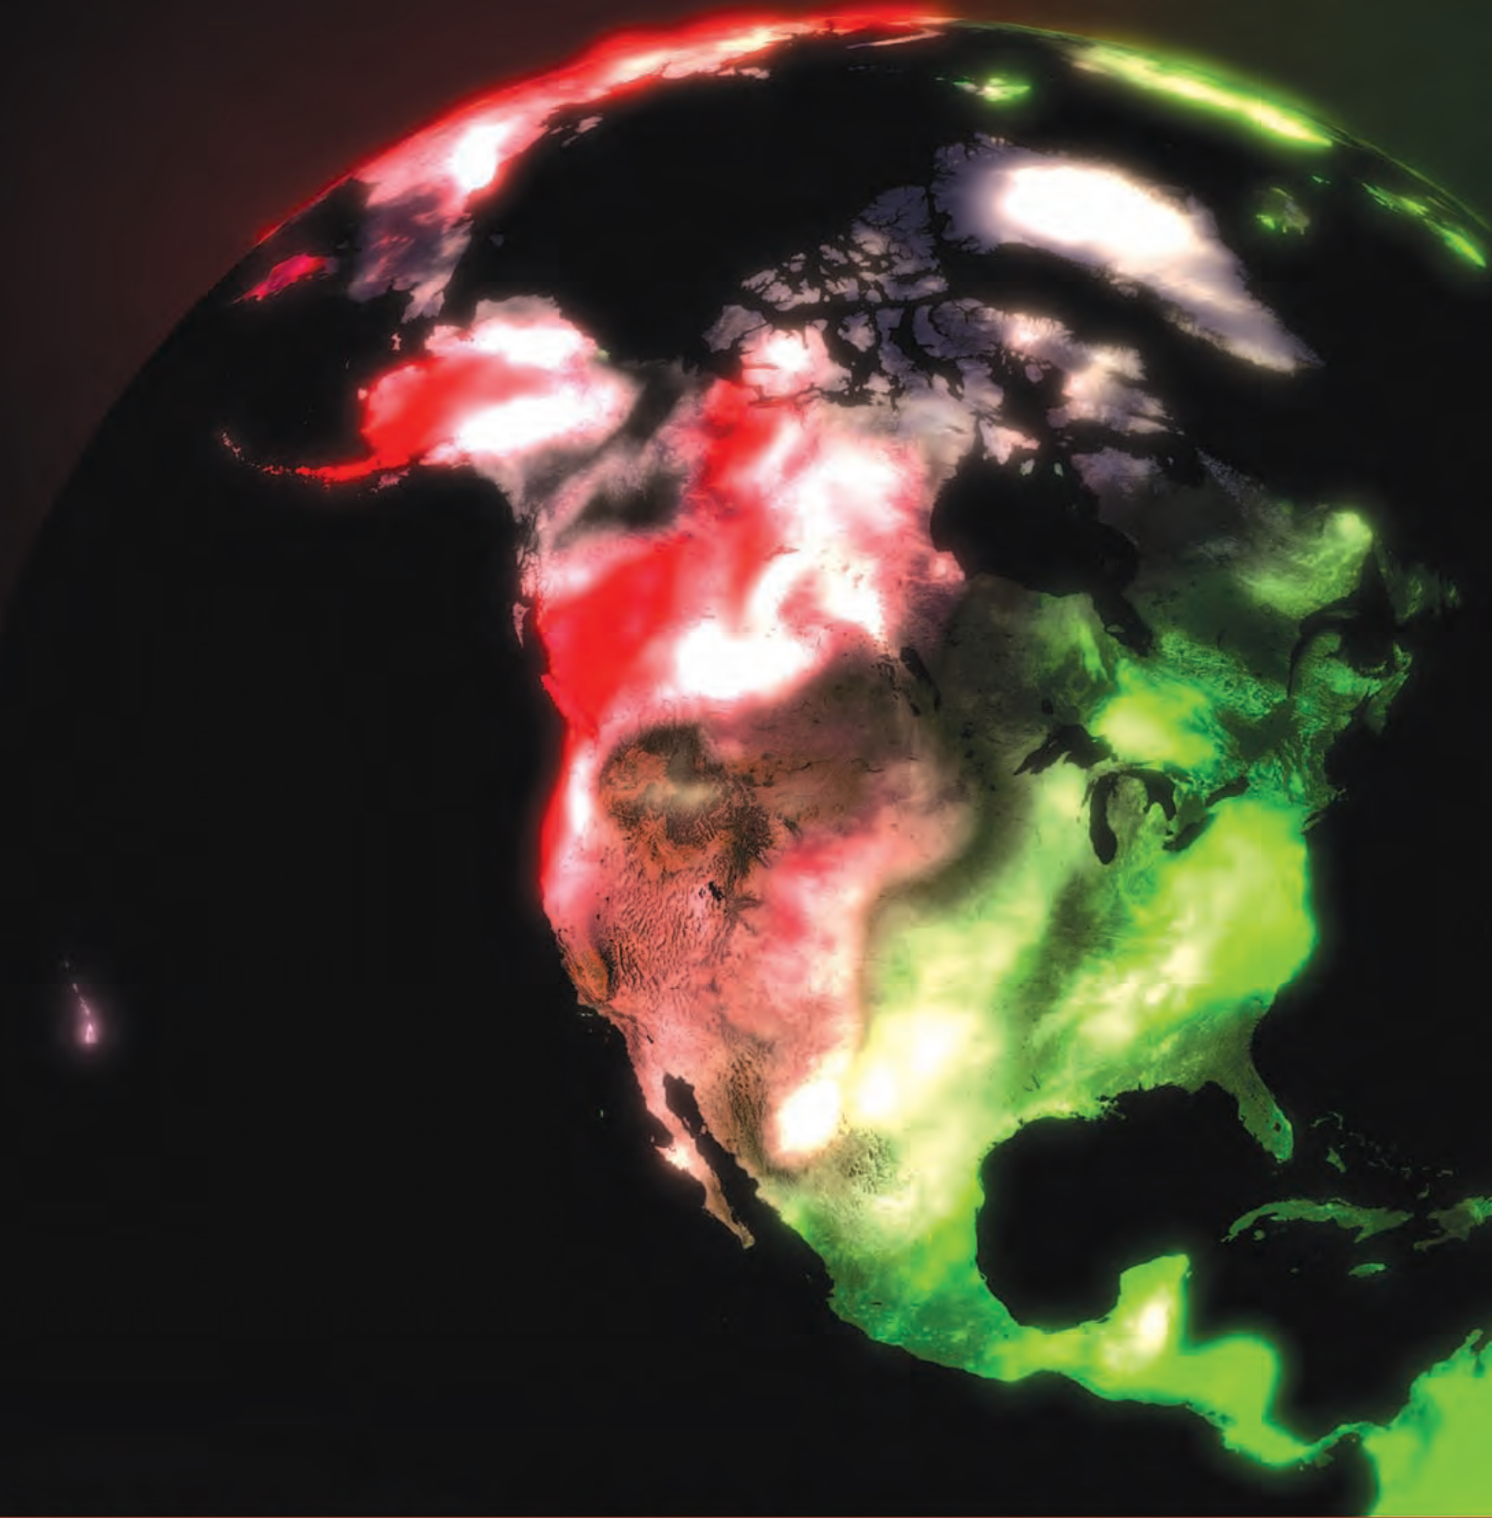

REAL WORLD SOLUTIONS FOR REAL WORLD PROBLEMS

Contents

About the Cover

The instantaneous net ecosystem exchange of CO_2 is shown projected onto the land surface from a Carbon-Land Model Intercomparison Project (C-LAMP) simulation during July 2004 as the sun rises over North America. Green represents an uptake by the biosphere because of photosynthesis where the sun has risen, while red represents a net flux to the atmosphere, dominated by heterotrophic respiration, where it is nighttime. *Scientific visualization by Jamison Daniel, ORNL/NCCS.*

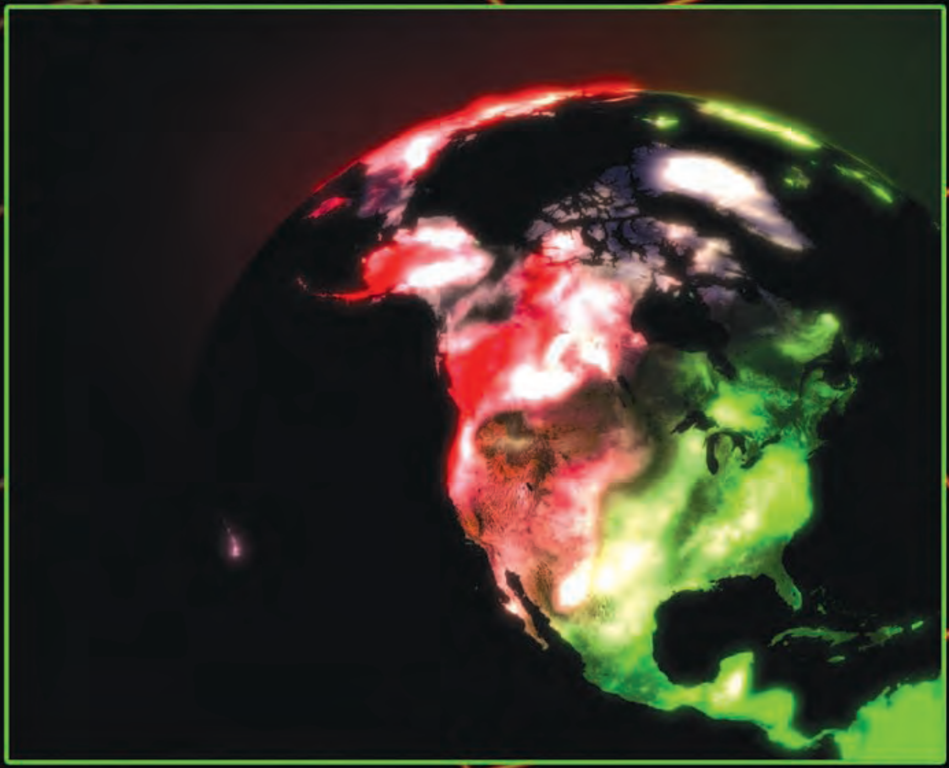

About the NCCS...

The National Center for Computational Sciences (NCCS) provides the most powerful computing resources in the world for open scientific research, research that will advance fundamental scientific knowledge about the world we live in and holds the promise to improve quality of life through scientific progress and technical innovation.

As a Department of Energy Leadership Computing Facility, the NCCS is poised to tackle some of our most pressing scientific and energy challenges.

As we now well know, the Earth is getting warmer, threatening everything from fisheries to arable land to currently settled coastlines. NCCS systems have helped to model the Earth's climate like never before, giving researchers a greater understanding of the mechanisms responsible and better preparing policymakers to tackle the various consequences of climate change.

One reason for this warming is our reliance on fossil fuels, a finite resource that still provides the basis for the majority of American energy production. This reliance also adds to the already large emissions of greenhouse gases into the air we breathe and cripples economic growth due to our nation's continued dependence on foreign fuel sources. Research at the NCCS is leading to a rapidly increasing understanding of fusion energy, a cleaner energy production method with minimal hazardous byproducts and virtually unlimited fuel sources. If realized, fusion energy will revolutionize the way in which the world meets its energy demands, making the air we breathe cleaner and reducing our dependence on finite resources.

Computational research into diesel engine combustion also promises to greatly enhance the efficiency of automobiles and freight vehicles, further cleaning up our air and increasing our energy independence and assurance.

These are just a few of the numerous areas of groundbreaking research enabled by the NCCS. Major advances in biology, astrophysics, and materials research are further examples of a very rich computational science portfolio supported by NCCS facilities and staff, advancing the spectrum of knowledge and innovation possible with computer simulation.

Landmark advances in scientific and technical understanding are the result of the NCCS's mission to provide capability computational science—simulations that take advantage of the majority of a system's potential. While other high-performance computer systems host large numbers of modest sized jobs, the NCCS seeks out and supports research that can only be tackled with the center's world-leading supercomputers, such as the Cray XT4 known as Jaguar. Ranked No. 2 on the Top 500 list of the world's fastest supercomputers in June 2007, Jaguar boasted a peak performance of more than 119 teraflops (119 trillion calculations per second).

The NCCS continues to aggressively expand its computing power. Steps are under way to expand the speed of Jaguar to 250 teraflops and to install a petaflop system, capable of a quadrillion calculations per second, in 2008. This magnitude of computing power provides researchers with a virtual laboratory, where simulation is the scientific tool of choice to confront problems out of the reach of theory and experiment. While simulation indeed takes place in a virtual world, it is increasingly providing the real world with the necessary solutions to its most pressing problems.

Besides scientific progress, America's competitiveness rests with its ability to provide technological expertise and innovation, a movement exemplified by the NCCS's ongoing mission: to provide researchers with state-of-the-art supercomputers capable of finding answers to complex problems that will enhance and improve our lives.

About ORNL

Oak Ridge National Laboratory (ORNL) is a multiprogram science and technology laboratory managed for the Department of Energy (DOE) by UT-Battelle, LLC. Scientists and engineers at ORNL conduct basic and applied research and development to create scientific knowledge and technological solutions that strengthen the nation's leadership in key areas of science; increase the availability of clean, abundant energy; restore and protect the environment; and contribute to national security.

ORNL pioneers the development of new energy sources, technologies, and materials and the advancement of knowledge in the biological, chemical, computational, engineering, environmental, physical, and social sciences.

James J. Hack, Director
National Center for Computational Sciences

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

TEL 865-241-7202 **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.

Introduction

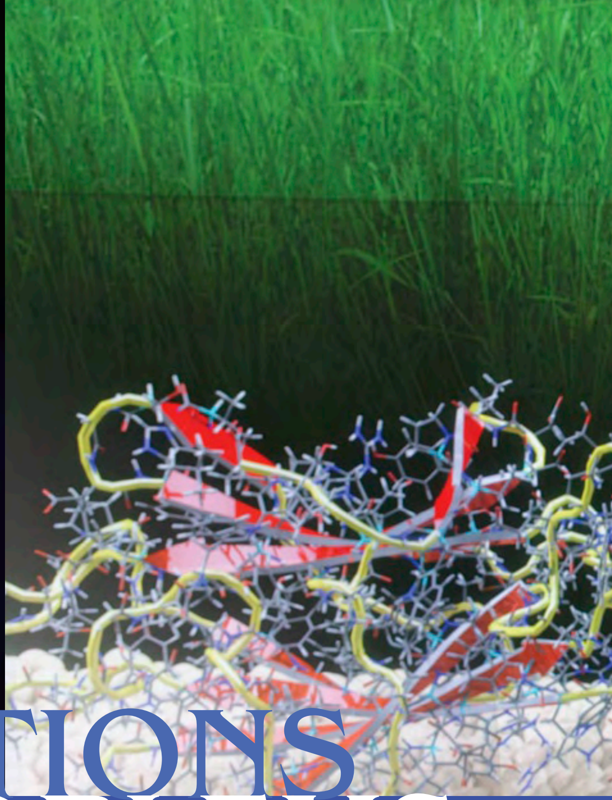

REAL SOLUTIONS for REAL PROBLEMS

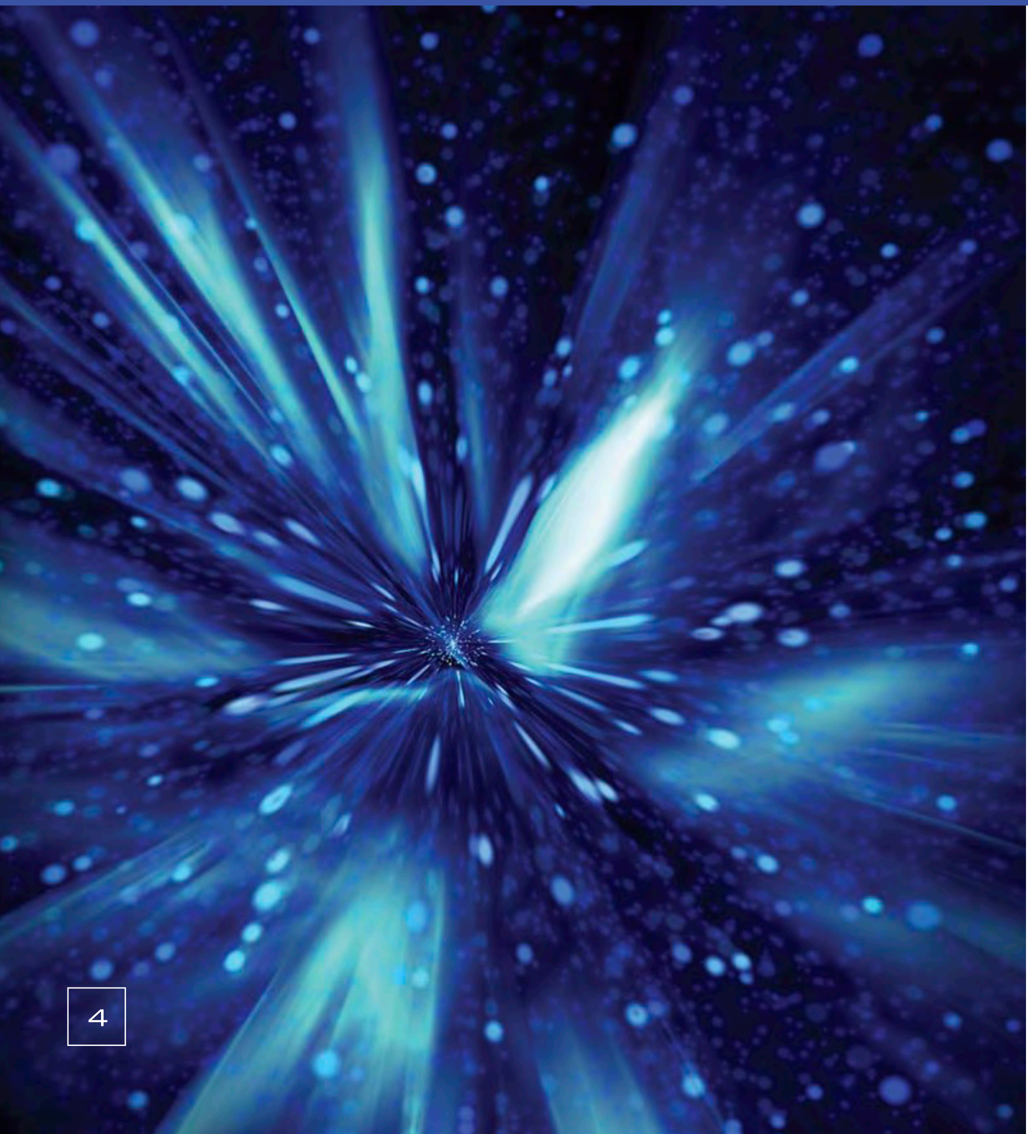

Looking back on the accomplishments of 2007 and forward to the achievements of 2008 and beyond, it is clear that computational science has become increasingly important for addressing real issues.

Energy assurance and climate studies, especially, have moved from the realm of the specialist, taking up residence on the front page as both fuel prices and global temperatures reach record highs. The world is confronted with serious and immediate problems that will require all of the scientific and technological resources at our disposal. Under the circumstances, we can be grateful for the explosive growth of computing power.

The challenges are obvious. Our reliance on foreign oil is particularly troublesome,

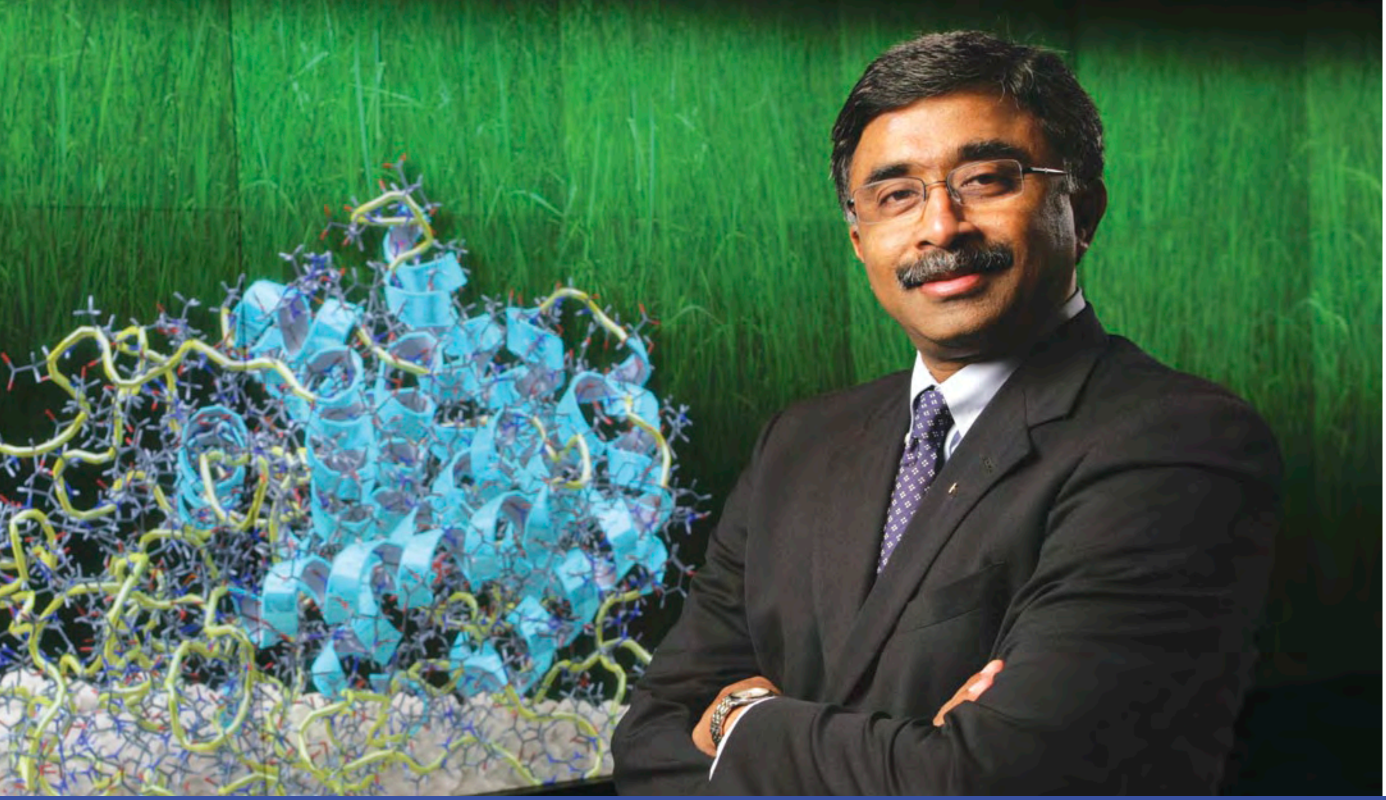

Associate Laboratory Director Thomas Zacharia leads ORNL's enterprises in computational science.

threatening both our economy and our environment. The challenges of limited energy supply and environmental damage will only increase as developing nations boost their own energy demand in order to improve their standard of living.

Solutions to these problems cannot come too soon. Fortunately, they will come sooner through the efforts of worldclass computational researchers working at the NCCS and other supercomputer centers. The breathtaking power of the world's most powerful computers—placed in the hands of computational physicists, chemists, biologists, and others—can inform and guide our response to climate change, speed the development of alternative energy sources, help find cleaner ways to make use of traditional energy sources, and transform our idea of what technologies are possible.

We at ORNL are proud to use our supercomputing resources to attack these problems from a variety of angles. Combustion and materials scientists are helping to make the most of conventional sources of energy, working to improve and maximize the efficiency of the combustion process and recycle waste heat in vehicles as electricity. Biologists are boosting the efficiency—and reducing the disruptiveness—of biofuels production, focusing on hardy, fast-growing, non-food plants to provide plentiful and inexpensive ethanol. Fusion scientists are pushing the limits of physics in their quest to turn nuclear fusion into a clean, inexpensive, almost limitless source of electricity.

Climate scientists, on the other hand, are working to explain and mitigate the environmental consequences of our energy use. ORNL provides a major component of the computing resources needed for the Community Climate System Model (CCSM), the world's leading framework for analyzing both the causes of and potential responses to climate change. This model provides the most sophisticated climate analysis tools in history, tools that are being continually refined and updated as we bring more powerful computing systems online. The talented computational scientists guiding this effort are providing policy-makers and the public alike the information needed to make responsible and possibly difficult choices. Geologists, too, are using NCCS supercomputers to improve the environment by developing the means to store massive volumes of greenhouse gases deep underground.

I am confident we will continue to face these challenges with creativity and perseverance. We will do our part by remaining at the cutting edge of computation, providing researchers with the world's most powerful systems and the computational expertise to efficiently use these powerful tools. These efforts will also help us to prepare for the unforeseen challenges of the future.—*By Thomas Zacharia*

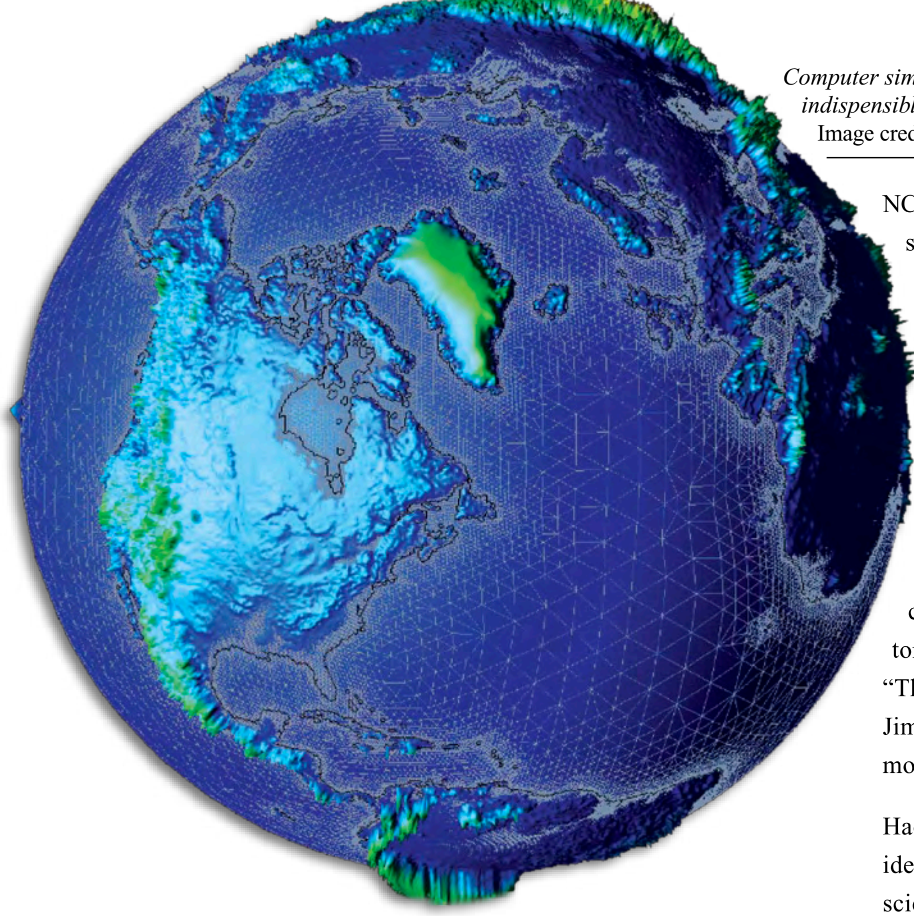

Computer simulation has become indispensible to climate studies.
Image credit: Ahmed Khamayseh, ORNL

NCCS Gets a New Director:
Meet James J. Hack

In 2007 James J. Hack, a senior scientist at the National Center for Atmospheric Research (NCAR) in Boulder, Colorado, joined the NCCS as its new director and heads the DOE Leadership Computing Facility (LCF).

The LCF, housed at ORNL, hosts some of the world's fastest machines, capable of carrying out trillions of operations each second (teraflops) and quickly approaching quadrillions of calculations each second (petaflops).

"We are thrilled that Jim Hack [has joined] us to lead America's premier open computing facility," said ORNL Director Thom Mason. "Jim is a global leader in climate research and has devoted his career to gaining insight into Earth's atmosphere, where dynamics are complex and a lot is at stake. He is well suited to lead the NCCS in addressing the grand scientific challenges of this century, which include climate but also extend to fields such as biology, chemistry, physics, and even computation itself. This research supports endeavors such as developing renewable energy and gaining a better understanding of our universe."

NCCS supercomputers provided more than one-third of the simulation data jointly contributed by DOE and the National Science Foundation (NSF) to the most recent assessment report of the United Nations' Intergovernmental Panel on Climate Change (IPCC), the group that shared the 2007 Nobel Peace Prize with former Vice President Al Gore. Hack was one of the principal developers of the model that was run on the NCCS supercomputers.

"Jim brings to the NCCS a depth of experience wresting great science from the world's most powerful machines," said Thomas Zacharia, ORNL associate laboratory director for computing and computational sciences. "The goal of supercomputing is deep insight, and with Jim's leadership the nation's top researchers will make the most of petascale computing."

Hack, with LCF Project Director Arthur "Buddy" Bland, identifies major high-performance computing needs from scientific and hardware perspectives and puts forth strategies to meet those needs as machines evolve to the petascale and beyond. Hack also leads the Climate Change Initiative at ORNL, directing a team of scientists and engineers across the laboratory in advancing the state of the art in Earth system discovery and policy through enhanced scientific understanding, Earth system modeling, and computational and observational programs.

"ORNL is in a unique position to substantially advance the most challenging of scientific problems, and I'm deeply honored to have been selected to play a leadership role in the laboratory's vision," Hack said. "I'm looking forward to being able to contribute to such an exceptional program in computational science and the important scientific insights it will reveal."

Hack headed the Climate Modeling Section at NCAR, a center sponsored by NSF, and served as deputy director of NCAR's Climate and Global Dynamics Division. He served as an adjunct professor in electrical and computer engineering at the University of Colorado at Boulder and is the author or coauthor of 98 scientific or technical publications.

After receiving his Ph.D. in atmospheric dynamics from Colorado State University in 1980, Hack became a research staff member at the IBM Thomas J. Watson Research Center, where he focused on mapping scientific algorithms to high-performance computing architectures.

He moved to NCAR in 1984, where he led development of the global atmospheric model currently called the Community Atmosphere Model. He has worked on all aspects of large-scale global modeling, including development and evaluation of numerical methods, development of analysis frameworks, and implementation of global models on high-performance computer systems.—*By Dawn Levy*

Meet the Management Team

The NCCS management team, from left: Doug Kothe, Jim Hack, Buddy Bland, Kathlyn Boudwin, and Jim Rogers.

The NCCS manages and supports major activities for the U.S. Department of Energy's Office of Science (SC), 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. Leadership Computing Facility Project Director Buddy Bland oversees installation and upgrades of the NCCS supercomputers. Kathlyn Boudwin is deputy director of the LCF project. As Director of Science, Doug Kothe guides the research teams that use the computing systems. Jim Rogers, director of operations, manages day-to-day undertakings and planning.

With input from an advisory committee and an operations council, the management team works closely with computer vendors and NCCS experts to bring new, high-end hardware and software into productive use. The interactions between government, industry, and academia result in a leadership-class resource for scientific users and guide development of future systems.

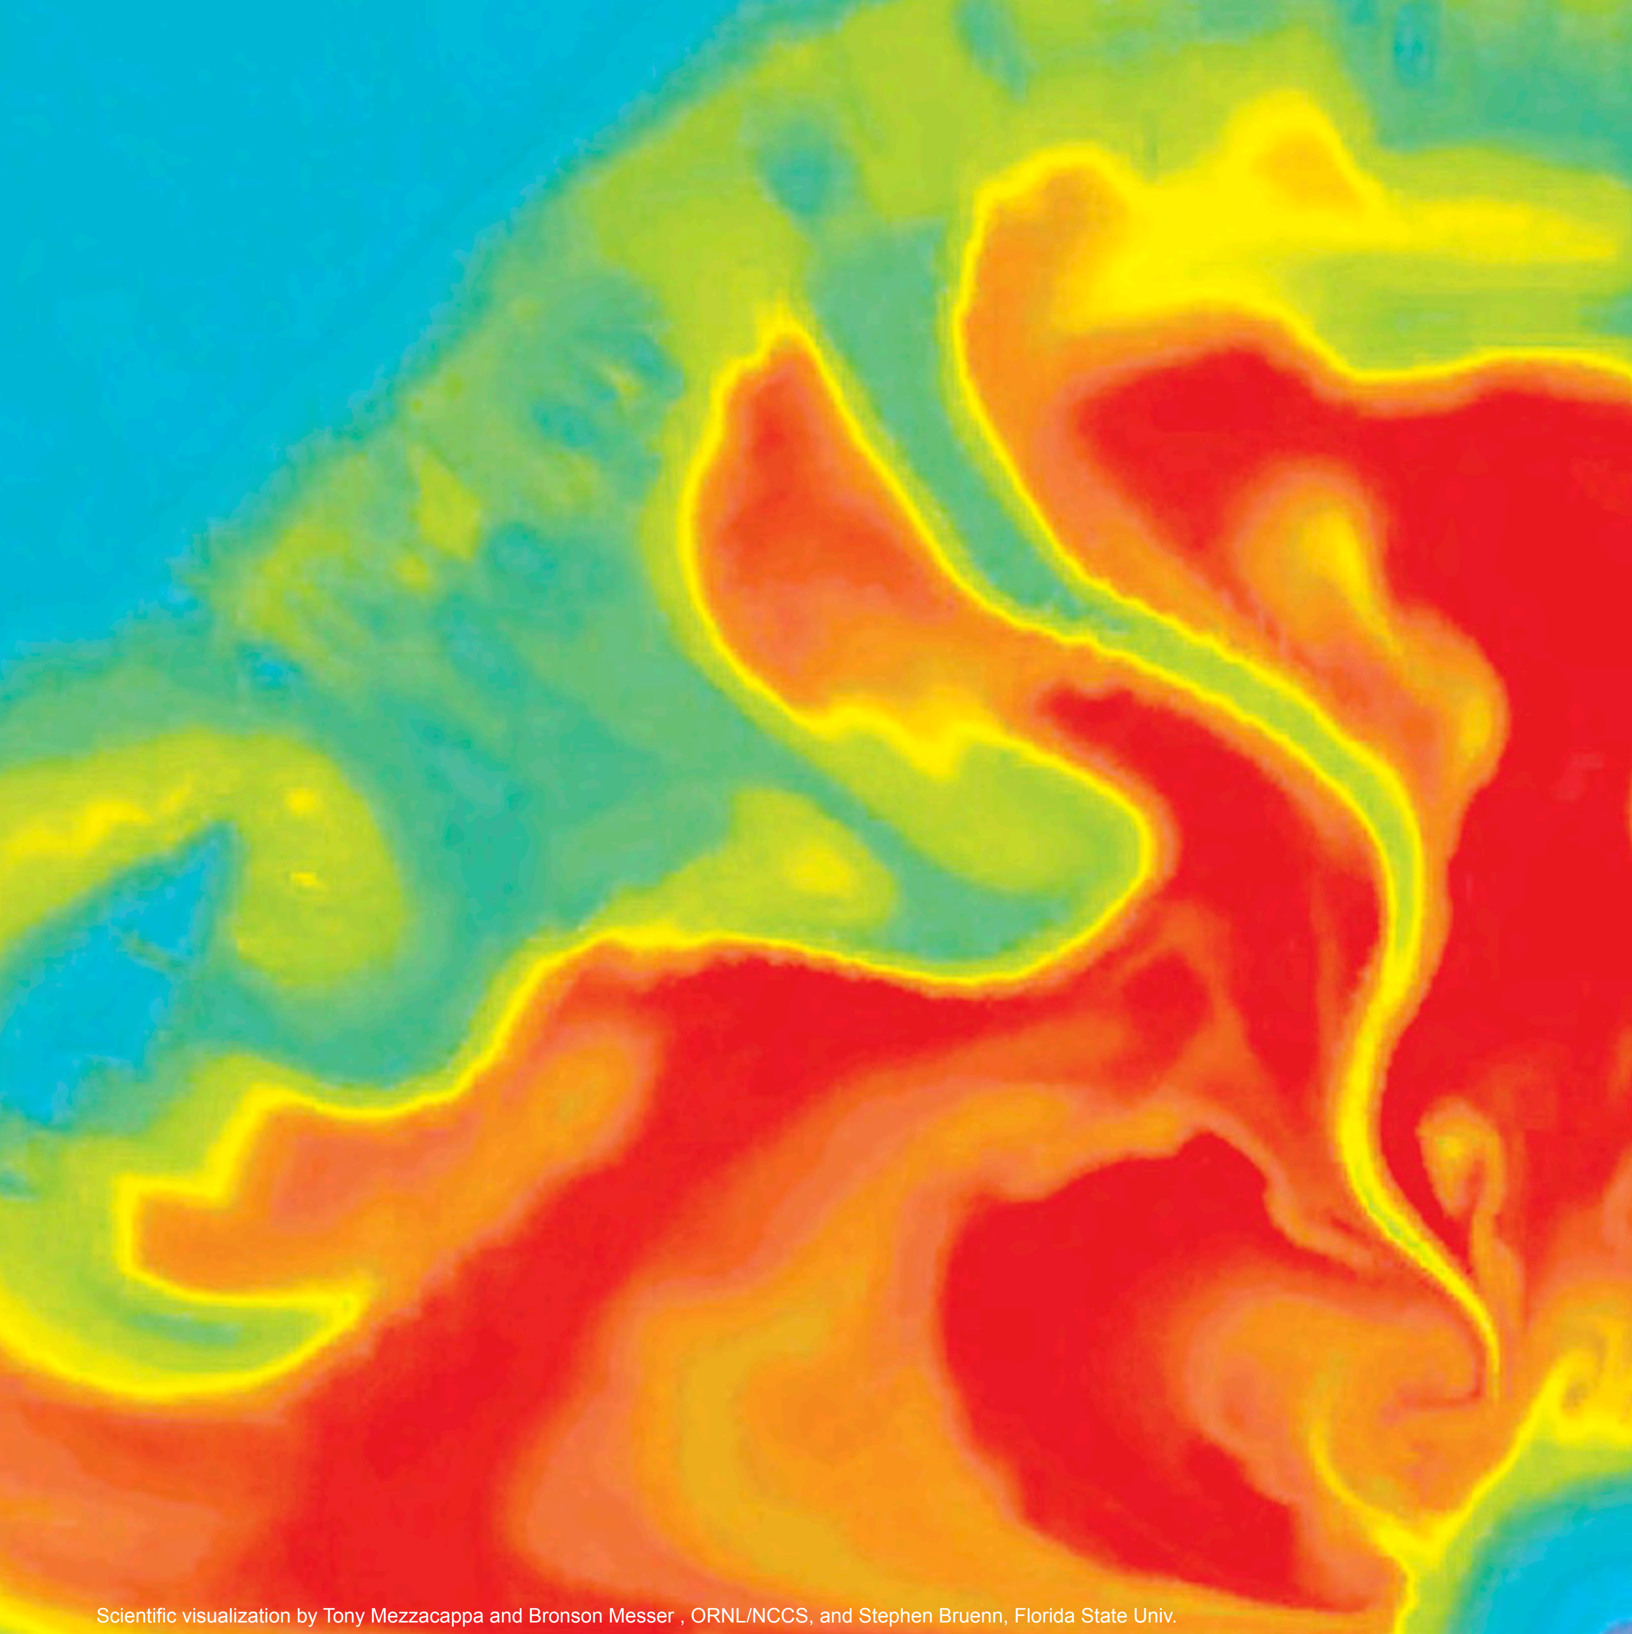

Scientific visualization by Tony Mezzacappa and Bronson Messer , ORNL/NCCS, and Stephen Bruenn, Florida State Univ.

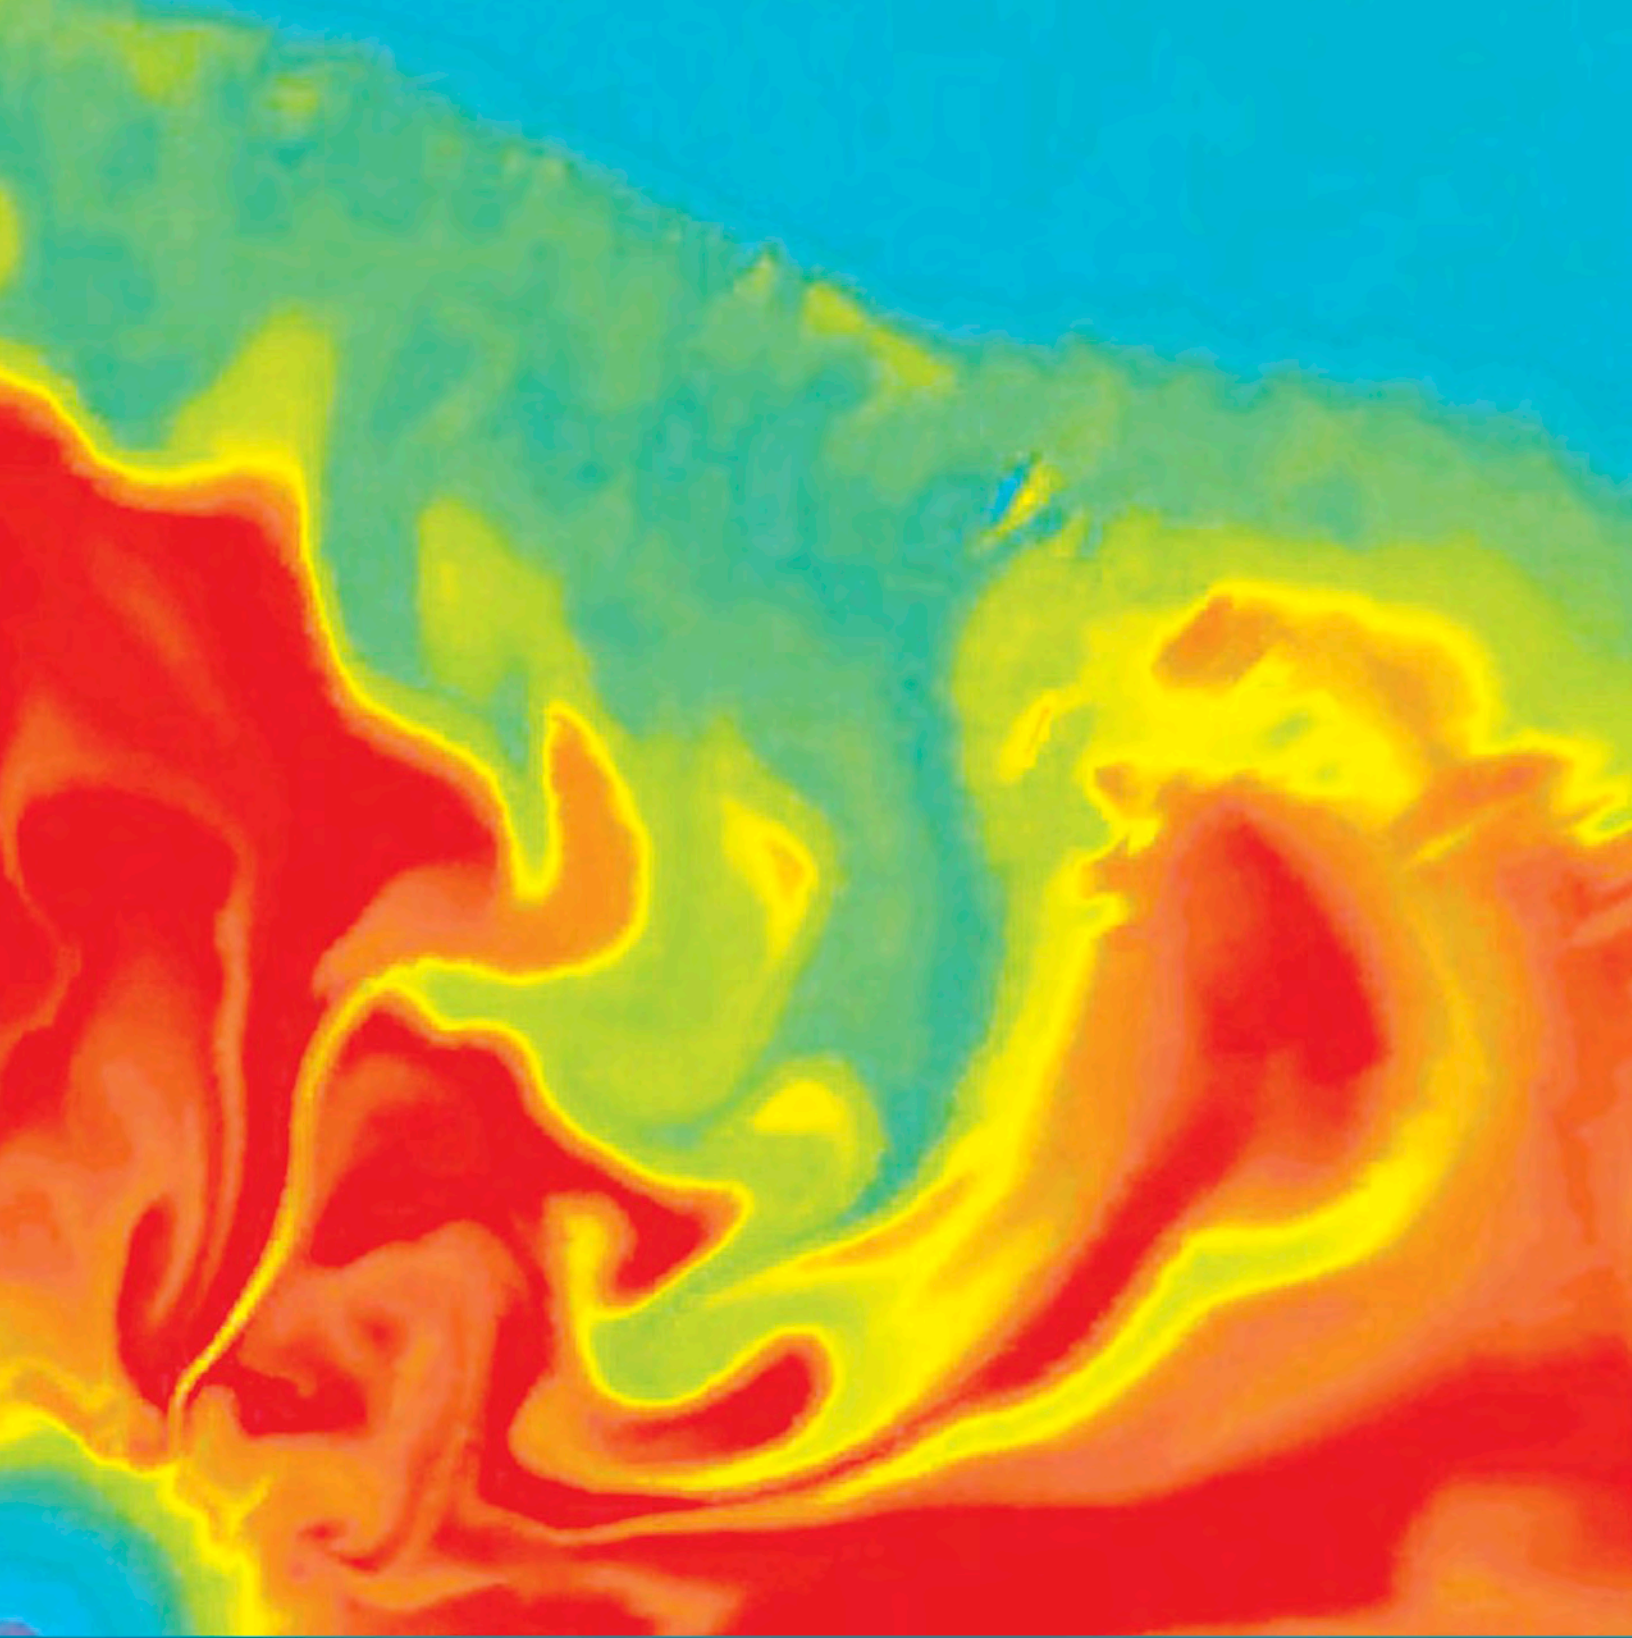

Science

Computation has come into its own as an invaluable tool in all areas of science, and we at the NCCS are proud of the role we have played in this revolution

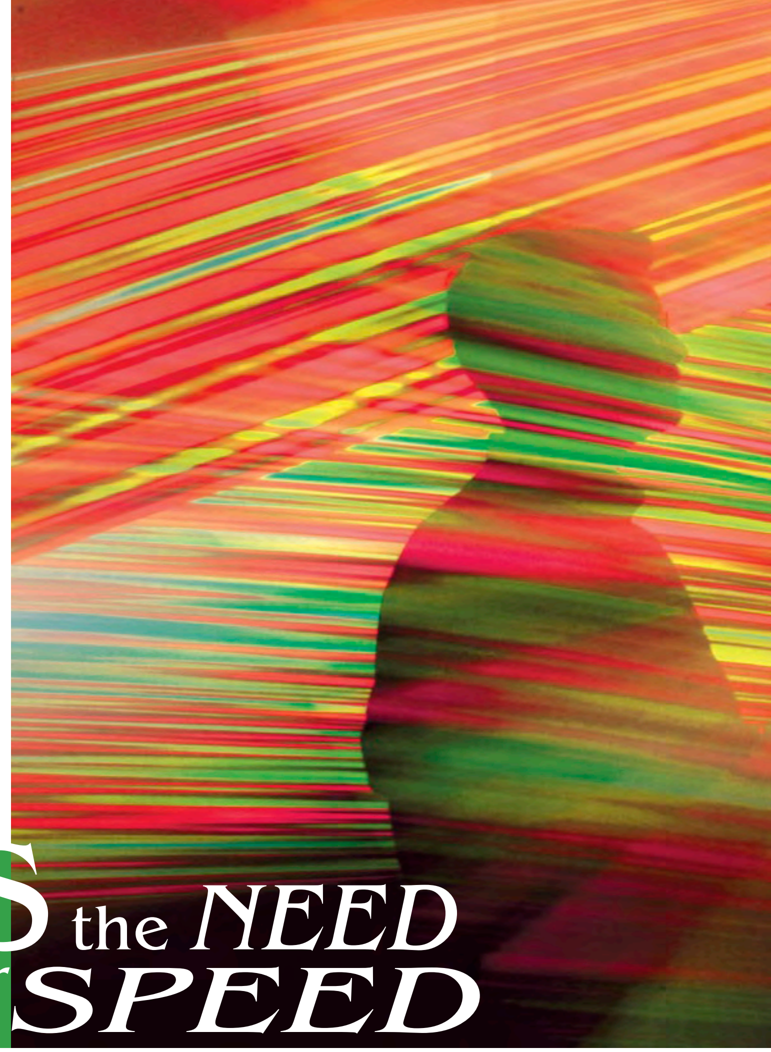

Science DRIVES the NEED for SPEED

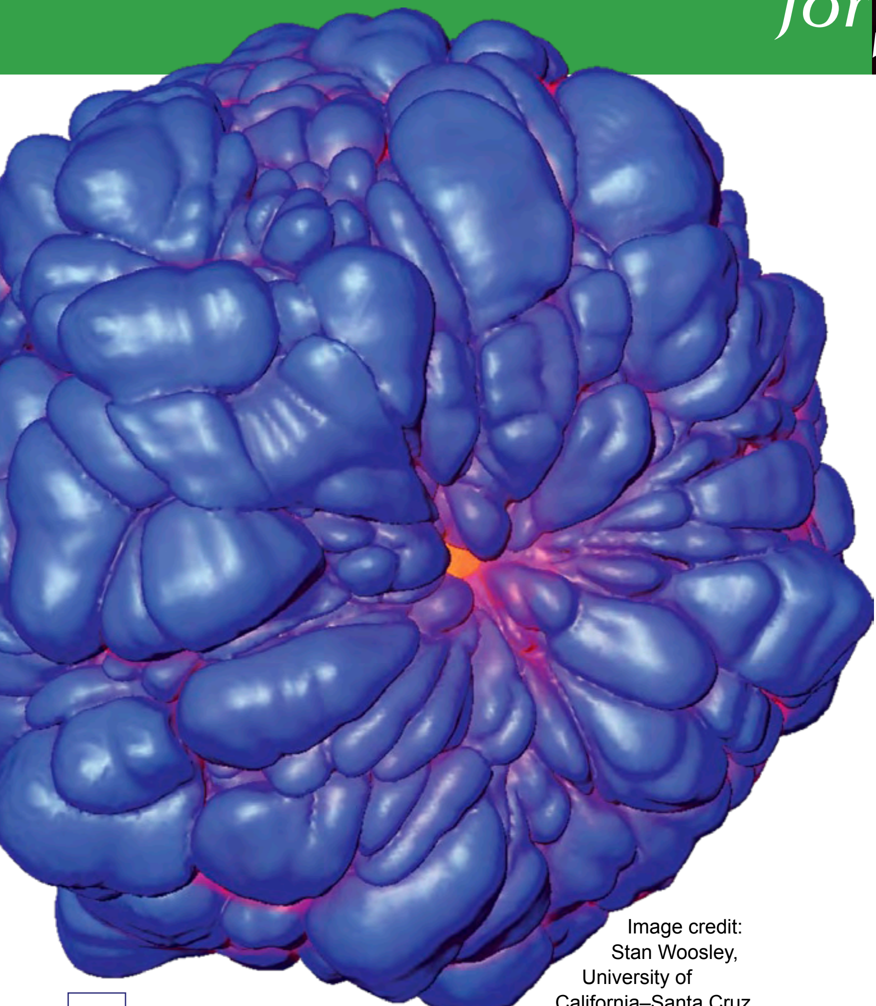

Image credit: Stan Woosley, University of California–Santa Cruz

As the nation's Leadership Computing Facility, we can say the promise of the original LCF proposal has been realized: Multicentury climate simulations—unheard of only a few years ago—are shedding new insight on climate change, three-dimensional simulations are driving design choices for the multinational ITER fusion reactor, and materials scientists have proved that a purely electronic model successfully describes high-temperature superconductors.

Our progress as a supercomputing center can be seen as we've grown within DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. INCITE focuses the nation's most powerful computing systems on the world's most challenging scientific problems. In 2007, the NCCS provided three-quarters of those resources, or 75 million processor hours. In 2008, that number will double to 145 million processor hours.

Terascale supercomputers such as ORNL's Jaguar system—capable of trillions of calculations each second—have

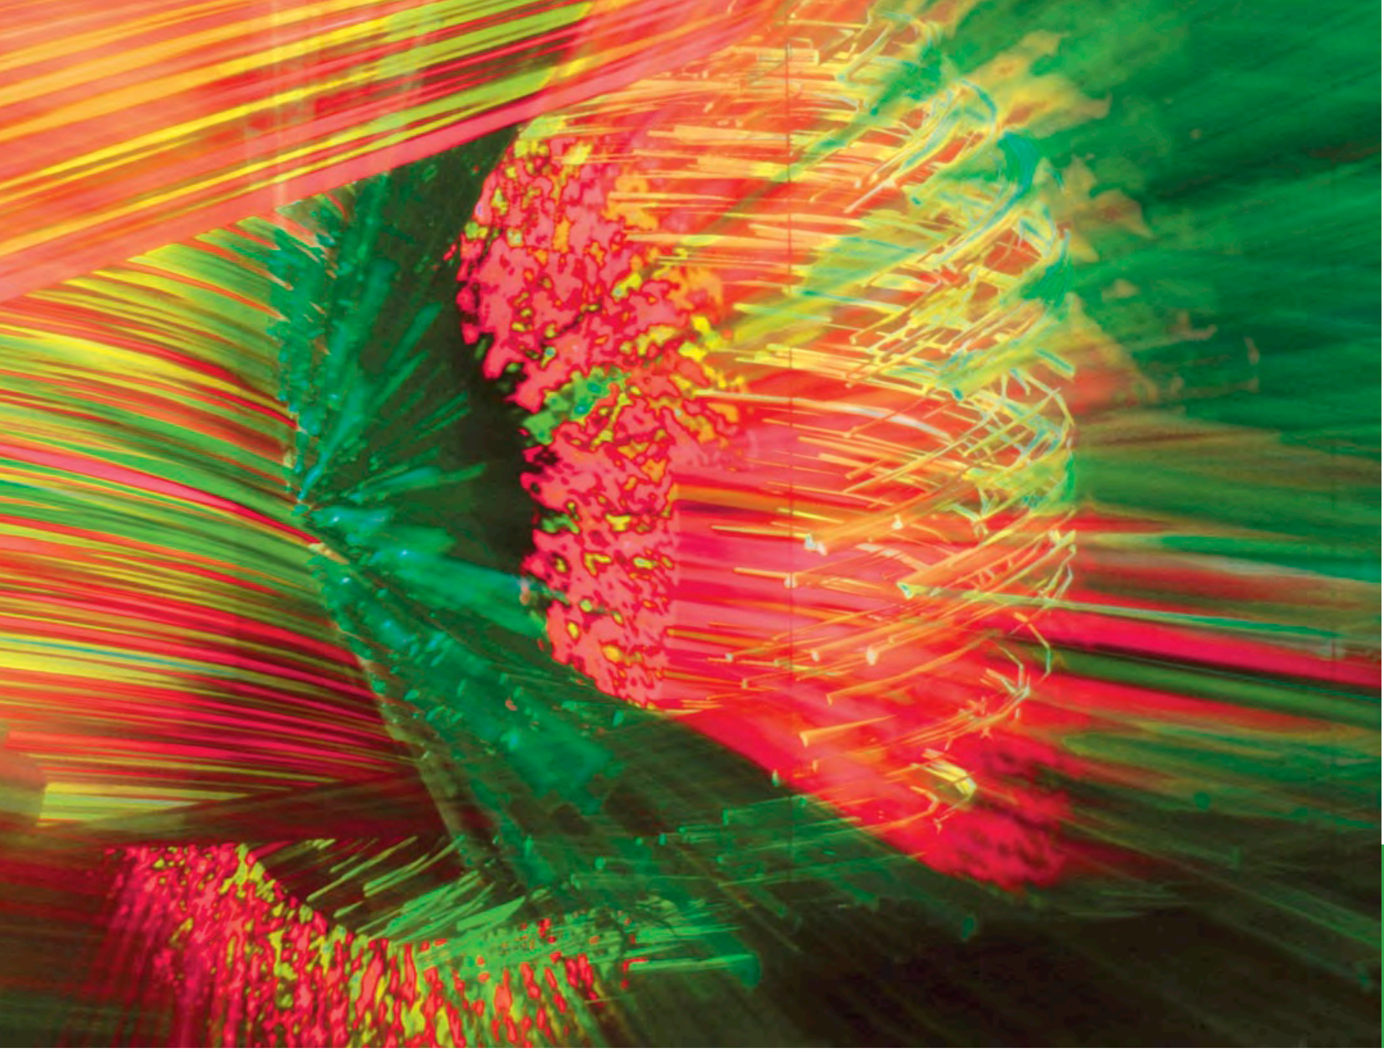

Scientific visualization by Scott Klasky, ORNL/NCCS

changed the world and our view of it. Petascale systems, a thousand times more powerful, will be arriving soon and promise a new revolution—a revolution the NCCS will launch in early 2009. Beyond that, we are making plans for the inevitable arrival of an exascale system within a decade. Capable of a million trillion—or 10^{18}—calculations each second, such a system will be truly transformational.

I cannot give adequate credit for all the exciting science produced at the NCCS in 2007. Nevertheless, I would like to focus on four broad areas, reviewing some of the successes we have enjoyed in materials science, climate studies, energy assurance, and fundamental science. Many of these advances rely on applications either developed by or contributed to by our own staff scientists. I will also touch on some of the advances we can look forward to in the era of petascale computing. For more information on these exciting achievements, please visit our Web site at *www.nccs.gov.*

When scientists understand how materials are put together at the atomic and molecular levels, they can determine how to manipulate those materials to do the seemingly impossible. One group working to develop materials that convert waste heat directly into electricity used NCCS supercomputers to simulate a supercell of more than 1,000 atoms—a fourfold to fivefold increase from previous simulations. Another group was able to resolve the molecular structure of successive layers of methane adsorbing to magnesium oxide—an achievement that promises to advance applications as diverse as fuel storage, manufacturing, and airport security.

As we look forward, petascale computing holds the potential to expand our understanding of high-temperature superconductors, explaining the differences in transition temperatures between superconducting materials. Electronics will benefit from an improved understanding of colossally magneto-resistive oxides and magnetic semiconductors and from new switching mechanisms in magnetic nanoparticles for ultrahigh-density storage. In addition, the world of medicine will see new insights into the physical–chemical factors and mechanisms that control damage to DNA.

The more we know about the Earth's climate, the better we understand how difficult it is to simulate the interrelation-

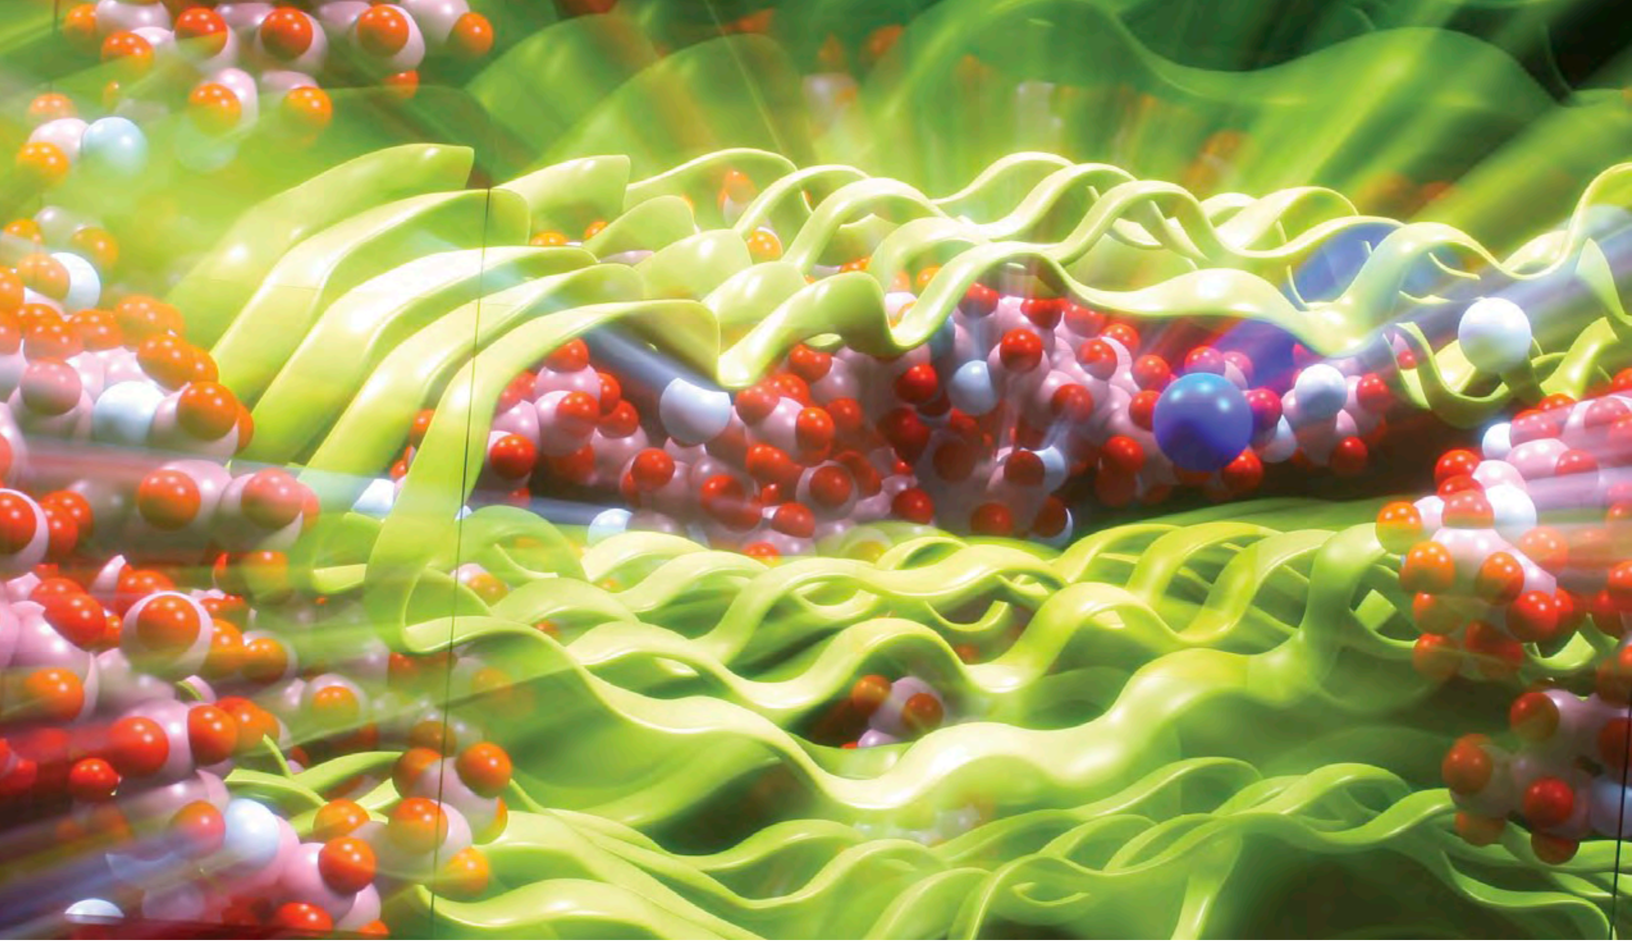

ship between the planet's atmosphere, its plant life, and its oceans. The CCSM is the nation's leading tool for evaluating our planet's climate in the past, present, and future. As home to the Climate Science Computational End Station Development and Grand Challenge Team, the NCCS is providing the computing power necessary to turn the CCSM into an even more predictive suite of scientific tools. Another group has used NCCS supercomputers to perform the most realistic global ocean simulation to date, following virtual tracers for a century to uncover the roles of eddies and water transport in the global ocean.

Petascale computing will greatly refine and focus these efforts, allowing climate scientists to better integrate models for the global ocean, sea ice, land, and atmosphere, as well as the journey of heat, water vapor, and greenhouse gases—such as carbon dioxide ($CO_{2)}$ and methane—through these systems. It will also enhance the ability of researchers to perform three-dimensional simulations of carbon sequestration, a process that will divert CO_2 from power plant emissions and inject it deep underground to dissolve in saline aquifers.

Energy assurance—encompassing production, distribution, and consumption—is the ability to guarantee the energy necessary to maintain and enhance our way of life in an economically viable and environmentally benign manner. Scientific computing increases our energy assurance by helping us make the most of our energy sources, both alternative and conventional. Fusion scientists are working out details of the 100-million-degree (centigrade) ITER reactor, increasing our understanding of such issues as ion and electron turbulence and radio-wave heating. Another project is strengthening our understanding of combustion, focusing, among other things, on flame stabilization in an ignitive environment relevant to fuel-efficient engines and turbines burning diesel, gasoline, and biofuels.

Energy assurance simulations become even more valuable at the petascale. Biologists will be able to simulate systems of several million atoms in their search for an efficient means of converting non-food plants such as switchgrass and poplar trees to ethanol. Chemists will be able to generate quantitative catalytic reaction rates and guide small system calibration. Combustion scientists will develop a simulation tool for designing new engines. Nuclear scientists will boost the efficiency, stability, and safety of fission reactors. And materials scientists and nanoscientists will begin to tackle the challenge of more efficient electrical energy storage systems by devising new designs able to accommodate a hundredfold increase in energy and power density.

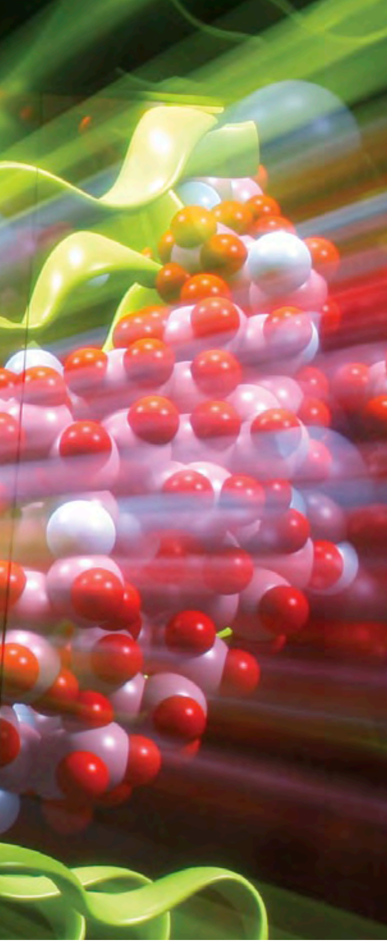

Scientific visualization by Sean Ahern, ORNL/NCCS

Supercomputer simulation at the terascale has greatly enhanced our understanding of how the universe is put together and how it works, at scales both too large and too small to be easily imagined. Astrophysicists closing in on the secrets of the core-collapse supernova have provided the first explanation for a pulsar's spin that matches observation and have provided exciting insights into the mechanism of the supernova explosion itself. Nuclear physicists, on the other hand, have provided the first-ever calculation of the calcium-40 atom using coupled-cluster theory and interaction, including triples correlations, which will enable nuclear structure research with exotic beams and reliable predictions of unknown nuclei crucial for astrophysical modeling.

At the petascale, scientists striving to understand the fundamental workings of the universe will be able to explain how stars of various masses go through core-collapse supernovas. They will provide a first-principles understanding of the mechanism for nuclear reactions, permitting a predictive capability for nuclear properties and scattering cross sections relevant to DOE programs. And they will compute the strong-force interaction between quarks and gluons so precisely that our knowledge will no longer be limited by theoretical uncertainties.

The challenges we face extend well beyond the capabilities of a single research team or computing center. Any effective response will require a focused, coordinated effort from a variety of institutions and a variety of disciplines. Nevertheless, the scientific computing community is well prepared to guide these efforts. Computing has become the point at which theory and experiment come together, capturing our knowledge and showing us a way forward. We have made an impressive start in many of these areas, and we are pressing ahead to provide the scientific community with tools of unprecedented power.

The challenges we face may appear daunting, but we welcome the opportunity to overcome them.—*By Doug Kothe*

A magnetohydrodynamic simulation performed with the GenASiS code explores what effects the instability of a supernova shock wave has on the magnetic fields in stellar cores during core collapse supernova explosions. Image Credit: Eirik Endeve, ORNL, and Ross Toedte, ORNL/NCCS.

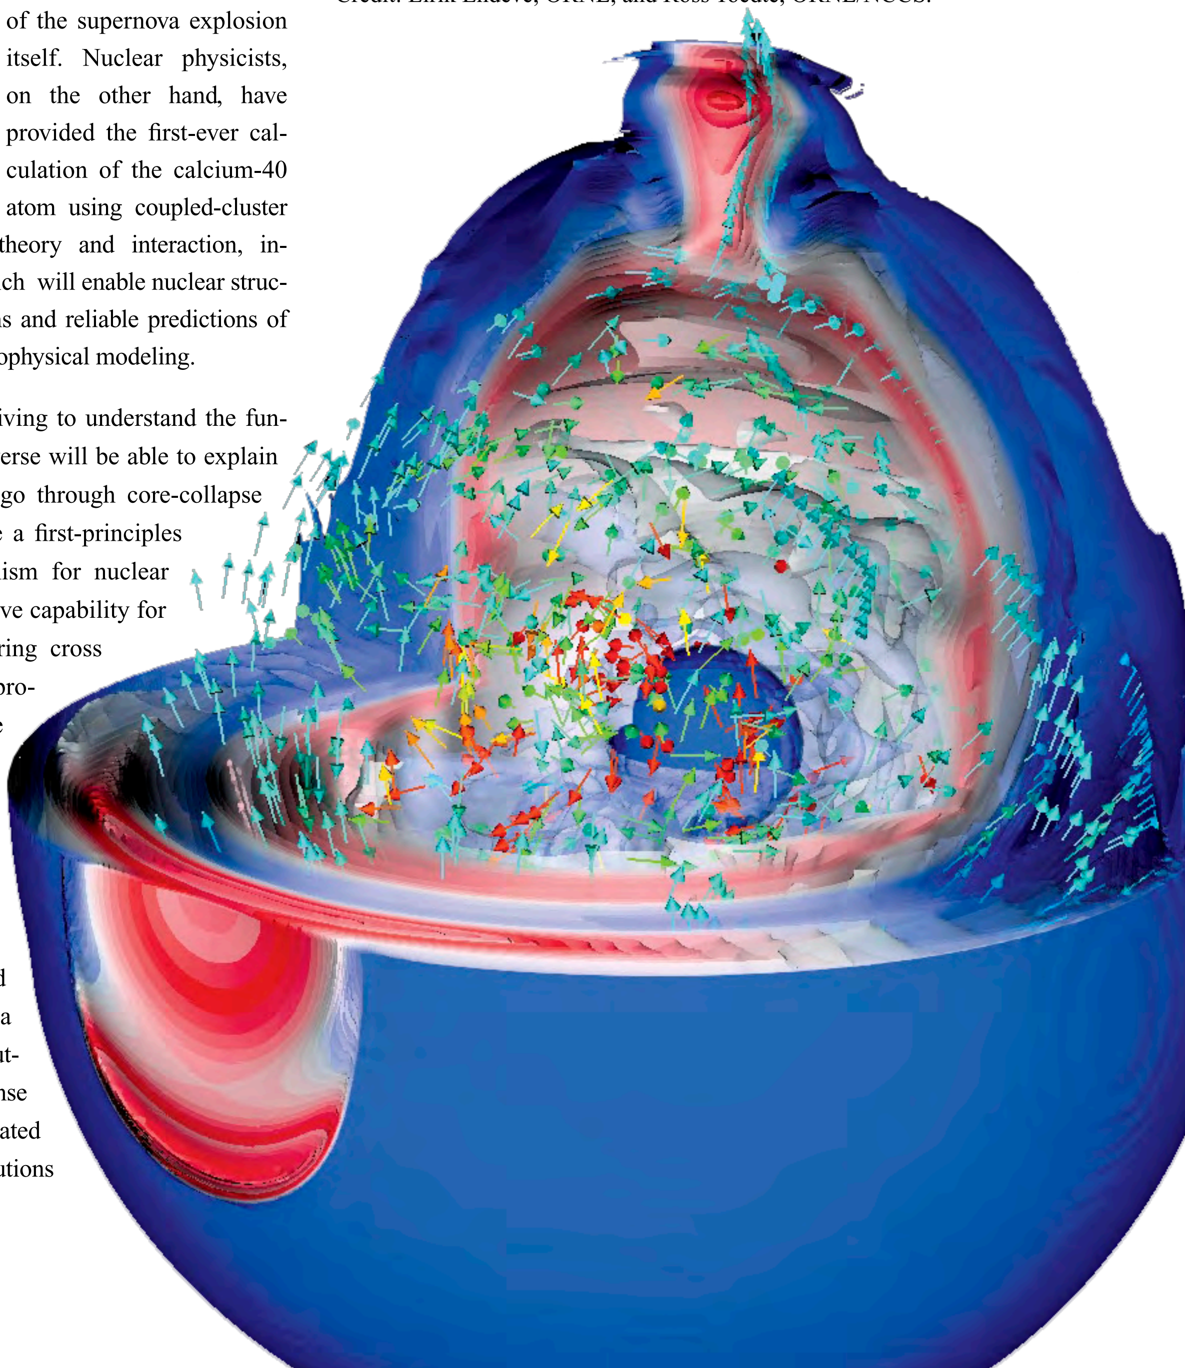

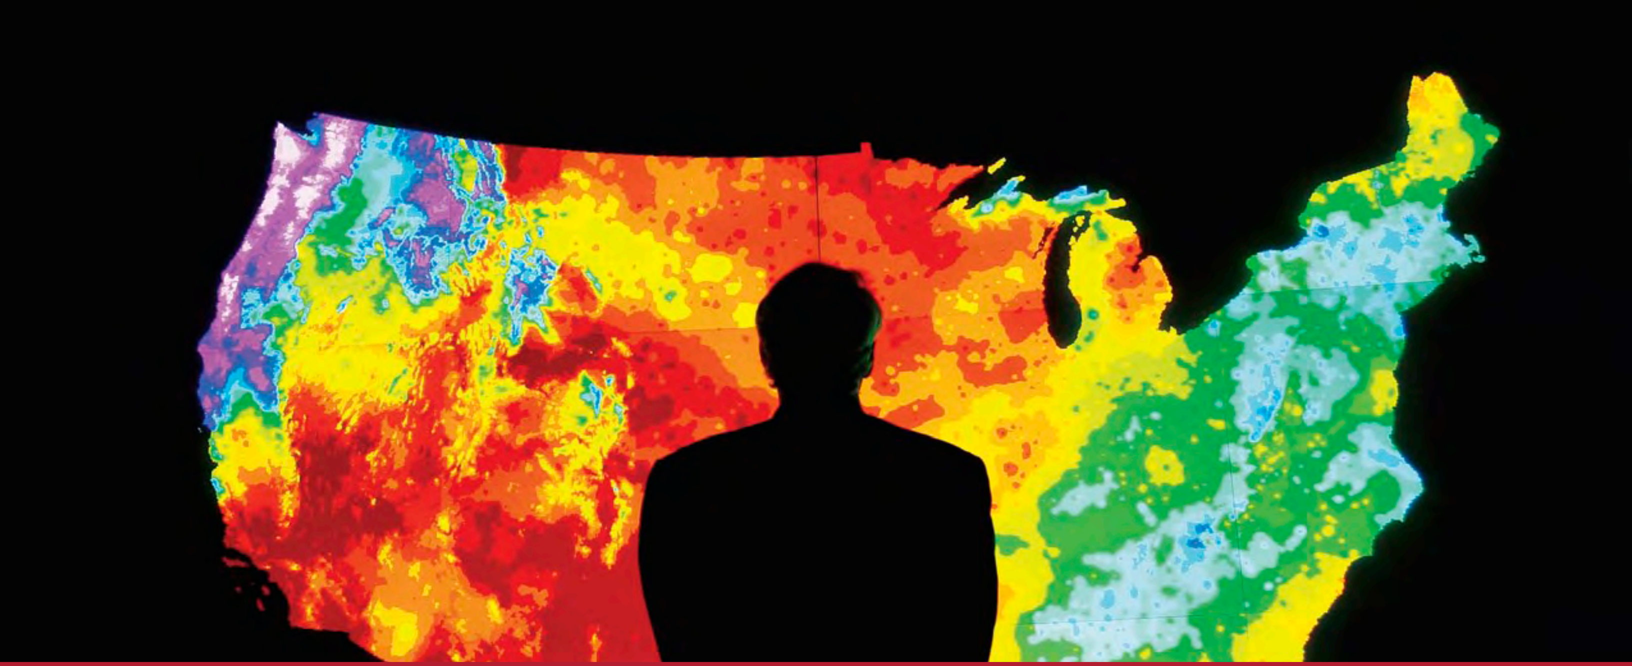

2007 Science Highlights

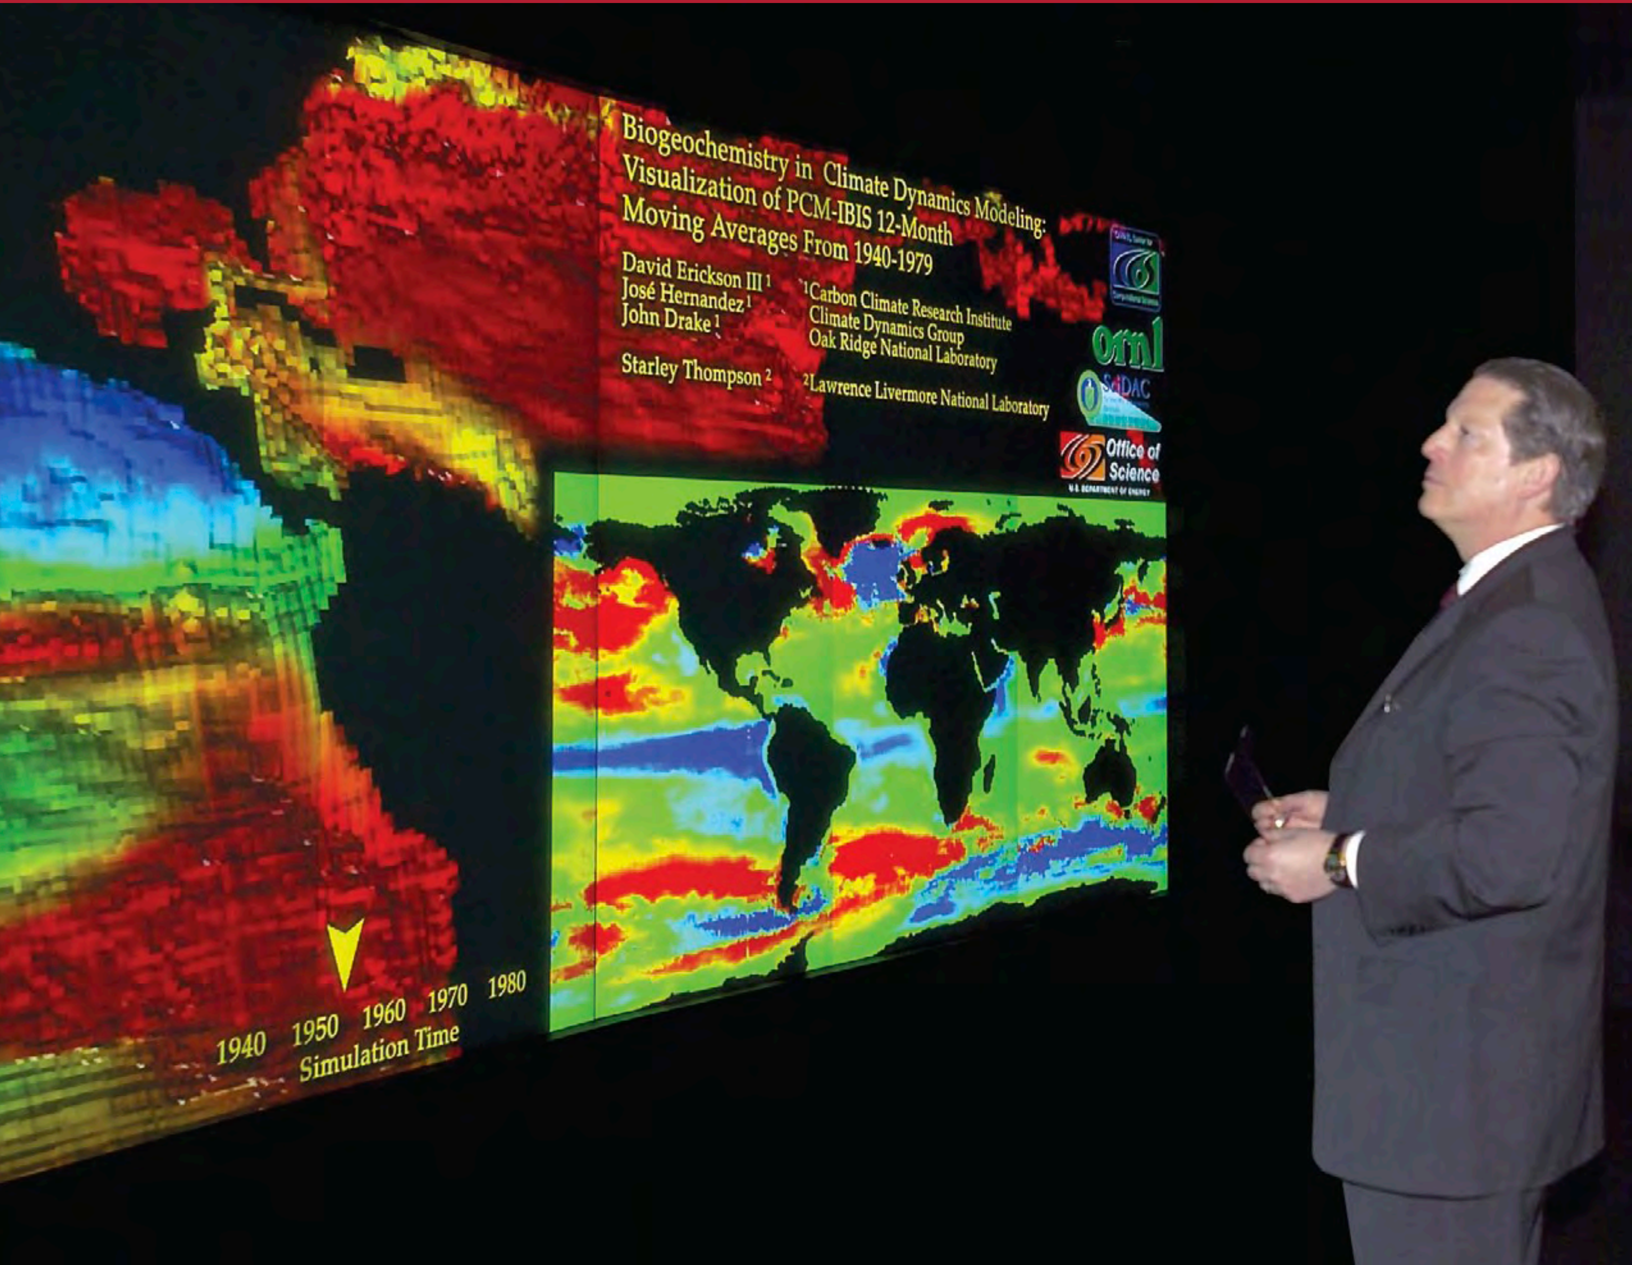

Software engineers turn to Pat Worley's team to fix problems that threaten simulations

Resolution Revolution

Everyone has questions about climate change. Is it a good idea to convert forests and food croplands to produce plants for biofuels? What technologies best capture and store carbon? How intense will hurricanes and heat waves get? Will release of methane trapped in permafrost accelerate climate change?

Finding answers increasingly depends on climate simulations, which set in motion a digitized world that mirrors our past and present and probes our future. That world would not turn without software applications such as the CCSM, a megamodel coupling four independent models whose codes describe Earth's atmosphere, oceans, lands, and sea ice.

Such simulation tools run on supercomputers like those of the NCCS at ORNL. Depending on those models is a global community of scientists representing, in the United States alone, organizations including the DOE, NCAR (funded by the NSF), National Oceanic and Atmospheric Administration, National Aeronautics and Space Administration, Environmental Protection Agency, and various national security agencies.

The Fourth Assessment Report (AR4) of the IPCC utilized simulations done in 2004 and 2005. Using the latest version of the CCSM, researchers carried out computations on resources at the NCCS, NCAR, National Energy Research Scientific Computing Center, and Japanese Earth Simulator. John Drake, a scientist working at the intersections of computer science, climate science, and applied mathematics, called the simulations "a watershed event for climate science and for the way in which we provide computational simulation results to the community." Drake leads ORNL's contribution to the Climate Science Computational End Station, which will support the IPCC's fifth assessment report, due in 2014.

"A fairly small team has been building the models and even a smaller team performs the simulations and posts the material on the Earth System Grid for others to retrieve," he said. "Very few sites in the world can field the kind of computational horsepower that the NCCS does and that various other large climate and weather centers have internationally. The fact that you can perform these simulations and then make the results very quickly available to the university researchers and people who don't have access to the machines or wherewithal to build the models multiplies the productivity of the science enterprise."

The proof is in the papers: In the months after simulation data was posted, scientists produced roughly 300 journal articles. The IPCC cited the studies in AR4, which concluded planetary warming during the 20th century was probably the result of human activities. In 2007 the IPCC shared the Nobel Peace Prize with former Vice President Al Gore for its efforts.

Former Vice President Al Gore viewed the EVEREST PowerWall during a 2007 visit to ORNL.

Weather versus climate

How can we simulate climate 100 years from now if we don't know the weather 100 days from now? "Climate is statistical, or average, weather," explained Drake. "It can tell you if hurricanes will be more likely, less likely, or stronger, but it can't tell you when they will occur."

With weather, small changes early on cause large changes later. "You can't do weather forecasting beyond 10 to 15 days because it's based on chaos theory," Drake said. "In Earth system modeling and climate studies, we're always aware of the effect of chaos—the butterfly flaps its wings, and that changes the path weather is going on, the fundamental dynamics of the atmosphere. Some people would throw up their hands and say you can't do this problem."

But if scientists study climate from a statistical standpoint—in essence sampling multiple flaps of the butterfly's wing—they can approach the problem. "An ensemble of paths is then averaged to get the most likely path," Drake said.

Just how well do statistics model reality? "To the extent that we can reproduce the paleoclimate record or recent historical record, if emission scenarios and forcings are accurate, then we believe the models are reasonably accurate," said NCCS Director James J. Hack, a renowned climate researcher who implements global models on high-performance computing systems. "But we know that they're not reliable on smaller -than-subcontinental space scales. In fact, on space scales similar to the North American continent, there's divergence in the models about what happens to precipitation over North America 100 years from now."

To answer questions about climate change at local levels, such as what's going to happen in the Tennessee Valley in a decade, scientists need higher-resolution models. "We want to employ numerical algorithms that can scale to use many, many, many more processors and keep the time to solution about the same," Hack said.

"To evaluate local or regional impacts of climate change, the computational requirements for climate modeling go way up," said ORNL computational scientist Patrick Worley. "The models are currently not able to efficiently exploit the computing resources that will be available in the near future. To take advantage of those, we need to modify some of the models fairly dramatically."

This image was produced from a Carbon-Land Model Intercomparison Project (C-LAMP) simulation performed as part of a SciDAC2 project on NCCS supercomputers. Scientific visualization by Jamison Daniel, ORNL/NCCS

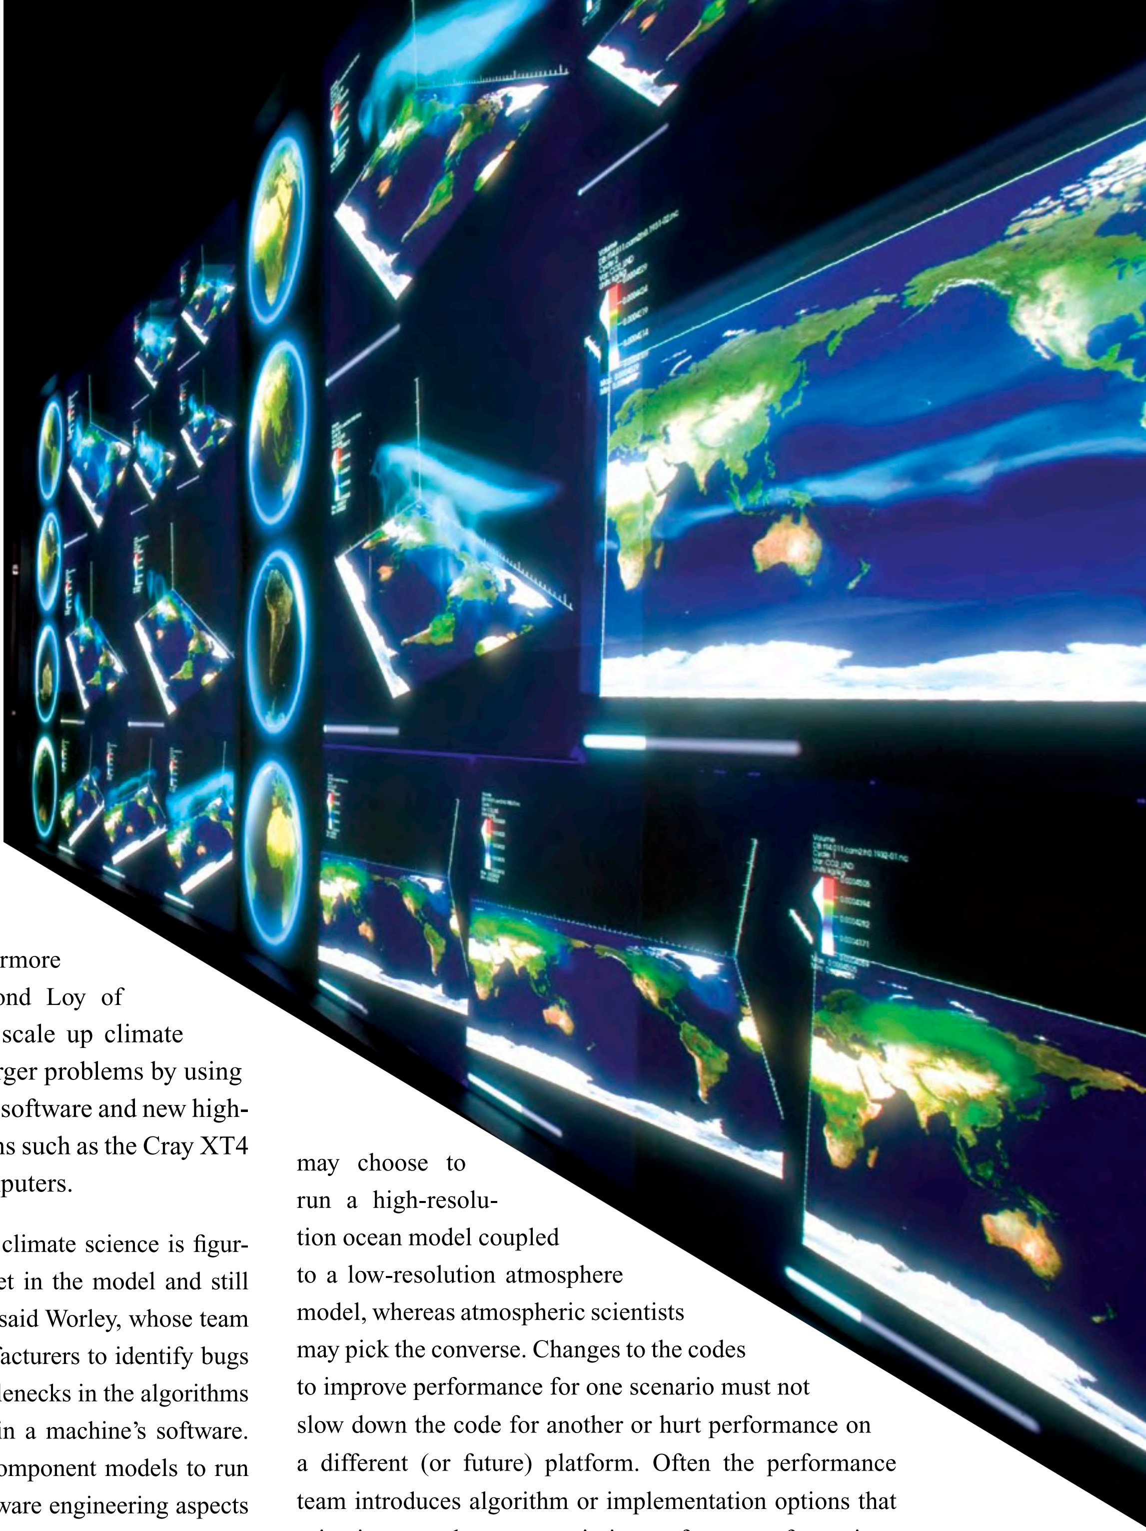

Simulations bring scientific insight. A 30-foot-wide display wall facilitates data analysis and exploration.

Performance police

When scientists want more accurate or more detailed simulations, they turn to modeling experts and software engineers who upgrade the capabilities of the simulation models. When the software engineers need help, they turn to Pat Worley. He leads a DOE Scientific Discovery through Advanced Computing (SciDAC) project with Arthur Mirin of Lawrence Livermore National Laboratory and Raymond Loy of Argonne National Laboratory to scale up climate codes, enabling them to solve larger problems by using more processors, and to evaluate software and new high-performance computing platforms such as the Cray XT4 and IBM Blue Gene/P supercomputers.

"An important practical aspect of climate science is figuring how much science you can get in the model and still get the simulations done in time," said Worley, whose team works with researchers and manufacturers to identify bugs in CCSM codes, performance bottlenecks in the algorithms used in the CCSM, and glitches in a machine's software. "Our contribution is getting the component models to run as efficiently as possible. The software engineering aspects of the code are always changing, and often the new code has unexpected performance issues. We monitor things. We're kind of the performance police."

Worley and his colleagues push codes to their limits. If a code runs slowly on 1,000 processors but quickly on 2,000, they might assign more processors to work on a problem. If, owing to algorithmic restrictions, the code can't use more than 1,000 processors, changing algorithms may be the only option to improve performance. Different science also imposes different performance requirements. Ocean scientists may choose to run a high-resolution ocean model coupled to a low-resolution atmosphere model, whereas atmospheric scientists may pick the converse. Changes to the codes to improve performance for one scenario must not slow down the code for another or hurt performance on a different (or future) platform. Often the performance team introduces algorithm or implementation options that scientists can choose to optimize performance for a given simulation run or on a particular computer system.

On Cray and IBM systems, the group has improved performance through both algorithmic and implementation efforts. Recent work improved performance 2.5-fold on benchmark problems on ORNL's Cray XT4. "With the improvements to the scalability of the CCSM software by Pat and his colleagues, along with the dramatic growth in the performance of Jaguar, the CCSM developers are seriously considering model resolutions and advanced

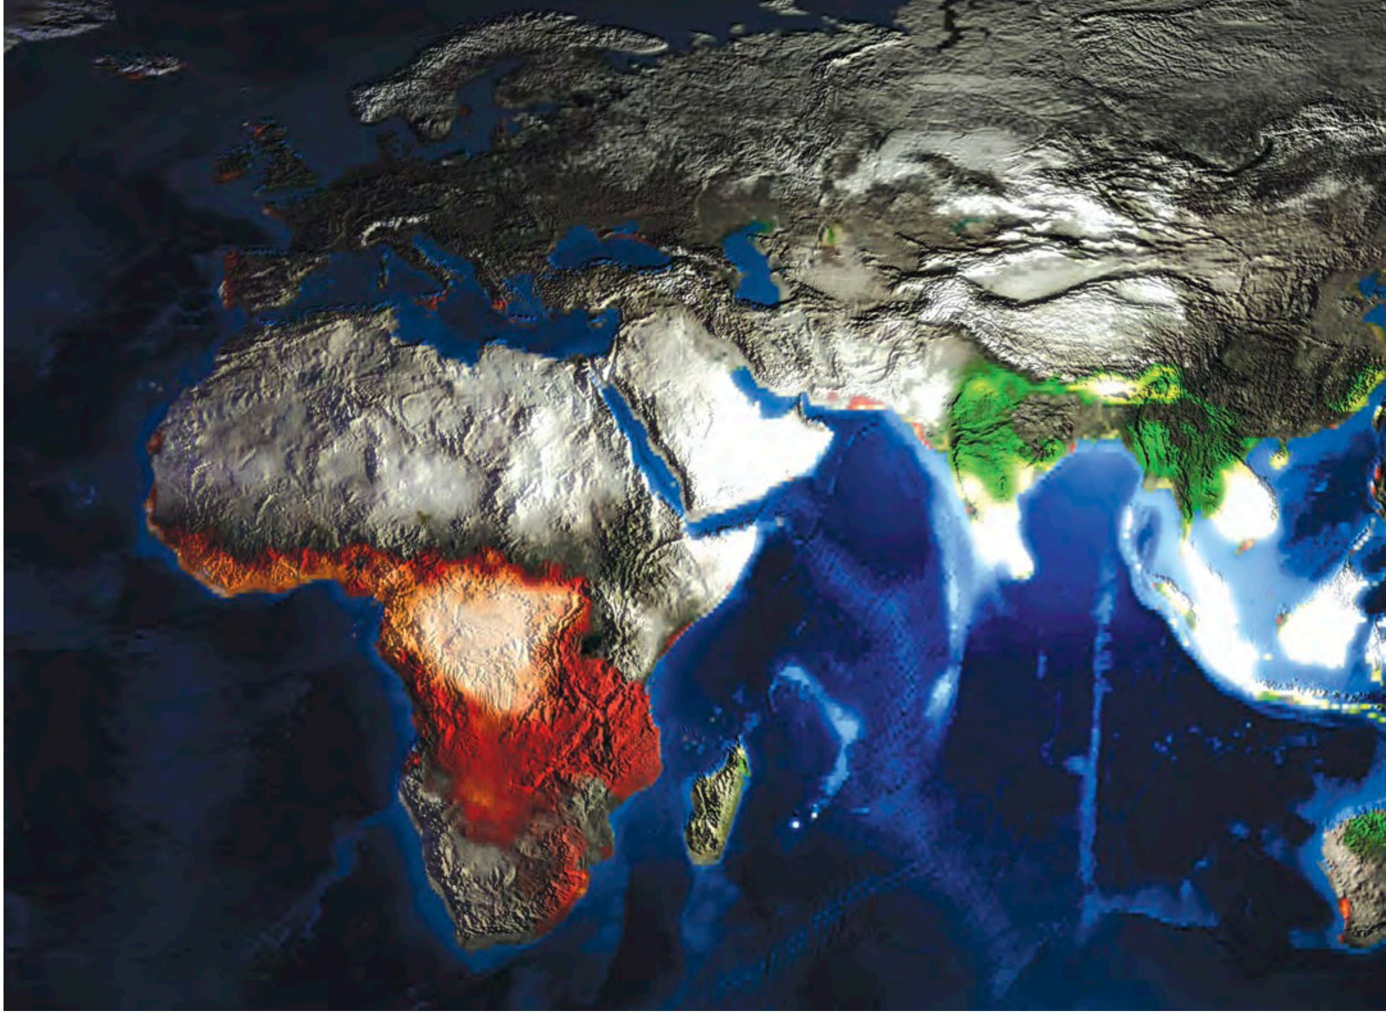

As the sun rises over Eastern Europe, the instantaneous net ecosystem exchange (NEE) of CO_2 is shown in the Eastern Hemisphere. Strong uptake is shown in green-to-white colors and is strongest in the tropics. A net release of CO_2 to the atmosphere is shown in red-to-white colors and is strongest over the Congo, where the sun is not shining. This image was produced from a C-LAMP simulation performed as part of a SciDAC2 project using NCCS supercomputers. Scientific visualization by Jamison Daniel.

physical processes that were not on the table before," said Trey White, who as NCCS liaison to the CCSM project helps the scientists make the most of the machines.

"Pat Worley's group has provided critical support in improving the scalability and performance of the CCSM across a wide range of architectures," said NCAR's Mariana Vertenstein, head of the engineering group responsible for CCSM's software development, support, and periodic community releases. "The CCSM project played a major role in the IPCC AR4 through an extensive series of modeling experiments and in fact resulted in the most extensive ensemble of any of the international global coupled models run for the IPCC AR4. This accomplishment could not have occurred without Pat's contributions."

Worley's team is currently working with a large multilab SciDAC project led by Drake with Phil Jones of Los Alamos National Laboratory to build a first-generation Earth System Model, which will extend the physical climate model by including chemical and ecological processes. The computer allocations are provided through the Climate Science Computational End Station, an INCITE program award led by NCAR's Warren Washington on Jaguar in the NCCS.

"For DOE, which is very concerned with the carbon cycle and with the impact of climate change on ecology and ecosystem services, this kind of Earth system model is really called for," Drake said. "We're trying to do whatever we can to get there as quickly as possible."—*By Dawn Levy*

Related Publications

Drake, J.B., P.W. Jones, M. Vertenstein, J.B. White III, and P.H. Worley. 2007. "Software Design for Petascale Climate Science." *Petascale Computing: Algorithms and Applications,* ed. D. Bader, Chapman & Hall/CRC Press, Taylor and Francis Group.

Mirin, A.A., and P.H. Worley. 2007. "Extending Scalability of the Community Atmosphere Model." *Journal of Physics: Conference Series,* **78**, 012082.

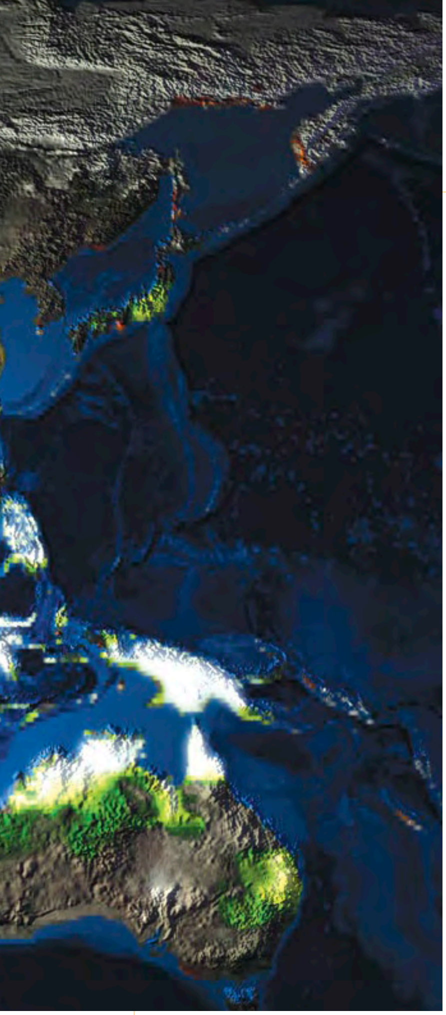

Better Answers to Climate Questions

In 2007, renowned climate modeler James J. Hack arrived at Oak Ridge to lead the national lab's new Climate Change Initiative and direct the NCCS. He came from the NCAR, an NSF–sponsored center in Colorado, where he led development of the model much of America's climate science community uses to simulate Earth's atmosphere. The Climate Change Initiative that he heads at ORNL aims to accelerate discoveries about Earth's climate system through lab-wide engagement of scientists and engineers from diverse directorates encompassing energy, environment, computing, and national security.

"The initiative is an attempt to leverage all that work and fill in the gaps so that we have a more comprehensive range of capabilities for answering questions about consequences of climate change," said Hack, who arrived at DOE's largest lab at the time when the atmospheric level of carbon dioxide was found to exceed the worst of four scenarios explored by the IPCC. His selection reflects the indispensability of high-performance computing, models, and simulations to assessing global climate change.

Capable of trillions of calculations per second, the supercomputers in Hack's purview have distinguished themselves in climate science. The NCCS's IBM pSeries Cheetah system provided more than one-third of the simulation data for the joint DOE/NSF contribution to the IPCC's fourth assessment report, which concluded planetary warming during the 20th century was likely due to human activities. In 2007, IPCC scientists shared the Nobel Peace Prize with former U.S. Vice President Al Gore for amassing and disseminating knowledge about human contributions to climate change.

"Now is the time to begin focusing on what the consequences are going to be to various stakeholders responsible for decisions and planning about energy, agriculture, water resources, land use, and other critical infrastructure needs that have long planning timescales," Hack said. "If we're able to demonstrate that the models have good predictability on regional space scales and we can get at things like how climate extremes change and where the thresholds are in the climate system—where you reach a certain tipping point that something significant changes with regard to the consequences—then the community at large can take that into account when they're planning for their growth."

Agencies conducting climate research, either independently or in partnership with DOE, include NSF, the National Aeronautics and Space Administration, the National Oceanic and Atmospheric Administration, and the Environmental Protection Agency.

"DOE needs to build its climate capabilities in partnership with these agencies," Hack said. "The piece that we're proposing to do has a very mission-oriented aspect—to stand up a capability to answer questions about climate change consequences. A big part of that is utilizing DOE's unique computational capabilities, including NCCS's terascale Jaguar system, which is already making important contributions to climate research."

The NCCS, DOE's Leadership Computing Facility, is well positioned to help researchers gain insights that may guide policymakers and planners in exploring options for addressing some of the greatest challenges of our time.— *By Dawn Levy*

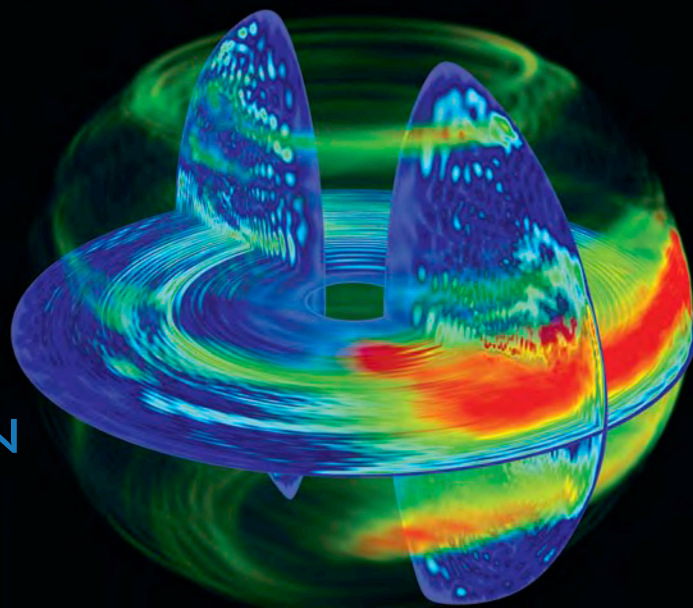

Researchers seek lingua franca so fusion codes can converse in coupled model

We Need to TALK

Projects explore radio frequency waves and magnetohydrodynamics in fusion plasmas

In the Bible, all people spoke the same language until God felled the Tower of Babel and produced diverse languages that prevented people from understanding each other. Today physicist Don Batchelor at ORNL is trying to solve a Babel-esque problem by creating a common language and a computational framework that will allow diverse software codes to communicate with each other in simulations of plasma—the hot, ionized gas that fuels nuclear fusion reactors.

Right now only some codes can share data. Different data formats and data names confound communication between many codes. Also challenging is getting codes to couple, or provide information to other codes at specific times to solve equations about the evolving state of plasma. Coupling is a tough task because a factor described by one code depends on another factor described by some other code.

"We've developed basically a *lingua franca* whereby different physics codes can talk to each other," Batchelor says. In 2007 he and a team of more than two dozen researchers at ten institutions used resources at the NCCS to make progress toward developing an integrated plasma simulator. The work was made possible with an allocation of 1.7 million processor hours on the center's Cray Jaguar XT supercomputer through DOE's INCITE program.

Getting codes to speak a common language is difficult but not impossible. Researchers in another field, climate science, have done it. They linked models describing different aspects of Earth's climate system and got the codes to communicate with each other in the CCSM coupled simulation tool.

"The climate community developed sophisticated models of the atmosphere, ocean, land, biomass, and sea ice—but you can't predict what the climate's going to do until you know how all of those things interact," Batchelor says. "We're in a similar situation in fusion. Over the years we have developed highly sophisticated models in different areas of plasma physics that are necessary to understand fusion experiments."

We need to talk

Controlling plasma is the key to getting cheap, clean energy from future commercial reactors. Construction is under way for ITER, an experimental fusion reactor expected to begin operations in Europe in 2016. The $13 billion, 30-year multinational megaproject will build and operate a full-scale experimental device to demonstrate the technical feasibility of fusion energy. Following ITER would be a full-scale commercial fusion plant by 2050 and large-scale fusion-power adoption over the ensuing 30 years.

Improved simulation capability is urgently needed to support ITER, Batchelor says. Simulations are critical to

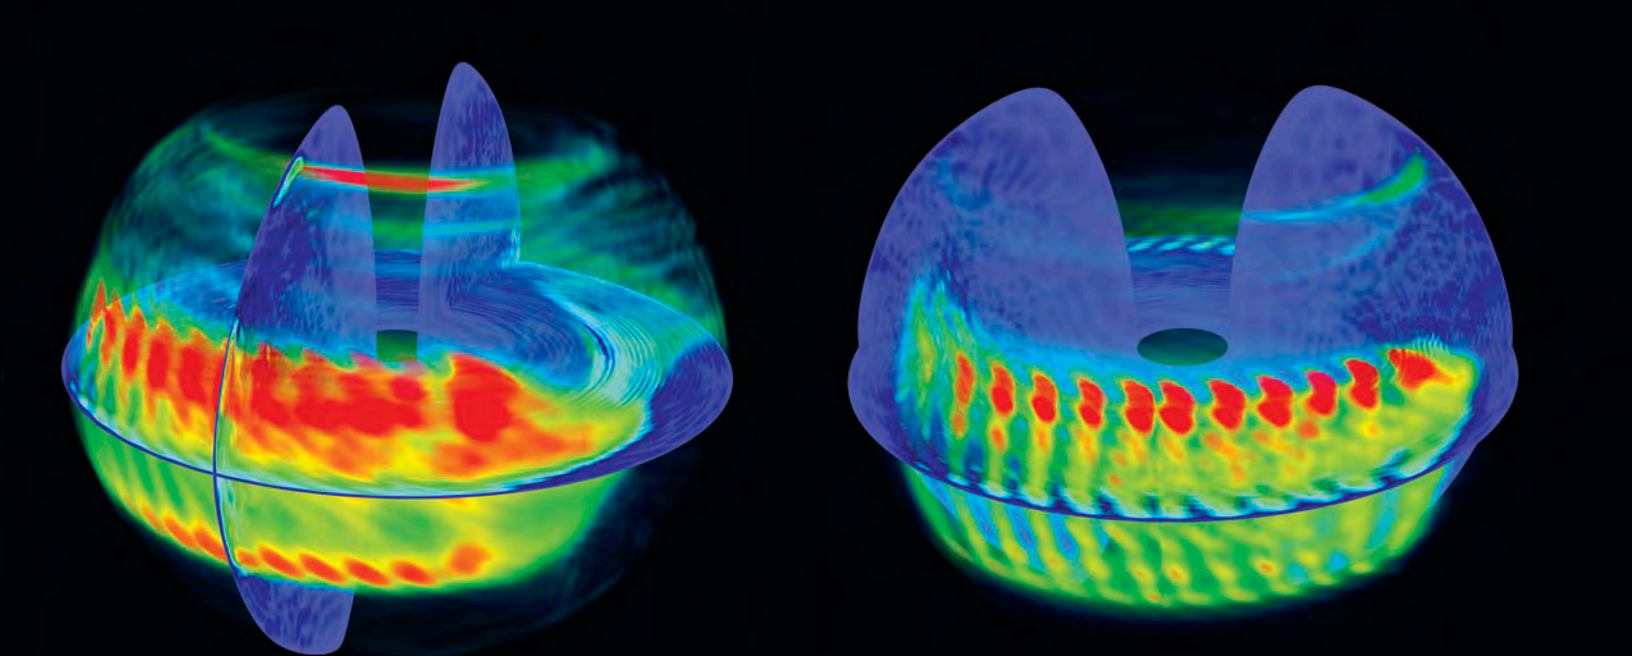

Fusion codes must communicate to control plasma in reactors. Image credit: Don Batchelor, Fred Jaeger, Sean Ahern

supporting both theory and experiment. Theorists use simulations to gain insight into complex equations that correctly describe the state of plasma but that are difficult to intuit without solving the equations and visualizing the data. Experimentalists use simulations to evaluate if the desired trial conditions are achievable with the available equipment, guide operation of equipment, and interpret measurements from experiments.

"Once you have a simulation which is validated against the experiment—that you have confidence in—then you can use it to predict the next device," Batchelor says.

An Esperanto of fusion software codes would allow Batchelor and collaborators to couple diverse phenomena taking place at multiple scales. The phenomena that they study range in speed from the fast oscillations of the radio waves that heat plasma (one cycle every 100-millionth of a second) to the slow interval it takes plasma to heat (about a second). They range in size from the gyroradius of an electron spiraling around a magnetic field line (about one-tenth of a millimeter) to the height of a fusion reactor (about 6 meters).

Says Batchelor, "When we come to take the next step in going from ITER to some sort of a demo reactor, there are going to be a lot of design decisions that have to be made, and simulations that are validated against experiments will be the way that we have to make these design decisions. That will ride on supercomputing."

Synchronized SWIMming

Scientists have developed theories and computer codes that deal with multiscale phenomena more or less in isolation. "Now the time has come to consider the interaction between these phenomena—which have been studied essentially as separate disciplines—because they feed back on each other," Batchelor says.

Two of the most important fusion codes simulate radio frequency (RF) waves, which heat and control plasma, and magnetohydrodynamics (MHD), the behavior of "fluid" that has a magnetic field and carries current. Researchers developed the RF and MHD codes with support from the SciDAC program, funded by the Office of Advanced Scientific Computing Research (ASCR). Batchelor's INCITE allocation supported three SciDAC projects. The first, the Center for Simulation of Wave Particle Interactions, developed RF codes including AORSA and TORIC, whereas the second, the Center for Extended MHD Modeling, developed MHD codes including M3D and NIMROD. The third, the Center for Simulation of Wave Interactions with Magnetohydrodynamics (SWIM), of which Batchelor is principal investigator, brought the RF and MHD codes from the two supporting projects into a framework to try to get the codes to talk to each other.

The goals of SWIM are to use coupled codes to improve our understanding of RF and MHD interactions in plasma, develop an integrated computational system for treating many physics phenomena, and serve as a prototype for the Fusion Simulation Project (FSP). The FSP aims to simulate the behavior of toroidal magnetic fusion devices called tokamaks on all important time and space scales and to account for the interactions of all relevant processes. It will contribute to ITER's experimental planning and performance optimization and design of devices beyond ITER.

A major element of SWIM is the development of a tool called the Integrated Plasma Simulator (IPS), a computational framework that will link physical factors beyond RF and MHD to make available an expanded, coupled model.

"The Integrated Plasma Simulator has the ability, we claim, to couple any physics code," Batchelor says. That's a lot of coordination—fusion simulations can take tens of thousands of processors working in parallel to compute heat sources and diffusion coefficients.

SWIM employs almost as many computer scientists and mathematicians as physicists and is funded equally through the Office of Fusion Energy Sciences (OFES) and ASCR. The United States is in a strong position to field a fusion simulation project, Batchelor says. "We have two offices at DOE who are interested in working with each other—OFES and ASCR—and we have experience working with each other through the SciDACs."

After ITER

Why study RF and MHD interactions? "The MHD stabilities put limits on the performance of the tokamak," Batchelor says. "They can reduce the confinement time, can cause the plasma essentially to blow away, even cause damage to the device structure." RF systems can control these instabilities.

Batchelor received an allocation of 300,000 hours on Jaguar in 2008 to make innovations to increase the codes' flexibility and performance. To finish the INCITE simulations, Jaguar will simulate RF waves energizing ions inside Alcator C-Mod, an experimental tokamak at the Massachusetts Institute of Technology that is powered by RF only. One of Batchelor's goals is to validate this aspect of simulations run using the IPS computational framework against experiments run on Alcator C-Mod.

If today's simulations support experiments, tomorrow's experiments may support simulations. "Looking ahead, you could consider that the purpose of the ITER device is to provide data to validate the FSP," Batchelor says. "If the purpose of developing all this simulation is to take the next

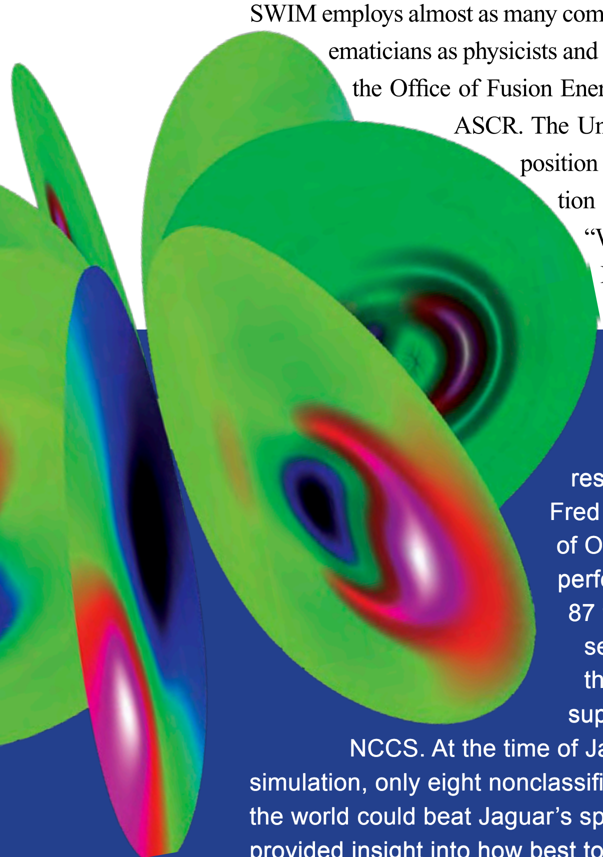

Tuner Melt

In 2007 a fusion research team led by Fred Jaeger and Lee Berry of ORNL achieved a performance of more than 87 trillion calculations per second, or teraflops, on the Cray XT4 Jaguar supercomputer at the NCCS. At the time of Jaeger and Berry's simulation, only eight nonclassified computers in the world could beat Jaguar's speed. The simulation provided insight into how best to heat an experimental reactor that is scheduled to begin operating in 2016. Formerly known as the International Tokamak Experimental Reactor, ITER is a multinational project that will be an important step toward developing commercial fusion power plants.

ITER will use antennas to launch radio waves carrying 20 megawatts of power into the reactor. That's the power equivalent of a million compact fluorescent light bulbs. These waves will heat the deuterium and tritium fuel to fusion temperatures—about ten times hotter than the surface of the sun. The deuterium and tritium form plasma, a state of matter created when gases become so hot that electrons get energized and fly off their atoms. In addition, the radio waves can drive currents that help confine the plasma. Jaeger's simulations will contribute to understanding how to make the most of the wave power in both heating and controlling the plasma.

"This run is the first two-dimensional simulation of mode conversion in ITER," Jaeger said of the 2007 simulation, which explored the conversion of fast electromagnetic waves to slow electrostatic waves. Before this run, mode conversion in ITER was simulated in only one dimension, although scientists could simulate mode conversion in two dimensions for smaller step—have a validated simulation that enables you to build a demo reactor or a commercial reactor—then all these experiments are to build a scientific basis for that, and a major factor for taking the next step beyond ITER is going to be simulation. On the one hand FSP will support ITER by helping it do its experiments. But the real value of ITER is going to be to validate the simulations so we can extrapolate with confidence to the next step."—*By Dawn Levy*

Related Publications

Jaeger, E.F., L.A. Berry, S.D. Ahern, R.F. Barrett, D.B. Batchelor, M.D. Carter, E.F. D'Azevedo, R.D. Moore, R.W. Harvey, J.R. Myra, D.A. D'Ippolito, R.J. Dumont, C.K. Phillips, H. Okuda, D.N. Smithe, P.T. Bonoli, J.C. Wright, and M. Choi. 2006. "Self-consistent full-wave and Fokker-Planck calculations for ion cyclotron heating in non-Maxwellian plasmas." *Physics of Plasmas* **13** (5).

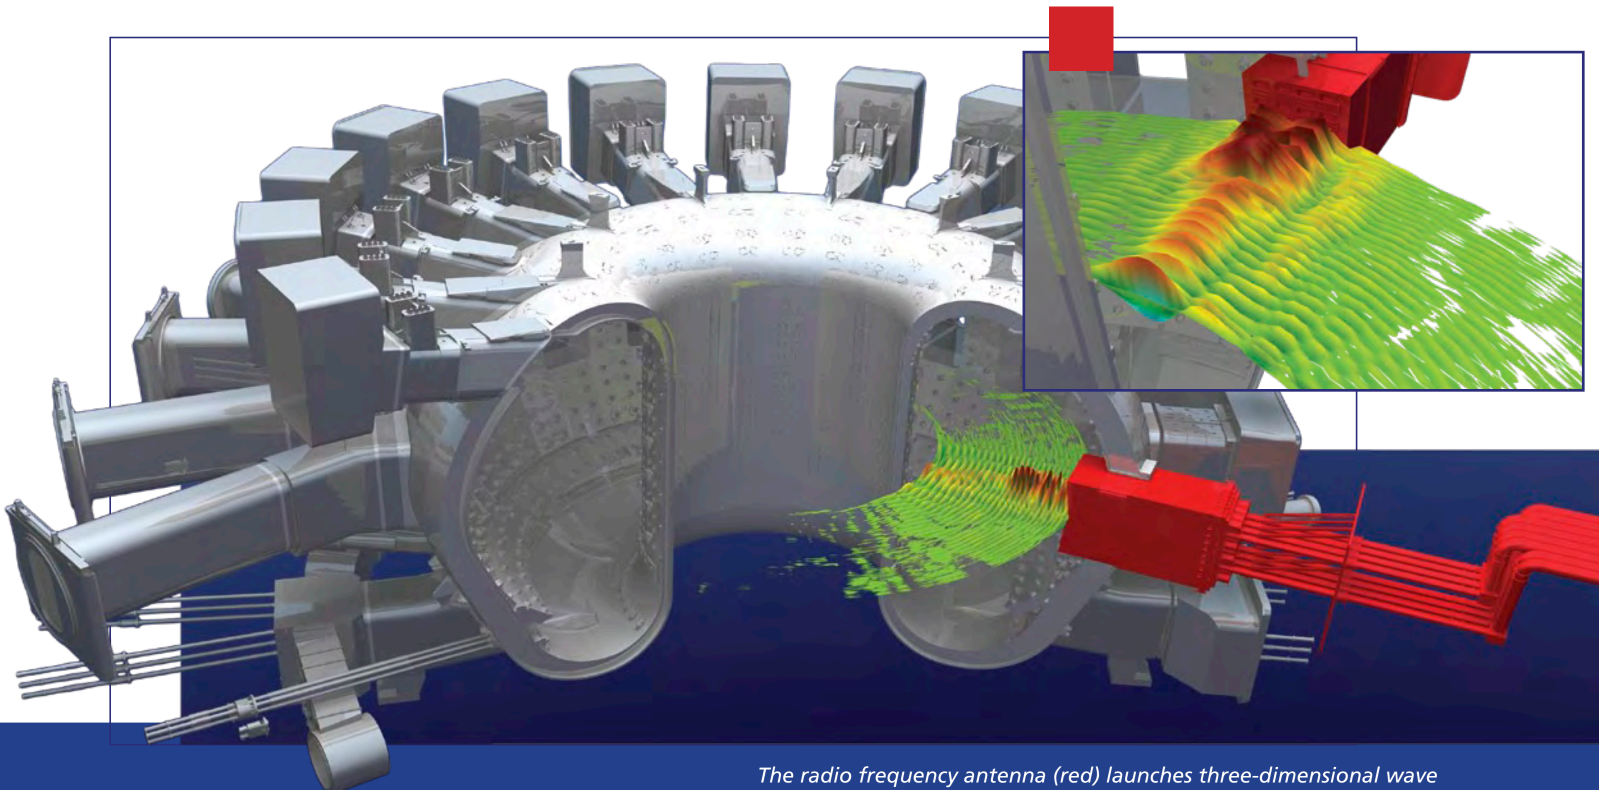

The radio frequency antenna (red) launches three-dimensional wave fields into the ITER plasma. The waves heat deuterium and tritium fuel to fusion temperatures about ten times hotter than the surface of the Sun. Image credit: Fred Jaeger and Sean Ahern/U.S. ITER Project Office.

tokamaks. "We need to know which types of waves are present because different waves can interact differently with the plasma."

Jaeger's team uses AORSA, a software code that solves Maxwell's equations for the electromagnetic wave fields in the plasma. For the 2007 simulation, the code employed 22,500 processor cores—98 percent of the machine's capacity—to calculate the interplay between radio waves and particles in the plasma as well as the current produced by the interaction. A mesh of 500 by 500 cells, or 250,000 individual cells—more than triple the resolution of earlier simulations—gave the team the ability to examine interactions in fine-grained detail.

Upon analyzing the energy distributions of the very-high-energy ions created when radio waves heat the plasma, the scientists found that, in some cases, these ions increased the fusion reaction rate. They learned the optimal frequency for driving current in the ITER plasma and identified the heat-loss channels that limit this current. In addition, comparing ITER with current tokamaks, they found stronger central focusing of radio waves in ITER, which bodes well for its ability to keep plasma hot enough for fusion.

The AORSA team is part of a SciDAC project known as the SciDAC Center for Simulation of Wave-Plasma Interactions. The team includes plasma scientists, computer scientists, and applied mathematicians from ORNL; the Massachusetts Institute of Technology; Princeton Plasma Physics Laboratory; General Atomics; CompX, Inc.; Tech-X Corporation; and Lodestar Research Corporation.
—By Dawn Levy and Leo Williams

TAP IT and TRAP IT

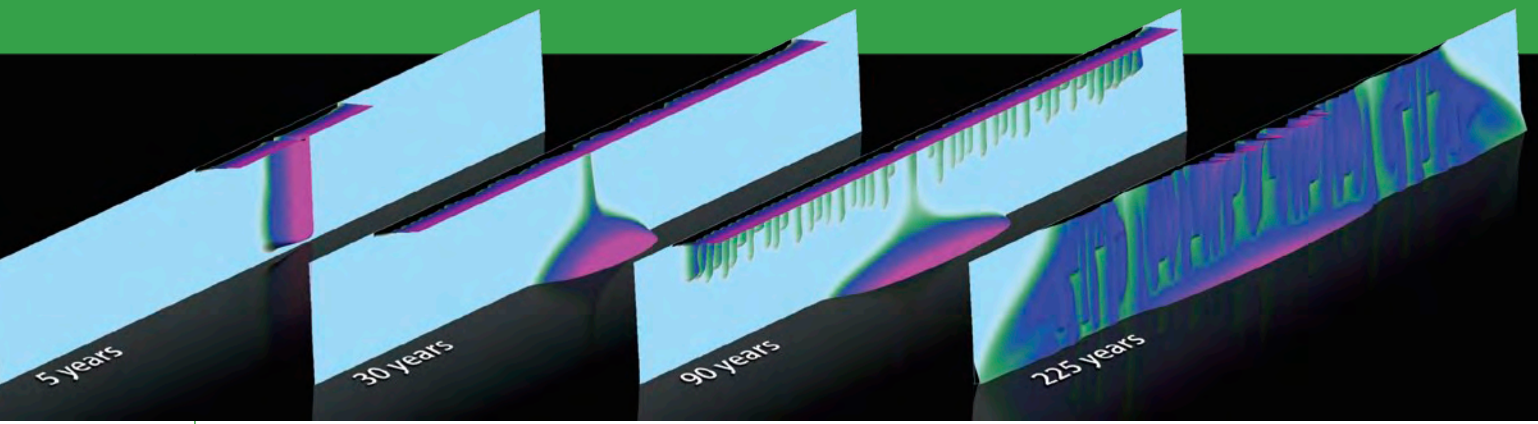

Carbon dioxide dissolving in a deep saline aquifer. Image credit: Peter Lichtner, Los Alamos National Laboratory, and Sean Ahern, ORNL/NCCS.

Not every response to global warming focuses on new energy sources. Even as we develop promising technologies such as solar power, biofuels, and nuclear energy, we face the prospect of being tethered for some time to the old energy sources—primarily fossil fuels such as coal and oil.

One proposal for mitigating the effect of coal power on the earth's climate involves separating CO_2 from power plant emissions and pumping it deep underground, where it can remain indefinitely dissolved in the groundwater or converted into a solid form of carbonate minerals.

A team of researchers led by Peter Lichtner of Los Alamos National Laboratory (LANL) is using ORNL's Jaguar supercomputer to simulate this process, known as carbon sequestration, searching for ways to maximize the benefits and avoid potential drawbacks. Using Jaguar, the team has been able to conduct the largest groundwater simulations ever seen.

Coal is very abundant in the United States, but coal power brings with it a variety of serious problems, one of which is that coal-fired power plants spew CO_2 into the air. Carbon dioxide is the most worrisome of the greenhouse gases; according to the IPCC, levels of this gas in the atmosphere are 35% higher than they were before the Industrial Revolution and are in fact higher than they've been in the last 650,000 years. Climate scientists believe it is no coincidence that we're also seeing a string of the warmest years since people began taking measurements more than 150 years ago.

Carbon sequestration, then, would be one way to reduce the amount of CO_2 being pumped into the atmosphere. The process being simulated by Lichtner's team involves taking CO_2 that has been separated from a power plant's emissions and injecting it nearby into a deep saline aquifer one to two kilometers below the surface. If all went according to plan, it would spread out under a layer of impermeable rock and get an opportunity to dissolve into the brine surrounding it.

The CO_2 would be pumped in a state known as a supercritical phase, which is present when it is kept above 50 degrees centigrade—120 degrees Fahrenheit—and over 100 times atmospheric pressure; it would be kept in that state by the heat and pressure naturally present deep underground. According to Lichtner, CO_2 in this phase is in some ways like a liquid and in some like a gas, but the primary benefit is that it avoids the rapid expansion that would go along with changes between the two phases.

Lichtner's team is investigating a process known as fingering that speeds the rate at which the CO_2 dissolves. Fingering grows out of the fact that while CO_2 in the supercritical phase is lighter than the surrounding brine, brine in which CO_2 has been dissolved is actually heavier than unsaturated brine. The result is a convection current, with fingers of the heavier, saturated brine sinking. This fingering in turn increases the surface area between the CO_2 and the brine and speeds the dissolution of the supercritical CO_2 into the brine.

The rate of dissolution is critical to the success of carbon sequestration. When it is first injected in the ground, the CO_2 pushes the brine out of place. Once the CO_2 dissolves, however, it adds little to the volume of the brine, which can then move back into place.

"The problem is that we're talking about injecting huge amounts of CO_2 by volume," Lichtner explained. "If you were injecting it into a deep saline aquifer, for example, you would initially have to displace the brine that was present, and then the question is, 'Where does that go?' If you pushed up that brine into the overlying aquifers you might contaminate, say, the drinking water for the whole Chicago metropolitan area. It's a race against time how rapidly this CO_2 will dissipate."

There are other hazards as well that must be thoroughly understood before large volumes of CO_2 can be pumped underground. If the CO_2 were to rise to the surface, that would create another substantial hazard. The process of dissolving CO_2 into groundwater is, in fact, known as carbonation; CO_2 rising rapidly to the surface could literally turn the groundwater into seltzer water.

"There are natural occurrences of CO_2 shooting out of the ground," Lichtner noted. "As you lower the pressure and temperature and the bubble of CO_2 you injected starts approaching the surface, you get a change of phase from supercritical to liquid to gas. Then it would occupy suddenly a much larger volume, and you might get a geyser forming like in Yellowstone [National Park].

"As long as the supercritical phase still exists, it presents a hazard to people living on the surface, because it could escape through fractures, abandoned boreholes, or boreholes that leak. And so the rate of dissipation is important to understand to know how rapidly you get rid of this supercritical phase."

A final issue that must be studied focuses not so much on the rate at which CO_2 dissolves, but rather on the changes this process brings to the aquifer itself. As Lichtner explained, CO_2 produces carbonic acid, which in turn lowers the pH of the brine. This could speed the reaction between the newly acidic brine and surrounding minerals and potentially release contaminants into the environment that would not be present otherwise.

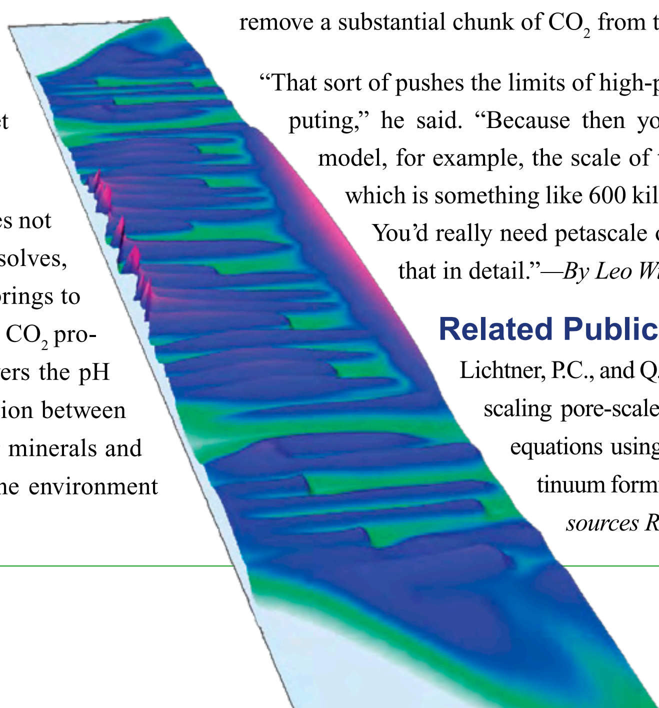

Lichtner's team is focusing its simulations on the Illinois Basin, a 60,000-square-mile area ranging across most of Illinois as well as eastern Indiana and Kentucky. This area relies heavily on coal power, but the size of the region also provides a daunting task to anyone who wants to model it computationally.

The team is simulating carbon sequestration using an application known as PFLOTRAN, which is built on the PETSc parallel libraries developed by a team led by Barry Smith at Argonne National Laboratory. Chuan Lu of the University of Utah developed the supercritical CO_2 implementation in PFLOTRAN while working with Lichtner as a postdoctoral researcher at LANL. Lichtner and his team have shown that PFLOTRAN can handle grids on the order of a billion cells—an unprecedentedly large number for a groundwater simulation. Nevertheless, each cell in such a simulation will be nearly 100 square meters, too large to comfortably analyze the fingering process, which takes place at the scale of tens of centimeters to tens of meters, depending on the properties of the aquifer.

Lichtner noted that his team is working both to improve the performance of PFLOTRAN and to prepare for the arrival of even more powerful supercomputers. To make PFLOTRAN more effective, for example, the team is working to evolve from the use of a structured grid, in which a quarter of the cells give no useful information, to an unstructured grid that can redistribute those cells where they will be of most value.

Beyond that, he said, the team is anticipating a new generation of supercomputers capable of speeds greater than 1,000 trillion calculations a second, or a petaflop. At that scale, Lichtner's team will approach the resources they need to guide the process of carbon sequestration, enabling us to remove a substantial chunk of CO_2 from the atmosphere.

"That sort of pushes the limits of high-performance computing," he said. "Because then you would want to model, for example, the scale of the Illinois Basin, which is something like 600 kilometers in length. You'd really need petascale or larger to model that in detail."—*By Leo Williams*

Related Publications

Lichtner, P.C., and Q. Kang. 2007. "Upscaling pore-scale reactive transport equations using a multiscale continuum formulation." *Water Resources Research* **43** (12).

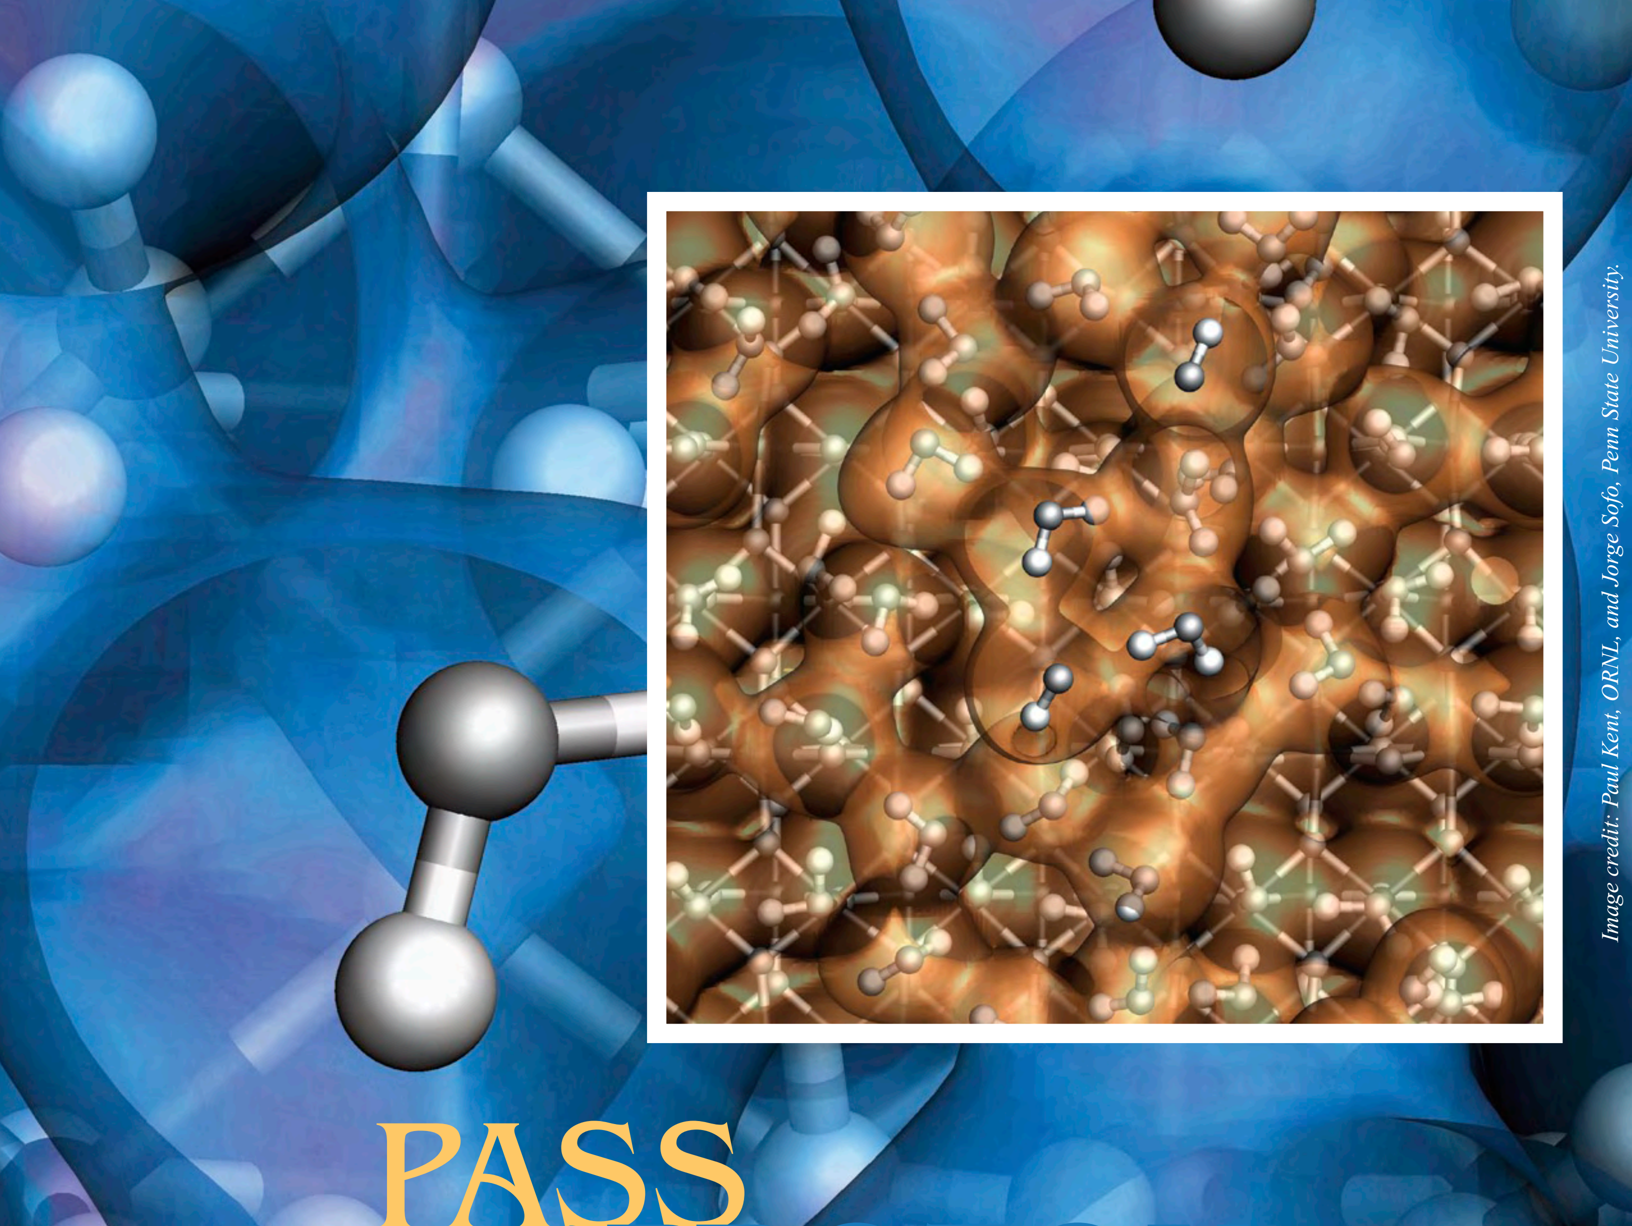

Image credit: Paul Kent, ORNL, and Jorge Sofo, Penn State University.

PASS the PROTON

Magic happens at the nanoscale. From antibacterial coatings to fuel-cell-powered cars, tomorrow's inspiring new technologies depend on today's researchers learning how to control the novel behavior of small-scale systems in a predictable way

A team of materials scientists led by Jorge Sofo of Penn State University and Thomas Schulthess of ORNL has used NCCS's Jaguar supercomputer to successfully simulate such a system. Working from fundamental principles of quantum mechanics, the team accurately simulated the behavior of water in the presence of the common catalyst titanium dioxide, a material routinely used in solar cells and hydrolysis. The work not only improves our understanding of a process that is already important in areas such as fuel cells and the geosciences; it also prepares the way for simulations of ever more complex systems.

"This whole simulation sets the stage for a lot more work on more complicated systems," explained Paul Kent, a member of the team who worked extensively on the computer application used in the research. "This is much more than a proof of concept because we've got a lot of science out of this, but the idea is obviously to move on to more complicated materials."

Specifically, the team simulated the process by which water passes protons from one molecule to another. While a molecule of water—H_2O—contains two hydrogen atoms and one oxygen atom, one of the hydrogens will occasionally break off; this leaves an unconnected hydrogen—the nucleus of which is a single proton—and a hydroxyl molecule (consisting of the leftover oxygen and single hydrogen). This proton can then be exchanged with other hydrogens in other water molecules, transporting the proton through the water.

The mechanism observed in the simulation, known as the Grotthuss mechanism, governs the way in which a proton is passed between hydroxyl atoms. More generally, the simulation shows that first-principles molecular dynamics—in this case performed using the Vienna Ab-Initio Simulation Package (VASP)—can be used to explore this process.

The process is very fast and localized. To capture the behavior of water molecules, the simulation tackled a system of 700 atoms in steps of a half femtosecond each. Equivalent to a half-quadrillionth of a second, a half femtosecond is to a second as a second is to 63 million years. To get enough useful information, the simulation proceeded through more than 20,000 time steps to get to 10 picoseconds, or 10 trillionths of a second of data, making it one of the largest such simulations undertaken to date.

The simulation put the water in contact with titanium dioxide for two primary reasons. First, the catalyst spurred the process enough for a 10-picosecond simulation to yield sufficient results. Second, technologies such as fuel cells typically use catalysts, albeit more expensive catalysts such as platinum, and the group aims to improve our understanding of the processes and thereby improve the technologies.

"We wouldn't be able to study the Grothuss mechanism in these simulations without the titanium dioxide," Kent explained. "The water is swimming around at room temperature, and every so often it will adsorb on the surface of the titanium dioxide briefly and a hydrogen will break off. And it will do a little interchange with other water molecules in the area, so they can swap protons. Without the titanium dioxide, the process is too infrequent for these simulations."

Titanium dioxide, also known as titania, is far less expensive than platinum. Used in a wide range of industrial applications, this common catalyst also makes white paint white, protects skin from ultraviolet radiation, and activates oxygen sensors.

The VASP application used to simulate the process solves the Schrodinger equation, which is, among other things, able to describe the chemical bonding between individual atoms. While VASP is a mature code, Kent has optimized it for use with Cray supercomputers such as Jaguar, improving the robustness of the application and its ability to scale to ever more processors.

Kent said the team sticks with the VASP code because it has proven very robust and effective over the years.

"We haven't changed the fundamental method used in the code at all," he explained. "In fact, we don't want to, because one of the reasons that people use it so much is there's a sense of trust and experience." Eventually, he said, more efficient linear scaling applications are likely to take over, but these are not yet ready for use in large science runs.

The computer specialists are able to compare their results with results from a team of experimentalists led by Dave Wesolowski of ORNL. Wesolowski's team evaluated the same system of water and titania molecules using neutron-scattering techniques at the Intense Pulsed Neutron Source at Argonne National Laboratory. It has looked at other oxide materials as well.

The collaboration illustrates the benefits of experiment and computer simulation working together. On the one hand, computational scientists need to know that their work matches systems in the real world, and Wesolowski's neutron scattering provided validation to Sofo's team. On the other hand, even the most advanced experimental techniques are limited in the amount of information they can provide, especially when looking at the scale of atoms and molecules.

"The concept that we're proving is that we can study this system with sufficient accuracy to get good agreement with experiment, where we've got data," Kent explained. "One of the cool things about this work is that the neutron scattering gives us a fingerprint for the dynamics of the water. And we can, in our simulations, go off and compute that fingerprint as well.

"But experimentalists can't see the individual atoms move around. They only get this fingerprint, which is an average of what's going on. Now maybe there will be some more sophisticated technique in the future. But the idea here is to have this close collaboration so that if we can verify that we've got the same dynamics that they're seeing, we can then go in and see the whys."

In the case of water and titania, neutron scattering is able to measure the vibrational modes of the water—in other words, how the hydrogen and oxygen atoms in a water molecule wiggle. The simulations, however, show in detail not only how individual molecules wiggle, for example the bending and stretching of the oxygen-hydrogen bonds, but also the collective response of all the waters in the system. This enables researchers to understand the necessary conditions for different processes such as local changes in the acidity of the water.

Because the simulations were able to match the neutron scattering results for vibrational modes at varying temperatures and densities, the researchers have reason to believe that other aspects of the system may be accurate as well, aspects that cannot be measured by neutron scattering.

"So we've got this cross reference," Kent noted, "and as closely as we're looking it seems to be accurate. The fact that the modes line up and are shifted appropriately because of the titania, it's confidence building. The cross reference is by no means everything, but it gives a hint that the dynamics in the simulation are pretty good and that we're justified in looking in more detail."

Kent stressed that these simulations are unlikely to lead to great technological innovation in the short term, but they are a necessary step along the way.

"For example, proton transport is critical for fuel cells," he noted. "This is clearly a major motivation for learning what we can about proton transport in water—on a surface that is well characterized and clean and where the neutron-scattering people can do measurements. Looking at a fuel cell membrane is much more complicated, so this is a prerequisite."—*By Leo Williams*

Related Publications

Macridin, A., M. Jarrell, T. Maier, P.R.C. Kent, and E.D'Azevedo. 2006. "Pseudogap and antiferromagnetic correlations in the Hubbard model." *Physical Review Letters* **97** (3).

Fuentes-Cabrera, M., X. Zhao, P.R.C. Kent, and B.G. Sumpter. 2007. "Electronic structure of xDNA." *Journal of Physical Chemistry B* **111** (30):9057-9061.

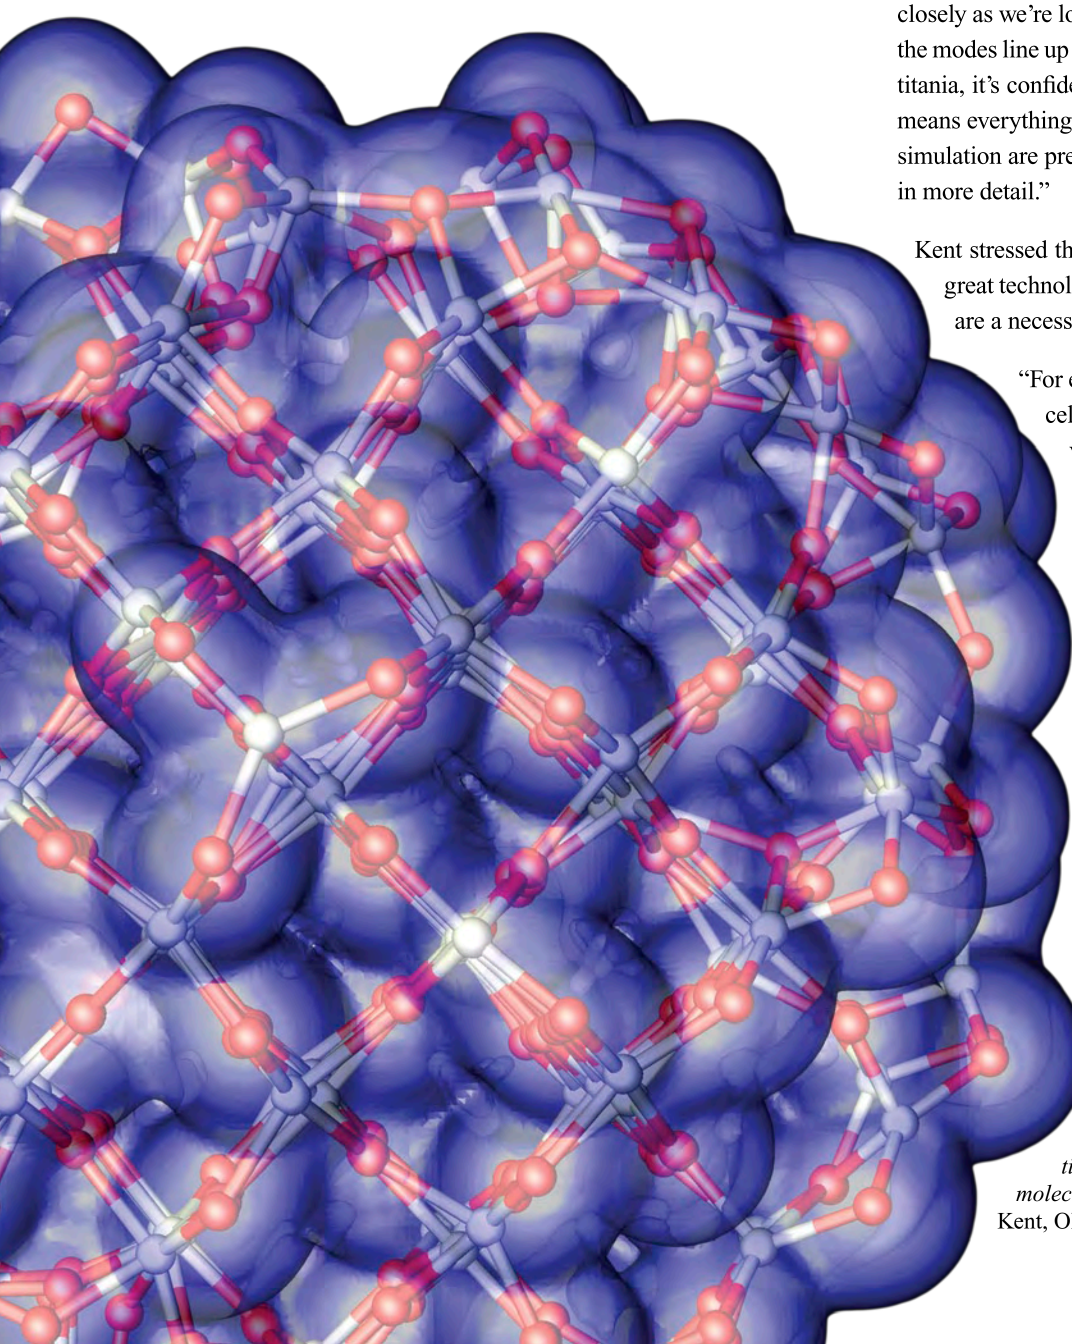

Structure and charge density of water on a titanium dioxide surface from large scale ab initio *molecular dynamics calculations.* Image credit: Paul Kent, ORNL, and Jorge Sofo, Penn State University.

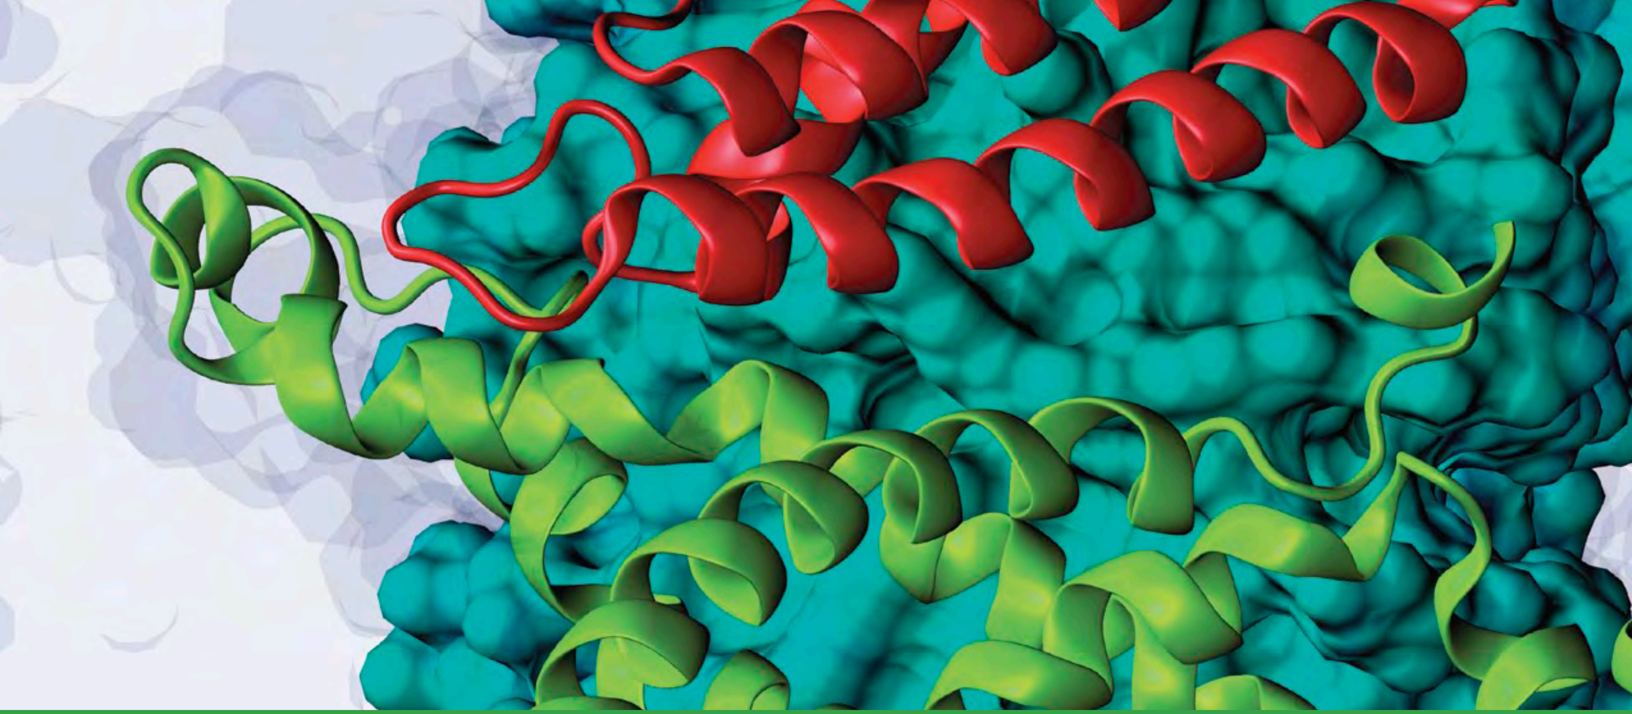

An OPEN and SHUT Case

We're not robots, but the human body has machine-like aspects. That's especially true at the scale of billionths of a meter, the size of atoms and small molecules. A great example of a molecular machine is a membrane protein that responds to spikes of electricity by changing shape to allow potassium ions to enter a cell. Scientists are now using NCCS resources to simulate the voltage-gated potassium channel in unprecedented detail.

"The study will serve as a future road map for simulating, visualizing, and elucidating the workings of molecular nanomachines," says Benoit Roux of Argonne National Laboratory and the University of Chicago. In essence, a voltage-activated ion channel is a nanoscale device acting as an electric switch, he says. With University of Illinois at Urbana-Champaign researchers Klaus Schulten and Emad Tajkhorshid, Roux uses the Leadership Computing Facility at ORNL to model the channel in its open and closed states and determine the gating charge driving the change in conformation between the two states.

If the switch operates normally, the potassium channel opens when activated and closes when resting. But if gating malfunctions—and it can go awry in various ways—cardiovascular or neurological disease can result. Muscle, heart, nerve, and secretory cells produce and respond to electrical signals, earning them the name "excitable cells." The important functions of potassium channels in excitable cells make them good drug targets. Other voltage-gated channels allow selective entry of sodium and calcium ions and are also promising targets.

The potassium channel is made of four identical protein subunits. Each contains segments of amino acids that cross the cell membrane six times like a switchback mountain trail. Two transmembrane segments from each subunit come together to form a pore through which only potassium ions may enter or leave.

The voltage sensor is the other important functional part of the potassium channel. Formed by four transmembrane segments surrounding the pore, the sensor responds to changes in electric potential. In a process called gating, the protein switches its shape to allow or block passage of ions across the cell membrane in response to a change in voltage.

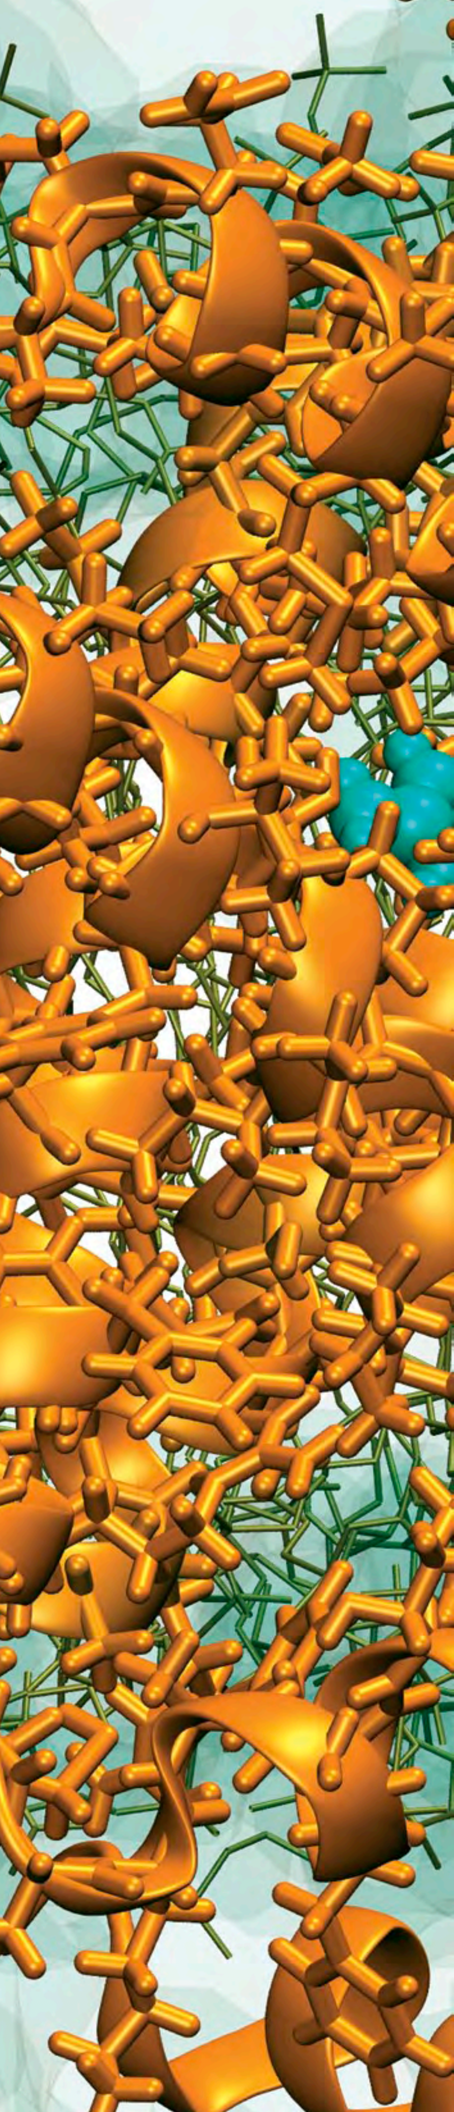

Voltage-gated ion channels allow different ions and charges to gather on each side of a cell's membrane. In a nerve cell, potassium is abundant inside, and sodium, an ion that works with potassium to propagate nerve impulses, is plentiful outside. At rest the cell's membrane potential—maintained by a predominant potassium ion conductance—is slightly negative (polarized). When a hormone, drug, or neurotransmitter binds to a receptor on the neuron, stimulating it, a chain reaction begins that discharges the membrane potential as a tiny electrical current. When the cell becomes excited, it depolarizes, and sodium channels open. Because ions and charges flow from areas of high to low concentration, sodium flows in, further depolarizing the local membrane. Nearby potassium channels respond by opening their gates, allowing potassium to flow down its electrochemical gradient to restore the resting membrane potential. The result? Channels open, inactivate, and close in sequence. The membrane potential changes quickly and transiently, propagating the signal down the nerve like a line of falling dominos.

Gating is the key

What makes the nerve impulse possible in the first place is cellular choosiness for specific ions, or voltage-activated gating. To understand its mechanism, the researchers first aim to determine the channel's structure in its open and closed states. X-ray crystallography, the best method since the late 1950s for determining the structures of proteins, is employed to analyze potassium channels obtained from rat brain tissue. Today, it provides an incomplete picture of the atomic-level structure of the open channel, and no X-ray structure exists for the closed channel. Preparing membrane proteins for X-ray diffraction studies is very difficult because the proteins have to be crystallized under conditions different from their home environment, the cell membrane.

Roux's team is using a computer program called Rosetta to predict the three-dimensional structure of the potassium-channel protein. The group found that simulations of the open and closed states are stable. Assessing stability is critical to supporting the model's validity. Collaborator Vladimir Yarov-Yarovoy, a research assistant professor at the University of Washington, recently adapted Rosetta to better predict the behavior of proteins embedded in membranes.

"Rosetta predicts protein structure starting with amino acid sequence information alone, without any starting template structure," Yarov-Yarovoy says. Strings of amino acids (primary structure) link through hydrogen bonds to form pleats and helices (secondary structure) that fold to form three-dimensional proteins (tertiary structure) that can associate with other proteins (quaternary structure). For a given sequence of amino acids, Rosetta conducts a large-scale search for three-dimensional protein conformations that are especially low in free energy and assumes the native state is the one with the least free energy.

The mechanism of voltage gating remains poorly understood. Currently, three models exist. A "transporter-like" model posits a small motion of four transmembrane segments that is focused in narrow, watery crevices. The two other models propose large movements of one transmembrane segment, which is less exposed to membrane lipids in the "sliding helix" model than it is in the "paddle" one. Finding out which model is correct will help provide a better understanding of these molecular machines and is the second major goal of Roux's project.

"The final truth will incorporate elements of each model," he says. "By constraining Rosetta's large-scale search with known experimental data, we were able to converge toward a well-defined model of the closed state." Because detailed structures are not yet known for the majority of membrane proteins, simulation guided by experimental constraints could aid efforts to obtain other structures.

In a step toward achieving their long-term goal of understanding how membrane-associated molecular protein machines function, the researchers simulated the motion of all atoms in the system using a molecular dynamics code for parallel processing that was developed in Schulten's lab. In 2002 a Gordon Bell Prize honored the code, called NAMD, as an outstanding achievement in high-performance computing. The code uses Newton's laws and an energy function to simulate protein behavior in steps on the order of one femtosecond, or quadrillionth of a second. By looking at how the potassium channel moves in tiny, ultrafast increments, researchers can build a biologically meaningful picture of its dynamics.

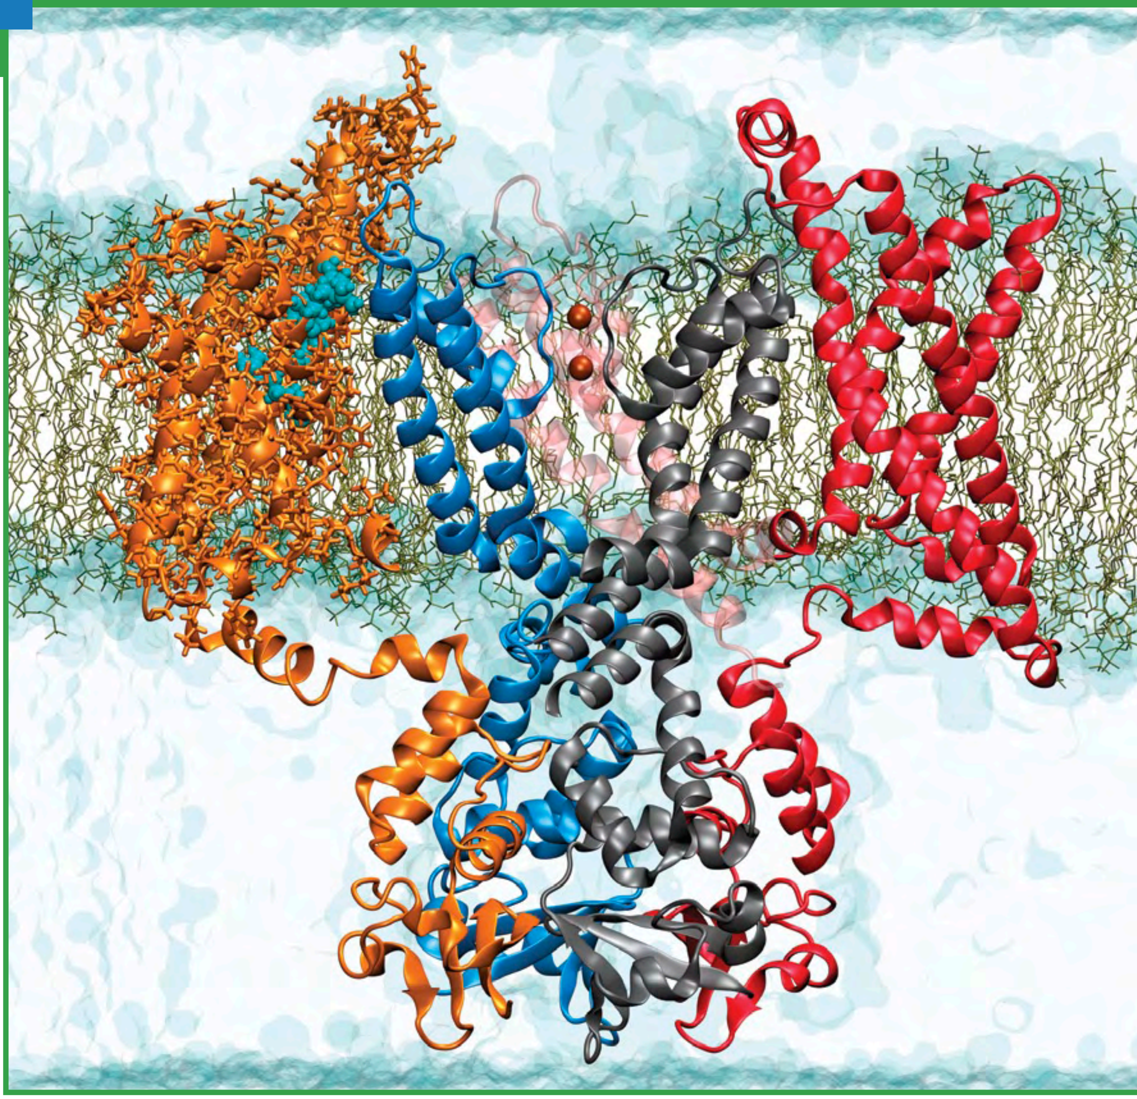

NCCS supercomputers illuminate the workings of the voltage-gated potassium channel, simulated here in its open state in a membrane environment. Image credit: Theoretical and Computational Biophysics Group, University of Illinois at Urbana-Champaign.

INCITE for molecular nanomachines

Roux's group received funding from the National Institutes of Health and an allocation of NCCS supercomputer time through the INCITE program, which supports computationally intensive research. In 2007 Roux and colleagues used about 2.5 million processor hours on the NCCS's Cray XT Jaguar supercomputer to model the behavior of systems with up to 350,000 atoms using NAMD. The models calculate the minimum energy state, an indicator of what shapes molecules would be most comfortable assuming. The molecular dynamics simulations, which can scale to monopolize the entire system, used more than 20 percent of the system's processing cores, according to NCCS liaison Sadaf Alam, who works closely with the researchers to assist them in achieving their scientific productivity goals on the NCCS platform through code optimization and scaling as the system hardware and software evolve. "We are in the process of unraveling the atomistic basis for the coupling between a voltage-gated channel and the transmembrane electric potential," Roux says.

To continue their studies, the researchers have received a 2008 INCITE grant of 3.5 million hours on Jaguar.

Unraveling Alzheimer's

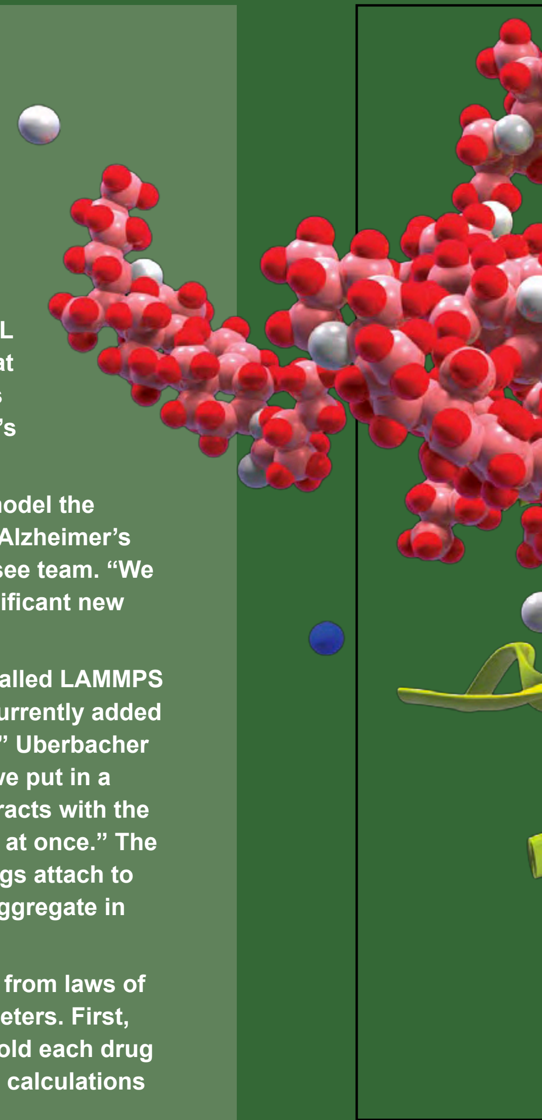

For the majority of people with Alzheimer's, the degenerative brain ailment is a two-protein disease. Amyloid protein outside neurons forms plaques; tau protein inside neurons forms neurofibrillary tangles. Current drugs delay symptoms but do not stop formation of plaques and fibrils.

In 2007, researchers Edward Uberbacher and Phil LoCascio at ORNL used 100,000 processor hours on the Cray XT4 Jaguar supercomputer at the NCCS to investigate the mechanisms by which a new class of drugs acts. The drugs, called caprospinols, may stop the growth of Alzheimer's fibrils and even disassemble them.

"This is the first time that molecular dynamics have been used to model the mechanisms these drugs use to interact with and reduce the growth of Alzheimer's fibrils," says Uberbacher, who leads a joint ORNL-University of Tennessee team. "We learned that there are several different ways these work and added significant new knowledge to the existing models."

Uberbacher, LoCascio, and colleagues employed a software code called LAMMPS to compute the molecular dynamics of 20 to 50 Alzheimer's drugs concurrently added to Alzheimer's fibrils. "This simulation is very much like an experiment," Uberbacher says. "Since the fibril is large and will require multiple drugs to affect, we put in a number of drug molecules and then track what each one does as it interacts with the fibril. This shows us lots of different possible interactions with the fibril at once." The researchers used this information to explore mechanisms by which drugs attach to and reconfigure small proteins called peptides bound in fibrils, which aggregate in the Alzheimer's brain as plaques.

The supercomputers performed *ab initio* calculations that originate from laws of physics and do not make assumptions such as model and fitting parameters. First, the supercomputers calculated the chemical bonds and energies that hold each drug molecule together. Then they performed standard molecular mechanics calculations to explore a drug's activity.

"Since we can do quantum mechanical *ab initio* calculations on 1,000 atoms or so, we can generate this knowledge in a way that is more accurate than what pharmaceutical companies usually do," explains LoCascio. "Hopefully this method will become more widespread in industry and lead to better drug design."

Results show promise. Some drugs bound to the growing ends of Alzheimer's fibrils, impeding further growth. A drug developed by researchers at Georgetown University and licensed by Samaritan Pharmaceuticals prevented an Alzheimer's peptide from changing to a conformation that would allow addition of peptides to a growing fibril. Another drug unraveled fibrils by causing their peptides to disassociate and assume conformations consistent with their free-solution form.

"This work provides an initial model for how peptide dissociation from the fibril can be made to occur," Uberbacher says. "The simulations performed on Jaguar are an important demonstration of a new paradigm for dynamic modeling of drug-protein interactions."

Based on knowledge gained in the 2007 simulations, subsequent simulations will investigate alternative structures within the same drug class. In addition, testing in mouse models of Alzheimer's disease and mouse-brain imaging related to drug testing are planned with the University of Tennessee.— *By Dawn Levy*

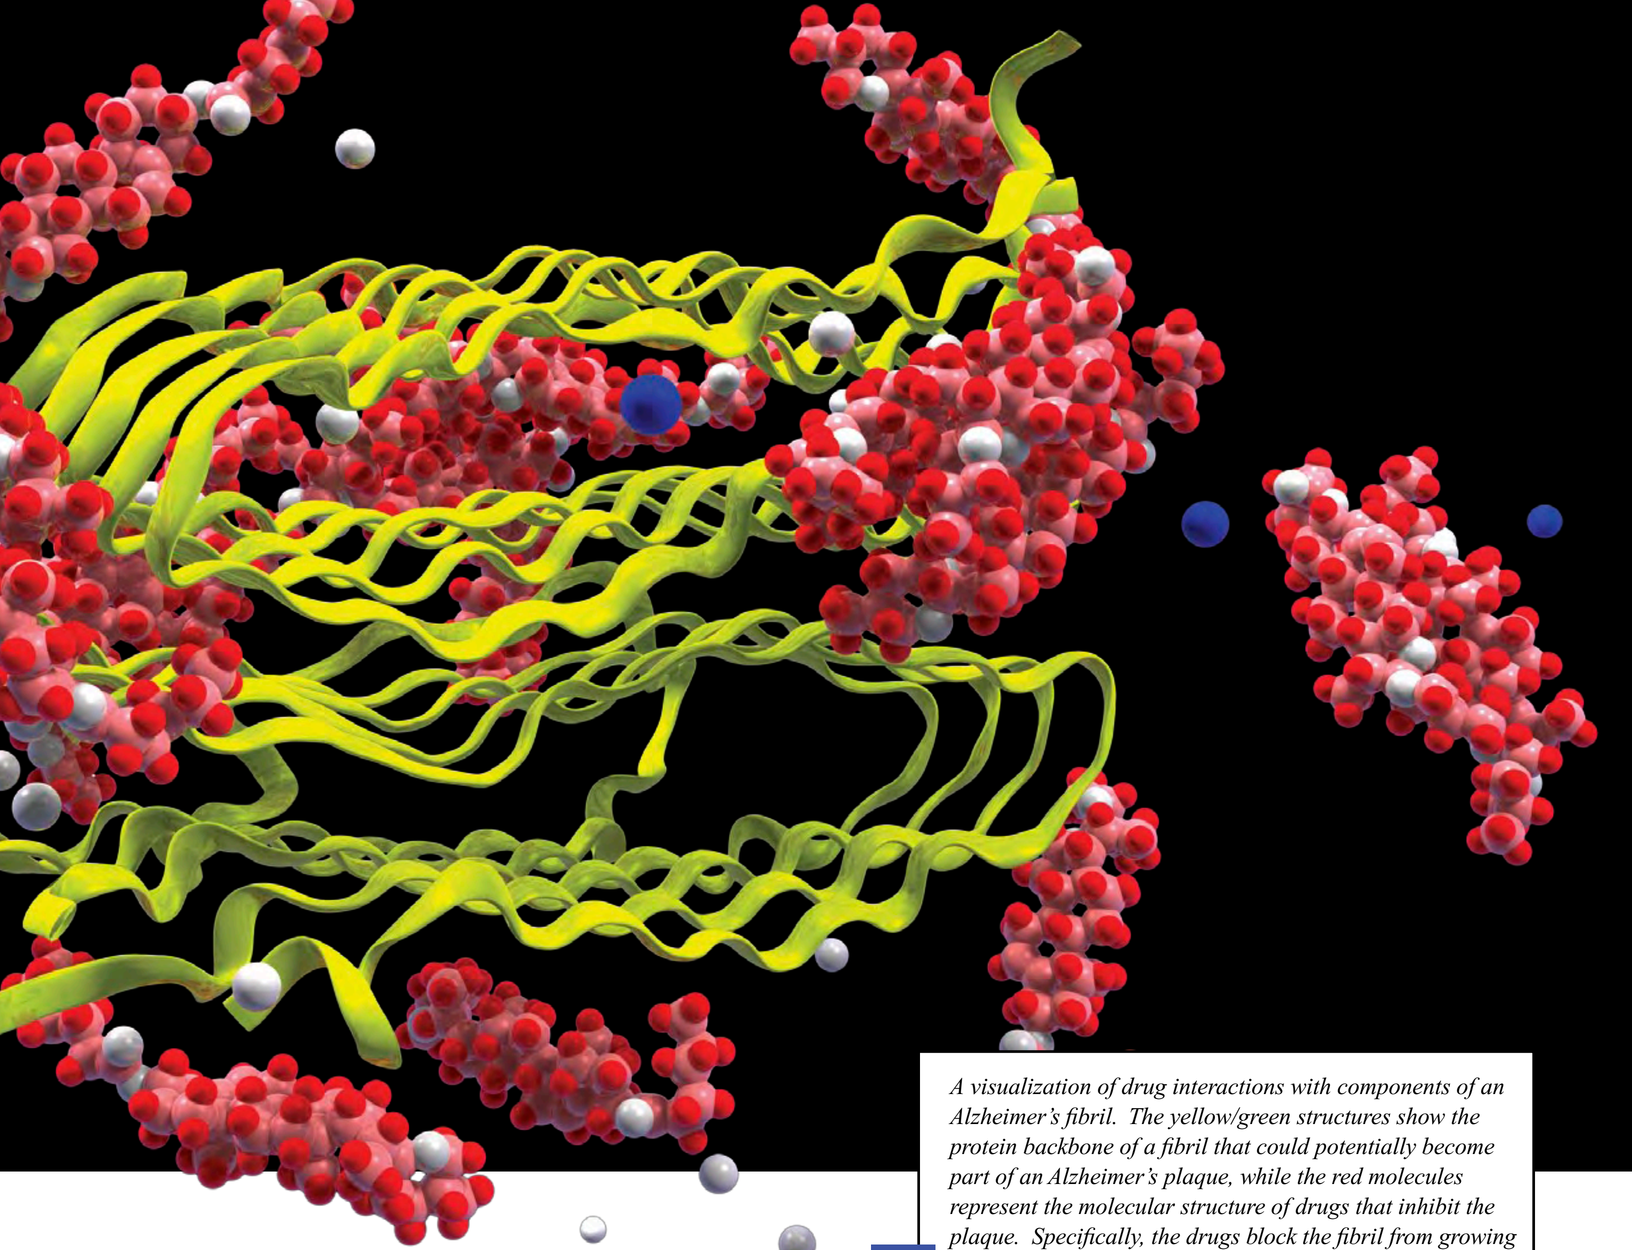

A visualization of drug interactions with components of an Alzheimer's fibril. The yellow/green structures show the protein backbone of a fibril that could potentially become part of an Alzheimer's plaque, while the red molecules represent the molecular structure of drugs that inhibit the plaque. Specifically, the drugs block the fibril from growing through the accumulation of additional material. The blue and white spheres represent salt ions. The drugs are also beginning to unravel parts of the fibril. Image credit: Ed Uberbacher, ORNL, and Sean Ahern, ORNL/NCCS

Simulating voltage-activated changes in the potassium channel will provide proof-of-principle that the functions of protein-based molecular machines can be simulated, Roux says. Researchers can then expand this computational strategy to study membrane proteins that transport molecules into the cell and those that pump substances out to mediate biological processes such as catalysis of biofuels, production of rare and unique compounds, and detoxification of organic waste products.

As a case in point, Roux points to a regulator of serotonin, a neurotransmitter with important roles in sleep, depression, and memory. "The serotonin transporter is *the* protein that is inhibited by Prozac and other selective serotonin reuptake inhibitors," he says. Understanding it and other similar membrane transport molecular nanomachines may open a floodgate of medical innovation.—*By Dawn Levy*

Related Publications

Cordero-Morales, J.F., L.G. Cuello, Y.X. Zhao, V. Jogini, D.M. Cortes, B. Roux, and E. Perozo. 2006. "Molecular determinants of gating at the potassium-channel selectivity filter." *Nature Structural & Molecular Biology* **13** (4):311-318.

Jogini, V., and B. Roux. 2007. "Dynamics of the Kv1.2 voltage-gated K+ channel in a membrane environment." *Biophysical Journal* **93** (9):3070-3082.

DIFFERENT FORCES RULE THE UNIVERSE AT DIFFERENT SCALES

KEEPING IT TOGETHER

As you look to the planets and the stars, gravity prevails, holding worlds and galaxies together. When you look around at the environment of your everyday life, electromagnetism takes charge, keeping you from falling through the floor and allowing you to walk across it.

As you look to the microscopic, the world is ruled by the strong nuclear force, usually known simply as the strong force. This is the force that holds quarks together to form protons and neutrons, the subatomic particles that make up more than 99 percent of the stars we see and the world we live in. A remnant of this force is also responsible for holding these particles together to form atomic nuclei, overcoming the natural aversion that positively charged protons have to one another and giving us the elements that make the world interesting and our lives possible.

A team led by David Dean of ORNL is exploring the nature of nuclei and of the remnant of the strong force that holds them together. By showing us the energy of an atomic nucleus and the energy necessary to pull it apart, the team is paving the way for safer, less expensive, and more efficient sources of energy. It is also giving us insight into the messy process by which stars create new elements.

The strong force is also known as the color force, and it is governed by a field of study known as quantum chromodynamics (QCD).“You only see a little bit of the strong force when you see the proton and neutron interacting,” Dean explained. “You only see the tail of the QCD force, and that’s the force that’s really working in the nucleus. The main QCD drivers are inside the protons and neutrons. In between, there’s a little bit of force. It’s strong, but it’s not nearly as strong as what goes on inside a proton or inside a neutron.”

Dean’s team is simulating nuclei on ORNL’s Jaguar supercomputer through the INCITE program. The team performed simulations with an allocation of 5 million rocessor hours in 2007, and it will move forward with an allocation of 7.5 million processor hours in 2008.

The team is pursuing its work with several techniques, one of which is known as coupled-cluster theory. This approach has two primary benefits. On the one hand, it is an *ab initio* —or first principles—approach, meaning properties of the nucleus are calculated without simplifying approximations of the physics involved. On the other hand, coupled-cluster methods scale more efficiently than other *ab initio* techniques, allowing Dean and his colleagues to pursue larger nuclei than would be otherwise possible.

These calculations become monumentally complex for several reasons, one of which is that scientists are still working to understand precisely how the strong force—which applies primarily to quarks and gluons—operates at the much larger scale of protons and neutrons. Another challenge is the fact that interactions between the protons and neutrons in a nucleus are strongly repulsive at short distances and mildly attractive at intermediate distances, meaning many, many basis states must be used to capture the behavior of the nucleus. Third, the project is simulating a range of nucleus types, including the weakly bound and very short

Aided by high-performance computing, physicists are exploring the nature of nuclei and the force that holds them together. Image credit: LeJean Hardin and David Dean, ORNL.

lived as well as the tightly bound and long lived. This range also greatly increases the number of basis states—and, in turn, calculations—that are necessary.

The team was able to use coupled-cluster theory to explore the nature of helium, the universe's second-lightest element with two protons. In particular, the researchers were able to add neutrons to the simulated atom and evaluate the nucleus as it grew, from helium-4 (with two protons and two neutrons) to helium-5, -6, -7, -8, and -9.The project, which the team documented in a paper in the journal *Physics Letters B*, was especially challenging because several of the nuclei—in particular helium-5, -7, and -9—are very short lived. While some configurations of protons and neutrons are stable, meaning they stick together, others hang together for only a miniscule fraction of a second. Known as resonance states, these configurations are very difficult to calculate.

"Let's say I did a collision," Dean explained. "So here's a particle that comes in and hangs around, but it just goes around the nucleus four or five times and then leaves. That would be a resonance state. If the particle came in and just stayed there forever, that would be a bound state.

"The helium-4 is well bound, but helium-5, -7, and -9 are unbound. That means they're in resonant states. You can't make them stick together in the laboratory; they just fall apart. But they have a width; they have a measurable lifetime."

Dean and his colleagues were also able to perform the first coupled-cluster analysis of a helium nucleus that clears a long-standing hurdle for nuclear analysis—the three-body force. To this point, analysis of these nuclei has simply summed the interaction of each pair of particles, leaving out the effect created by more than two particles at a time. Instead of simply calculating the effect of every particle on every other particle, though, Dean's simulations moved forward to calculate the interaction of multiple particles at the same time. Dean likens the three-body force to the interactions among the Earth, sun, and moon, which together create tides.

Quantum Spin Doctors Dissect Exotic States Of Matter

When German physicist Max Planck created quantum theory in 1900, he was not trying to revolutionize the world. He was just trying to provide a theoretical foundation for the way a heated object radiates energy and, thereby, to improve the efficiency of light bulbs.

Modern scientists, too, cannot know which research will change the world, but history suggests the next scientific revolution will focus on complex systems rather than isolated bits of matter. If so, it may stem from computer simulations such as those performed by Tommaso Roscilde of France's École Normale Supérieure de Lyon (and recently of Germany's Max-Planck Institute). A team led by Roscilde is using ORNL's Cray XT4 Jaguar supercomputer to explore the quantum mechanical phenomena that give us superconductors and superfluids.

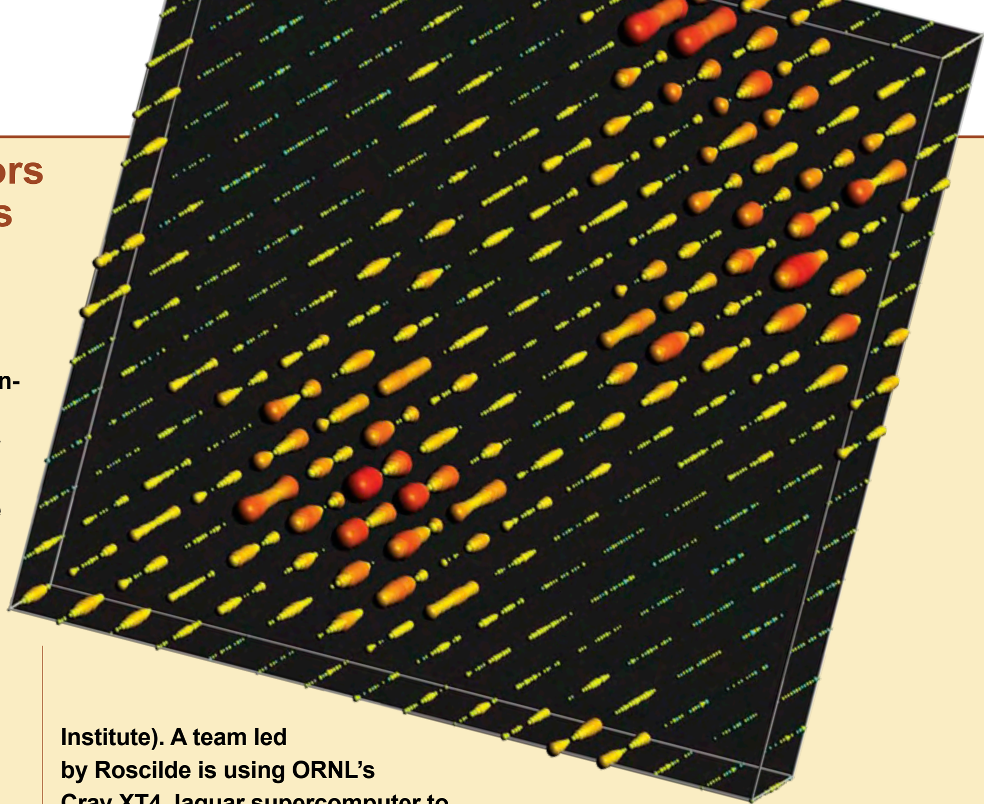

A simulated lattice of atoms lets researchers probe the properties of exotic materials. Image Credit: Tommaso Roscilde, École Normale Supérieure de Lyon, and Ross Toedte, ORNL/NCCS.

Dean noted that calculations of the three-body force, while difficult, are increasingly important to researchers as they move on to heavier nuclei."In theory the nucleons are treated like point particles" having no height, width, or depth, he explained. "But they're not point particles; they have a 1 fermi radius [1×10^{-15} meters, or a quadrillionth of a meter]. And in the nucleus, they can actually deform a little bit when they interact. They can excite. Those things become a three-body force. So this three-body force takes three nucleons and says they're going to talk to each other somehow.

"I think it was the first time ever coupled-cluster theory was used to look at these widths. It was the first time ever it was applied to weakly bound nuclei. So that was a big step in the method, because that's one of the things we wanted to do."

The team is also pursuing the coupled-cluster approach with heavier nuclei. While it may not be able to tackle nuclei of the heaviest atoms, Dean is confident that the approach will be successful in analyzing medium-mass nuclei, perhaps up to mass 100, giving researchers information they can use to make a reasonable assessment of even heavier nuclei.

This assessment will be made using a technique known as density functional theory. Dean explained that coupled-cluster theory is only one in a range of techniques used by researchers to delve into the nature of nuclei. For the lightest nuclei, he said, researchers use an approach called configuration interaction, which is even more rigorous than coupled-cluster theory but is also more cumbersome in scaling to larger nuclei. And for nuclei too heavy for coupled-cluster theory, they use density functional theory, which uses an approximation known as a density functional to help uncover the secrets of the nuclei.

By applying coupled-cluster theory to ever-heavier nuclei, Dean and his colleagues are providing the information needed to pursue these heavy nuclei.

"Density functional theory can be used to calculate a broad range of thousands of nuclei," he explained, "but it needs to have a landmark. It needs to have an anchor in reality. The coupled cluster will help you calibrate the density functional theory. That's one of the overarching goals.

"I'm not sure we want to calculate thousands of nuclei in coupled-cluster theory, but we want to calculate key isotopes, and to me that means we probably want to do the nickels [with 28 protons] and the tins [with 50 protons] in an ab initio fashion, out to nickel-78 and to tin-132 and beyond. Those are the nuclei that we're aiming for in the next few years. They are important markers for the nuclear energy density functional, the nuclear DFT [density functional theory]. So we calculate something ab initio, and the DFT

Roscilde and his teammates—Stephan Haas of the University of Southern California and Rong Yu of ORNL—are using Jaguar through the INCITE program. A grant for 800,000 processor hours in 2007 allowed the team to simulate a lattice of atoms in a quantum magnet to examine two extraordinary quantum phases, or states of matter.

In the first phase, called Bose-Einstein condensation, atoms throughout the material occupy the same state, with the same momentum, range of probable locations, and spin. In a quantum system, this is the closest they can get to being in the same place at the same time. By simulating the introduction of impurities into the material, the team is also able to create the second phase, known as Bose glass. In a Bose glass, the impurities force the condensation into separate islands throughout the lattice, with atoms sharing the same state only with other nearby atoms.

Roscilde's team is using a technique called Quantum Monte Carlo to simulate disorder in the quantum magnet and thereby create Bose glass. The team hopes its efforts will allow its collaborators—Vivien Zapf and Marcelo Jaime at the National High Magnetic Field Laboratory at Los Alamos National Laboratory—to perform the first experimental confirmation of Bose glass.

The work is at the cutting edge of condensed-matter science.

"I find that in this particular instance of a study of a solid-state system, you're really trying to tailor matter to a level of control that was unthinkable a few decades ago or even a few years ago," Roscilde said. "What you have is a system where, in principle, you can tune the system among completely exotic phases that have no analog in classical systems."

Roscilde, like Planck a century before, cannot say whether his work will produce benefits beyond a deeper understanding of the universe. He noted that a device based on quantum systems would likely be very different from existing technologies.

"You have to think hard what to make out of these systems," he said. "It's not just that once you know them, you know what to make out of them. You have to totally think of new functionalities." *—By Leo Williams*

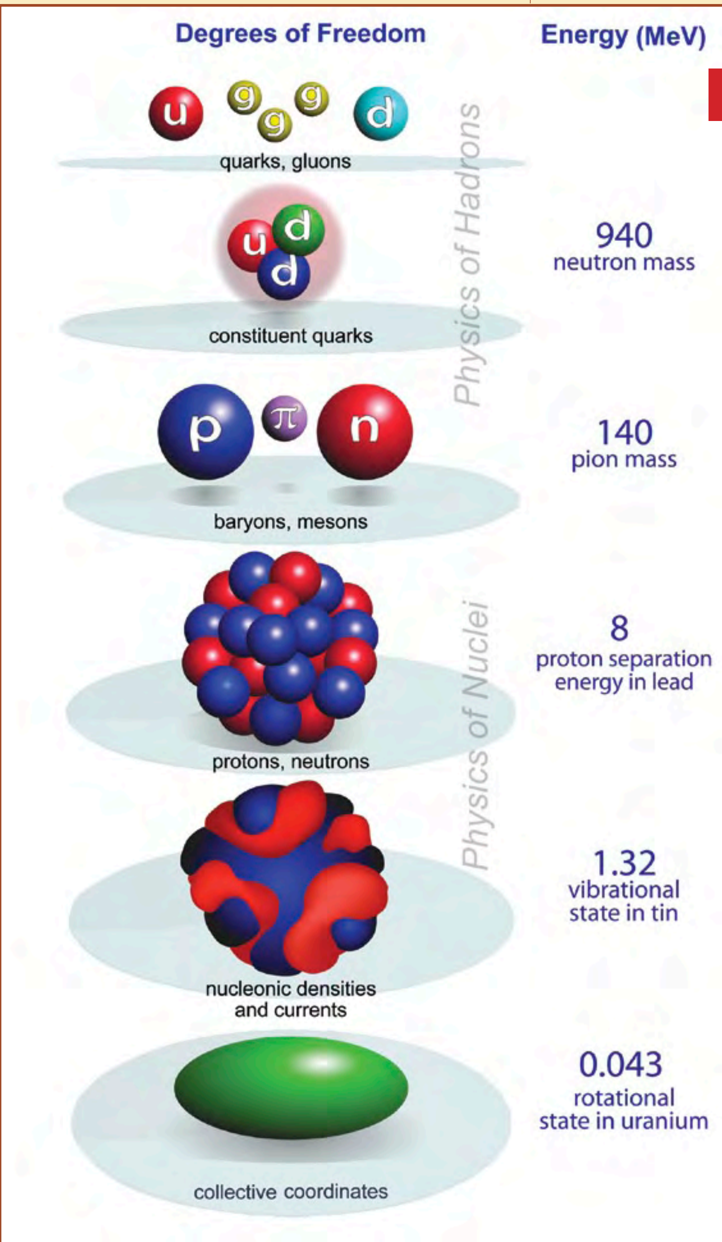

folks calculate something from density functional theory, and we can compare, and make a statement about the robustness of the density functional theory in those nuclei."

With that information, scientists will know better both how nuclei stick together—which will be key in creating future generations of fusion reactors—and how they come apart—which is important for designing safer and more efficient fission reactors. They will also understand better how stars create and distribute the building blocks of the universe.

"From the standpoint of energy independence, you get a better idea of how atoms fuse and how they split apart and the kind of energy that's created," Dean explained. "From a more esoteric point of view, you're adding a piece of the puzzle on the nature of the universe. This is how we exist."*—By Leo Williams*

Related Publications

Hagen, G., D.J. Dean, M. Hjorth-Jensen, and T. Papenbrock. 2007. "Complex coupled-cluster approach to an ab-initio description of open quantum systems." *Physics Letters B* **656** (4-5):169-173.

Gour, J.R., P. Piecuch, M. Hjorth-Jensen, M. Wloch, and D.J. Dean. 2006. "Coupled-cluster calculations for valence systems around O-16." *Physical Review C* **74** (2).

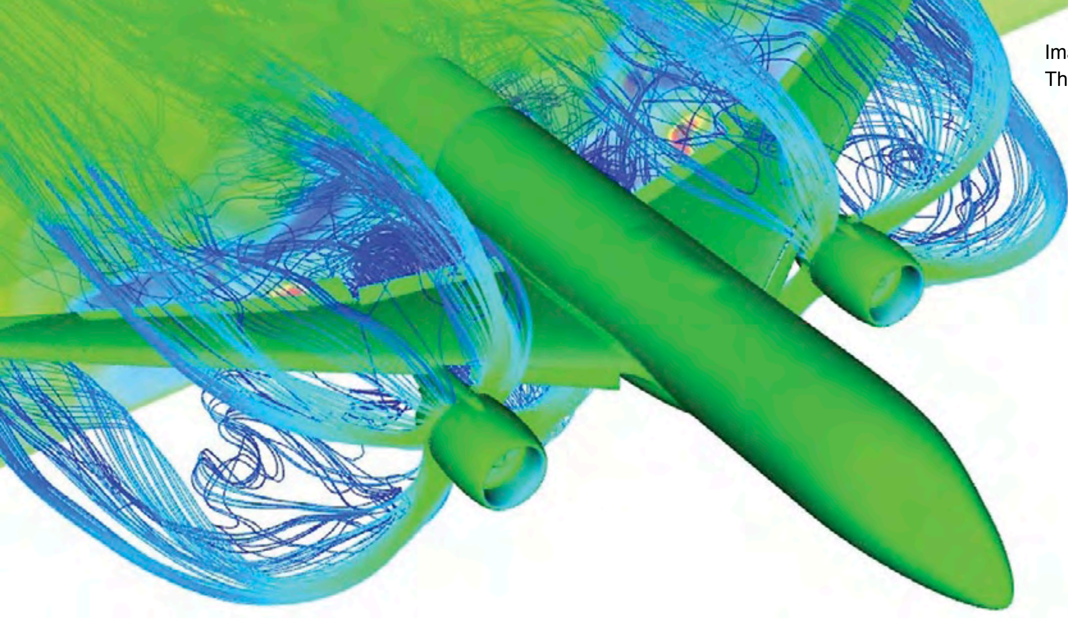

Image courtesy of
The Boeing Company

IN ADDITION TO TACKLING TODAY'S MOST CHALLENGING SCIENTIFIC DILEMMAS, THE NCCS HAS PARTNERED WITH AMERICA'S INDUSTRY LEADERS TO FOSTER INNOVATION AND FIND PRACTICAL, REAL-WORLD SOLUTIONS TO PRESSING REAL-WORLD PROBLEMS

TAKING CARE of BUSINESS

Boeing, General Motors (GM), and the National Energy Technology Laboratory (NETL) have all used the NCCS's computing systems to address a range of important issues. From designing safer aircraft to sequestering carbon from coal-burning power plants, these industrial partnerships promise to keep American industry competitive and benefit the American public through technological innovation.

Take the world's leading aerospace contractor, Boeing, for example. The aviation giant is under intense pressure to manufacture airplanes in less time with less money, all the while without sacrificing safety. A great many of the world's aircraft are aging and inefficient, and expanding economic powers like China and India have put unprecedented demand on the global fleet. Leading manufacturers like Boeing are expected to help fill the void. Through simulation, the computing muscle available at the NCCS has allowed Boeing to more rapidly and inexpensively test airplane components, such as wing surfaces under stress, and phenomena, such as the flow of air around the wings. "Boeing relies on partnerships such as INCITE to be able to access these computers in order to do proof-of-concept development and validation," says Moeljo Hong, a Boeing researcher and principal investigator of the project at the NCCS.

Researchers at GM are likewise using NCCS resources, in this case to make our air cleaner and our automobiles more efficient. By modeling various materials on NCCS supercomputers, GM engineers are exploring the possibilities for converting the wasted heat from automobile exhaust into usable electricity. Exhaust consumes 60 percent of the energy generated by an engine and is largely responsible for the increasingly problematic air pollution in many urban centers around the world. The potential of this research to greatly improve our automobiles, thus reducing the amount of oil our nation consumes, and clean up our air in the process is indeed promising. In a testament to the collaborative success between the NCCS and GM, the automotive giant recently performed the largest-ever simulation of a 1,000-plus-atom supercell—a feat unattainable at any other available facility.

GM isn't the only organization to recognize the potential of high-performance computing for making our world cleaner and greener; researchers from NETL are also using NCCS systems to help us all breathe a little easier. By simulating coal gasification, a process in which coal is pulverized (thus separating fuel-rich hydrocarbons from pollutants), scientists hope to build a new generation of eco-friendly power plants. Given that coal represents the major source of generated power in America and many other countries, and that its byproducts are known pollutants, this research has the potential to impact a range of industrial and environmental issues. As generating power from coal becomes more efficient, the cost of that power will come down, benefitting industry financially. And by making coal-generated power cleaner, the research benefits the entire world by significantly reducing industry's pollution footprint.This new generation of coal-fired power plants is envisioned to emit no nitrogen oxides, mercury, or sulfur and to trap most carbon dioxide. Until we are able to viably produce alternative sources of energy to meet the world's growing demand, coal will remain a necessary source of abundant, affordable power. However, as a result of the collaboration between the NCCS and NETL, it doesn't have to continue to pollute the air we breathe.

increasing investment in research, strengthening science and mathematics programs throughout the nation's education system, and building a scientific and academic infrastructure conducive to innovation and technological advancement.

Essentially, the America COMPETES Act aims to maintain America's global leadership position in the realms of science and technology through more investment in research and development and through the development of more robust educational initiatives to encourage interest in science and mathematics in elementary, junior high, and high schools.

The unique partnership between government and industry provided by the INCITE program is mutually beneficial. Industry receives time to conduct research on some of the world's most advanced supercomputers and, owing to the open science aspect of the INCITE program, those companies must share their findings with other industry leaders. This symbiotic platform for shared research will inevitably benefit Americans in countless ways and keep the United States in the lead in an increasingly competitive global scientific arena.—*By Scott Jones*

NCCS Partners with Industry

Previously, both the GM and NETL allocations were Director's Discretion projects, meaning that they showed early potential to scale well to large systems such as the NCCS's Cray XT4 Jaguar supecomputer. During this time, NCCS staff helped port and scale the codes and assisted the projects toward maximizing their research potential. In 2007, both GM and NETL were awarded INCITE allocations on Jaguar, demonstrating the ability of the NCCS to help researchers optimize their codes and perform the best simulation science possible.

Besides the obvious benefits that come with safer aircraft, more efficient automobiles, and cleaner energy, all of these projects help spur American innovation and keep the country competitive in the global economy. For this reason the industry-sponsored research at the NCCS exemplifies the goals outlined in the recent legislation titled "America COMPETES" (Creating Opportunities to Meaningfully Promote Excellence in Technology), which focuses on three major objectives:

Discussing how supercomputers can boost America's industrial competiveness, NCCS Visualization Task Leader Sean Ahern, far left, addresses, from left, DOE Manager of the Oak Ridge Site Office Gerald Boyd, ORNL Director Thom Mason, U.S. Representative Zach Wamp, and U.S. Senator Lamar Alexander.

Operations

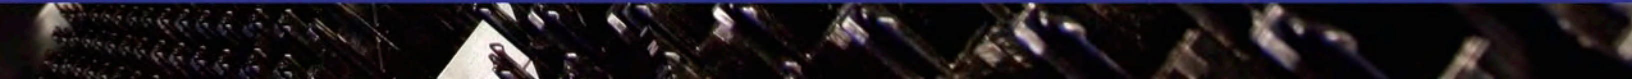

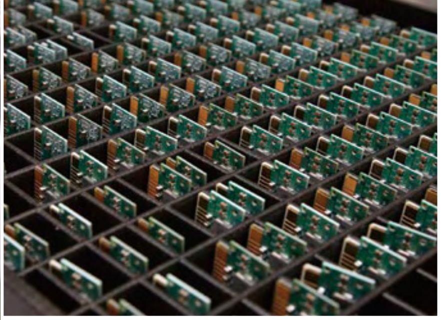

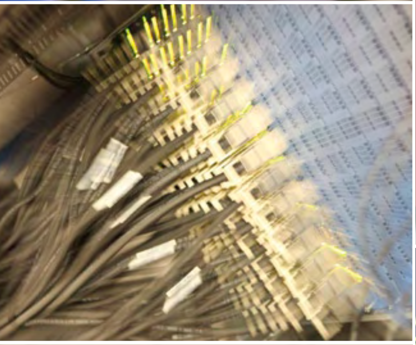

RAMPING UP!

The growth curve of super-powered computers at NCCS shows no sign of slowing. There is commensurate demand for all the infrastructure required to keep the big machines operating and manage their information output.

A year ago, NCCS operating staff faced the challenge of making room for a new petaflop computer to be installed in 2008 in ORNL's Computer Science Building (CSB). Now they're expecting twins: The NSF awarded a second near-petaflop machine to the University of Tennessee (UT) and its partners, including ORNL, in September 2007. The two massive systems will share a room on the first floor of the CSB. The DOE computer is scheduled to be installed in late 2008. The NSF machine is being built in stages and will reach its full capacity in 2009.

In addition to preparations for the new petaflop machines, NCCS completed upgrades of the Cray XT4 Jaguar that quadrupled the speed of that computer during 2007 and added an IBM Blue Gene/P to the computer collection. The demands of the expanded and new supercomputers for power; cooling; and data transport, storage, and management continue to push the limits.

Jaguar Redux

To accommodate the computational science community's appetite for speed, Jaguar was upgraded twice during 2007. In March, acceptance testing was completed for a new 68-cabinet, 65-teraflop Cray XT4 installed in the computer room on the second floor of the CSB. Cabinets from the existing Cray XT3 system in the CSB's first-floor computer room were then moved to the second floor and connected to the XT4 system, and the two systems were merged. The combined XT3/XT4 system consisted of 124 cabinets with dual-core AMD Opteron 2.6-gigahertz processors, giving it more than 23,016 computational cores. Its operating speed was 119 teraflops. Memory was doubled along with speed, maintaining a memory balance of 2 gigabytes of memory per core. The combined system offered a total 46 terabytes of memory and 750 terabytes of disk storage.

The upgrade to 119 teraflops made Jaguar the fastest open computer in the world on June 2007 and the second-fastest of all computers.

In December 2007, Jaguar's 124 cabinets were partitioned into two sections—one with 84 cabinets and one with 40. The 84-cabinet partition was upgraded to bring Jaguar to 263 teraflops. All 7,832 of its compute sockets were upgraded to 2.1-gigahertz AMD Opteron quad-core processors, doubling the number of cores per socket. The memory was doubled to preserve the memory balance. The system has 63 terabytes of aggregate memory and 750 terabytes of disk space. The change to this quad-core Opteron processor also increases the per-core computational capability from two flops per cycle to four. All things considered, the total system capability increased by a factor of 3.5 in a smaller physical footprint.

During a High-Performance LINPACK (HPL) run in early 2007 that demonstrated benchmark performance in excess of 100 teraflops, the machine consumed 2.2 megawatts of power. Fortunately, the AMD Opteron cores maintain the same power profile for the quad-core processors as for the dual-core, so the electrical power and cooling demand for the machine didn't change when it went to 263 teraflops. That means it's providing 3.5 times as much computing power at about the same operating cost.


Jaguar XT4 supercomputer

Testing and acceptance for the 263-teraflop Jaguar was scheduled to be complete early in 2008, at which time it was to return to production.

The increased power and speed will mean larger allocations for researchers. Several projects with allocations on Jaguar in 2008 will have at least 10 million hours, and one project will have 18 million hours. In 2007, by contrast, 10 million hours was the largest allocation granted.

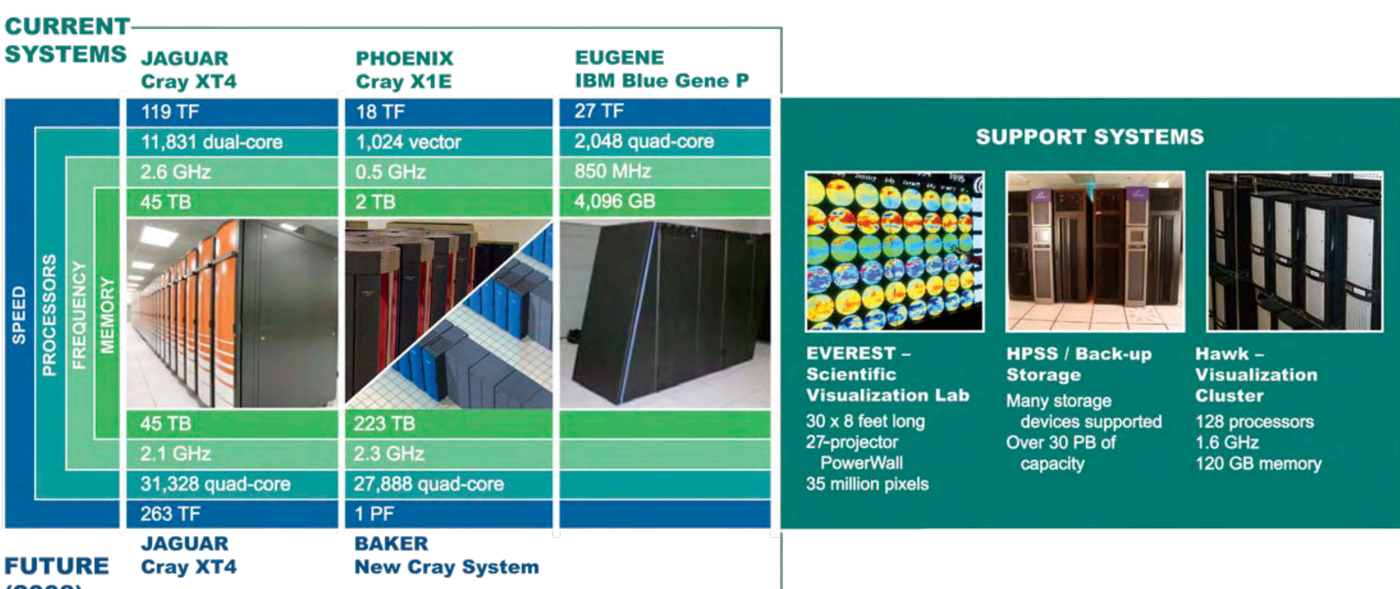

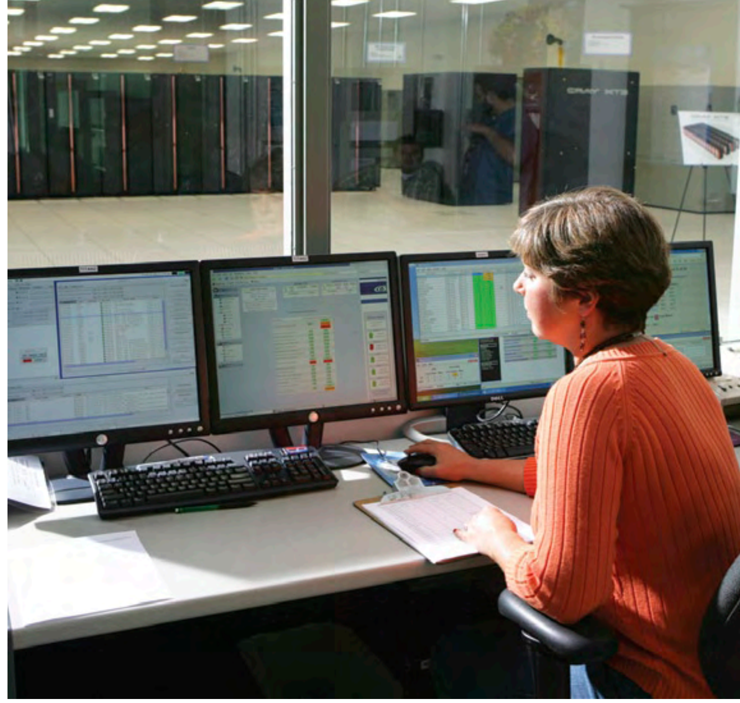

Compute Node Linux

NCCS worked with Cray throughout 2007 to accelerate the development and testing of a new operating system, Compute Node Linux (CNL), created for XT systems with high core counts. Catamount, the operating system used on Jaguar since its installation, was designed to support single- or dual-core processors and is inadequate for the quad-core processors in the upgraded Jaguar. CNL provides greater performance and stability than Catamount and can scale to 100,000 nodes.

The joint effort in developing CNL paid benefits to both Cray and the NCCS. The NCCS allotted Cray large blocks of computing time throughout 2007 to run CNL on the NCCS machines to develop it and put it into production quickly. As a result, CNL reached stability much more quickly than it would have if Cray had developed it using only in-house resources.

The collaboration made CNL available to NCCS for installation on Jaguar during the process of partitioning and upgrading it. When Jaguar was partitioned, CNL was installed on the smaller system to stabilize it and test its performance before it was installed on the 263-teraflop Jaguar. Both systems were running CNL by late 2007. Another benefit to DOE was that because of the development work at NCCS, the National Energy Research Scientific Computing Center (NERSC) was able to field its new 100-plus-cabinet

Spider Will Free Users From Web of Time-Consuming Chores

NCCS users will find it easier to manage the data files from their calculations with the advent of a new center-wide shared file system that saves all files to one location.

A key initiative during 2007 was beginning the installation of a Lustre-based file system termed "Spider." It will replace multiple "islands" of file systems in various locations on the NCCS network with a single scalable file system that eventually will serve all the NCCS systems. It will connect to the InfiniBand and Ethernet internal networks. Because all simulation data will eventually reside on Spider, users will not need to transfer files among multiple computers and data management systems.

By the end of 2007, a small version of Spider with 5 gigabytes per second of aggregate bandwidth and 650 terabytes of usable space was mounted on the Lens analysis cluster, the EVEREST visualization cluster, the Smoky End-to-End cluster, Jaguar CNL, and the Ewok development cluster. In 2008, NCCS staff will deploy two systems, each with 650 terabytes of space and 15 gigabytes per second of bandwidth. These early systems will augment the Jaguar "scratch" file system or other local scratch systems.

By the end of 2008, Spider will have been expanded to support the petaflop computer. That computer will use Spider exclusively and will have no local scratch files. At that point, Spider will provide 10 petabytes of storage space and over 200 gigabytes per second of bandwidth and will be mounted on all major NCCS systems.

Deployment of Spider will be a welcome development for researchers. Having a single repository of simulation data will increase their productivity, allowing them to spend more time pursuing their research goals. By simplifying the use of the data analysis and visualization tools, it may encourage more researchers to take advantage of them and thus increase the value of their data.—*By Deborah Counce*

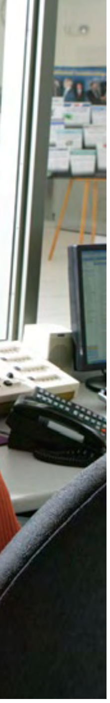

supercomputer with CNL installed instead of beginning with Catamount and switching to CNL.

Blue Gene

In October 2007, the NCCS installed a 27-teraflop IBM Blue Gene/P computer.

For ORNL staff who do not have INCITE allocations, the Blue Gene provides access to a system with thousands of processors. It has two cabinets, each with 1,024 sockets holding four 850-megahertz PowerPC 450 processors. The four processors on each socket share 2 gigabytes of memory for a system total of 4 terabytes. The sockets are interconnected by three high-performance networks that support three different types of interprocessor communication: a three-dimensional torus for point-to-point message passing, a global collective network that supports one-to-many communication, and a global interrupt network that enables low-latency barriers and interrupts. The low-latency network benefits applications for which communication between processors is important.

The IBM Blue Gene/P computer

Preparing for Petaflops

In late 2008, NCCS will install a 1 petaflop Cray XT5. It will incorporate 24 XT5 blade modules, each with eight processors, per cabinet. The system will use quad-core AMD Opteron processors and deliver a peak performance of over 1 petaflop. It will have 2 gigabytes of memory per core and an aggregate of more than twice as much memory as any other computer on the TOP500 list of supercomputers.

The cabinets for the petaflop system will use a Cray-designed R134-a refrigerant and chilled water cooling system. This is necessitated by the estimated peak power demand of more than 6.5 megawatts for this one system. Removing the heat generated by the 6.5 megawatts of electrical power flowing through the system requires substantial upgrades to the cooling system in the building. The CSB had three existing 1,200-ton chillers. Two more 1,500-ton chillers were added during 2007, bringing the total cooling capacity to 6,600 tons of chilled water. The additional 3,000 tons of cooling will be brought online in July 2008.

A 12-inch chilled water distribution line is being added under the raised floor of the computer room to circulate the chilled water. It will connect to 48 240-kilowatt-capacity Liebert XDP pumping units in the computer room. R134a refrigerant will be pumped to an evaporator located at the top of each XT5 cabinet. A high-speed fan at the bottom of each cabinet will also pull 3,200 cubic feet per minute of air through the cabinet to move heat from the processors to the evaporator, where the circulating refrigerant will absorb the waste heat. Chilled water, supplied from the CSB Central Energy Plant, will absorb the heat from the refrigerant and remove it.

The escalating power requirements of the NCCS lineup also require a more robust system for guaranteeing uninterrupted power and backup power to critical systems. An additional 1 megawatt of rotary uninterruptible power supply (UPS) was installed to supplement the existing 500 kilowatts of UPS provided by a bank of batteries. The computers themselves are not on UPS because they require more power than UPS could supply; however, the disk subsystems are on UPS to eliminate failure of the file system and the potential loss of data. A 2,200-hp generator was added to the existing backup power system to provide another 1.5 megawatts of emergency power in case of an outage. The generator power system allows essential systems to continue operation if power is interrupted.

Keeping pace with the power needs of the supercomputer center is an endless task. The peak power demand of the upgraded Jaguar is about 2.5 megawatts. When the DOE petaflop computer begins operating, it will have a peak power consumption of more than 6.5 megawatts, more than double that of Jaguar. Work was under way for most of 2007 and will continue during 2008 to supply the additional power needed to keep all those machines running.

A large-scale electrical and mechanical systems project was carried out in March 2007 to prepare the CSB for the required power and chilled water upgrades. More than 150 people, including ORNL staff and electrical and mechanical contractors, worked around the clock for 36 hours over a weekend to install electrical gear and water lines. The massive upgrade project and an associated electrical and mechanical outage required three months of planning and preparation.

Four 2.5/3.3 megavolt ampere (MVA) transformers and 400-amp switchboards will be installed during the first half of 2008 to supply power to the NCCS, and a fifth will be installed to supply power to additional equipment required for the increased computing load. The transformers will be installed inside the CSB to supply power to the DOE petascale computer.

During summer 2008, the Tennessee Valley Authority (TVA) will install a fourth 161 kilovolt (kV)/13.8-kV transformer and a new switchgear bay at its primary ORNL power substation to increase the substation transformer capacity from 210 to 280 MVA. The new transformer is scheduled to be energized by December. TVA will install a new 161-kV power line to the substation and two new 13.8-kV distribution feeds from the substation to the NCCS area during spring and summer and energize them by late summer. The power upgrades will boost the power available for the supercomputer systems from 7.3 to 20.5 MVA and the power for NCCS support systems from 8 to 11.3 MVA.

Going Green

Given the enormous energy requirements of ORNL's leadership computing systems, the NCCS is committed to operating these systems as efficiently as possible. As a result, the NCCS is among the most energy-efficient computing centers in the country.

The CSB was designed with energy efficiency in mind. Vapor barriers in its walls minimize the need to dehumidify the computer rooms. Its high-efficiency chillers are known for their energy stinginess. And the center follows a variety of strategies to keep hot air separate from chilled air, making the chillers even more efficient. As a result, the NCCS is among a very small number of data centers worldwide that have received the Leadership in Energy and Environmental Design, or LEED, certification from the U.S. Green Building Council.

As the NCCS prepares for the new petascale system, it is going even further to save energy and protect the environment. The new system will use 480-volt power instead of the standard 208-volt power, thereby reducing energy loss and saving several hundred thousand

Chilled water pipes in the NCCS's mechanical room provide cooling to the center's machines.

dollars each year in electricity bills. The system will be liquid-cooled, meaning heat will be removed from the computer room before it is mixed with cold air in the room. Even packing for the new cabinets will be minimized, with the packing for each truckload going back to be used for the next truckload, rather than going to a landfill.

Networking

A major issue for users of high-performance computers is the capability to move large data files from one system or location to another in a reasonable amount of time. For the NCCS, with its many remote users, accurate, high-speed data

HPSS: Supercomputers Get More Closet Space

The minutely detailed simulations conducted on the NCCS computers require ongoing high-speed, reliable storage and retrieval of huge amounts of data. To keep abreast of the demands of the world's most powerful collection of computing resources, the NCCS added new storage capacity in 2007 and took steps to enhance the security of the system against equipment failure.

Two Storage Tek SL8500 storage silos were added to increase data storage space for data archival. The new silos contain 32 T10000 titanium tape drives. They were installed in Building 4500 North at ORNL, in space formerly used as a computer room. Two Powderhorn storage silos remain on the second floor of the CSB, and two older silos were removed from service.

Data from simulations are first written to disks by high-speed data movers and then migrated to tape drives. Smaller files remain on disk for fast retrieval, whereas large files are removed from disk storage after they are transferred to tape.

Building in redundancy to the High-Performance Storage System (HPSS) is a current focus. NCCS staff are setting up the metadata servers as two mirrored systems so that all stored data will still be available if part of the system should fail. The core processes that make up HPSS are switching to Linux and to a different brand of server for better performance and redundancy.

New disk and tape resources are being implemented to position HPSS to handle larger files and larger file sets that come in at faster rates. "We're adding resources, changing striping, and generally reconfiguring the existing system so that data sets come in faster and go to more appropriate slots," said HPSS administrator Stan White. "As the files get bigger, the requirement is faster throughput which includes storing data in larger chunks. We're going for density—the largest amount of data in the smallest amount of space." Most storage systems have some empty space, but the NCCS is trying to minimize unused capacity, White said. Even a tiny percentage of empty space adds up in a system as large as HPSS.

At the end of 2007, HPSS was storing about 2 petabytes of data files. The amount stored has been doubling every year, but it's uncertain how adding two petascale systems will affect the volume. The capacity is expected to be 10 petabytes in 2008 and 18 petabytes by 2009.

About 2.5 terabytes of data were coming in to HPSS daily in 2007, White said. "We have had 9-terabytes days—a year earlier, we couldn't have handled that volume." He expects to be storing 40 terabytes a day in the foreseeable future. An application scheduled to run soon is expected to produce 24 terabytes of data per day, and that's in addition to several other simulations that will be running at the same time.

The average rate for writing data to HPSS is 174 megabytes per second and the average read rate is 235 megabytes per second, said White. The maximum rates, under ideal conditions, are 478 megabytes per second for writing data to the system and 663 megabytes per second for reading.—*By Deborah Counce*

transfer is essential—network capability must keep pace with computing capability. The NCCS is working to put high-throughput networks in place among the different systems in the center and between the NCCS and other research institutions to ensure that moving data from the petascale computers doesn't become a roadblock to scientific discovery.

Upgraded connections to ESNet and Internet2 dramatically boosted the speed of data transfer to and from the NCCS in 2007. A single 1-gigabit-per-second connection to the networks was replaced with two 10-gigabit-per-second connections. The several-fold increase in network speed enables researchers to move the larger data sets from the ever more powerful computers to and from their sites more easily.

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

At the end of 2007, InfiniBand connected the Lens analytic cluster, the Ewok development cluster, and the Smoky end-to-end cluster. Eventually, all major NCCS systems will be connected through InfiniBand, including HPSS. Initially the InfiniBand network will support a parallel file system with an aggregate performance of 100 gigabytes per second. Ultimately it will grow to support more than 200 gigabytes per second required for the petaflop machines.

Installation of InfiniBand infrastructure during 2007 included installing switches, running cables between systems, and extensive testing. One of the challenges of deploying the network in the NCCS is that the various systems it will support are widely separated physically. Connecting all the NCCS systems requires the use of newly developed 100-meter InfiniBand cables.

Ross Toedte, left, and Mark Dobbs convene at EVEREST to discuss a materials science visualization.

Visualization

The visualization and analysis team at the NCCS began installing new hardware in late 2007 to replace the venerable visualization cluster Hawk. The new system will enable NCCS users to analyze larger amounts of data at higher speeds and view it with higher-quality visualizations. This is especially important for users remote to ORNL, as remote visualization is often their only realistic method of analyzing large data sets.

Hawk will be replaced by three new clusters, each designed for a particular task. The Lens cluster will provide analysis and remote visualization with 512 processor cores, 2 terabytes of memory, 64 high-performance graphics cards, and a high-speed interconnect. Distributed on 32 nodes, Lens will serve most of the analysis needs of NCCS users. Lens is expected to be able to handle datasets 20 times larger and analyze them 10 times faster than Hawk. It is scheduled to enter production in May of 2008.

An indispensable research tool at the NCCS is the visualization facility that enables researchers to see and analyze the patterns in the data sets produced by their simulations.

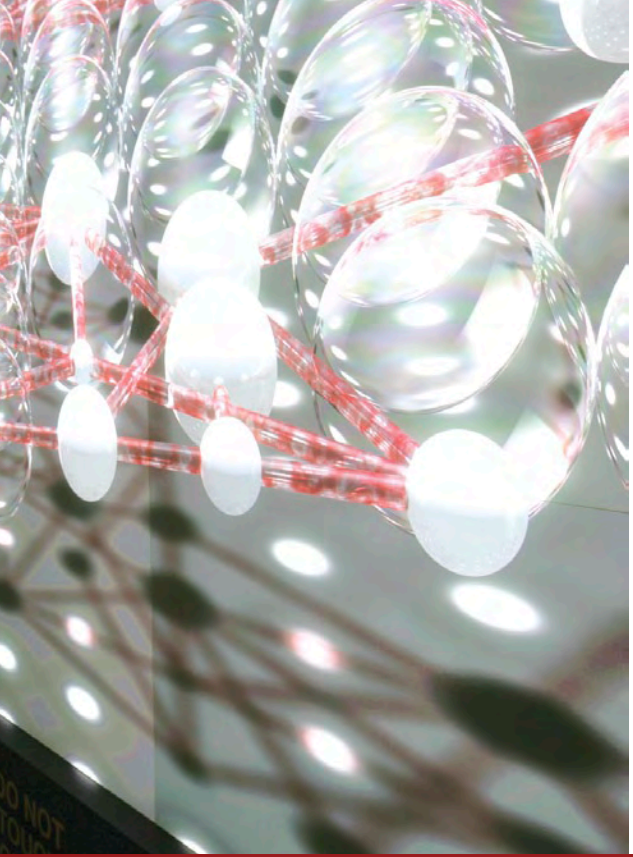

It features a 35-megapixel, 30-by-8-foot PowerWall called EVEREST (Exploratory Visualization Environment for RE-search in Science and Technology) on which 27 projectors present a bright, highly detailed display of data.

The PowerWall will be driven by the dedicated EVEREST cluster. Its 14 nodes with two graphics cards each will be directly connected to the PowerWall projectors. Each node has a high-performance disk that provides data rates capable of streaming data to the PowerWall at rapid rates.

The Orb cluster is a new parallel Lustre file system that will be used for the EVEREST PowerWall. It can speed data to Everest fast enough that full-wall movies can be displayed at 30 frames per second, compared with 2 frames per second on the Hawk cluster.—*By Deborah Counce*

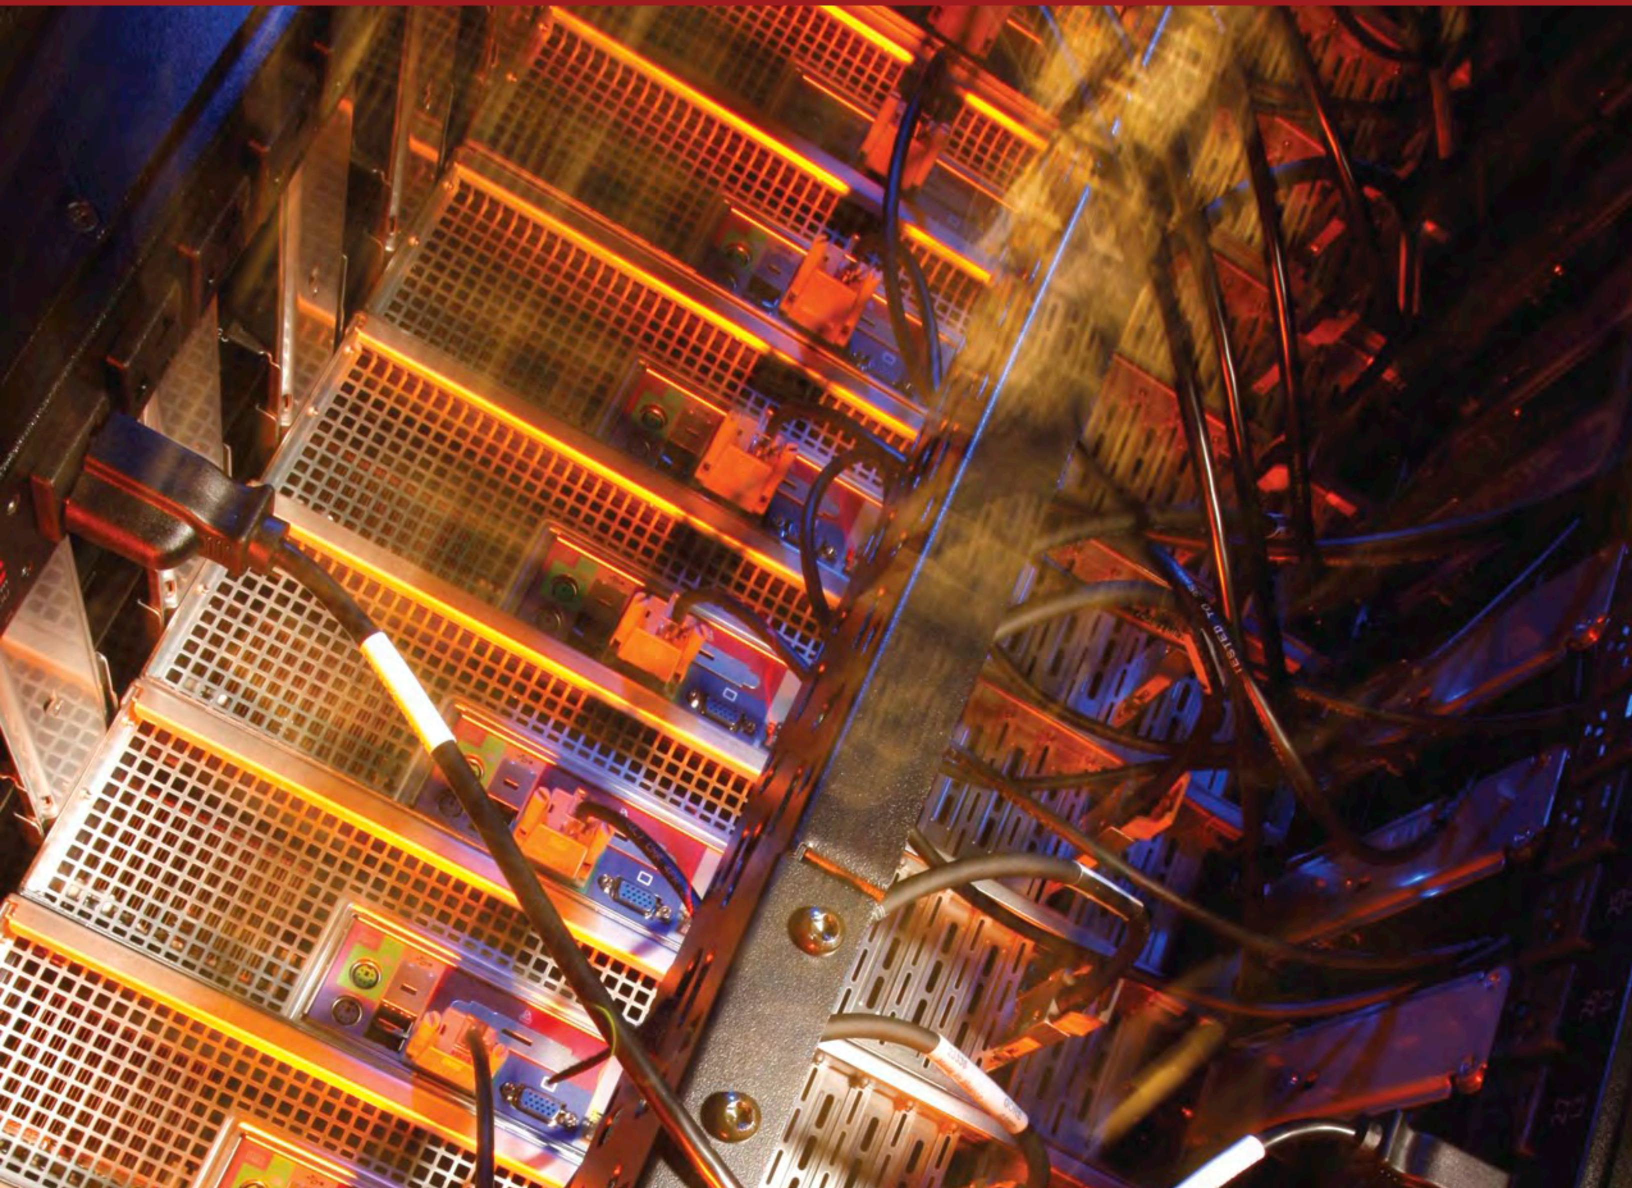

ENABLING BREAKTHROUGH SCIENCE

When ORNL astrophysicist Anthony Mezzacappa needs help honing the Chimera code, the application his team uses to simulate core-collapse supernovas, Scientific Computing Group member Bronson Messer is there.

Roselyne Barreto and Bronson Messer in the Scientific Computing group work with researchers such as Kwan-Liu Ma of the University of California, who created the visualization shown, to wrest the best science from NCCS supercomputers.

Messer is a Ph.D. astrophysicist. As the group's liaison to Mezzacappa's project, he contributes his scientific computing expertise from the perspective of a domain expert.

"We're lucky," Mezzacappa said. "I think the overall setup, having a liaison who's a domain scientist working with the teams, it's a wonderful idea that works, and in our case we really are lucky."

Messer bolsters Mezzacappa's group in a variety of ways. He coordinated efforts to improve Chimera's data input, allowing thousands of processors to write to a single file. He also works closely with the group to maximize Chimera's efficiency. Considering the magnitude of these simulations—Mezzacappa's project had an allocation of 7 million processor hours in 2007 and 16 million in 2008—even a seemingly small improvement can make a big difference.

"Each of our runs is a month to a month and a half," Mezzacappa explained. "If Bronson comes in and tells me, 'I did X, Y, and Z with so and so, and we just shaved 25 percent off of the run time,' that's a full week gone now because of that optimization."

Messer and his colleagues exemplify how the NCCS helps projects make the most of their time on ORNL's state-of-the-art supercomputers. Coming from a variety of fields—including materials science, quantum mechanics, fusion science, and molecular biology—group members not only provide their own expertise; they also offer projects entry into a deep and unique repository of knowledge and expertise at the center.

"In a lot of respects, I could do what I do for this project if I sat in Tony's group," Messer explained, "but because I'm a Scientific Computing Group member, I know what people are doing in the center, and I know who to go ask if something goes awry. I know who to ask for help—whether that be immediate help or other things."

"It is without question one of the very best aspects of the entire organization for computational science," Mezzacappa said. "It could very well be the best thing they have done in setting up the LCF [Leadership Computing Facility] and the NCCS—to have this Scientific Computing Group with a domain-expert liaison. That would be the last thing I would ever change in the center."—*By Leo Williams*

NCCS STAFF

It takes extraordinary skill and commitment to effectively support breakthrough science while at the same time preparing for the new systems that will keep the NCCS at the forefront of computational science. Each of the four NCCS groups—Scientific Computing, High-Performance Computing Operations, User Assistance and Outreach, and Technology Integration—added staff in 2007 as the center continued to grow, accumulating resources that will help bring the center and researchers into the era of petascale computing.

Scientific Computing Group

The Scientific Computing Group (SCG) is a collection of research scientists and experts in scientific computing, some science domains, data management, analysis, and visualization who work directly with users to help them run their applications productively. Many of the research scientists with time allocations at the NCCS are not experts in computation and need intensive assistance in developing codes for high-performance machines. SCG research liaisons have expertise both in the research fields using the NCCS and in designing and optimizing code specifically for massively parallel computers. Other members of the group ensure that scientists can organize, move, store and retrieve, and analyze their data and then see the patterns through visualization and the progress of their simulations via a computational workflow.

"Most of our users aren't computer scientists. They're experts in biology, chemistry, physics—not in writing code to run efficiently on tens of thousands of processors. A big part of our job is filling that gap, supplying the computational and computer science expertise that enables them to make the fullest use of the NCCS computers to solve their scientific problems.

We know researchers want to spend their time doing research, not doing administrative chores. We are automating as many of the routine, repetitive data management tasks as we can so the users can easily do their computations and also concentrate on analyzing the results of their simulations.

A scientist can't look at trillions of separate data points on a spreadsheet and see the patterns. That's what visualization does—it presents a picture of what the data show; it lets the researchers see what's happening in their simulations."

—Ricky Kendall, Group Leader

User Assistance and Outreach Group

> This is the most powerful computational resource that our users will ever access. We make their transition to running simulations as seamless as possible, and celebrate their achievements by effective communications with research communities and the public.
> —Julia White, Group Leader

The User Assistance and Outreach (UAO) Group promotes and manages front-line support for NCCS users and coordinates internal and external communications activities for the center. It assists the users with technical issues, such as porting and troubleshooting their codes, and organizational tasks such as managing their accounts and allocations. UAO staff serve as the voice of the users at in-house meetings and maintain avenues of communication between the NCCS and users, such as mail lists, the NCCS web site, workshops, and the annual users meeting. UAO also publishes printed and electronic media aimed at the scientific community and the public to foster greater understanding of the role—and impact—of the center and its users.

High-Performance Computing Operations Group

The High-Performance Computing (HPC) Operations Group keeps the NCCS computers up and running 7/24/365. When a new system arrives at the NCCS, HPC Operations staff execute the rigorous acceptance testing necessary to bring it into full operation. All systems are monitored around the clock to ensure optimal operation and to spot potential problems or any suspicious activity that might indicate a cyber security breach. The staff also handle configuration management during the lifetime of each system, tracking every change and upgrade.

> "Our goal is to run a world-class center for high-performance computing. We are aware of the enormous value of our resources to our users, so we strive for excellence in system performance, availability, and service."
> —Ann Baker, Group Leader

Technology Integration Group

Members of the Technology Integration Group provide, identify, develop, and integrate new technologies into the NCCS to improve efficiency and support user requirements. They are responsible for keeping the entire infrastructure of internal and external networks, file systems, data storage, cyber security, and application management tools at the NCCS at the forefront of available technology. Their competencies include Infiniband, MPI, MPI-IO, the HPSS, system programming, parallel file systems, Linux, Linux kernel, and a variety of tools. The Technology Integration group provides the expertise to implement new infrastructures such as Spider and Infinband at the NCCS.

"We develop and maintain the infrastructure that connects the different parts of the NCCS—the file systems, the storage systems, the networks that allow users to move their data from one place to another. We manage the data superhighways within the NCCS and between it and other institutions.

We're out in front of the technology. Because we're doing things no one has done before, the tools our users need may not exist until we invent them."
— Shane Canon, Group Leader

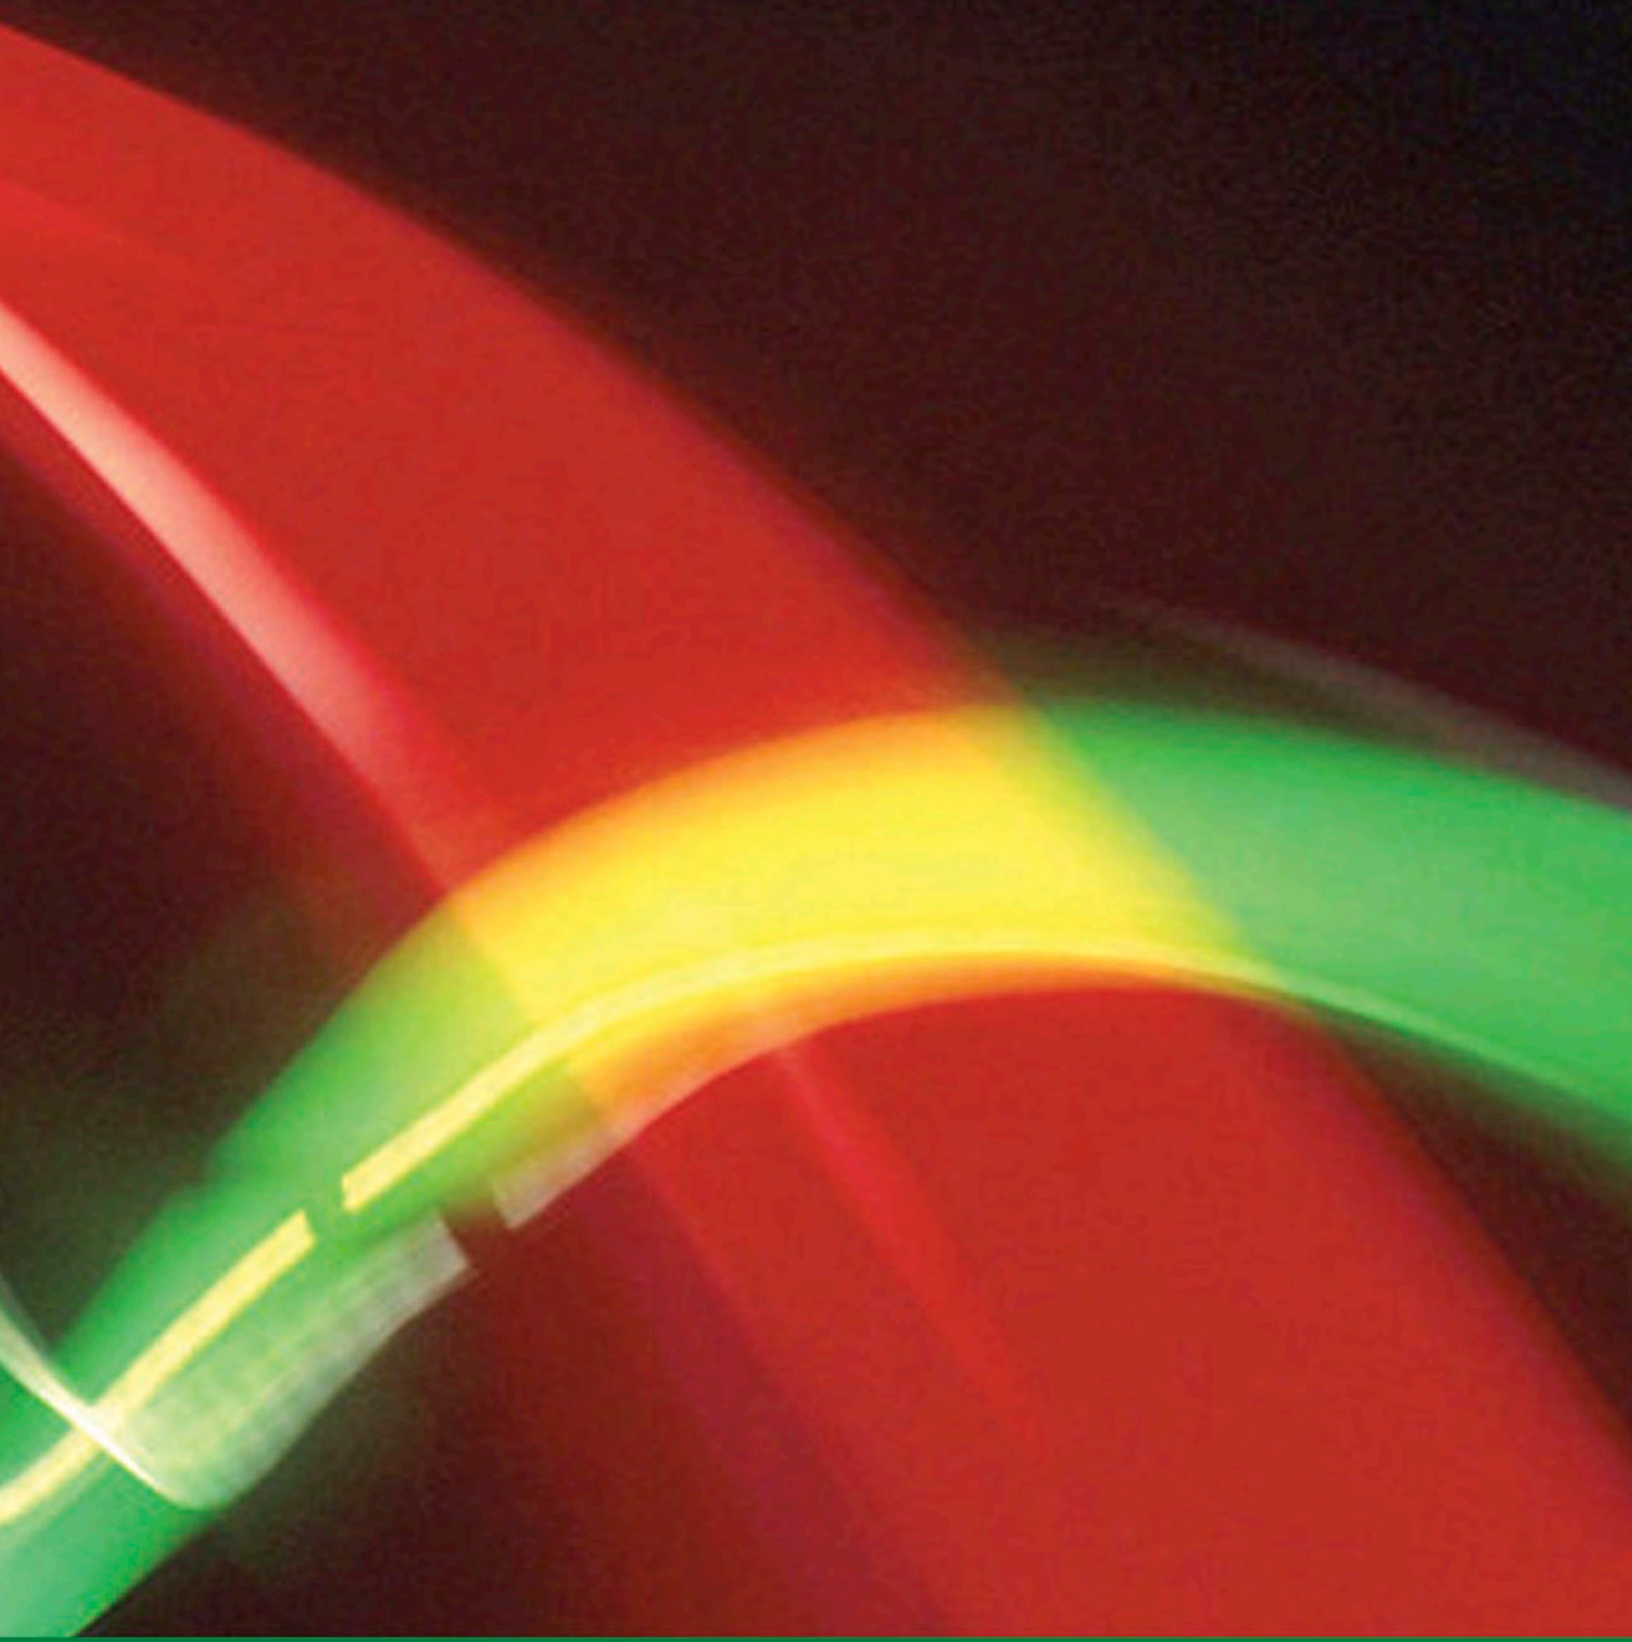

Looking Forward

Petascale Frontier...

The era of petascale computing is upon us. With this new frontier comes unexplored territory in every area of computational science and simulation. The potential of petascale systems (capable of 1,000 trillion calculations per second) to resolve some of science's most urgent issues is profound. For example, it will now be possible to simulate the Earth's climate with unparalleled precision, helping us understand our impact on the planet and more effectively battle global warming. Our newfound ability to model the exploding stars that litter the universe with the building blocks of life will reveal some of the universe's deepest secrets. Fusion energy, with its promise of a cleaner, more abundant source of electricity, will come one step closer to reality. And researchers will be able to model resistance-free superconductors like never before, propelling technologies such as magnetic levitation trains and MRI machines. The possibilities are virtually limitless, and the NCCS will lead the way by providing the systems that enable this research. The new Cray petascale system will be the first of its kind for open scientific research, featuring liquid-cooled cabinets and quad-core Barcelona 2.3-gigahertz processors. This revolutionary computing power will give scientists the tools they need to tackle science's greatest challenges and improve the world we live in.—*By Scott Jones*

Image courtesy of CRAY, Inc

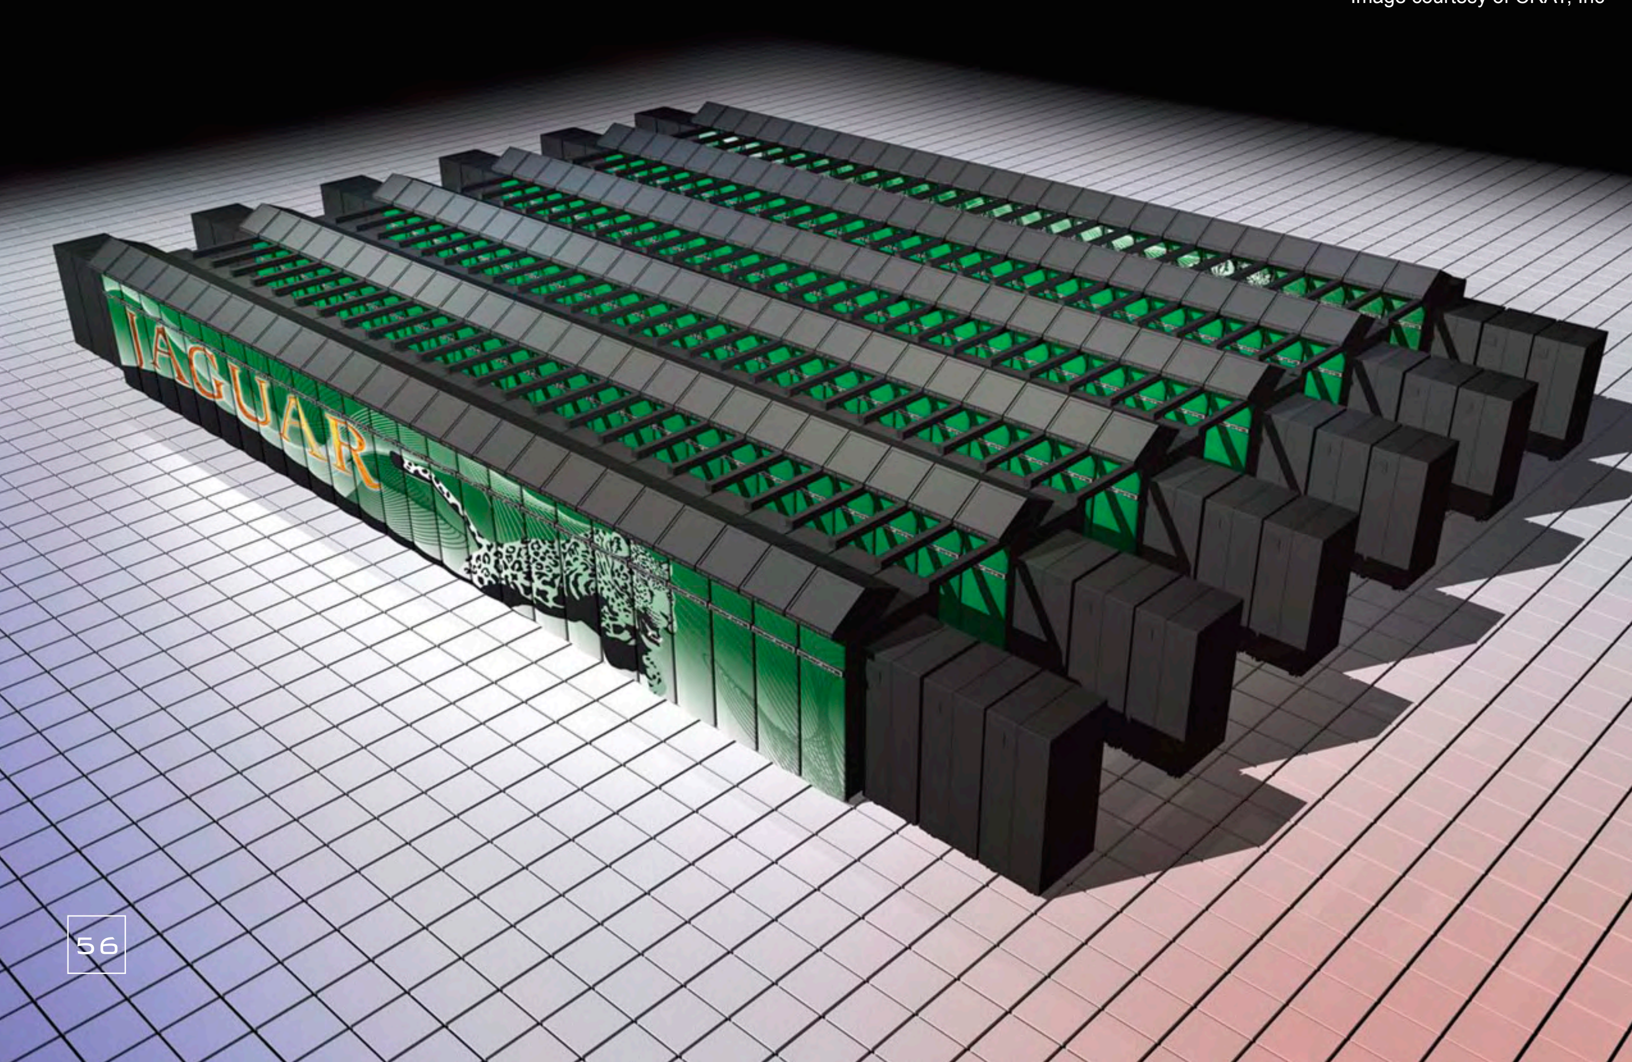

2008 INCITE PROGRAM

Office of Science INCITE Advancing America's Science and Industrial Competitiveness

Thirty projects have been awarded more than 145 million processing hours on supercomputers at ORNL by DOE's 2008 INCITE program.

Through INCITE, researchers from industry, academia, and government research facilities receive access to computing power at the NCCS to study climate change, fusion energy, nanoscience, materials, chemistry, astrophysics, and other areas. The 2008 INCITE program showcases numerous projects that explore cleaner, alternative energy sources and the effects of existing ones on our planet.

For example, computational biologists aim to discover a more efficient process for cellulose-to-ethanol conversion, which could reduce our dependence on oil, a finite fossil fuel. In a similar automotive vein, a team of mechanical engineers led by Jacqueline Chen of Sandia National Laboratories is modeling combustion in diesel engines in hopes of creating more efficient, cleaner-burning engines and boilers. Fusion researchers are using the NCCS's leadership computing systems to model heat and particle loss in a fusion reactor (such as the upcoming ITER reactor in France), a device that, if perfected, could revolutionize the way the world meets its energy demand.

All of these projects have the potential to greatly alter many of our most harmful energy-related habits, which in turn could alter the Earth's climate. And just as climate change is receiving more and more attention worldwide, it will likewise have an increased presence in the INCITE portfolio. Leading researchers from NCAR and other institutions will continue to advance tools that predict climate change with unprecedented accuracy. Other related projects will dissect the role of ocean currents in regulating climate, examine the relationship between carbon dioxide and abrupt climate change, and explore the feasibility of storing carbon dioxide underground.

The 2008 INCITE program will also continue to reflect the critical research role played by American industry. General Motors will explore materials at the nanoscale, searching for ways to convert a vehicle's waste heat into usable electricity. Researchers from the Boeing Company will continue using NCCS systems to create next-generation tools for designing aircraft. Physicists from General Atomics will continue their examination of turbulence in fusion tokamaks, further advancing the promise of fusion energy.

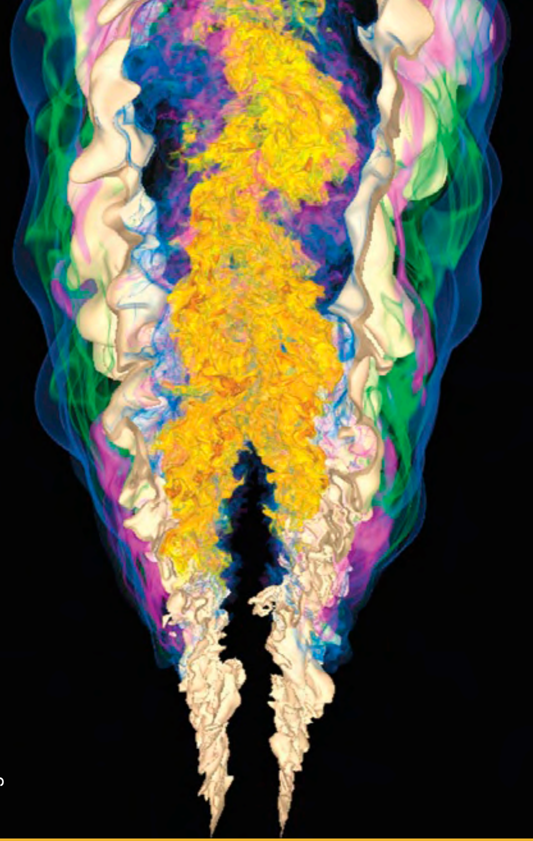

Image credit: K. L. Ma of SciDAC Ultrascale Visualization Institute

The 2008 allocations will showcase an aggressive program of upgrades to the Cray XT4 Jaguar supercomputer. When the current upgrade project is complete, Jaguar's 31,000-plus processing cores will be capable of 263 trillion calculations a second, or 263 teraflops—more than double its peak performance a year ago.

The INCITE program and recent hardware upgrades will give researchers an invaluable opportunity to continue pushing the boundaries of knowledge, and their efforts promise to improve both our lives and our understanding of the world in which we live. Furthermore, they have the potential to add to the impressive list of computational scientific achievements that took place at ORNL in 2007, when astrophysicists from ORNL and North Carolina State University released the first explanation for the spin of a pulsar that matches observation, publishing their findings in the preeminent journal *Nature*.

To read about all the 2008 INCITE awards, go to the DOE Office of Science home page at *www.science.doe.gov.*
—By Scott Jones

2008 INCITE

ASTROPHYSICS

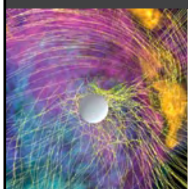

Multidimensional Simulations of Core Collapse Supernovae

Oak Ridge National Laboratory
Anthony Mezzacappa, Jaguar: 16,000,000 hours

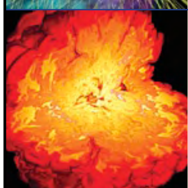

First Principles Models of Type Ia Supernovae

University of California, Santa Cruz
Stan Woosley, Jaguar: 3,500,000 hours

Numerical Relativity Simulations of Binary Black Holes and Gravitational Radiation

National Aeronautics and Space Administration
Joan Centrella, Jaguar: 1,000,000 hours

CLIMATE

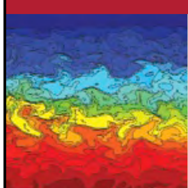

The Role of Eddies in the Meridional Overturning Circulation

University of California, San Diego
Paola Cessi, Phoenix: 486,000 hours

Assessing Global Climate Response of the NCAR-CCSM3: CO_2 Sensitivity and Abrupt Climate Change

University of Wisconsin, Madison
Zhengyu Liu, Phoenix: 420,000 hours

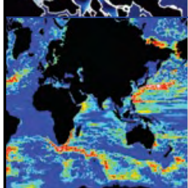

Eulerian and Lagrangian Studies of Turbulent Transport in the Global Ocean

ASC/Alliance Flash Center, University of Chicago
Synte Peacock, Jaguar: 3,163,000 hours

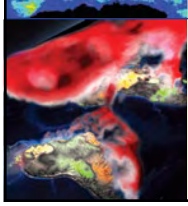

Climate-Science Computational End Station Development and Grand Challenge Team

National Center for Atmospheric Research
Warren Washington, Jaguar: 15,718,000 hours

COMPUTER SCIENCE

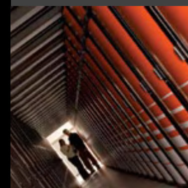

Performance Evaluation and Analysis Consortium End Station

Oak Ridge National Laboratory
Patrick Worley, Jaguar: 4,000,000 hours

FUSION

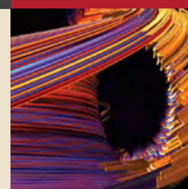

Verification and Validation of Petascale Simulation of Turbulent Transport in Fusion Plasmas

University of California, San Diego
Patrick Diamond, Jaguar: 8,000,000 hours

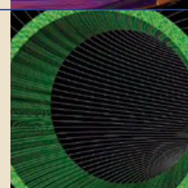

Fluctuation Spectra and Anomalous Heating in Magnetized Plasma Turbulence

University of Maryland
William Dorland, Jaguar: 4,000,000 hours

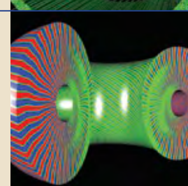

Gyrokinetic Steady-State Transport Simulations

General Atomics
Jeff Candy, Jaguar: 1,500,000 hours

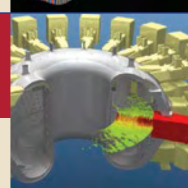

High-Power Electromagnetic Wave Heating in the ITER Burning Plasma

Oak Ridge National Laboratory
E. Fred Jaeger, Jaguar: 1,000,000 hours

ENGINEERING

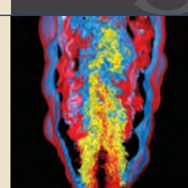

High-Fidelity Simulations for Clean and Efficient Combustion of Alternative Fuels

Sandia National Laboratories
Jacqueline Chen, Jaguar: 18,000,000 hours

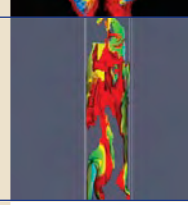

Clean and Efficient Coal Gasifier Designs Using Large-Scale Simulations

National Energy Technology Laboratory
Madhava Syamlal, Jaguar: 3,000,000 hours

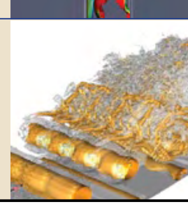

Landmark Direct Numerical Simulations of Separation and Transition for Aerospace-Relevant Wall-Bounded Shear Flows

University of Arizona
Hermann Fasel, Phoenix: 400,000 hours

High-Performance Computing Projects

PHYSICS

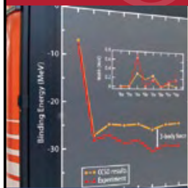

Computational Nuclear Structure

Oak Ridge National Laboratory
David Dean, Jaguar: 7,500,000 hours

Petascale Computing for Terascale Particle Accelerator: International Linear Collider Design and Modeling

Stanford Linear Accelerator Center
Lie-Quan Lee, Jaguar: 4,500,000 hours

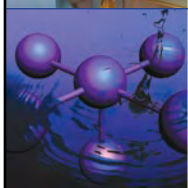

Lattice QCD

University of California, Santa Barbara
Robert Sugar, Jaguar: 7,100,000 hours

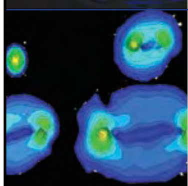

Computational Atomic and Molecular Physics for Advances in Astrophysics, Chemical Sciences, and Fusion Energy Sciences

Auburn University
Michael Pindzola, Phoenix: 2,000,000 hours

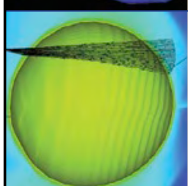

Modeling Heliospheric Phenomena with an Adaptive, MHD-Boltzmann Code

University of California, Riverside
Nikolai Pogorelov, Jaguar: 850,000 hours

MATERIALS

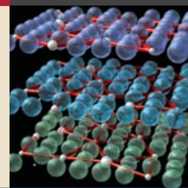

Predictive and Accurate Monte Carlo-Based Simulations for Mott Insulators, Cuprate Superconductors, and Nanoscale Systems

Oak Ridge National Laboratory
Thomas Schulthess, Jaguar: 10,000,000 hours

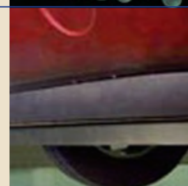

Electronic, Lattice, and Mechanical Properties of Novel Nano-Structured Bulk Materials

GM R&D Center
Jihui Yang, Jaguar: 10,000,000 hours

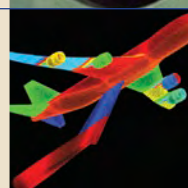

Development and Correlations of Large-Scale Computational Tools for Flight Vehicles

Boeing Company
Moeljo Hong, Jaguar: 100,000 hours
Phoenix: 300,000 hours

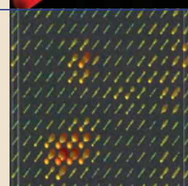

Bose-Einstein Condensation vs. Quantum Localization in Quantum Magnets

Max-Planck Gesellschaft
Tommaso Roscilde, Jaguar: 1,200,000 hours

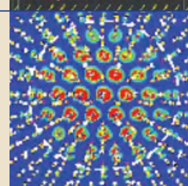

Linear-Scale Electronic Structure Calculations for Nanostructures

Lawrence Berkeley National Laboratory
Lin-Wang Wang, Jaguar: 2,100,000 hours

CHEMISTRY

Molecular Simulation of Complex Chemical Systems

Pacific Northwest National Laboratory
Christopher Mundy, Jaguar: 750,000 hours

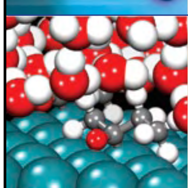

An Integrated Approach to the Rational Design of Chemical Catalysts

Oak Ridge National Laboratory
Robert Harrison, Jaguar: 10,000,000 hours

BIOLOGY

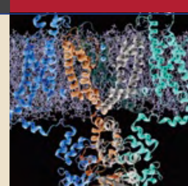

Cellulosic Ethanol: Physical Basis of Recalcitrance to Hydrolysis of Lignocellulosic Biomass

Oak Ridge National Laboratory
Jeremy Smith, Jaguar: 3,500,000 hours

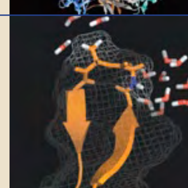

Gating Mechanism of Membrane Proteins

Argonne National Laboratory and University of Chicago
Benoit Roux, Jaguar: 3,500,000 hours

GEOSCIENCES

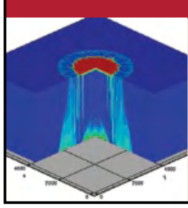

Modeling Reactive Flows in Porous Media

Los Alamos National Laboratory
Peter Lichtner, Jaguar: 1,800,000 hours

Appendices

The NICS is housed in the JICS facility on the ORNL campus.

The NCCS collaborates with and receives advice from a number of affiliate groups, including the Joint Institute for Computational Sciences (JICS), an external advisory committee, and an operations council.

NCCS Partnerships

Sean Ahern of NCCS confers with attendees at the NICS dedication ceremony.

Joint Institute for Computational Sciences

JICS was established in 1991 through an agreement with UT's Science Alliance (a UT Center of Excellence) and ORNL. Its purpose was to advance scientific discovery and the knowledge of computational modeling and simulation by making full use of the capabilities of the computers at ORNL and educating the next generation of researchers in using computational techniques to address scientific problems. From this beginning, JICS has expanded to include seven regional core universities.

When UT-Battelle, LLC, won the contract to manage ORNL, the relationship between ORNL and JICS expanded. The state of Tennessee built a 52,000-square-foot building on the ORNL campus across the street from the NCCS to house JICS. The new building, completed in the summer of 2002, has office space for 70 staff members as well as lab and classroom space. It provides five incubator suites, each about 1,600 square feet, that can be used as classrooms, conference rooms, or computer laboratories by ORNL staff, visiting faculty, and graduate students working together on computational research projects.

The year 2007 was a big year for JICS, as it welcomed the National Institute for Computational Sciences (NICS). Formed through a $65 million NSF grant to UT and its partners, NICS is housed in the JICS facility on the ORNL campus and will soon deliver a petascale computer system that promises to make substantial contributions in many scientific arenas, including climate change, astrophysics, alternative energy, and biomolecular science.

The arrival of NICS and its flagship computer, Kraken, will make ORNL the leading scientific computing center in the United States. Kraken, a Cray XT4 currently housed in the Computational Sciences Building at ORNL, came online in early 2008 and will eventually evolve into a petascale system.

At that point, ORNL, in collaboration with UT and NICS, will have more computing power than any other high-performance computing center as the NCCS likewise introduces a petascale system. This will result in two petascale machines under one roof. Along with other supercomputers at ORNL, the two petascale systems will contribute to over 3 petaflops of computing muscle available to researchers, allowing them to tackle tomorrow's most daunting scientific challenges.

External Advisory Committee

The NCCS Advisory Committee is composed of 12 to 18 distinguished scientists from academia, national laboratories, industry, and other research institutions. The committee provides advice to the NCCS director in the areas of computational science, computer science, applied mathematics, operation of a national user facility, and interagency communication and coordination.

Operations Council

The NCCS is one of 18 major user facilities at ORNL. The mission of the Operations Council is to ensure that the center operates in a safe, secure, and effective manner. The council is chaired by the director of operations and meets weekly to discuss current operational status, concerns, activities, and future direction.

The Research Alliance in Math and Science

The Research Alliance in Math and Science (RAMS) program provides opportunities for students and faculty in scientific and technical disciplines at U.S. colleges and universities to gain experience by collaborating in research with ORNL scientists. The program's goals are to improve U.S. competitiveness in research and to increase the representation of individuals from minority groups among holders of advanced degrees in science, mathematics, engineering, and technology. RAMS is based on the belief that cooperation between national labs and universities is the best way to build a well-qualified and diverse 21st-century workforce.

RAMS brings graduate and undergraduate students and faculty to ORNL each summer for an internship of 10 to 12 weeks. Each student is assigned a research mentor with whom to work on a project of interest to the student, the student's professor, and the ORNL researcher. Students are required to maintain daily journals of activities and experience, attend weekly technical seminars, attend workshops to build their skills, write summary papers discussing their research projects, give oral presentations of research results, and present the projects in poster sessions and meetings at ORNL and national conferences. Academic credit for the internship can be arranged through the college or university.

The development and expansion of educational relationships with historically black colleges and universities and other minority educational institutions are conducted through the Computing and Computational Sciences Directorate at ORNL. The RAMS program is sponsored by DOE's Office of Advanced Scientific Computing Research.—*By Scott Jones*

RAMS students prepare for a tour of the ORNL campus.

NCCS Outreach

In addition to facilitating today's top science, the NCCS is dedicated to providing the public with a robust, multidimensional outreach program that both encourages and educates.

An integral part of the outreach program is the NCCS's dedication to and interaction with a broad range of students, from the high school to university levels.

Education

The NCCS was involved in multiple educational initiatives aimed at expanding the horizons of current researchers and tomorrow's computational scientists. One such event was the Day of Science, an annual event sponsored by DOE and ORNL. It has grown to be one of the largest DOE educational events and in 2007 included more than 1,200 students from 125 colleges and universities.

The 2007 Day of Science was three times larger than in any previous year, giving future scientists a glimpse at the numerous career opportunities available in computational science. While many other nations are witnessing increases in their numbers of young scientists, America is witnessing a steep decline, a trend the Day of Science hopes to counter. Participants were shown first-hand the latest developments in green energy, nanoscience, fusion, and astrophysics.

ORNL also participated in a series of distance lectures for students at Georgia Tech University. Those attending were invited to learn about high-performance computing (HPC) via video from experts at the NCCS and ORNL, including Associate Laboratory Director for Computing and Computational Sciences Thomas Zacharia and LCF Project Director Buddy Bland, among others.

At the 2007 Day of Science, ORNL's Debbie McCoy speaks with students about opportunities in computational sciences.

In June, about 40 graduate students and undergraduates—pursuing careers in math, chemistry, physics, biology, and computer science—gathered at JICS to get a taste of leadership computing on the NCCS's IBM Cheetah system. "A Crash Course in Supercomputing," taught by Rebecca Hartman-Baker of ORNL's Computer Science and Mathematics Division, introduced students to the UNIX operating system and tools for writing the highly parallel programs required by modern high-performance computers.

The NCCS also hosts a seminar series that brings speakers to ORNL to share their research and expertise and interact with NCCS staff. The series provides an opportunity for collaboration building and helps to communicate NCCS objectives and accomplishments throughout ORNL. The NCCS hosted a total of eight seminars in the series in 2007, including Paul Bonoli of the Massachusetts Institute of Technology, who spoke on "Simulating Wave–Particle Interactions in Fusion Plasmas: Challenges and Successes," and Doug Scalapino of the University of California–Santa Barbara, whose talk was titled "The Name of the Rose—Spin Fluctuations." Other presenters included Jacqueline Chen of Sandia National Laboratories and John A. Turner of Los Alamos National Laboratory.

Finally, the NCCS sponsors interns from local high schools and area universities during the academic year and throughout the summer. Through these internships, future researchers get an up-close look at careers in science and mathematics and a head start in paving the way for the next generation of scientific progress.

Attendees at the NCCS users meeting discuss research during poster presentations.

Meetings and Workshops

As part of its outreach curriculum, the NCCS hosts and/or participates in a series of workshops aimed at educating the wider HPC community and fostering interaction and cooperation, such as the Cray Users Group and SciDAC 2007. Numerous other meetings and workshops are included in the outreach curriculum and are listed below.

Cray Technical Workshop

February 26–28

The first Cray Technical Workshop–USA, sponsored by the NCCS and the Cray User Group, took place in Nashville, Tennessee.

Attendees from the United States and Europe representing both current and prospective users of Cray XT3 and XT4 systems came from universities, national laboratories, and industry.

Cray users also gave a number of presentations. Topics included benchmarking and performance analysis of the XT3 and XT4 systems, performance analyses of applications in a wide range of scientific and engineering areas, and strategies for optimizing input and output on the XT3 and XT4.

NCCS Users Meeting

March 27–29

More than 70 NCCS users and others interested in the center gathered at ORNL to discuss the center's direction, network with staff and one another, and hone their skills.

The second annual NCCS Users Meeting allowed researchers to discuss the work they are performing on the center's state-of-the-art systems and outline their needs as the NCCS moves to petascale systems and beyond.

Users also had an opportunity to gather for a meeting of the NCCS Users Group. The group elected ORNL climate researcher John Drake as its chairman for the coming year and adopted an updated charter, which can be viewed on the NCCS Web site (*www.nccs.gov*).

SC07

November 10–16

Once again ORNL enjoyed tremendous success at SC07 (short for Supercomputing 2007), the premier international symposium for HPC, networking, storage, and analysis. The laboratory's booth featured all-electronic content displaying the latest breakthroughs in alternative energy solutions, astrophysics, climate modeling, fusion energy, overviewed educational opportunities, and previewed the future of supercomputing: petascale science.

The ORNL booth, which featured several keynote speakers, enjoyed strong attendance and introduced NICS, a collaboration established between UT and ORNL in response to the university's recent $65 million award from the NSF.

Part of the ORNL Booth at SC07

Ambassador Howard H. Baker, Jr., left, signs the Cray X1E computer as former Lab Director Jeff Wadsworth and Associate Lab Director Thomas Zacharia observe.

ORNL's Becky Verastegui was the general chair of the conference, and several ORNL researchers served in key committee positions. NCCS staff members also hosted a "birds-of-a-feather" session.

NCCS Scaling to Petaflops Workshop

July 30–August 1

Experts from the NCCS and Cray worked directly with researchers during a 3-day workshop focused on ORNL's upcoming multicore systems.

The NCCS Scaling to Petaflops Workshop, intended to help researchers use the center's enormous computing resources to deliver scientific breakthroughs, was held July 30 through August 1 at ORNL. During this hands-on event, NCCS and Cray experts met individually with computational scientists to ensure they will be able to take full advantage of upcoming Cray systems using quad-core processors.

The workshop helped them identify areas that might be bottlenecks on the upcoming architectures, as well as new algorithms that could aid in scaling applications to higher node counts. NCCS and Cray staff explained the benefits of multithreading application programming interfaces such as OpenMP and POSIX Threads, which can be used to alleviate the limited memory and network injection bandwidth for each node. They also helped users investigate ways the applications can be modified to achieve optimal I/O performance for running larger core counts.

TeraGrid Workshop

December 11–12

A contingent from ORNL attended and participated in a petascale preparation workshop titled "Building PetaScale Applications and Software Environments on TeraGrid."

The workshop, held at Arizona State University's Fulton School of Engineering, featured a series of lectures and presentations on algorithms, techniques, and performance tools.

Ricky Kendall, group leader for the scientific computing group at the NCCS and acting group leader for NICS, discussed the new NICS facility at ORNL. Bronson Messer of ORNL presented information on parallel compilers and languages.

Tours

The NCCS continued to host a variety of individual and group tours in 2007. More than 300 groups toured the facility and saw first-hand the machines that are enabling today's groundbreaking science. Some of the many dignitaries to visit the center included former Vice President Al Gore; Senator Lamar Alexander; Congressman Zach Wamp; Dennis Spurgeon, assistant secretary for nuclear energy; former senator Howard Baker; the Japan Atomic Energy Agency; ITER Director General Kaname Ikeda; a delegation from the European Parliament; and Senator Bob Corker.—*By Scott Jones*

Jim Hack, left, and Thomas Zacharia lead a tour of local high-school students.

NCCS Statistics

NCCS Usage by Program

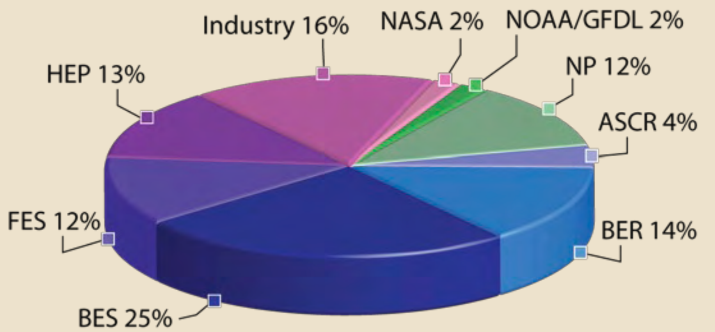

NCCS Usage by Discipline

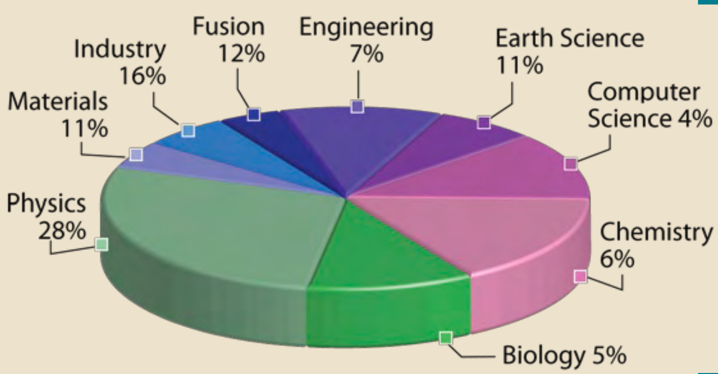

Active Users by Sponsor

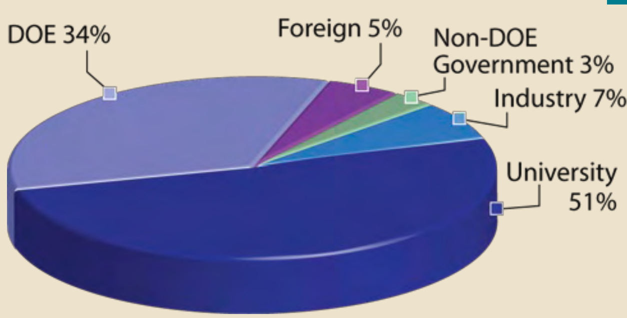

Active Users by Discipline

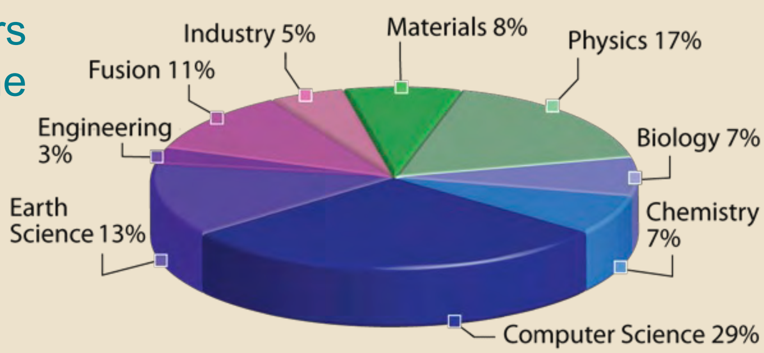

The charts on this page present a statistical overview of the NCCS in 2007: the research conducted here, the scientists conducting the research, and the organizations supporting the research. These charts show clearly that leadership computing at the NCCS spans a large range of scientific disciplines and research organizations.

Acronym Glossary

ASCR—Advanced Scientific Computing Research

BER—Biological and Environmental Research

BES—Basic Energy Sciences

FES—Fusion Energy Sciences

HEP—High Energy Physics

NASA—National Aeronautics and Space Administration

NOAA/GFDL—National Oceanic and Atmospheric Administration's Geophysical Fluid Dynamics Laboratory at Princeton University

NP—Nuclear Physics

Using the resources of the NCCS, researchers continue to produce numerous scientific breakthroughs.Listed below are a small sampling of the more than 300 publications from 2007, grouped by related discipline, that highlight a portion of the work being achieved through the combination of talented researchers, leadership-class systems, and the dedicated staff of the NCCS. For the complete list, please see *www.nccs.gov*.

Selected User Publications

Astrophysics

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

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." *The Astrophysical Journal* **660** (2): 1344–1356.

Chemistry

Yanai, T., R.J. Harrison, T. Nakajima, Y. Ishikawa, and K. Hira, 2007. "New implementation of molecular double point-group symmetry in four-component relativistic Gaussian-type spinors." *International Journal of Quantum Chemistry* **107** (6): 1382–1389.

Climate

Teng, H.Y., W.M. Washington, G.A. Meehl, L.E. Buja, and G.W. Strand, 2006. "Twenty-first century Arctic climate change in the the CCSM3 IPCC scenario simulations." *Climate Dynamics* **26** (6): 601–616.

Abram, N.J., M.K. Gagan, Z.Y. Liu, W.S. Hantoro, M.T. McCulloch, and B.W. Suwargadi, 2007. "Seasonal characteristcs of the Indian Ocean Dipole during the Holocene epoch." *Nature* **445** (7125): 299–302.

Combustion

Chen, J.H., E.R. Hawkes, R. Sankaran, S.D. Mason, and H.G. Im, 2006. "Direct numerical simulation of ignition front propagation in constant volume with temperature inhomogeneities – I. Fundamental analysis and diagnostics." *Combustion and Flame* **145** (1–2):128–144.

Hawkes, E.R., R. Sankaran, J.C. Sutherland, and J.H. Chen, 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

Nevins, W.M., S.E. Parker, Y. Chen, J. Candy, A. Dimit, W. Dorland, G.W. Hammett, and F. Jenko, 2007. "Verification of gyrokinetic delta f simulations of electron temperature gradient turbulence." *Physics of Plasmas* **14** (8).

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).

High Energy Physics

Edwards, R.G., et al., 2006. "Nucleon axial charge in full lattice QCD." *Physical Review Letters* **96** (5).

Colgan, J., M.S. Pindzola, and F. Robicheaux, 2007. "Triple differential cross sections for the double photoionization of H-2." *Physical Review Letters* **98** (15).

Materials

Schulthess, T.C., W.M. Temmerman, Z. Szotek, A. Svane, and L. Petit, 2007. "First-principles electronic structure of Mn-doped GaAs, GaP, and GaN semiconductors." *Journal of Physics-Condensed Matter* **19** (16).

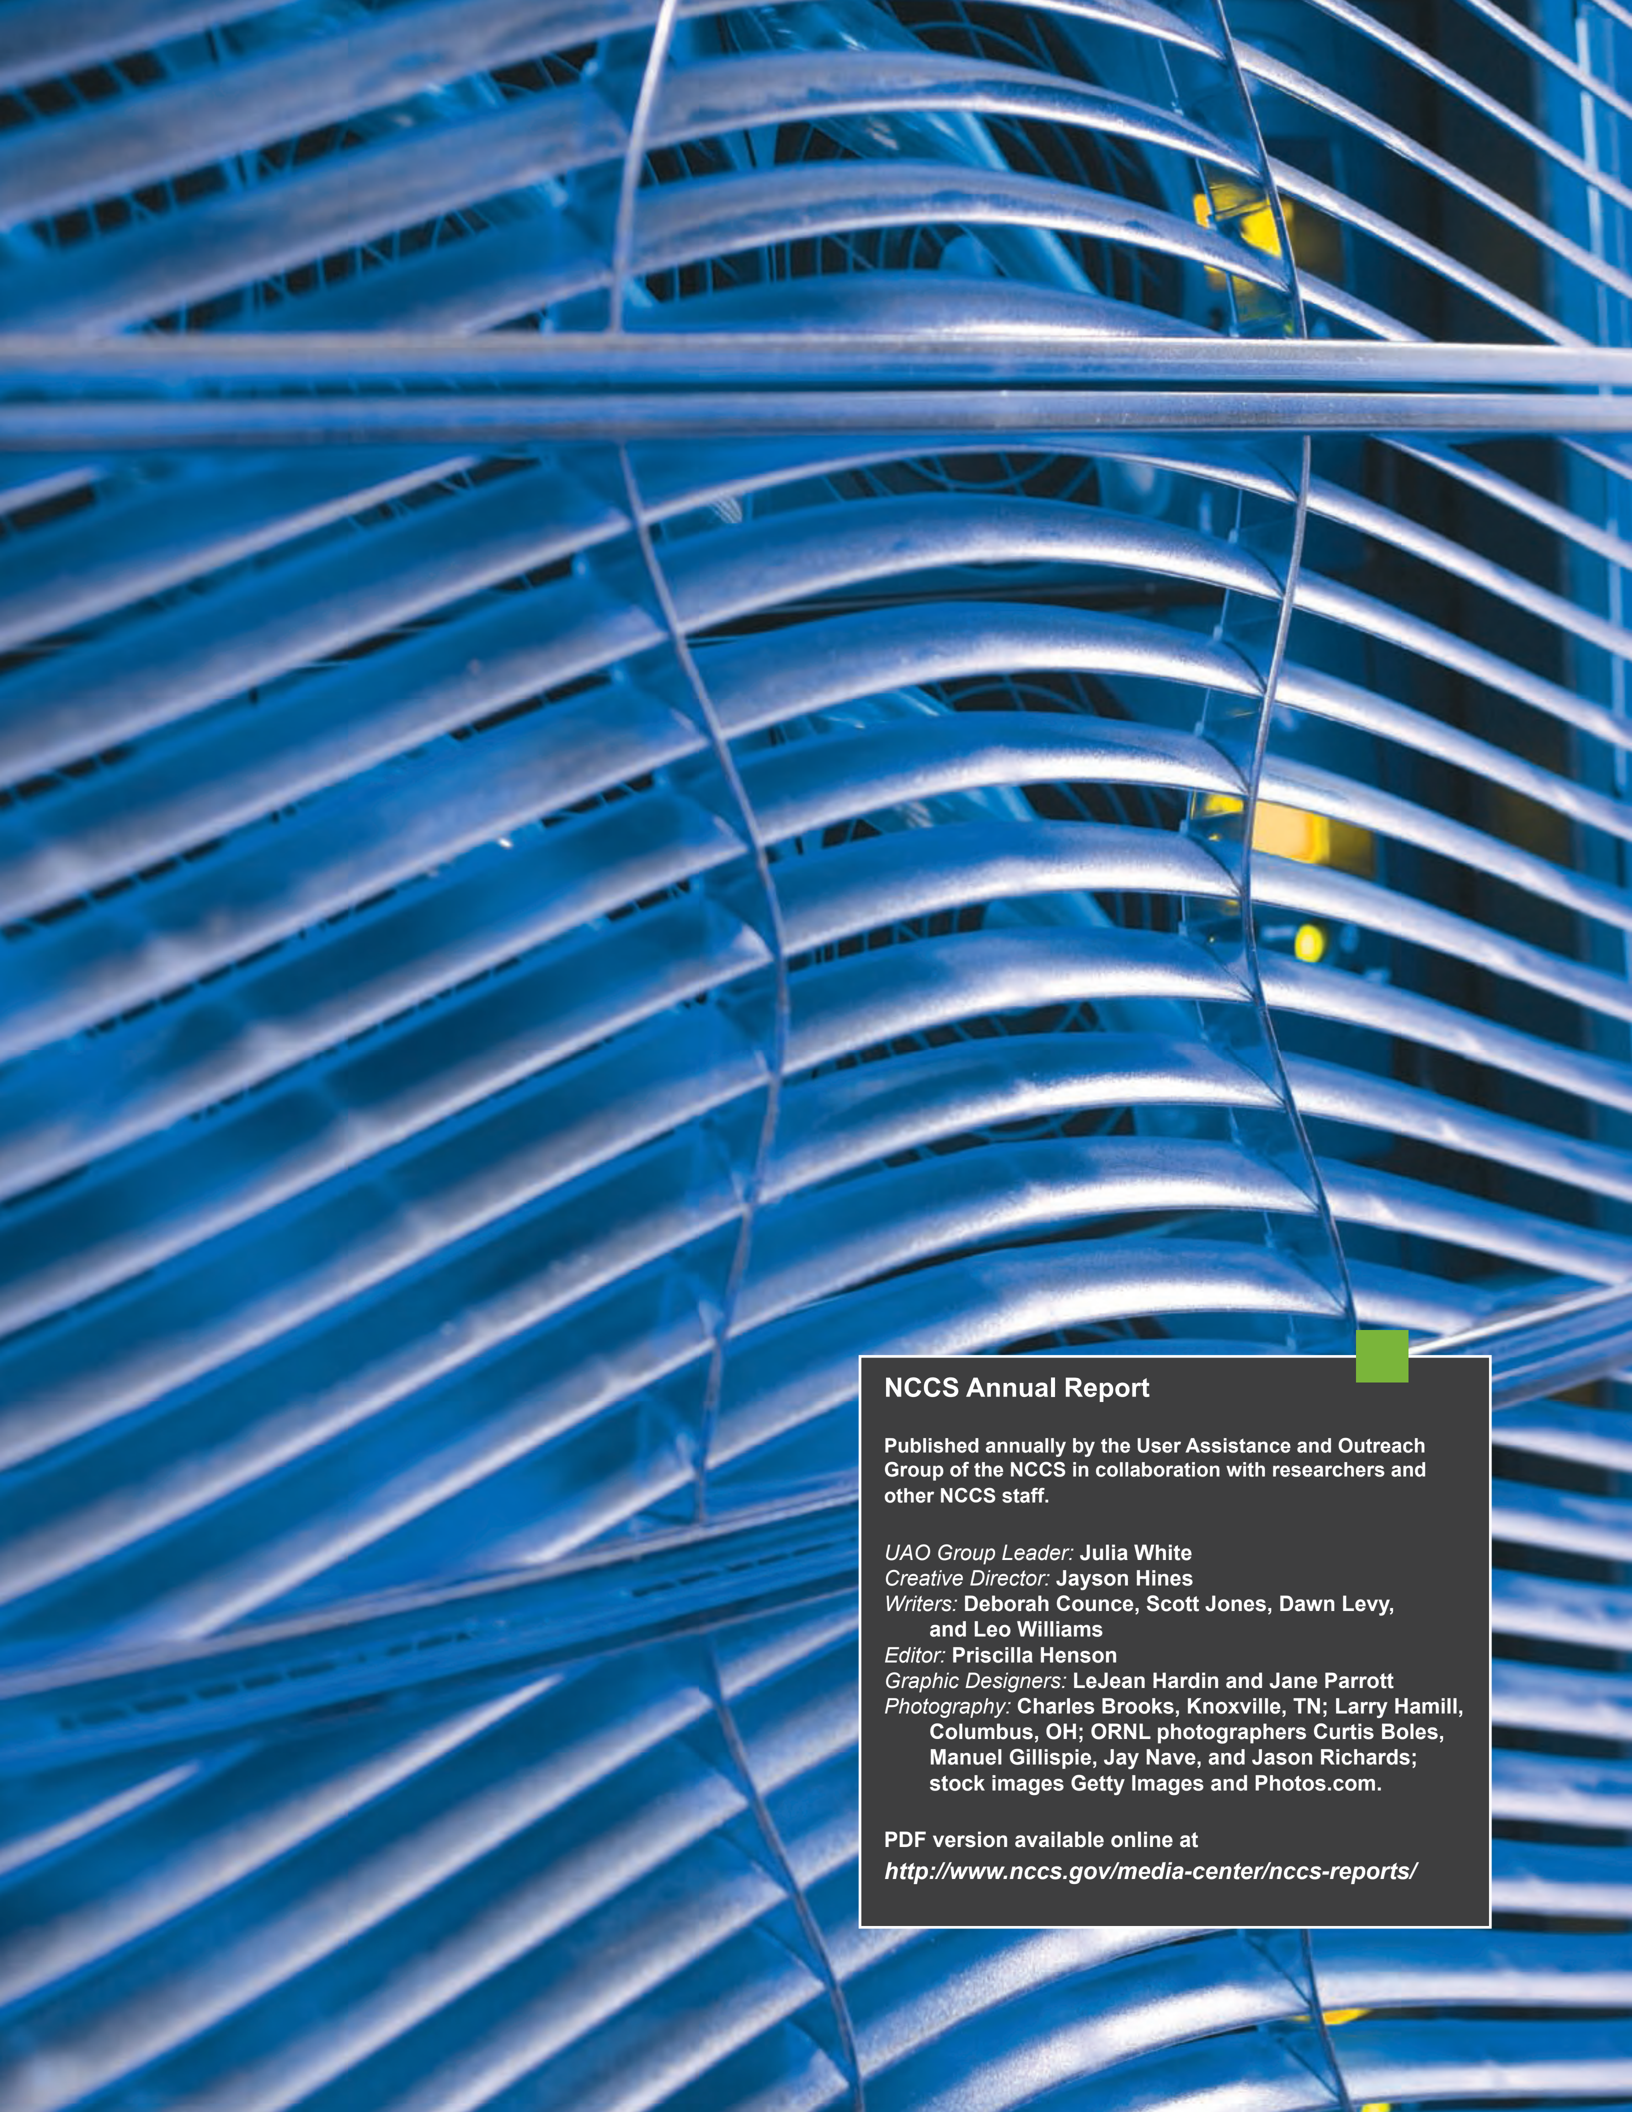

NCCS Annual Report

Published annually by the User Assistance and Outreach Group of the NCCS in collaboration with researchers and other NCCS staff.

UAO Group Leader: **Julia White**
Creative Director: **Jayson Hines**
Writers: **Deborah Counce, Scott Jones, Dawn Levy, and Leo Williams**
Editor: **Priscilla Henson**
Graphic Designers: **LeJean Hardin and Jane Parrott**
Photography: **Charles Brooks, Knoxville, TN; Larry Hamill, Columbus, OH; ORNL photographers Curtis Boles, Manuel Gillispie, Jay Nave, and Jason Richards; stock images Getty Images and Photos.com.**

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

Oak Ridge National Laboratory

ORNL 2008-G00459/jcp