ANNUAL REPORT 2006

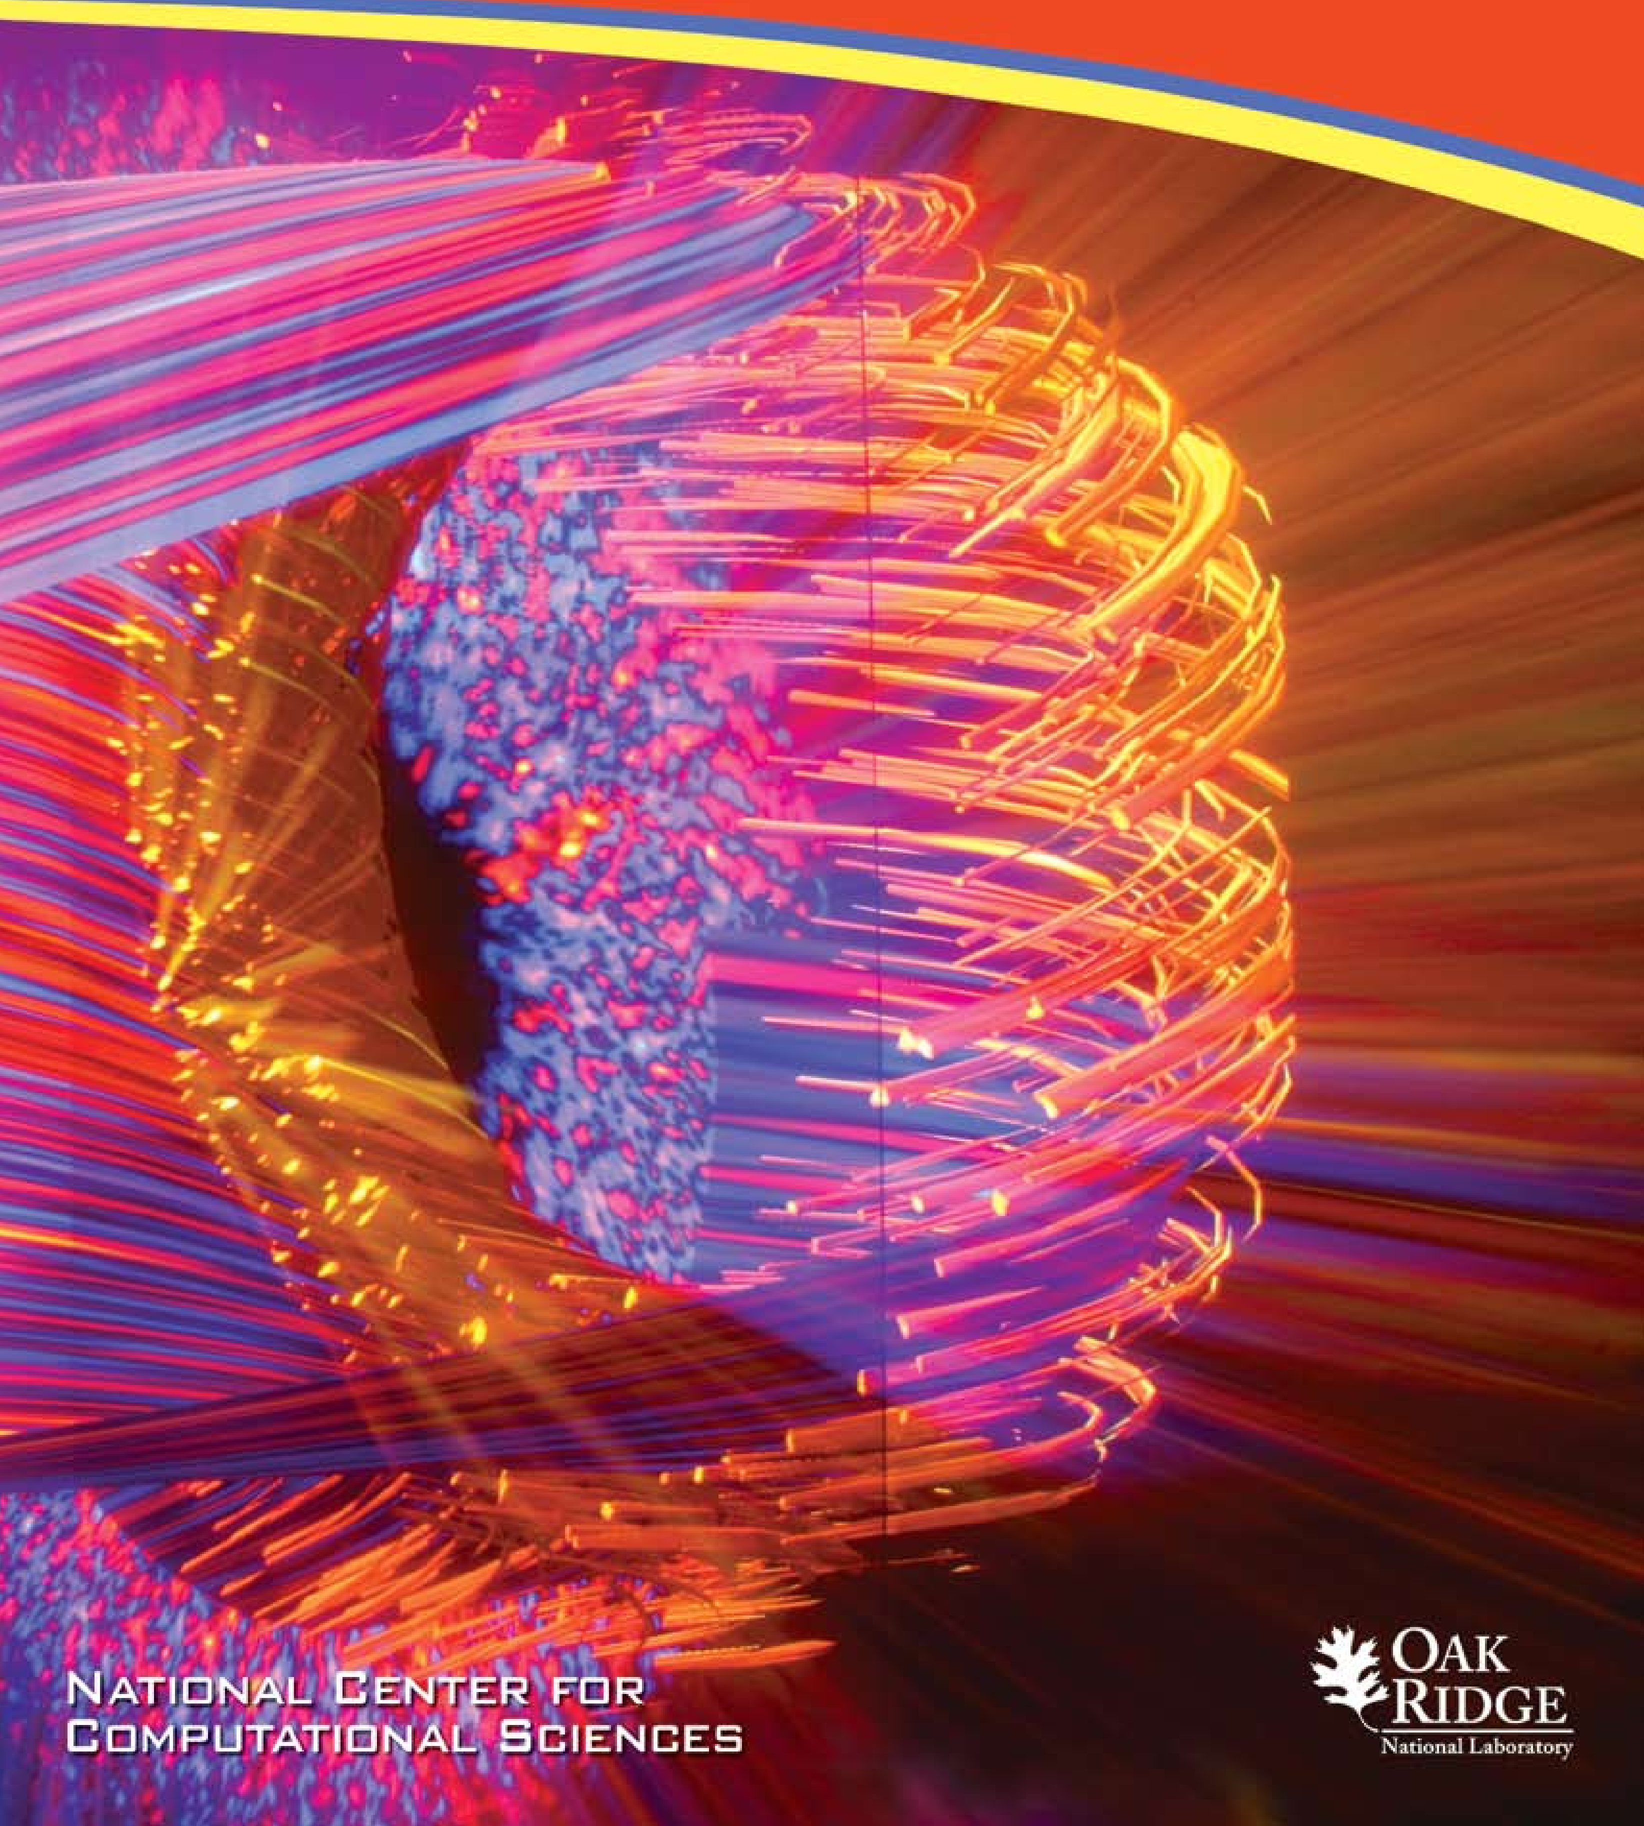

NATIONAL CENTER FOR
COMPUTATIONAL SCIENCES

OAK RIDGE
National Laboratory

CONTENTS

National Center for Computational Sciences

Located at 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 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.

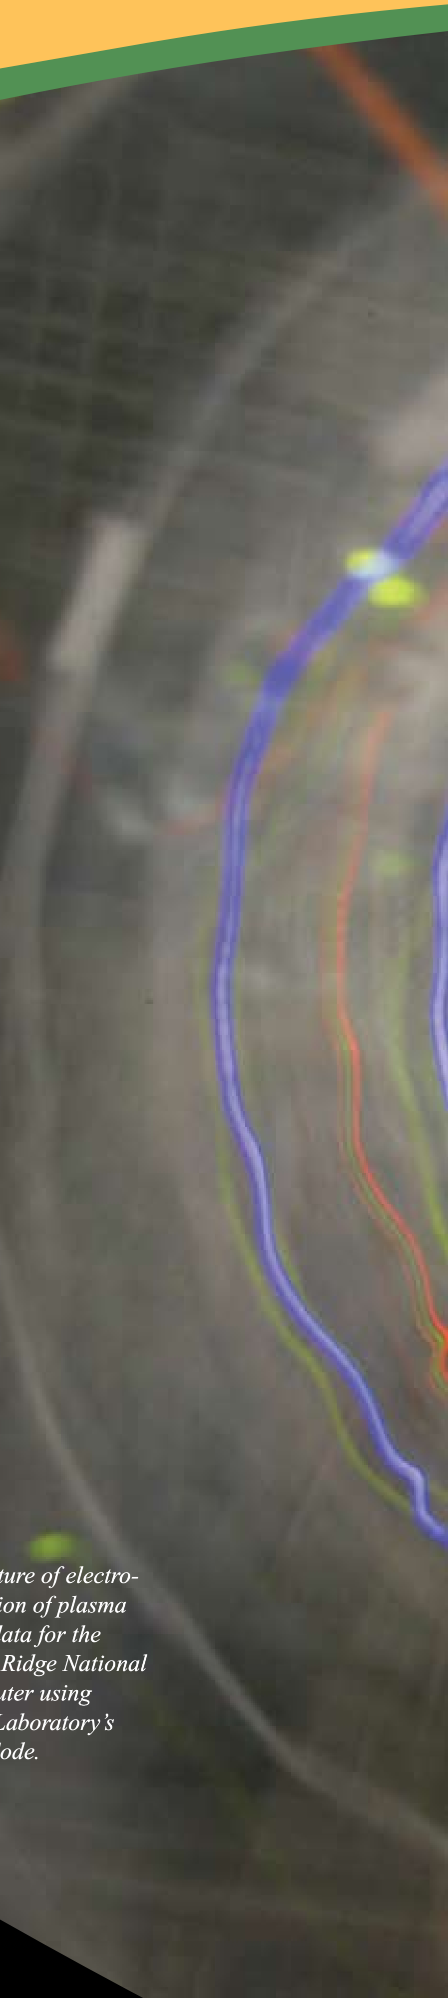

Front Cover: *This visualization shows the structure of electrostatic potential in a three-dimensional simulation of plasma microturbulence in a fusion reactor. The data for the visualization were generated on Oak Ridge National Laboratory's Jaguar supercomputer using Princeton Plasma Physics Laboratory's Gyrokinetic Toroidal Code.*

NCCS OVERVIEW

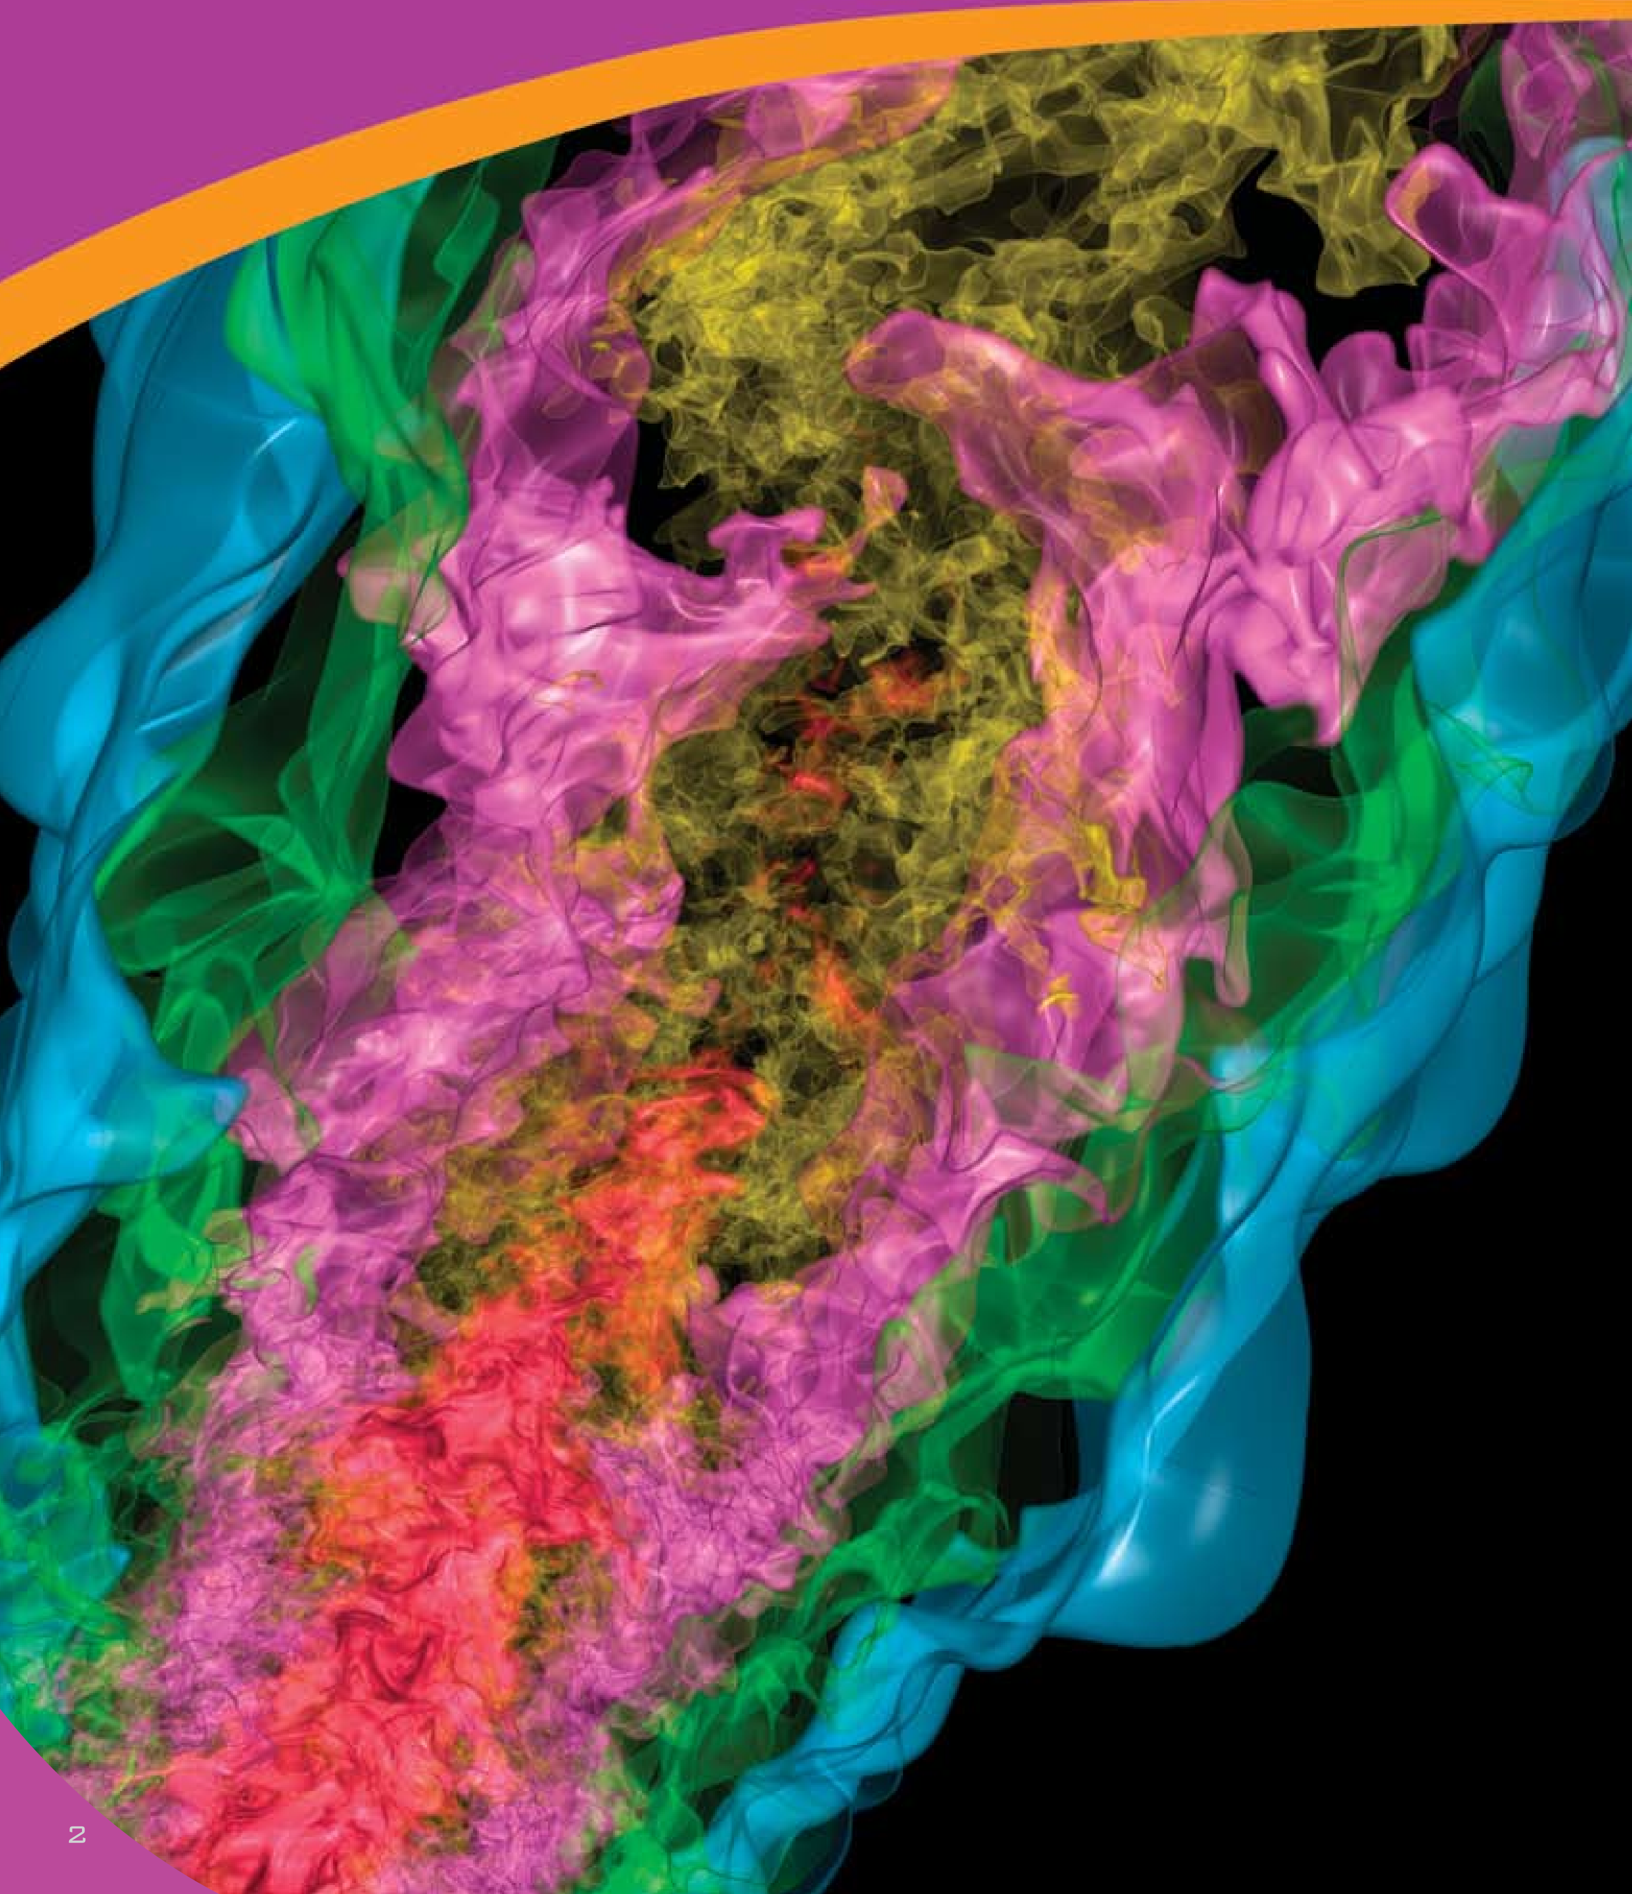

FUTURE DIRECTION OF NCCS

Supercomputers have become an integral part of our everyday lives.

It is difficult to find a jetliner or even a highway that was not designed at least in part with high-performance computing systems. In the morning, weather reports created through the courtesy of supercomputers let us know whether we need a coat or an umbrella; in the evening, containers designed with supercomputers hold the milk or detergent we bring home from the store.

The benefits of open scientific computing are also beyond question. Supercomputers at Oak Ridge National Laboratory (ORNL) and elsewhere are broadening our knowledge of the universe at a vast array of scales, from the interactions of atoms and molecules to the forces that bind galaxies. Computer simulation is accelerating the advances made in the more traditional scientific areas of theory and experiment, with these simulations often providing knowledge available in no other way.

It will come as no surprise if the next energy breakthrough or scientific revolution is built on computer simulation. There has been an explosion in the power of supercomputers, and its effects are being felt throughout the scientific world. ORNL's Jaguar Cray XT4, capable of performing more than 50 trillion calculations a second and ranked as the 10th most powerful system in the world, is 176 times more powerful than the world's fastest supercomputer of a decade ago.

Nevertheless, within 2 years the laboratory will be operating a computer 20 times again more powerful, able to perform 1,000 trillion calculations in a second. ORNL was the first institution to order a computing system able to operate at this level, known in the field as a "petaflop." But we have no illusion that this level of power will be enough. The universe is very, very complex, and scientists will need far more power to accurately simulate the forces that act on it, from the microscopic to the galactic and beyond.

Researchers in fusion science, materials science, climate prediction, and astrophysics have made it very clear how they will use the computers now in our plans, as well as the computers we have not yet even conceived. Fusion scientists will take us ever closer to a practical, virtually unlimited energy source. Materials scientists will design materials with a predetermined set of properties. Climate scientists will give us the information we need to make choices regarding how we live our lives and maintain our planet. And astrophysicists will continue to unlock the mysteries of the universe.

By design ORNL hosts only computing projects capable of producing groundbreaking science. We choose a small number of promising projects led by prominent scientists, and we give them an enormous amount of computer time to realize that promise. In doing so we are dedicated to changing the world for the better. Don't be surprised if the research that leads to that next great scientific advance is done here.

Thomas Zacharia, Associate Laboratory Director

LETTER FROM THE LCF PROJECT DIRECTOR

Buddy Bland, LCF Project Director

The NCCS's mission is to provide the most powerful open scientific computing resources in the world to a user community of scientists from government, academia, and industry and to enable breakthrough science using these systems.

The Leadership Computing Facility (LCF) continues to deliver on that mission and aggressively upgrade its computing systems. Last summer we more than doubled the computing power of our premier Cray XT3 system, known as Jaguar, to a peak performance of more than 54 trillion calculations a second (54 teraflops). The move pushed Jaguar to No. 10 on the TOP500 list of the world's fastest supercomputers, and the system is the most powerful supercomputer in the United States for open scientific applications.

As a result of these upgrades, our users have been able to leverage an amazing amount of computing power. We were able to deliver more than 37 million processor hours on our leadership systems in 2006, and our users were able to harness that power to deliver exciting scientific accomplishments, including the following:

- At the scale of stars, a team led by ORNL's Tony Mezzacappa used NCCS supercomputers to discover the first plausible explanation for a pulsar's spin that fits the observations made by astronomers. The team's achievement was featured in the preeminent science journal *Nature*.
- At the scale of molecules, a team of ORNL researchers used NCCS systems to advance our understanding of adsorption, the process by which one molecule attaches to another. Adsorption is used in the manufacture of thousands of products, and the team described its accomplishment in a cover article for the prestigious *Journal of Physical Chemistry C*, a publication of the American Chemical Society.
- At the scale of our planet, a team led by Warren Washington of the National Center for Atmospheric Research (NCAR) used NCCS systems to simulate the effect of carbon dioxide emissions on global climate, highlighting with increasing clarity the choices we must make to keep the earth a hospitable place to live.

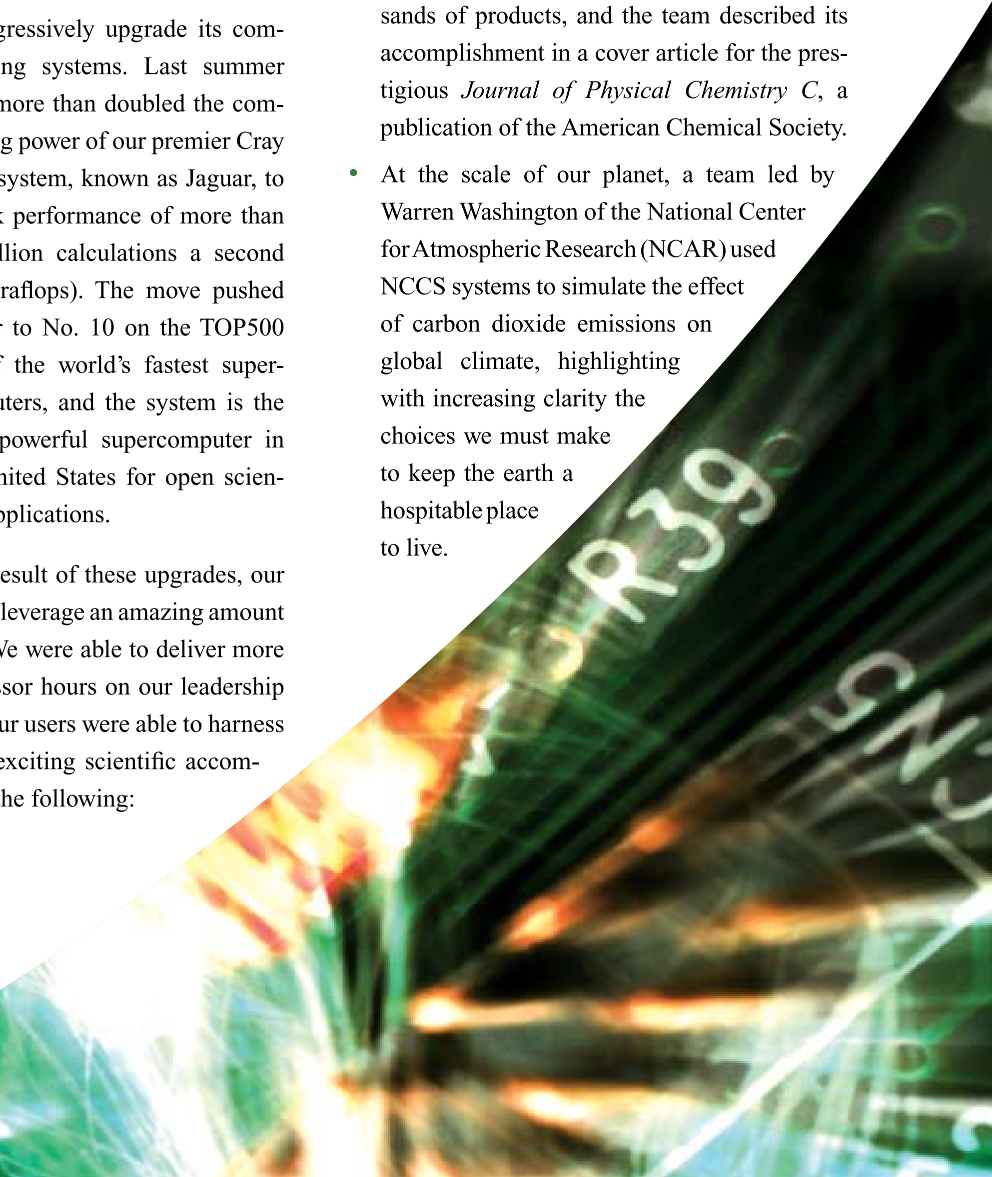

ORACLE vs. Jaguar: The 8-hour Day

ORNL's first big computer, the Oak Ridge Automatic Computer and Logistical Engine, was the best of its kind in its time, which was the 1950s. It could do 100 man-years of computing in eight hours.

We asked Bronson Messer, a computational scientist specializing in astrophysics, how that compares with the Cray XT4 Jaguar. "Hmmm," responds Bronson. "The simplest thing to do is to think that statistics for ORACLE were based on a man [or woman] doing a calculation every second. If you multiply that out, ORACLE did about 100,000 operations a second, or 100 kflops. The upgraded Jaguar will be capable of 119 teraflops: 119 x 10^{12} / 100 x 10^3 ~10^9, i.e. a billion times faster. So, the translation might be: Jaguar will be able to do 100 billion man-years of computing in eight hours."

Bronson also provides an astronomical angle: "There are about 100 billion galaxies visible to us right now in the observable universe. So, if you had one person in each galaxy in the observable universe computing for a year, Jaguar would beat them to an answer by day two."

These are just a few of the accomplishments we can point to at the NCCS, where researchers are expanding the boundaries of knowledge in a wide range of disciplines, from fusion energy to biology to particle physics to combustion.

The pace of our upgrade program will continue to accelerate in 2007 and beyond. As a result, we fully expect the pace of scientific advances being made here to increase as well. As of this writing, we have upgraded our Jaguar system to a peak performance of 119 teraflops, and by the end of 2007 we will have installed a quad-core upgrade to the system that will boost its peak performance to 250 teraflops. Because of these efforts, we are on track to more than double the computing resources we offer our researchers in 2007, to more than 75 million processor hours.

To keep up with these stunning increases in computing power, we have also expanded our staff and pushed forward upgrades to our network and storage systems. Each of our four groups—Scientific Computing, User Assistance and Outreach (UAO), High-Performance Computing Operations (HPC Ops), and Technology Integration (TechInt)—added staff in 2006, and our organization will continue to expand in 2007.

We have upgraded our connections to the Energy Sciences Network and Internet2, allowing our users to move up to 20 billion bits of information each second into and out of the center. We have also increased the storage in our High-Performance Storage System (HPSS) by more than half, from just under 750 trillion bytes (terabytes) at the end of 2005 to more than 1,250 terabytes as of this writing. Development continues on the sitewide high-performance Lustre file system known as Spider. Once in production, Spider will provide users with centralized, high-bandwidth access from each of the center's high-performance computing systems.

We here at the NCCS are very excited about the advances we have made, and we are very enthusiastic about the achievements yet to come. The faith we have placed in our researchers is paying off richly, and we are quite proud of our contribution to the nation's scientific achievement.

SCIENCE PERSPECTIVES

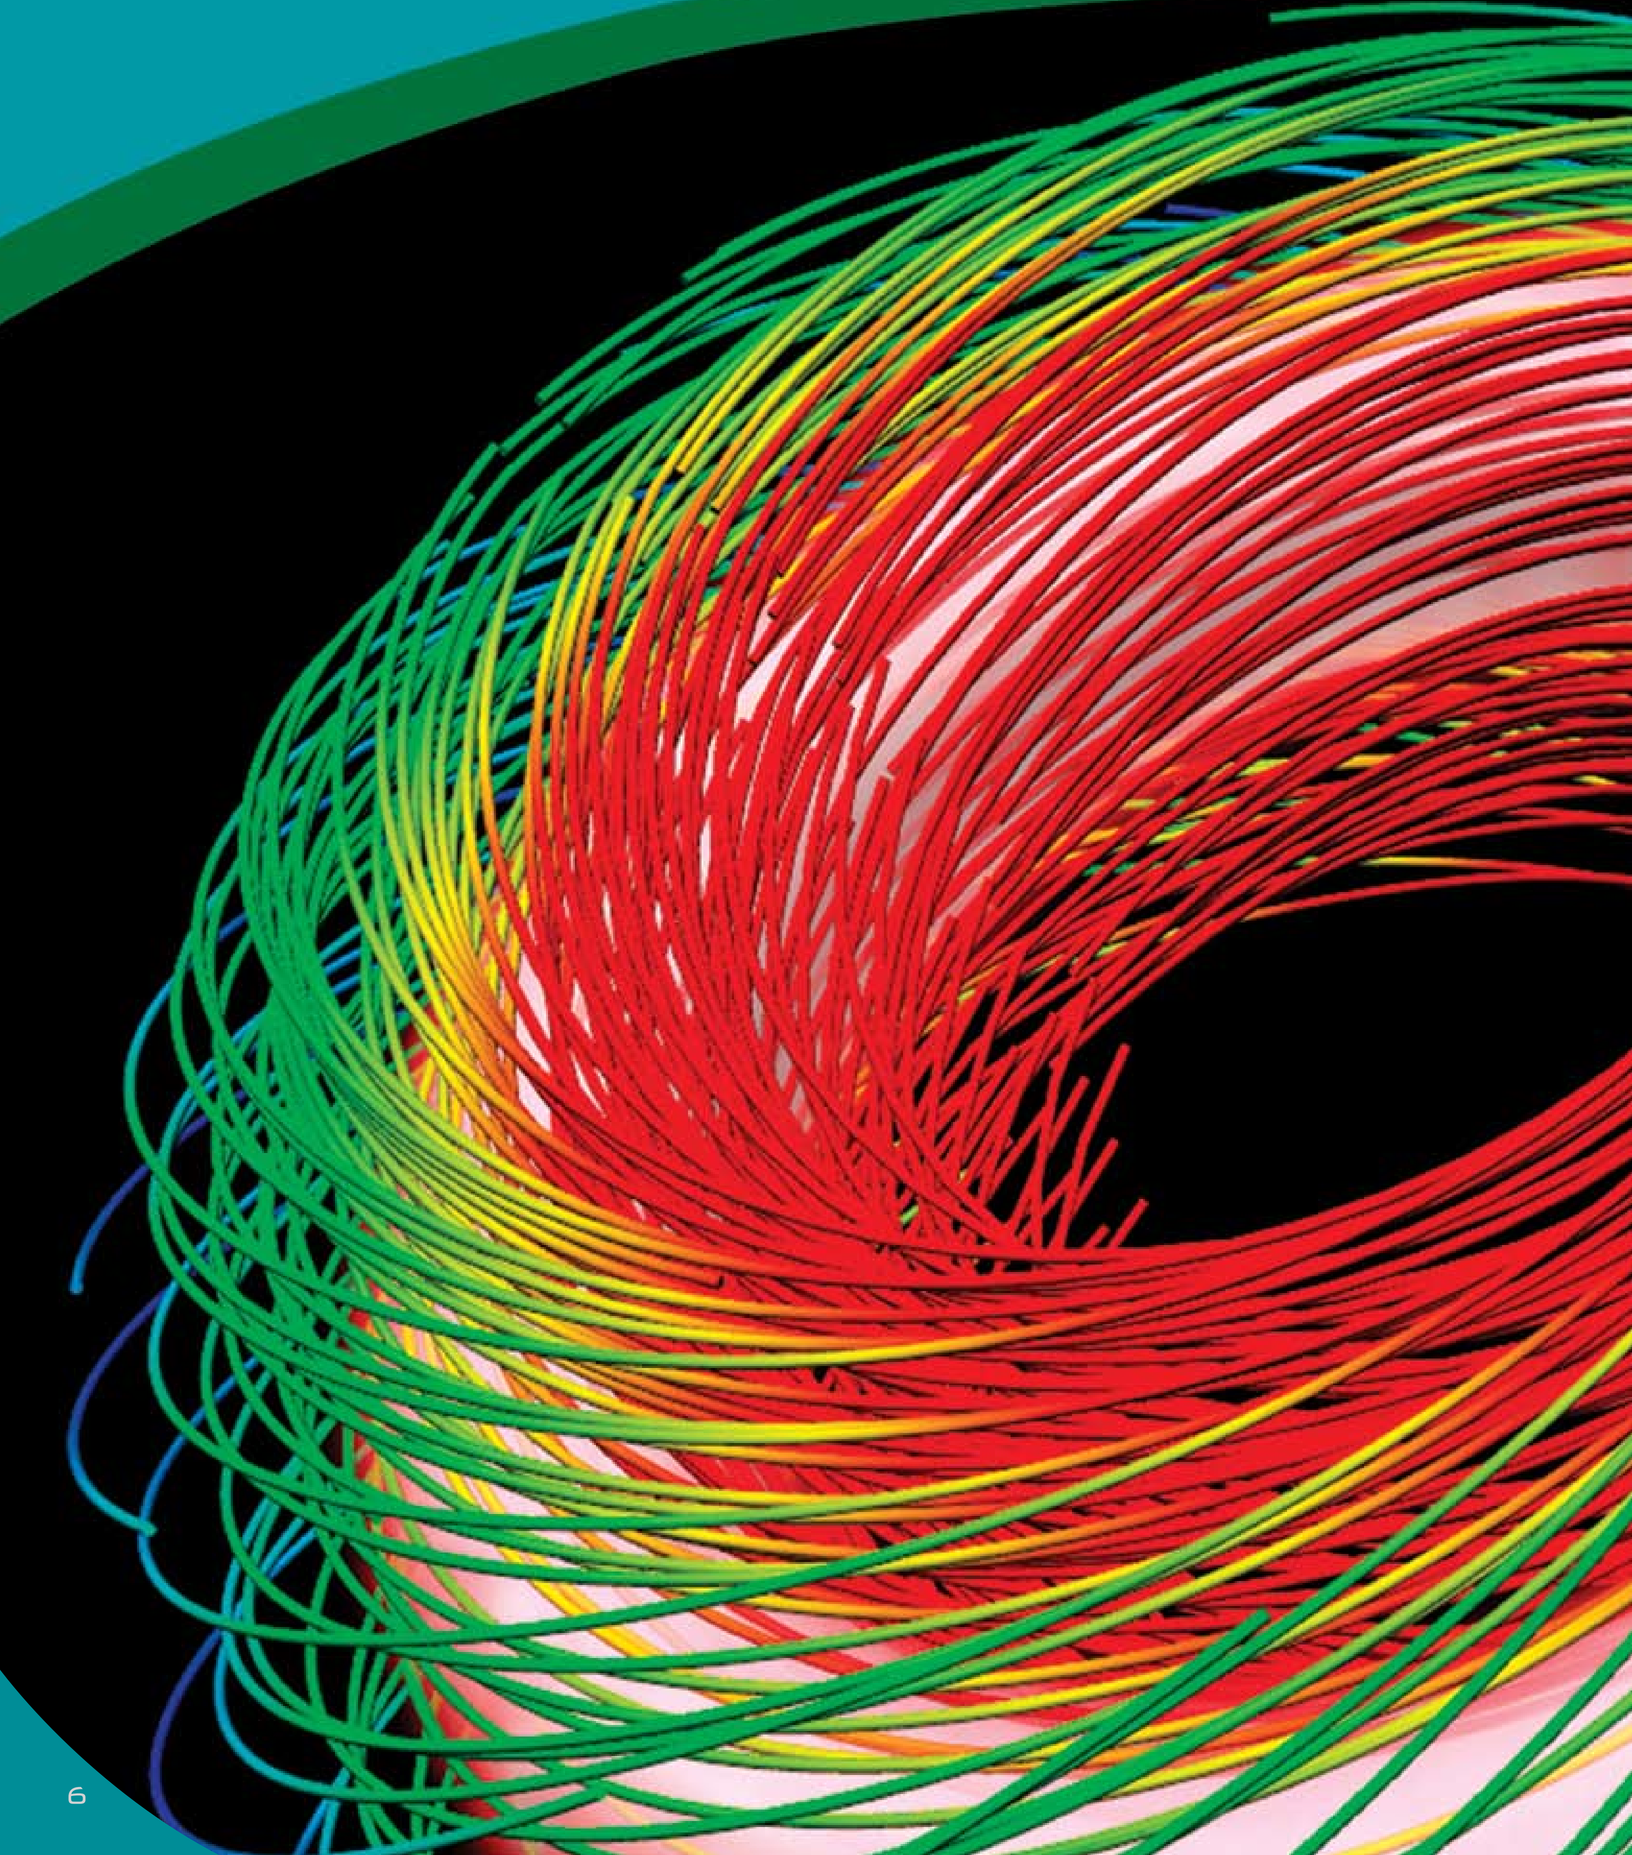

LETTER FROM THE DIRECTOR OF SCIENCE

From a science perspective, 2006 was a tremendous year for ORNL's NCCS.

We now have simulation capabilities across the sciences, from biology to high-energy physics. By offering unique, state-of-the-art facilities and support, we have attracted the leaders in each of these fields, and through this combination of talent and technology, we have expanded the boundaries of knowledge.

Being our first full year in operation, 2006 also provided a range of learning opportunities for us. We were able to discover what works and what doesn't. We were also able to use the year to expand our relationship with our users, learning how best to meet their needs and enable the groundbreaking science for which the center was created.

The NCCS hosted 22 major projects in 2006, with more than 50 scientific application codes running on our leadership systems. This exciting portfolio kept our support staff—as well as our Scientific Computing Group (SCG) and its scientific liaisons—busy working with researchers from across the United States and around the world.

To ensure that the codes running on our systems effectively produce breakthrough science, we established a new process—led by our Applications Requirements Council—that analyzes specific applications and matches them with NCCS supercomputers. We are confident that this process, and decisions resulting from it, will lead to more productive and higher quality science through an optimal use of the center's leadership resources.

I cannot do justice to the breadth and volume of fantastic science results we achieved in 2006, but I would like to provide a sampling. Maybe it will illustrate why we are very excited about the computational science research undertaken at the NCCS this past year.

Fusion

To realize the promise of fusion energy, scientists must overcome profound technical challenges. The plasma in a fusion reactor, which reaches temperatures of 180 million degrees Fahrenheit, can be unstable both system-wide and at the scale of electron interactions, and it behaves differently at the core of the reactor than at its edges.

Leading researchers used our Jaguar and Phoenix supercomputers to push the boundaries of fusion science, with projects simulating several key challenges of fusion reactors. One simulated the interplay of ions and electrons, another the effect of turbulence across the entire reactor, and a third the heating of fusion plasma.

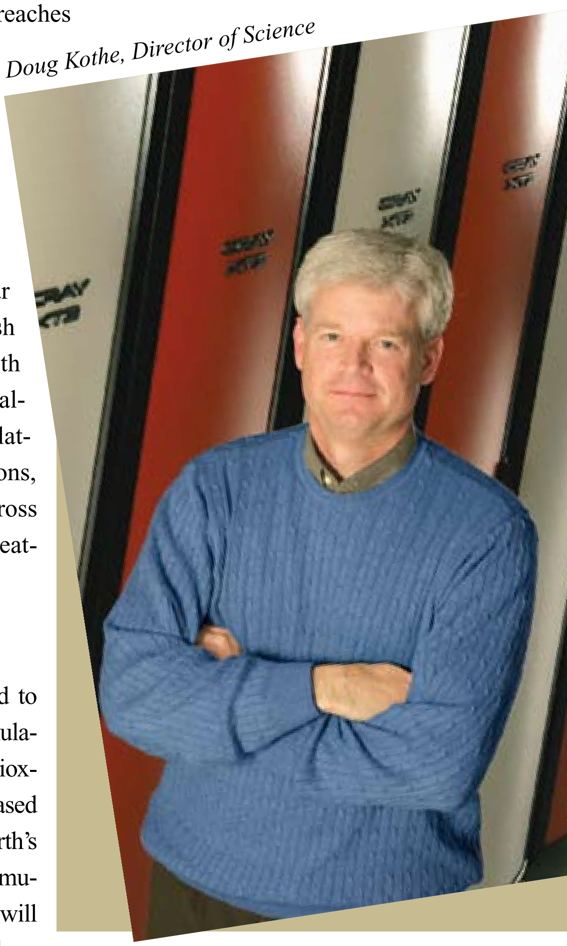
Doug Kothe, Director of Science

Climate

A team led by the NCAR continued to develop increasingly detailed simulations that show the effect of carbon dioxide—especially carbon dioxide released by industrial activities—on the earth's climate. The next version of the Community Climate System Model (CCSM) will also include more specific models of how carbon dioxide is absorbed by the earth's forests and oceans.

Because the stakes of this science are so high—with profound implications for the world's economies as well as its climate—these simulations must be accurate. Our Jaguar and Phoenix supercomputers gave the leading scientists involved the best possible resources to ensure they were producing information with unprecedented certainty.

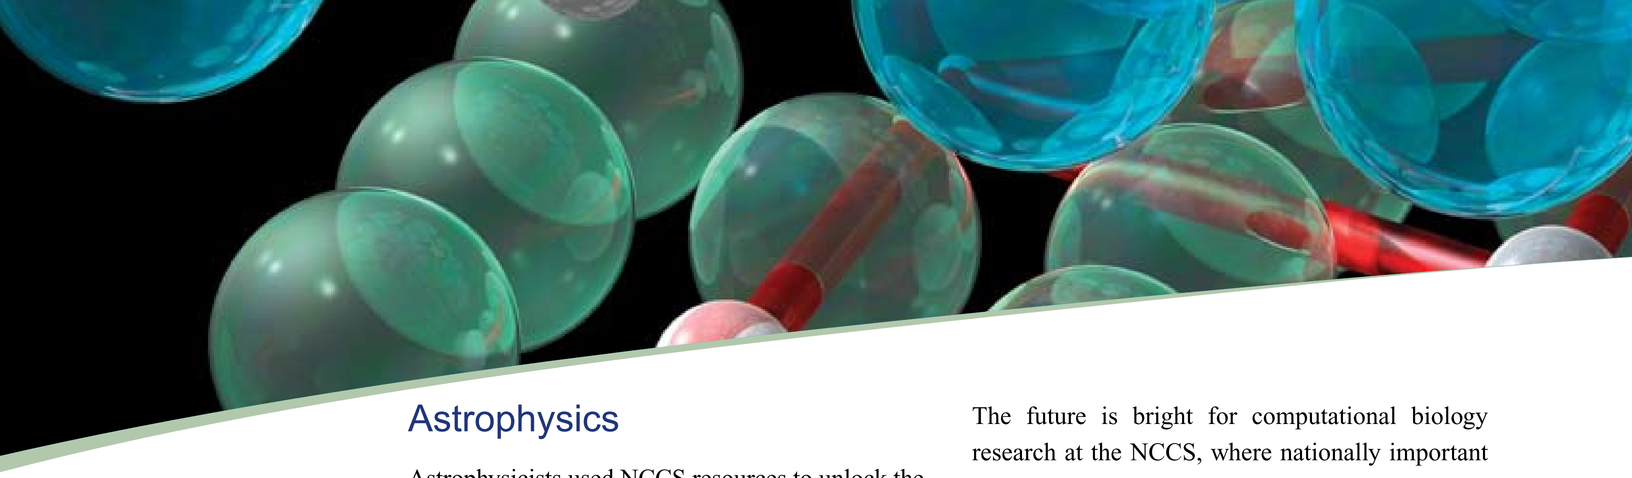

Astrophysics

Astrophysicists used NCCS resources to unlock the secrets of exploding stars. One team used Jaguar and Phoenix to explain the mechanism for core-collapse supernovas, the stellar explosions that provided most of the elements on earth. Because it had access to NCCS resources, the team was able to add more physics—and more of the right physics; it is going a long way toward putting to bed the mystery of the core-collapse supernova explosion mechanism.

Another team used Jaguar to explain the mechanism behind Type Ia supernovas. Besides being the largest thermonuclear explosions known, they are also key to understanding the accelerating rate at which our universe is expanding.

I'm optimistic that we're going to see some major astrophysics headlines in the coming year.

Biology and Bioenergy

One project I find especially fascinating used biomolecular dynamics simulations on Jaguar to probe the function of molecular motors in everyday living processes. One such "motor," the F1ATP-synthase, rotates around its shaft a thousand times a second, producing three molecules of ATP (adenosine triphosphate, which supplies energy for our every action) each time, for a total of 3,000 molecules per second.

Other computational biology research moved toward a better understanding of protein structure and dynamics. This work could lead to the design of better drugs for curing diseases and the production of better and novel enzymes for industrial uses such as the low-cost conversion of cellulose into ethanol.

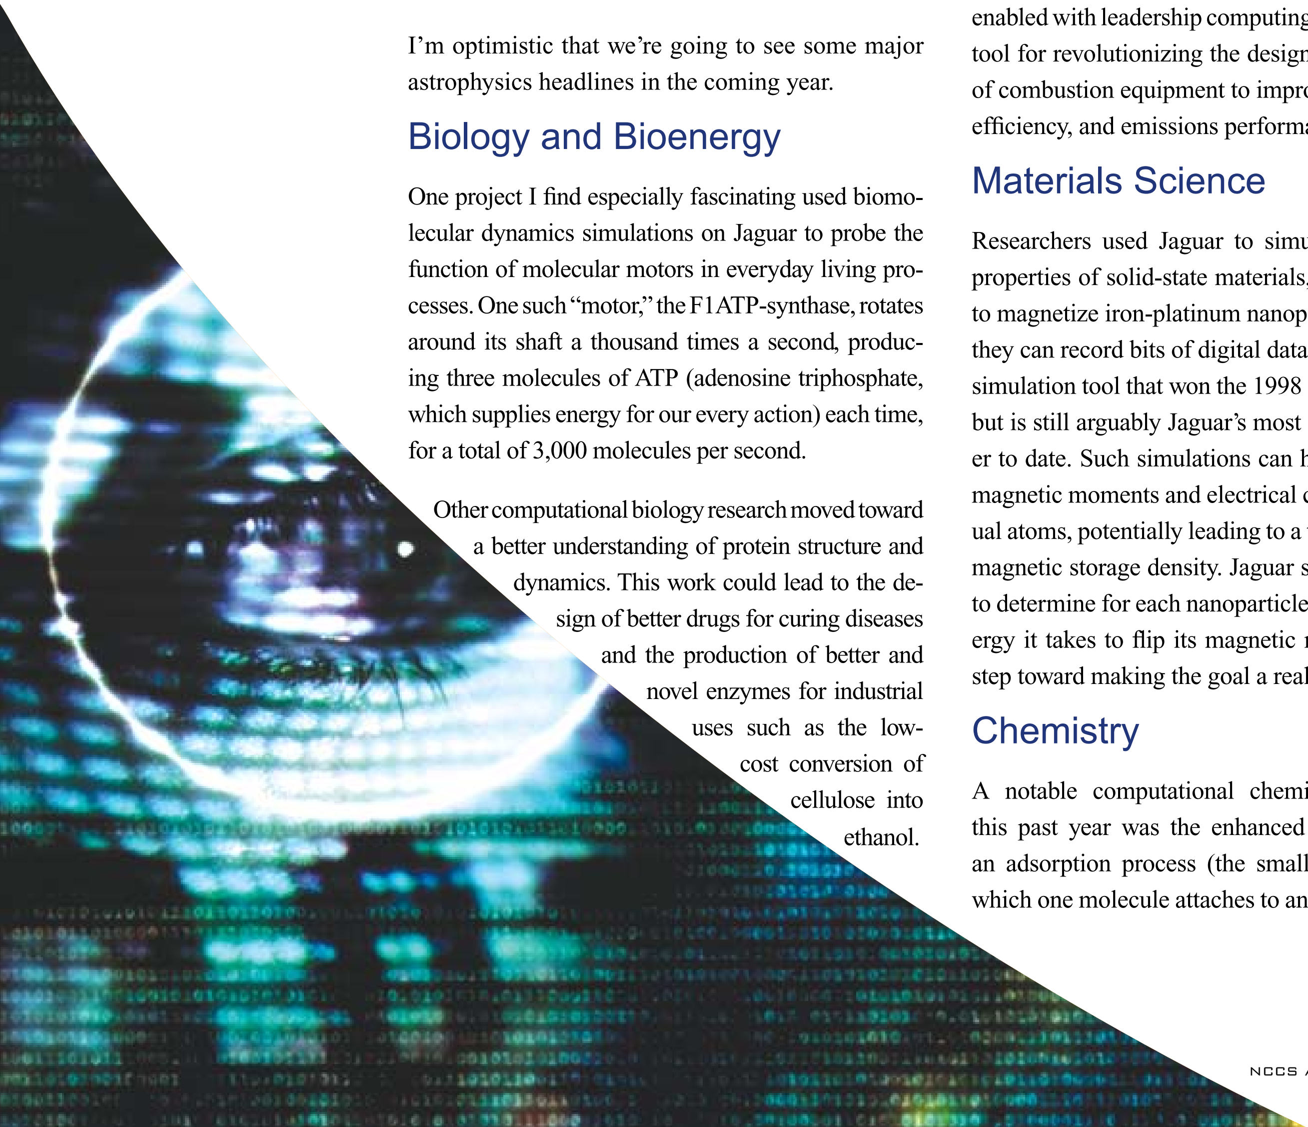

The future is bright for computational biology research at the NCCS, where nationally important initiatives such as bioenergy will rely on high-fidelity computational biology simulations on leadership systems to help guide the way.

Combustion

One of the most scalable and efficient consumers of our leadership systems has been a very impressive simulation tool that probes turbulent combustion processes via direct numerical simulation. This tool helped to clarify the extinction and reignition mechanisms of premixed and non-premixed fuel-air mixtures injected into next-generation combustion devices, such as internal combustion engines, power turbines, and industrial furnaces. Predictive models enabled with leadership computing are an unmatched tool for revolutionizing the design and performance of combustion equipment to improve reliability, fuel efficiency, and emissions performance.

Materials Science

Researchers used Jaguar to simulate the quantum properties of solid-state materials, looking for ways to magnetize iron-platinum nanoparticles reliably so they can record bits of digital data. This work used a simulation tool that won the 1998 Gordon Bell Prize but is still arguably Jaguar's most effective performer to date. Such simulations can help determine the magnetic moments and electrical charges of individual atoms, potentially leading to a tenfold increase in magnetic storage density. Jaguar simulations helped to determine for each nanoparticle the amount of energy it takes to flip its magnetic moment, a crucial step toward making the goal a reality.

Chemistry

A notable computational chemistry achievement this past year was the enhanced understanding of an adsorption process (the small-scale process by which one molecule attaches to another), namely the

molecular structure of successive layers of methane as they attach to magnesium oxide. This interaction of molecules with a surface is a fundamental problem found in a variety of important catalysis applications. Catalysts touch many aspects of our everyday lives, being directly involved in the synthesis of an estimated one in five industrial products as well as being used to produce gasoline and develop fuel cells.

Aerodynamics

In one of our industrial projects with the Boeing Company, leadership simulations were used to computationally investigate what happens to an airplane wing when a flap is suddenly deployed. Such "flutter analyses" have historically not included simulation of both the structural wing response (the wing flaps) and the airflow around the wing. Simulation of both simultaneously requires coupling of two complex codes in a nonlinear fashion: one for the aerodynamics and one for the structural response of the wing. Such accurate models can help prevent unnecessary maintenance and promote more aggressive designs in the future.

Computer Science

A performance end station project (focused solely on computer science) allowed potential future users to experiment with their codes, evaluate the performance of our leadership systems, and use the systems to develop tools for the high-performance-computing community. This effort is needed before current and future scientific applications can effectively use the 250 teraflops and 1 petaflop leadership systems due to be deployed at the NCCS within the next 2 years.

Nuclear Physics

Using a "coupled cluster" approach (made popular in quantum chemistry), whereby a set of coupled, nonlinear algebraic equations is used to describe the interactions of nuclear particles, researchers used Jaguar to explore the roles of both two- and three-body forces in intermediate-mass nuclei (e.g., oxygen-16). Such a project has never been successfully attempted for nuclei this massive. The ambitious goal is to put together a theoretical nuclear model that accurately predicts the mass, shape, size, and basic energy spectrum of intermediate-mass elements such as oxygen and calcium. Given the nuclear properties revealed by these models, crucial energy-security thrusts such as nuclear fusion and fission can be understood and predicted with more fidelity.

Accelerator Physics

Researchers continued to make impressive progress in the computational design of low-loss cavities for the International Linear Collider (ILC), a proposed electron-positron collider that will help answer questions about what the universe is made of and how it works. The ILC cavity design must be optimized, reaching the highest possible energy for beam acceleration while maintaining stability and focus.

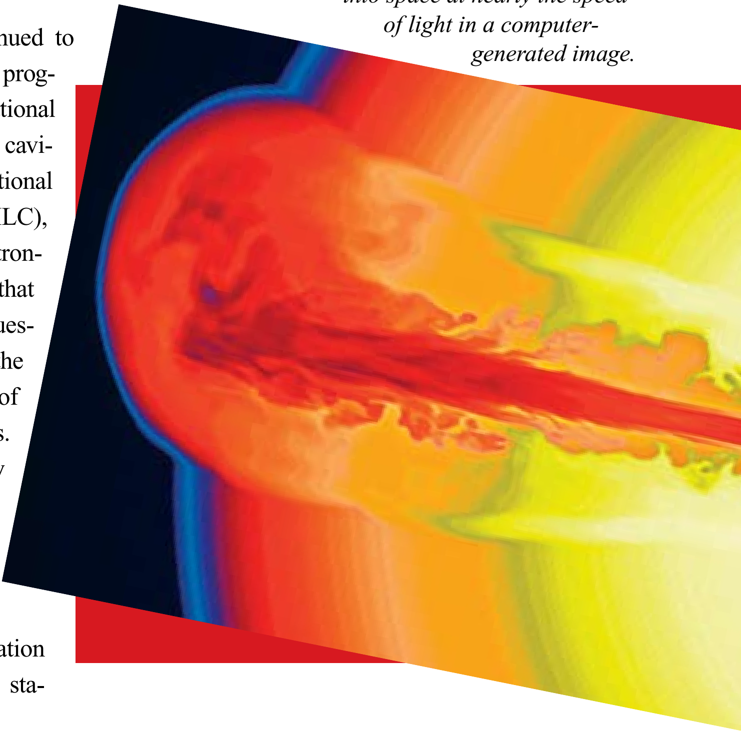

A superheated jet of plasma, the last gasp of a dying star, fires into space at nearly the speed of light in a computer-generated image.

These projects and others being carried out at the NCCS show beyond a doubt that scientific computer simulation has effectively established itself as a tool for accelerating discoveries and breakthroughs. Given the breathtaking power of the systems we have been able to provide and the world-class talent of the researchers using these systems, I fully expect work being performed here to be at the center of many of the world's most important scientific advances in the near future.

YOU'RE GETTING WARMER

The earth is getting warmer. Over the last century, the average temperature at the planet's surface has risen about 1.5 degrees Fahrenheit, and without intervention there appears no reason to believe the trend is going to stop.

While the change might seem modest, the potential consequences could be dramatic. Melting glaciers and thermal expansion are contributing to a rise in the sea level. Typhoons and hurricanes are likely to become more severe. Some areas are getting wetter, while others can look forward to long-term droughts. Even growing seasons are lengthening.

The issue came to the forefront early this year with the release of the Fourth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC), sponsored by the United Nations. Scientists from more than 100 countries contributed to the report, which concluded that global warming is, in fact, beyond serious question, with 11 of the last 12 years ranking among the warmest in the last one and a half centuries.

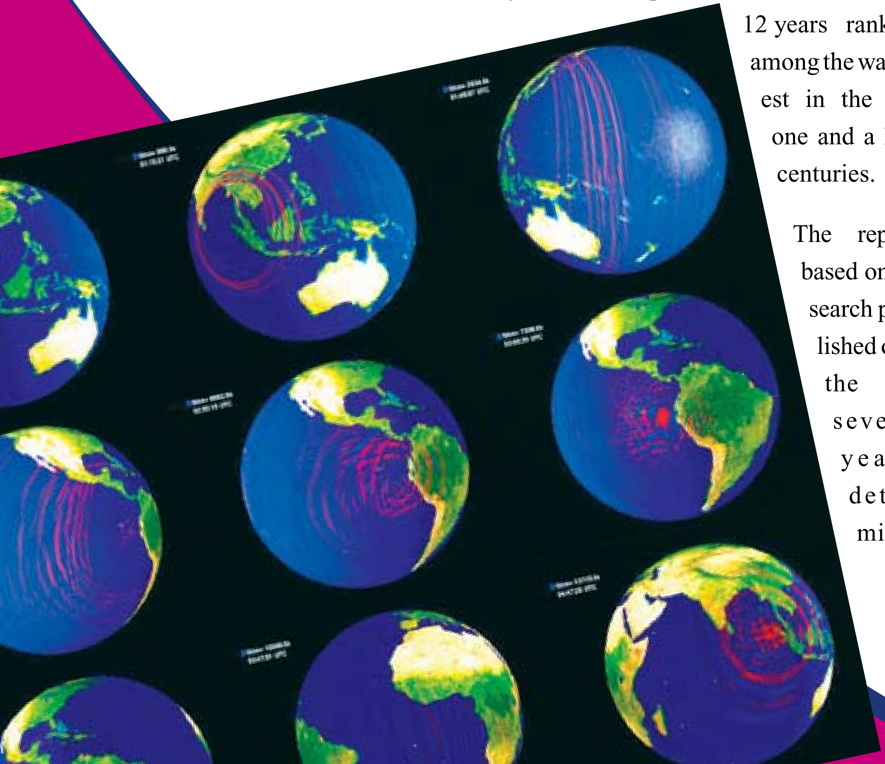

SPECFEM3D, an NSF and ORNL 2011 climate model.

The report, based on research published over the last several years, determined that there is more than a 90 percent chance that climate warming over the last 50 years has been the result of human activity. The primary culprits are chemicals known as greenhouse gases because they cause the atmosphere to trap more of the sun's heat than it would otherwise. The most important greenhouse gas is carbon dioxide, which is increased through the burning of fossil fuels by power plants, industry, and vehicles. Two other important greenhouse gases—methane and nitrous oxide—are linked mostly to agriculture.

A project headed by Warren Washington of NCAR is using the formidable computing resources of the NCCS to increase the confidence in global climate analysis. Through its research, the project will provide information needed by the public and policy makers alike to address our changing climate.

The project, known as the Climate-Science Computational End Station Development and Grand Challenge Team, is working to push climate simulation to a new level, both through the efforts of team members and by providing the scientific community as a whole with improved analytical tools.

"There will be intense research over the next decade or so to try to mitigate the effects of global warming through the possible use of increased renewable energy sources and possibly nuclear energy," Washington explained. "Advanced climate models can provide various scenarios of possible future climate change that will help guide policy makers in those decisions."

The challenge faced by climate scientists is no longer about whether the climate is getting warmer;

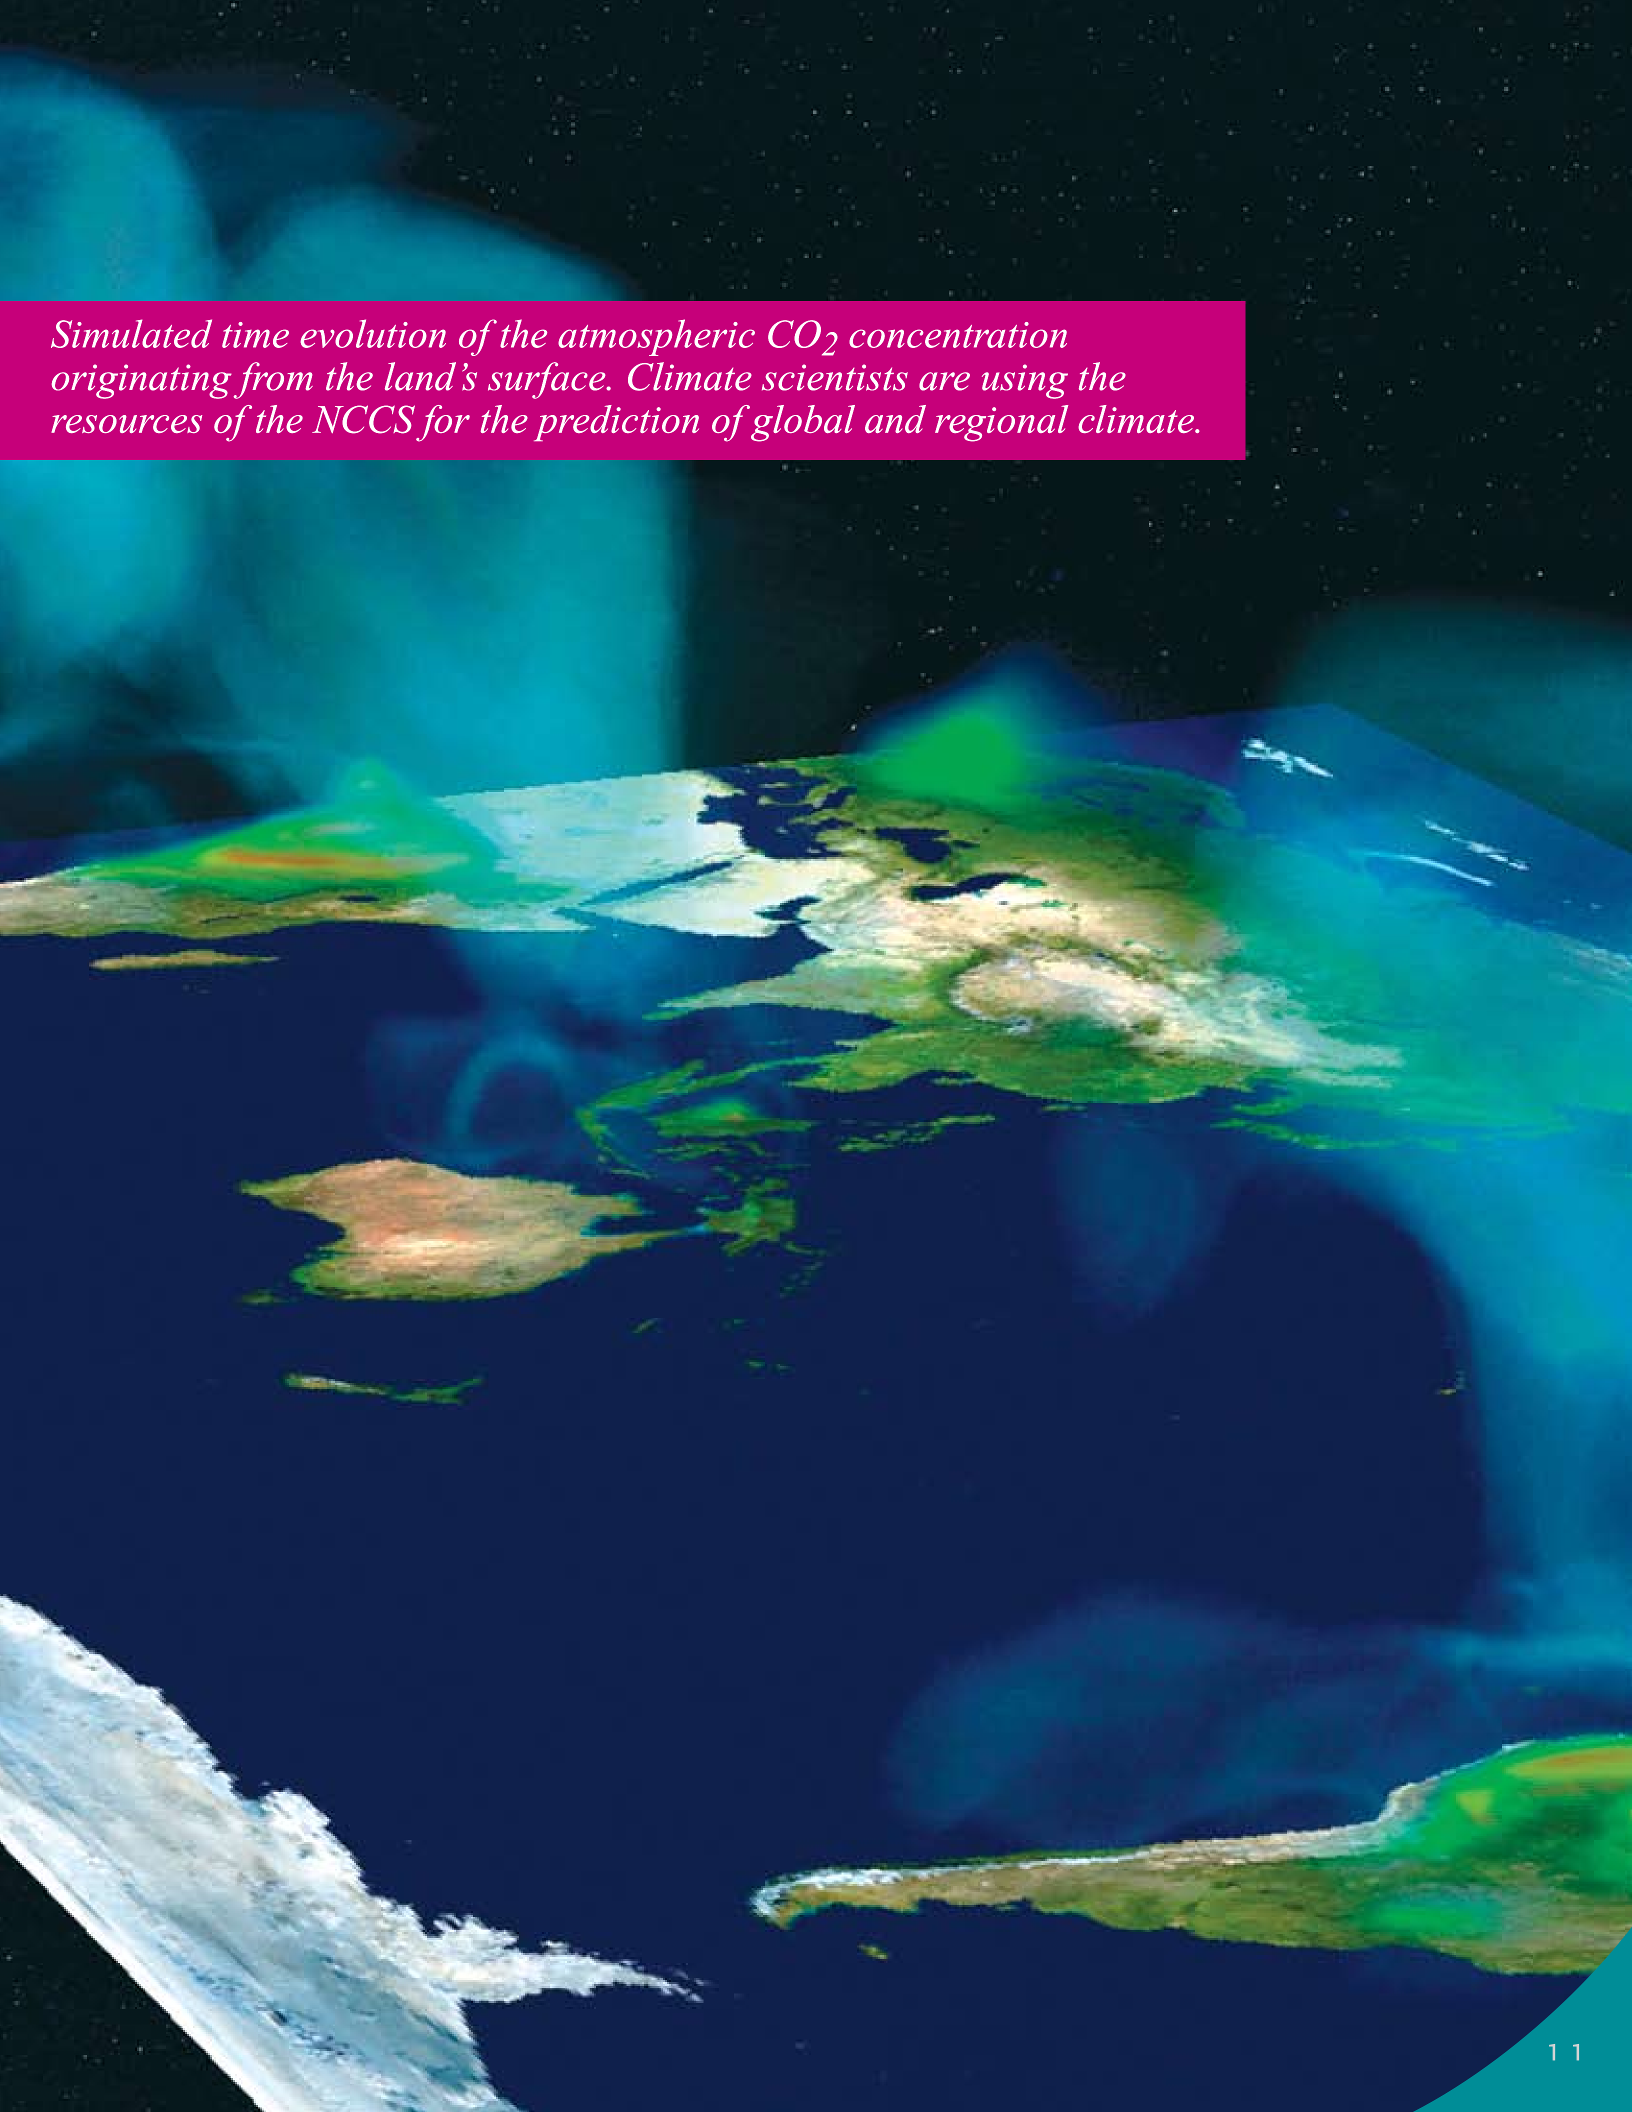

Simulated time evolution of the atmospheric CO_2 concentration originating from the land's surface. Climate scientists are using the resources of the NCCS for the prediction of global and regional climate.

Climate model simulation displayed on the EVEREST PowerWall.

that fact has been established by decades of measurement and research. Instead, scientists are working to determine exactly why the climate is changing, what the consequences are, and what, if anything, we can do about it.

"The warming discussed in the IPCC report, that's data," said ORNL's John Drake, chief computational scientist for the project. "That's observational data. There's very little that you'd even say is theory in that.

"So then, how do you characterize that? Here's where the theory starts to come in and where computing starts to play a role. Is it caused by human fossil-fuel burning, etc., or is this a natural variability within the climate system?"

Washington's team includes leading climate researchers from government and academia. Besides the scientists at NCAR, project members include the Department of Energy (DOE), the National Science Foundation, and the National Aeronautics and Space Administration. Researchers working on the project come from Oak Ridge, Lawrence Livermore, Lawrence Berkeley, Los Alamos, and Pacific Northwest national laboratories as well as universities around the country.

The group has already made substantial contributions to the field. It is responsible for the CCSM, which effectively simulates the effect of increasing greenhouse gases on temperature across the planet. The team also made the largest contribution of simulation data to the recent IPCC report.

At the time of those simulations, the project used the NCCS IBM pSeries Cheetah supercomputer, taking up more than half of that system for nearly a year and making up 30–40 percent of all IPCC simulations. Moving into 2007, the team is using the far more powerful Cray XT4 Jaguar and Cray X1E Phoenix systems at the NCCS.

The project needs the increased computing power and more to improve the CCSM. While existing models allow researchers to accurately simulate the impact of increased greenhouse-gas concentrations on climate across the globe, the scientists are committed to simulating such impacts on much smaller regions.

"By improving the scale and coming closer to the actual data, you can say what's the impact on a smaller area," Drake explained. "We'd like to be able to say, for example, what the different effects are between this side of the Appalachians and the other side of the Appalachians. That's a little fuzzy currently because we don't really resolve scales enough to tell what's the difference between Knoxville [Tennessee] and Asheville [North Carolina]. We'd like to be able to do that; clearly, that would be important information for people to have."

In addition, the researchers are committed to simulating the entire carbon cycle. The job is formidable. Instead of determining the increased levels of carbon dioxide in the global atmosphere, newer versions of the model will calculate emissions where they occur—at factories, power plants, and urban settings—and model the journey of carbon dioxide in the atmosphere. They will include the exchange of carbon dioxide between the atmosphere and the ocean. They will include the trees that take up carbon dioxide in the spring and summer and release it back in the fall and winter. And they will simulate the effect of deforestation in places such as the Amazon, where trees are increasingly unavailable to take carbon dioxide out of the atmosphere.

The job will be difficult. As the model becomes increasingly complex, each new element will result in new problems that must be found and corrected.

"This is a hard problem," Drake explained. "It's not going to be solved in 5 years, or even probably in 10 years. We will introduce errors that then need to be found and squeezed out of the theory."

The outcome will provide benefits beyond those to the dozens of leading scientists working on the project because the tools being developed at the NCCS will also be available to the scientific community in general.

"No one can develop a climate model on their own anymore," Drake explained. "It's just way too complicated."

Ocean surface height variability from ORNL Jaguar simulation.

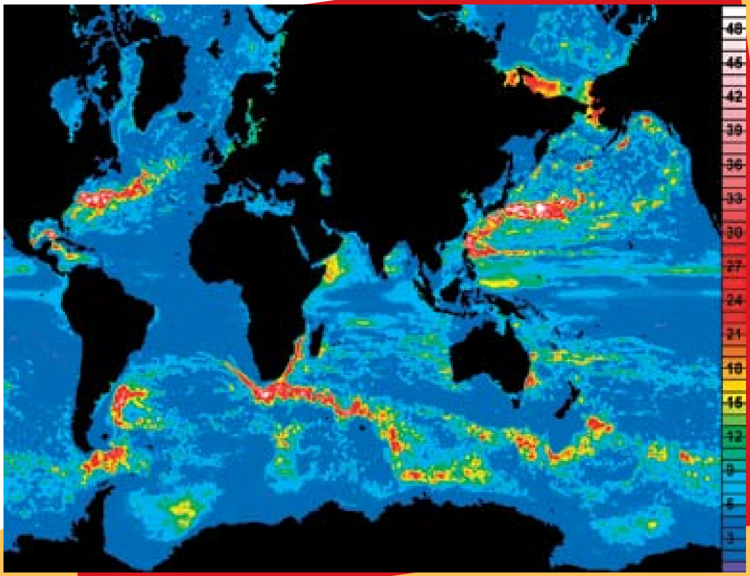

Current Events

The ocean is the 800-pound gorilla of climate studies.

A change in ocean currents could dramatically influence which regions on earth will be frigid and which will be temperate, which will be jungle and which will be desert. But while the ocean is undeniably important to climate, its role is not well understood.

A team led by Synte Peacock of the University of Chicago is using the immense computing power of the NCCS to run the most fine-grained, global-scale simulations ever of how the world's oceans work. In doing so the team will cast light on the currents and processes at work in the oceans as well as the possible fate of chemicals and gases—including carbon dioxide and other greenhouse gases—released into them.

The team is using the NCCS's Cray XT4 Jaguar supercomputer to perform the first-ever 100-year simulation of the ocean at a fine enough scale to include the relatively small, circular currents known as eddies. The task is daunting. The earth is enveloped by 319 million cubic miles of ocean, which covers nearly three-quarters of the planet to an average depth of more than 12,000 feet.

"There is a lot we don't understand," Peacock explained. "If you release carbon dioxide, what will happen to it? It will mix, get diluted, get transported into the interior ocean, and eventually make it back to the surface, but how long will it take, and what pathways will it take?"

The group will perform a series of virtual tracer studies, following both particles and dissolved material as the simulations push them around the ocean. These experiments will give scientists an idea of how something might be transported from the surface to the ocean depth and around the globe and of how long the journey might take. The job will be complicated by the difference in timescales between the surface (with a turnover on the order of decades) and the deep ocean (where the same process can take centuries).

The inclusion of ocean eddies is a key element in the project. Ranging in diameter from about 60 miles near the equator to about 6 miles near the poles, eddies play a key role in the dynamics of the ocean. Nevertheless, until recently researchers lacked the computing power to directly simulate eddies on a global scale.

"There is evidence that eddies can actually change the average flow of the oceans," Peacock noted. "Once you get to the resolutions you need to include the eddies, you get features that don't exist otherwise in a coarser-mesh model."

The team's ocean model will eventually be incorporated into the global CCSM. The information it provides will help us become more knowledgeable and responsible guardians of the planet.

Fueling the Future

Imagine a world with a virtually unlimited source of clean, safe energy that has almost no by-products or waste and zero risk of a nuclear meltdown.

While it might sound like a pipe dream, fusion reactors could one day turn this dream into a reality. Their potential is being explored by computational physicists like Jeff Candy of General Atomics, who is using the enormous resources of ORNL's NCCS to model heat and particle loss in a fusion reactor.

Candy's revolutionary GYRO code simulates plasma in a doughnut-shaped tokamak reactor. Plasma is an ionized gas that serves as the reactor's fuel; the plasma creates energy when hydrogen atoms within it collide, creating high-energy alpha particles and neutrons.

According to Candy, the plasma fuel in a tokamak is turbulent, creating small eddies that are very difficult to simulate. Due to this plasma turbulence, a tokamak lacks "perfect containment." As alpha particles heat the plasma, turbulence carries some of the heat away and eventually throws out the particles.

Once an alpha particle loses its energy, its usefulness is over. Candy likens the spent alpha particle to water in gasoline; it must be eliminated for the machine to run smoothly. Likewise, some heat loss in a fusion reactor is necessary for operational efficiency. The key, says Candy, is to know the rate at which the plasma is leaking heat and particles because while some loss is necessary, too much can be detrimental.

"Turbulence is necessary for a tokamak reactor," Candy said, adding that his GYRO code computes "optimal turbulence," finding the perfect balance of heat and alpha-particle production and loss. Managing turbulence affects potential tokamak reactors in a number of ways, including the cost, which will no doubt be a key factor in bringing fusion reactors to reality.

Faster! Faster!

Princeton Plasma Physics Laboratory (PPPL) fusion researchers using Jaguar set a speed record in 2006 for running the Gyrokinetic Toroidal Code (GTC), topping the previous record held by Japan's Earth Simulator supercomputer.

Stephane Ethier of PPPL succeeded in running GTC on 10,386 of Jaguar's 10,424 processing cores. The run advanced 5.4 billion particles per step per second, a 13 percent improvement over the 4.8 billion particles per step per second on the Earth Simulator, which ran the code on 4,096 processors. The three-dimensional GTC was developed at PPPL to study the dynamics of turbulence and the associated transport of particles driven by variations in temperature and density within a reactor system.

Ethier was particularly pleased with the efficiency of the runs on the dual-core processors of the new Cray XT4. GTC demonstrated more than 95 percent efficiency on the second processor of each dual-core node in the runs.

PPPL chief scientist William Tang said the run on Jaguar reached an extremely high statistical resolution. The resolution of fusion simulations will continue to climb as computing systems reach petascale and higher levels. High-resolution calculations with low noise help build an understanding of turbulent behavior in plasmas on realistic time scales that are characteristic of experimental observations.

The substantial contributions of NCCS staff to reaching the milestone, particularly Scott Klasky and Don Maxwell, were noted by the PPPL researchers.

The achievement is another step toward accurately simulating plasma behavior in fusion reactors such as ITER. The design and operation of fusion reactors must take turbulence into account because it is thought to be the primary mechanism by which particles and energy leave the magnetic field of a tokamak reactor, causing the system to lose energy.

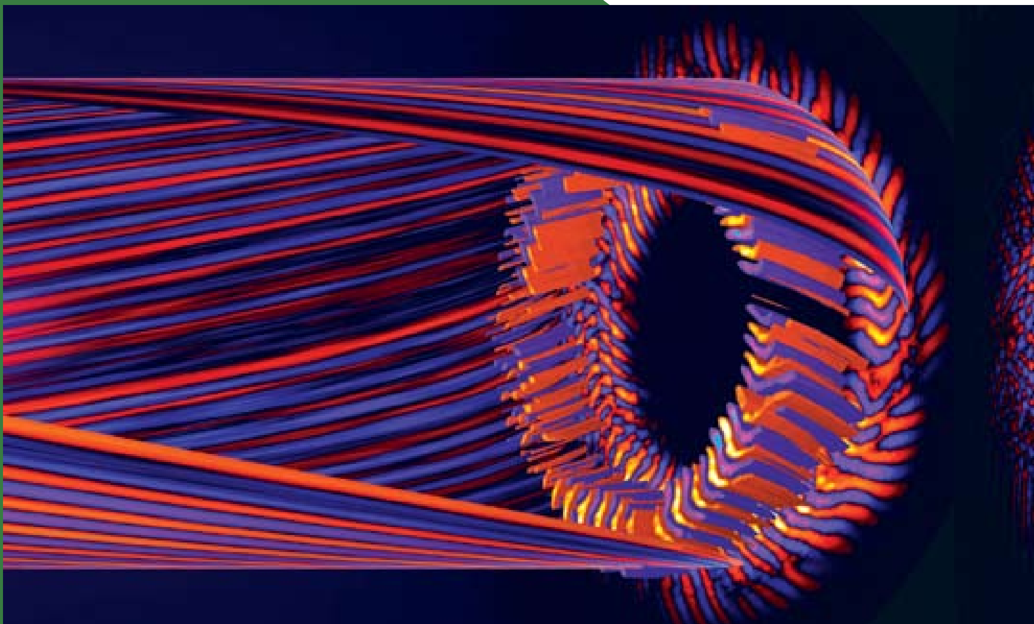

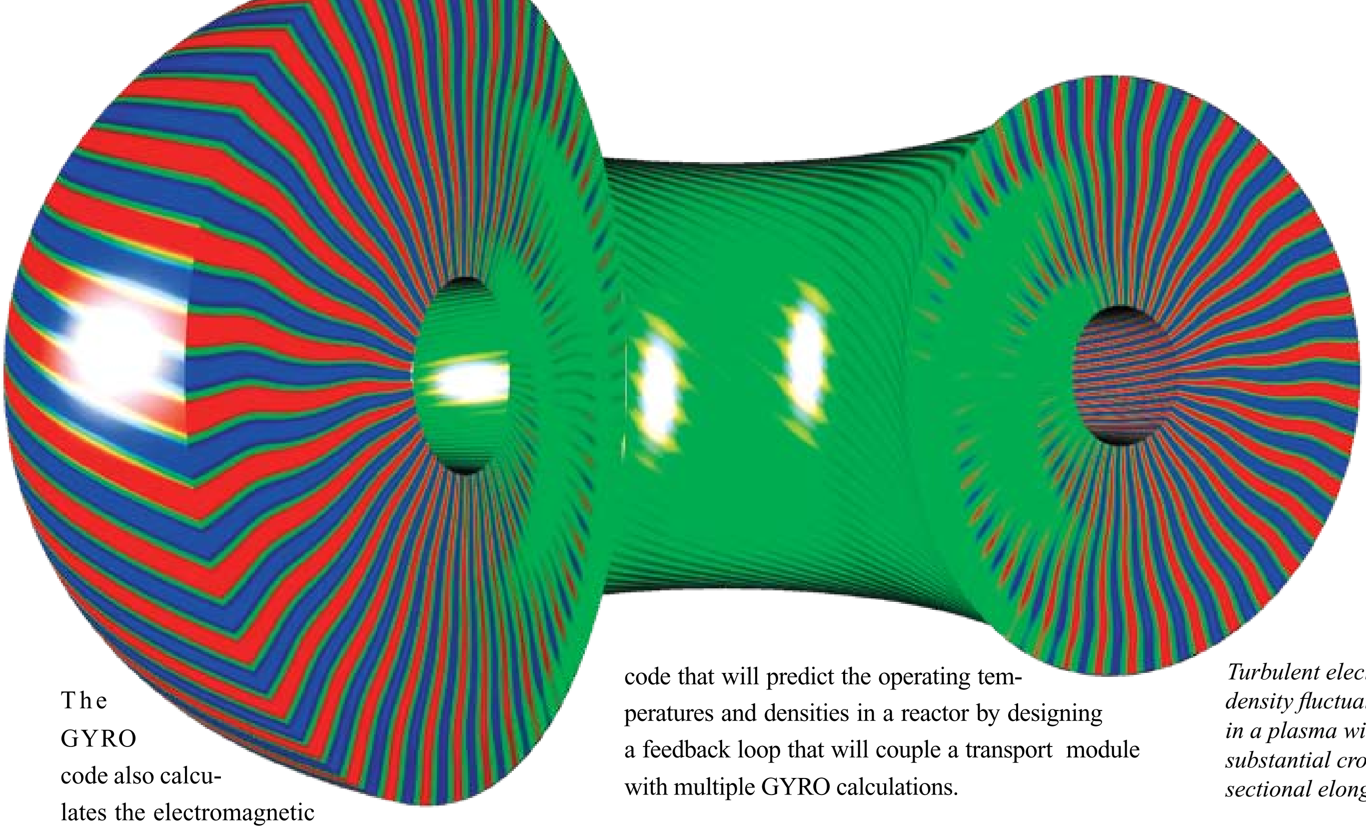

Turbulent electron density fluctuations in a plasma with substantial cross-sectional elongation.

The GYRO code also calculates the electromagnetic fluctuations in a plasma, a challenge that frustrated computational physicists for years. According to Candy, an algorithm was eventually developed that allows electromagnetic simulation in a "clean, accurate way." This algorithm, he says, is the only one known that can correctly calculate electromagnetic fluctuations.

Candy has an allocation at the NCCS to study turbulence simultaneously at the atomic scale of ions and the much smaller, subatomic scale of electrons. The job takes enormous computing resources because the eddies created by ion turbulence are dozens of times larger than those created by electron turbulence. Candy has used the GYRO code to model the two together and gain a "real, theoretical, broad understanding of ion-electron turbulence." Results from these simulations, carried out on the NCCS's Cray X1E Phoenix supercomputer, were presented at the International Atomic Energy Agency's 21st Fusion Energy Conference in Chengdu, China.

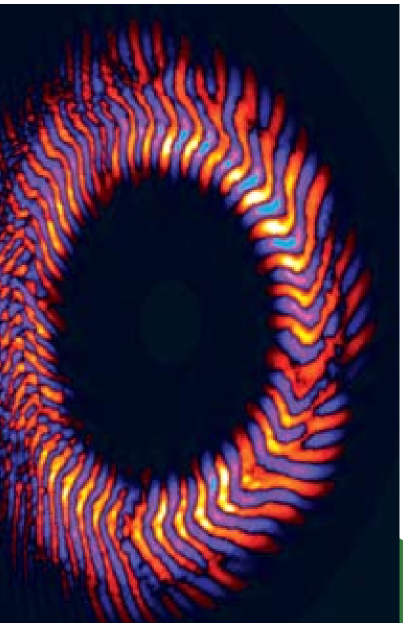

Candy also has a new project at the NCCS for 2007. In it his team will develop a new type of code that will predict the operating temperatures and densities in a reactor by designing a feedback loop that will couple a transport module with multiple GYRO calculations.

"Gyrokinetic simulations are a bit like a black box," Candy explained. "You tell it what the densities and temperatures are, and it will tell you how violent the turbulence is. By introducing feedback schemes, you find the unique temperatures and densities for which heat loss balances heat production."

Fortunately, he said, the project is able to take full advantage of the growing computing power offered at the NCCS. While the simulations will take up to 100 times the processing power of previous work, the GYRO code is able to take advantage of the thousands of processors available on the center's state-of-the-art supercomputers.

Using this new feedback technique, Candy and collaborators will attempt to predict the fusion performance of the ITER tokamak reactor, an international collaboration located in France. ITER, whose U.S. headquarters is located at ORNL, is currently a top priority of DOE's Office of Science. It is geared toward reaching the fusion energy break-even point, getting more energy out of the reactor than goes in.

Candy's work at the NCCS could eventually lead to a better understanding of tokamak plasmas and some day help deliver clean, safe energy to a world in need. With projects such as ITER already under way, the environmental side effects of conventional power production could one day be a thing of the past.

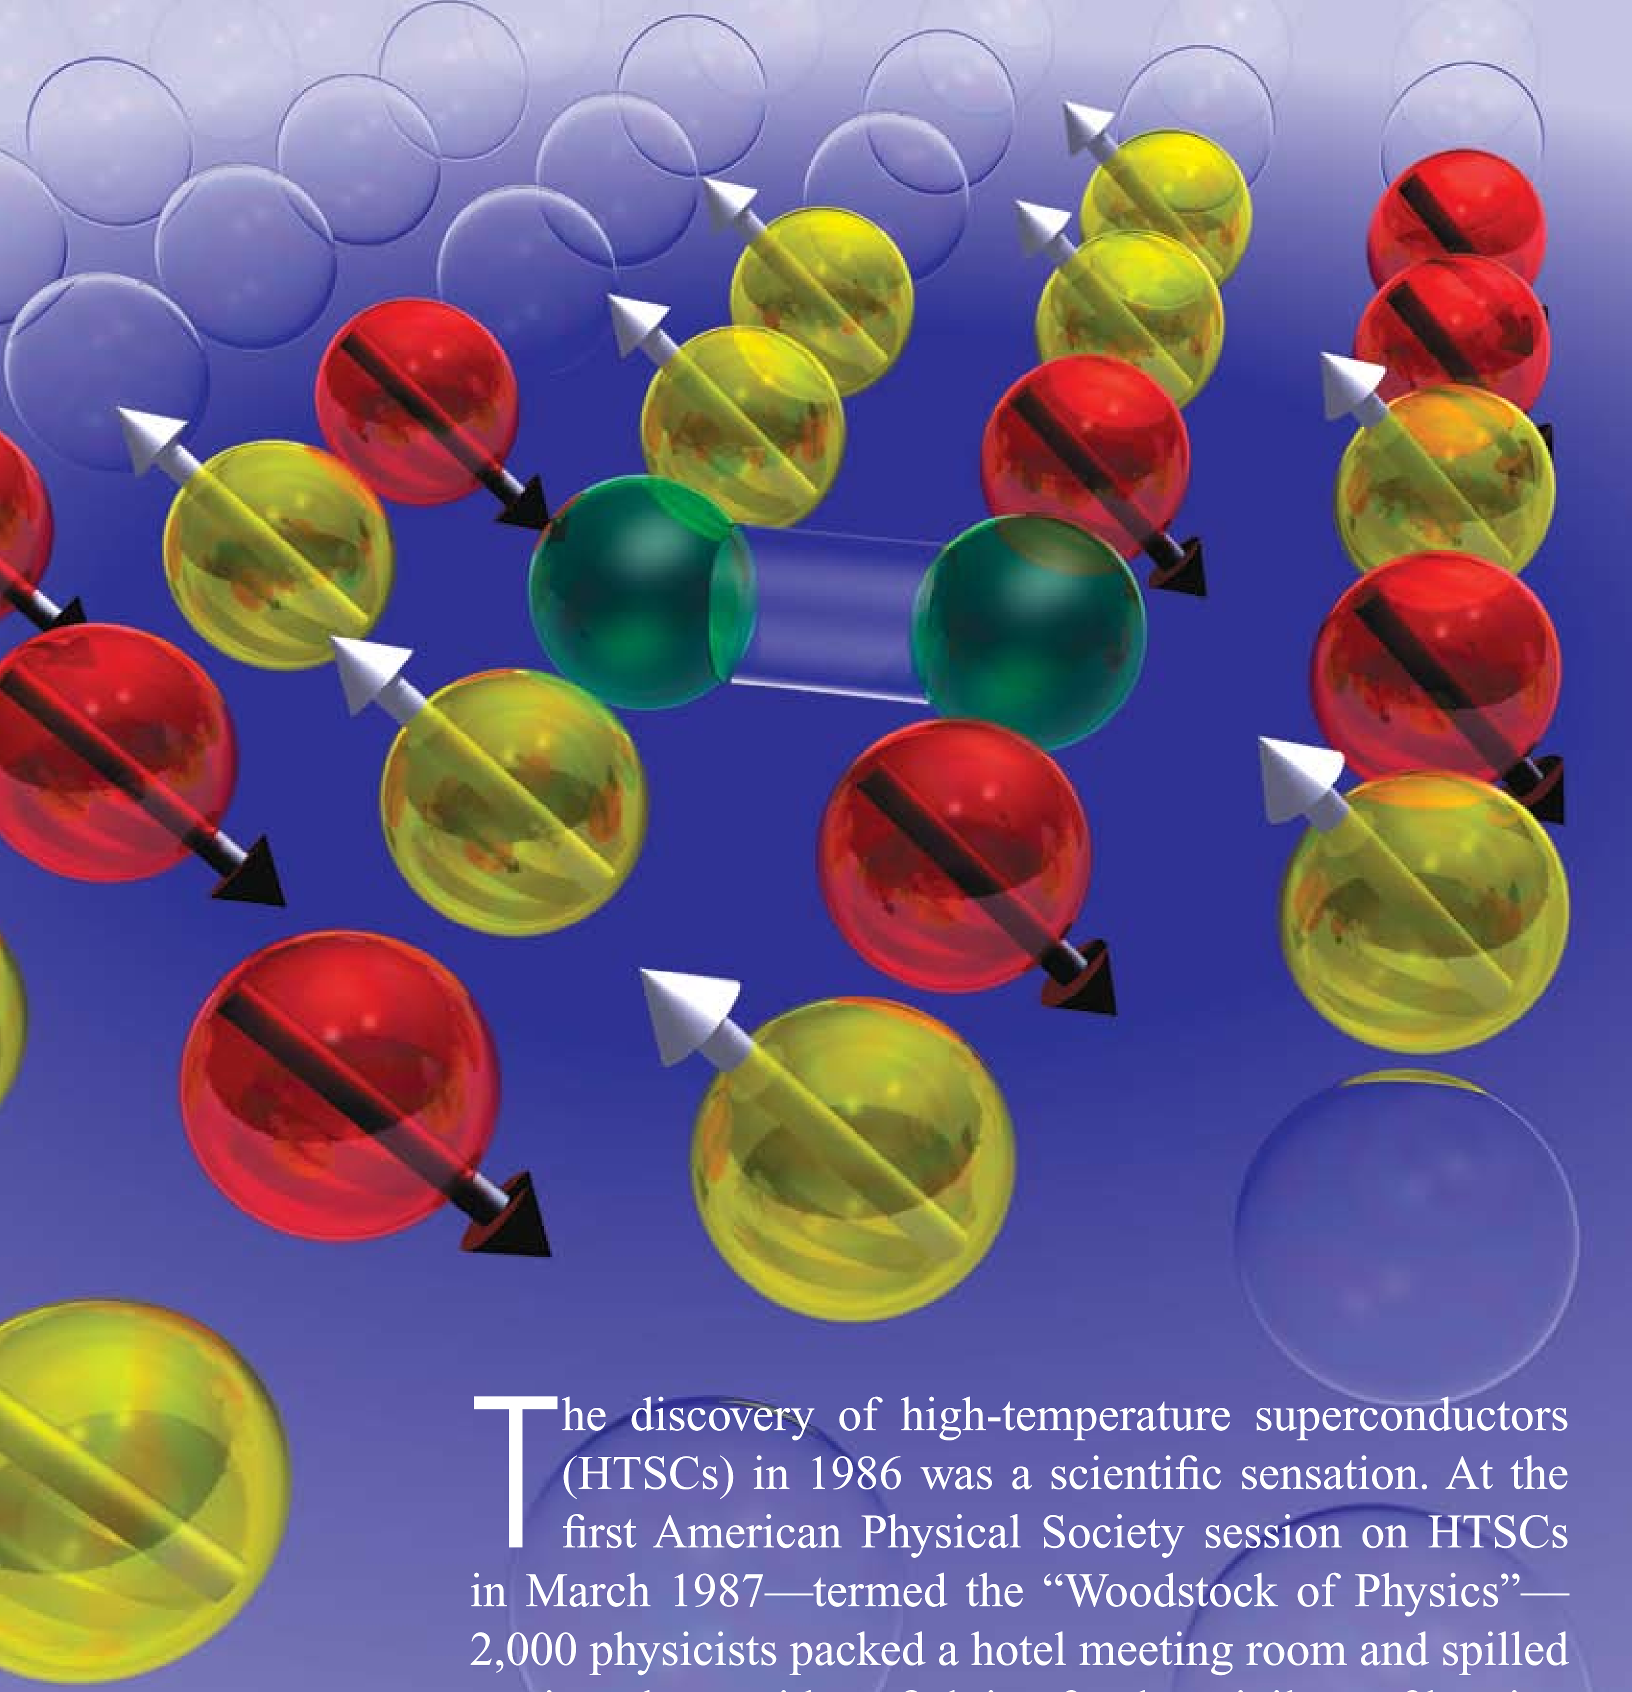

The discovery of high-temperature superconductors (HTSCs) in 1986 was a scientific sensation. At the first American Physical Society session on HTSCs in March 1987—termed the "Woodstock of Physics"—2,000 physicists packed a hotel meeting room and spilled out into the corridors, fighting for the privilege of hearing the technical papers firsthand. The session lasted all night. It was perhaps the only meeting of physicists in history ever characterized as "a riot."

RESISTANCE IS FUTILE

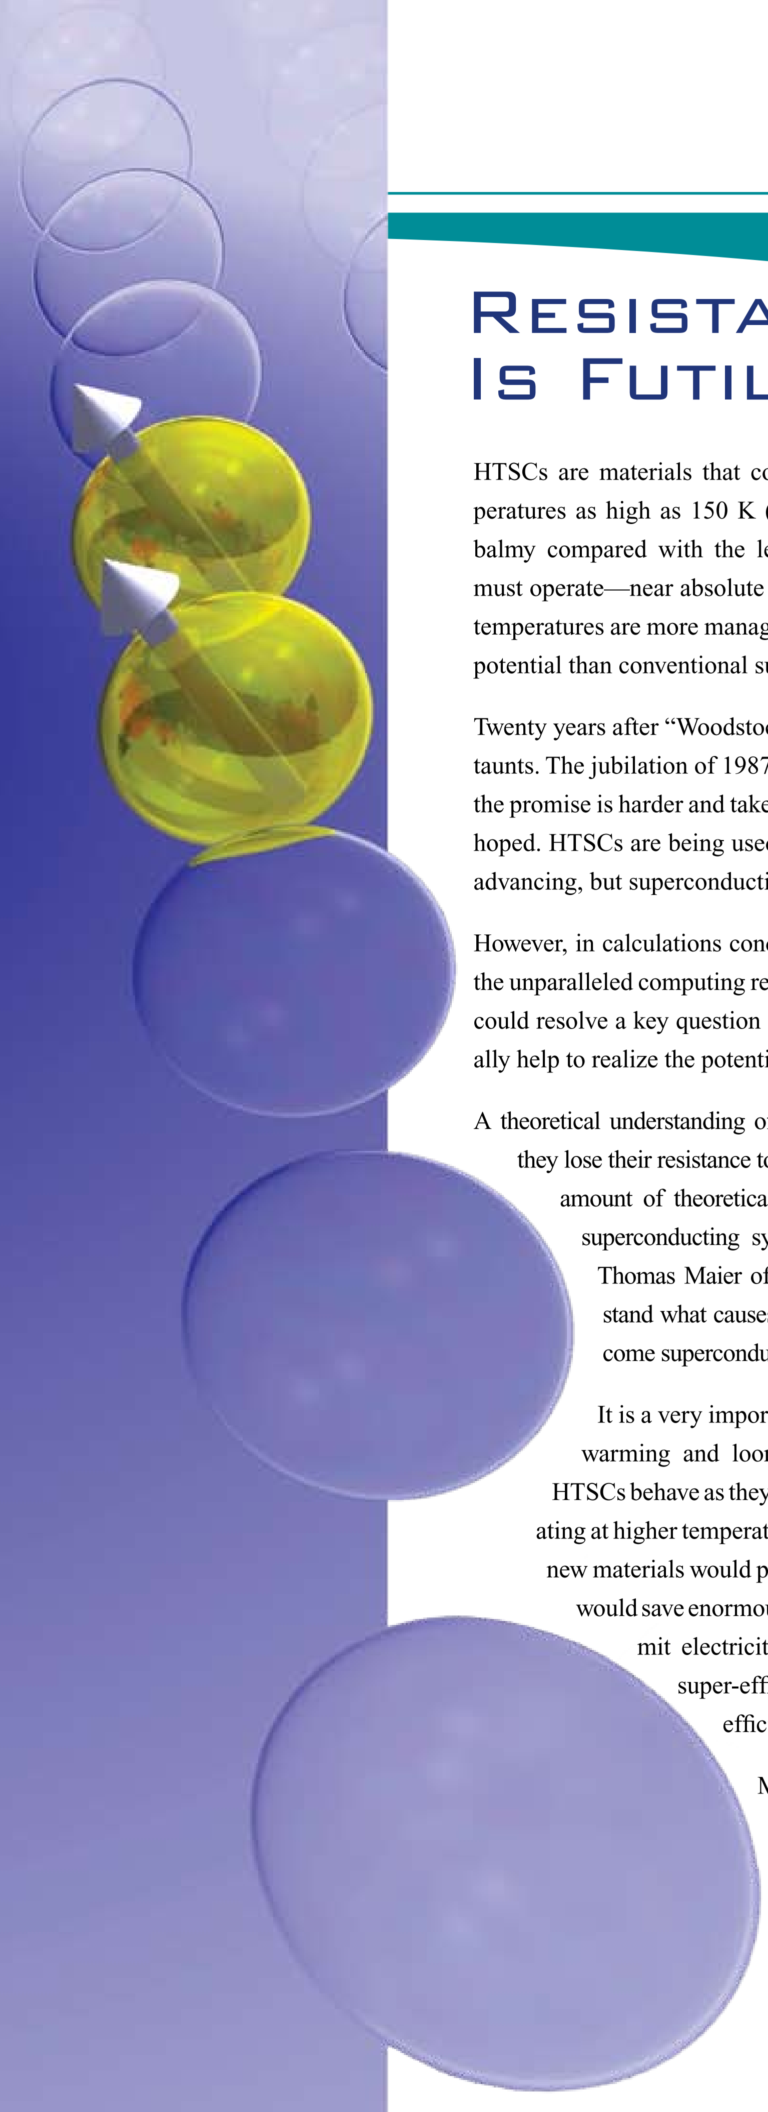

HTSCs are materials that conduct electricity without resistance at temperatures as high as 150 K (roughly −190°F). That's really cold, but it's balmy compared with the level at which conventional superconductors must operate—near absolute zero, or below −400°F. Because these higher temperatures are more manageable, HTSCs have much more technological potential than conventional superconductors.

Twenty years after "Woodstock," the promise of HTSCs tantalizes but also taunts. The jubilation of 1987 came up against a practical reality: fulfilling the promise is harder and takes longer than those starry-eyed physicists had hoped. HTSCs are being used in some applications, and the technology is advancing, but superconductivity is far from being a part of daily life.

However, in calculations conducted during 2006, a team of scientists used the unparalleled computing resources at NCCS to identify a mechanism that could resolve a key question about how HTSC materials work and eventually help to realize the potential of these highly energy-efficient materials.

A theoretical understanding of the most basic question about HTSCs—why they lose their resistance to electricity—eludes us. "There has been a huge amount of theoretical and experimental work on high-temperature superconducting systems, but no complete understanding," said Thomas Maier of ORNL. "One especially would like to understand what causes the pairing interaction, why these systems become superconducting."

It is a very important question, particularly in a world of global warming and looming energy shortages. Understanding why HTSCs behave as they do would open the door to new materials operating at higher temperatures, perhaps even at room temperature. These new materials would pave the way for breakthrough applications that would save enormous amounts of energy—power cables that transmit electricity without losses, practical electric vehicles, super-efficient high-speed trains, and leaps in energy efficiency for all sorts of electrical machinery.

Maier is part of a team working at the NCCS to develop a theoretical description of HTSCs.

Simulations of embedded atom clusters revealed that spin fluctuations cause electrons to form a superconducting state in the Hubbard model of cuprate superconductors.

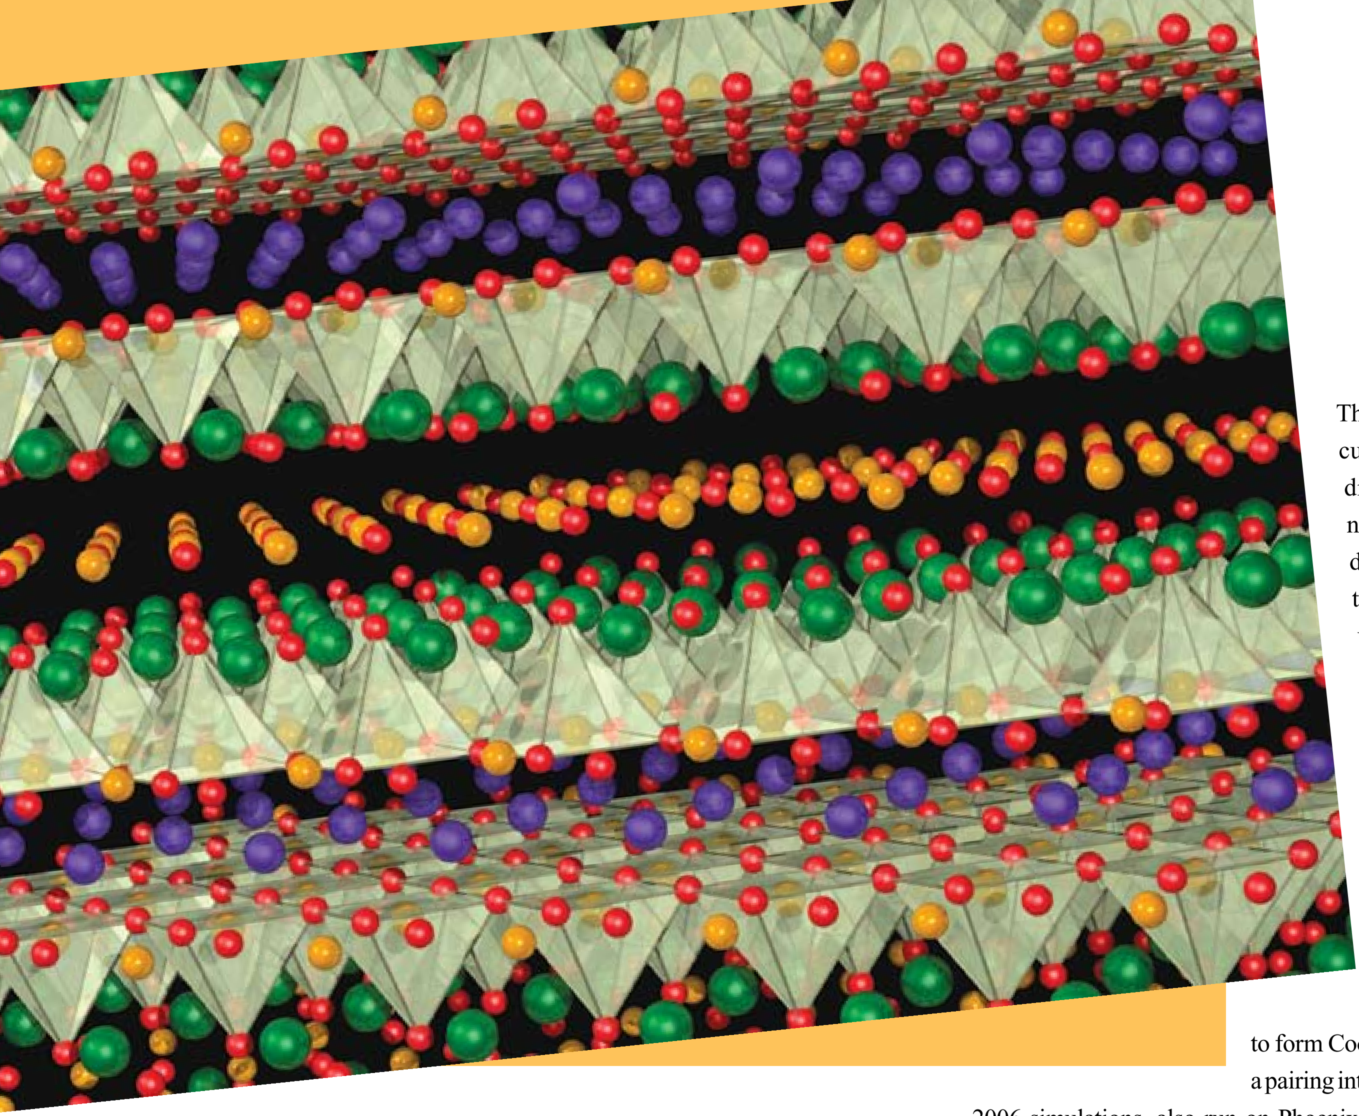

Crystal structure of the high-temperature cuprate superconductor $YBa_2Cu_3O_7$. The superconducting copper-oxide planes are modeled with a two-dimensional Hubbard model.

One mystery they are trying to solve is why the electrons in some materials bond to form pairs called "Cooper pairs," which settle into a state in which they conduct electricity without resistance. "The ultimate reason a material becomes superconducting is that the electrons pair into Cooper pairs," Maier explained. "If you drive a current through the system, the superfluid phase formed by the Cooper pairs doesn't resist it."

As a basis for their calculations, the scientists are studying the two-dimensional Hubbard model, believed by most physicists to provide an appropriate framework for describing the physics underlying HTSCs. The team resolved a key question in 2005 with simulations showing that superconductivity emerges as a result of strong interactions among electrons. Those calculations, conducted on the Cray X1E Phoenix supercomputer, were the first solution of the Hubbard model ever to include a large enough atom cluster to provide confidence in the results.

The 2006 calculations addressed the next step in developing the theory: uncovering the mechanism that underlies the Cooper pairing. Scientists know that some force causes electrons to form Cooper pairs via a pairing interaction. The 2006 simulations, also run on Phoenix, were aimed at determining what that force is. They revealed that the interaction behind Cooper pairing is driven by a mechanism called spin fluctuation (a magnetic effect associated with the rotation of electrons).

"All the structure in this pairing interaction comes from the spin fluctuation contribution. We have shown that in the Hubbard model commonly believed to be a description of high-temperature superconductors, it is spin fluctuation that mediates the pairing that leads to superconductivity," said Maier.

All of the known HTSCs are cuprates, materials with copper-oxide planes separated by atoms of other elements. Theorists have long speculated that spin fluctuations could lead to the type of symmetry found in Cooper pairs in the cuprates, Maier said. But the calculations performed at the NCCS were the first conducted at an adequate scale to analyze a microscopic model for the source of the pairing interaction and confirm the theorists' understanding.

Doug Scalapino of the University of California–Santa Barbara was one of the first to propose that spin fluctuations underlie the pairing mechanism in HTSCs. Scalapino, a key figure in the HTSC research community and a contributor to the NCCS project, said the new findings help relate the Hubbard model to lab experiments on cuprates, particularly neutron scattering and angle-resolved photoemission scattering (ARPES). The NCCS result shows ways to address the pairing mechanism by using data obtained from actual materials and these techniques.

Researchers don't know how to directly measure the pairing mechanism in materials, Scalapino explained, but the Hubbard model used to reveal the pairing mechanism can also be used to obtain neutron scattering and ARPES results. Once that is done, "we can find the particular combination, the certain way of putting those together, that does a good job of predicting what we know the pairing interaction is," said Scalapino. Experimentalists then can use that formula to do the same thing for the actual material—use neutron scattering and ARPES data obtained directly from materials to calculate the pairing interaction in real materials.

The team will build on the 2005 and 2006 breakthroughs to answer other critical questions. For example, what causes different cuprates to have different transition temperatures (i.e., to become superconducting at different temperatures)? What mechanism besides spin fluctuations might enhance the pairing interaction and affect transition temperatures? One possibility suggested by experiment is inhomogeneities in a material, said Maier. In 2007 his team will add inhomogeneities to their model and see how they affect the pairing interaction and the transition temperature.

The calculations for the inhomogeneities will probably be conducted on NCCS's Cray XT4 Jaguar supercomputer instead of the Cray X1E Phoenix because the problem will demand a large number of processors, whereas the previous simulations required a smaller number of faster processors.

"The interest ultimately is in how to increase the transition temperature to room temperature," Maier said. "If we can understand why a material is superconducting at 150 K, then we can ask what must we do to raise that temperature. It's not a guarantee that once we understand it, we can raise the transition temperature, but it's a big step."

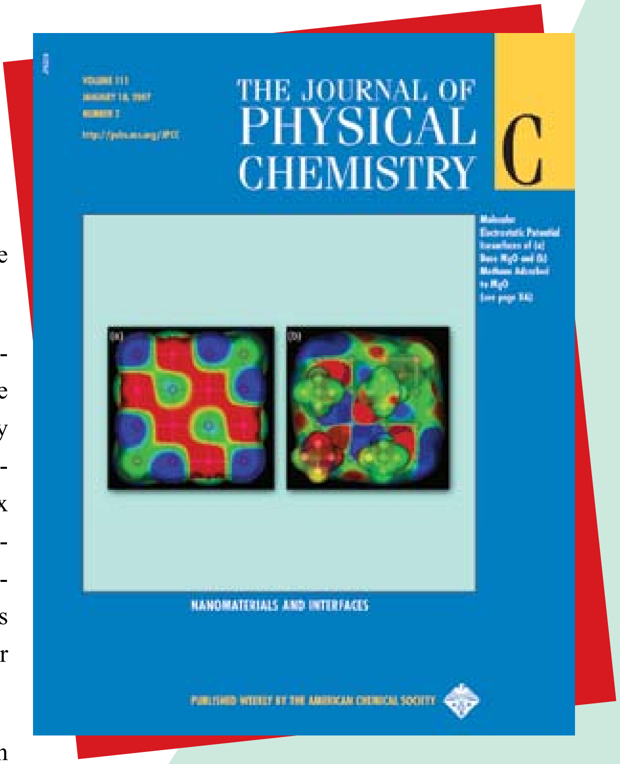

Getting Attached

Adsorption, the small-scale process by which one molecule attaches to another, plays a decidedly large-scale role in our lives. It is used to make thousands of common products—from margarine to gasoline—and it is the process by which chemical sensors keep us safe at airports and fuel cells will likely store hydrogen in the high-tech future.

Given the importance of this process, it makes sense that we would want to understand it as well as possible. Four Oak Ridge researchers using high-performance computers at the NCCS have significantly enhanced our understanding of adsorption, an achievement that landed them on the cover of a prestigious chemistry journal. Michael Drummond, Bobby Sumpter, William Shelton, and John Larese are the authors of "Electronic Structure Investigation of Surface-Adsorbate and Adsorbate-Adsorbate Interactions in Multilayers of CH_4 on MgO(100)." The paper was featured as a cover article in the *Journal of Physical Chemistry C*, a publication of the American Chemical Society that focuses on nanomaterials and surface science.

Sumpter noted that the interaction of molecules with a surface is a fundamental problem facing scientists, but it relates to a variety of important applications, including catalysis and chemical sensors.

"It's very important to understand what drives the adsorption of gases, for example, on surfaces," he explained. "In catalysis, the first thing that happens is a molecule adsorbs on the surface. With sensors, if you're trying to sense some gas—which is a big deal in airports—the first thing that happens is you've got to get the molecule on the material that's part of the sensor."

Sumpter said the collaboration began with a question from Larese, a neutron scientist with ORNL's Chemical Sciences Division, and became a good example of computer simulation complementing experimental results.

"That's typical. Theory and computation can do two things. We can try to understand an experiment, or we can actually come up with a material and tell experimentalists they should look into it."

Picking Up Good Vibrations

Proteins are highly efficient biological machines, but some "tinkering" could make them better.

That conviction drives the use of scientific simulation to provide a more detailed understanding of how biomolecules such as proteins and enzymes operate. (An enzyme is a type of protein that catalyzes or speeds up a chemical reaction.) Ultimately, the research, conducted on the NCCS's Jaguar supercomputer, is expected to pave the way for designing enzymes that operate more efficiently and harnessing them to improve the efficiency of chemical processes.

"The implications of the fundamental understanding of biomolecular systems will be felt in many fields, including health and medicine, energy, and the environment," said Pratul Agarwal of ORNL, principal investigator for the project.

The pioneering biomolecular simulations have produced major new insights into protein structure, dynamics, and functions. A key finding in 2006 was that the shape or design of an enzyme in nature is closely connected to the type of its internal motions or "vibrations." In the past enzymes were viewed as static structures, but computational studies reveal that they are "dynamic entities constantly undergoing conformational fluctuations," Agarwal said. He compared the motions inside proteins at the molecular level to the constant fluttering of trees in a breeze.

The new insights about protein vibrations were important in several aspects of the project's investigations during 2006. For example, studies of enzyme shapes in organisms ranging in complexity from simple *E. coli* bacteria to human beings revealed that the vibration patterns in the enzymes have been retained in their entirety as organisms evolved into more complex forms, Agarwal said. It appears that over many evolutions, these simple-looking molecules have been optimized not only for structure, but also for internal dynamics and the role of those dynamics in protein function. The implications of this finding are significant for many chemical and industrial processes.

"Once we understand the nature of the vibrations, we can use these principles to design new enzymes for industrial processes," Agarwal explained. "We can start manipulating the vibrations to conduct the kind of chemistry we want on a much faster time scale." Clarifying the role of vibrations could also open doors to the design of new pharmaceuticals. Earlier simulations on Jaguar provided insights into how protein vibrations speed up the rates of chemical reactions inside enzymes. A more full understanding will indicate to scientists which amino acids in proteins they can change or mutate to make the chemical reactions more efficient, he said.

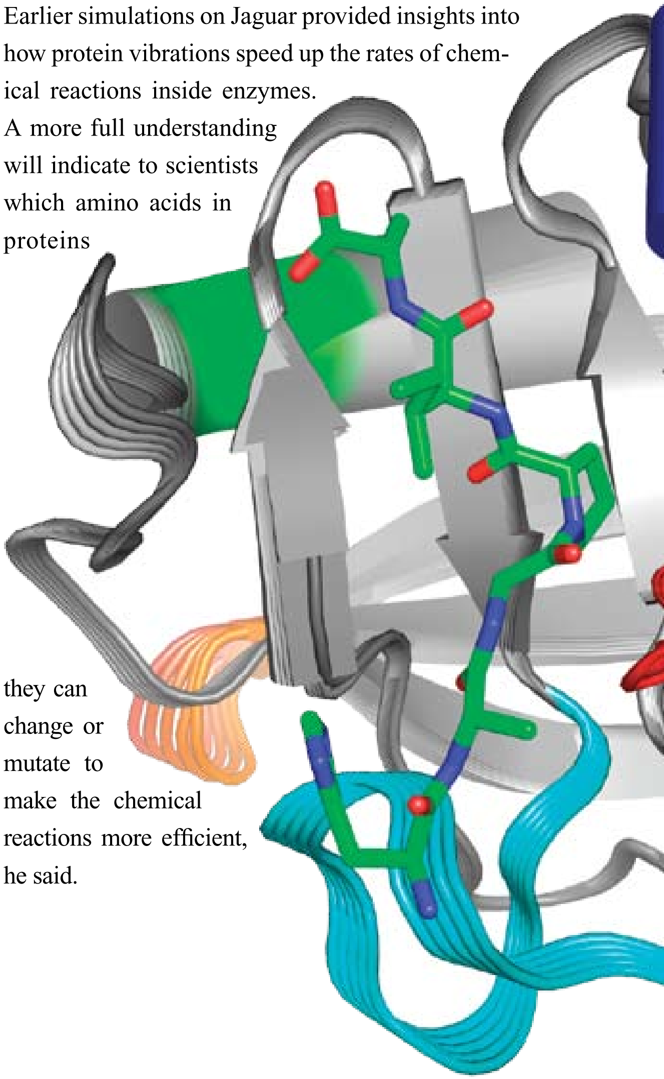

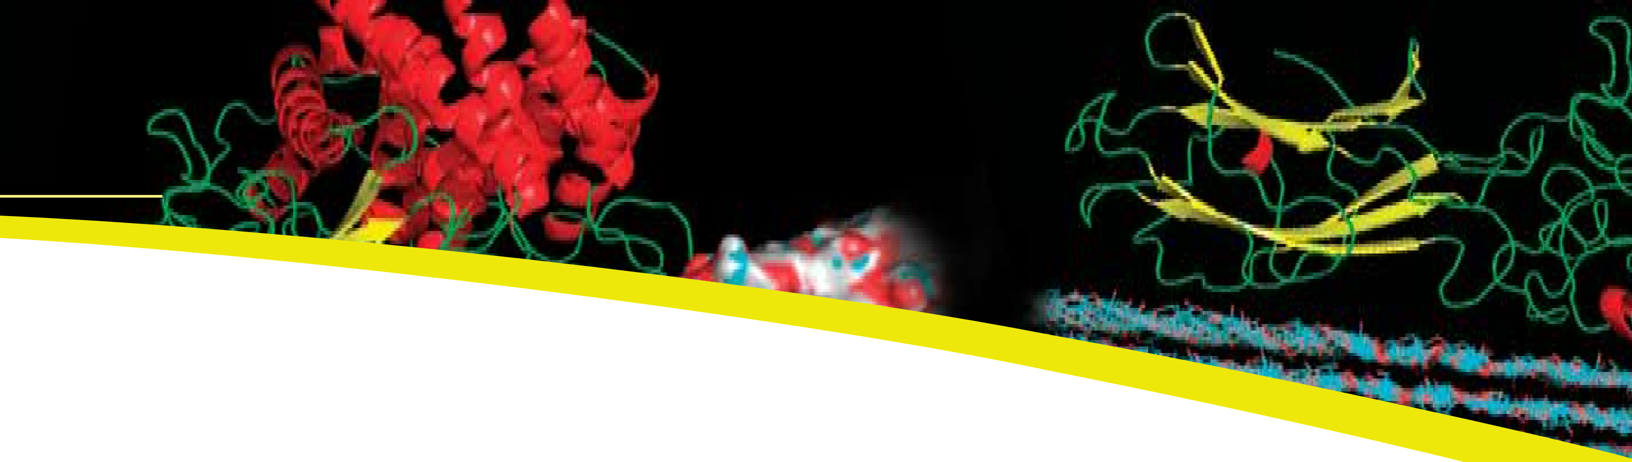

Another focus of the 2006 research was the internal dynamics of the cellulase enzyme, which converts the cellulose in biomass into sugars. This model has been created based on the crystal structure of Cel9A cellulose from *Thermonospora fusca*, which is available in the protein data bank. “Once we understand the mechanism of this enzyme, we will be able to design better or more efficient enzymes so we can improve the process for producing ethanol at low cost,” Agarwal said.

Cellulase function involves two different activities of interest. First, Agarwal explained, the enzyme binds to the cellulose surface and splits off individual cellulose chains, which involves breaking hydrogen bonds. Second, linkages in the cellulose chains are hydrolyzed to release glucose or slightly more complex sugar units, which can be converted to ethanol. Releasing sugars from the cellulose chains is what currently limits the efficiency and cost-effectiveness of ethanol production. More efficient cellulases are needed. Once the mechanisms for splitting off cellulose chains and for releasing sugars are understood, Agarwal said, scientists will have an idea of the amino acids that can be optimized through protein engineering for more efficient conversion of cellulose to sugar.

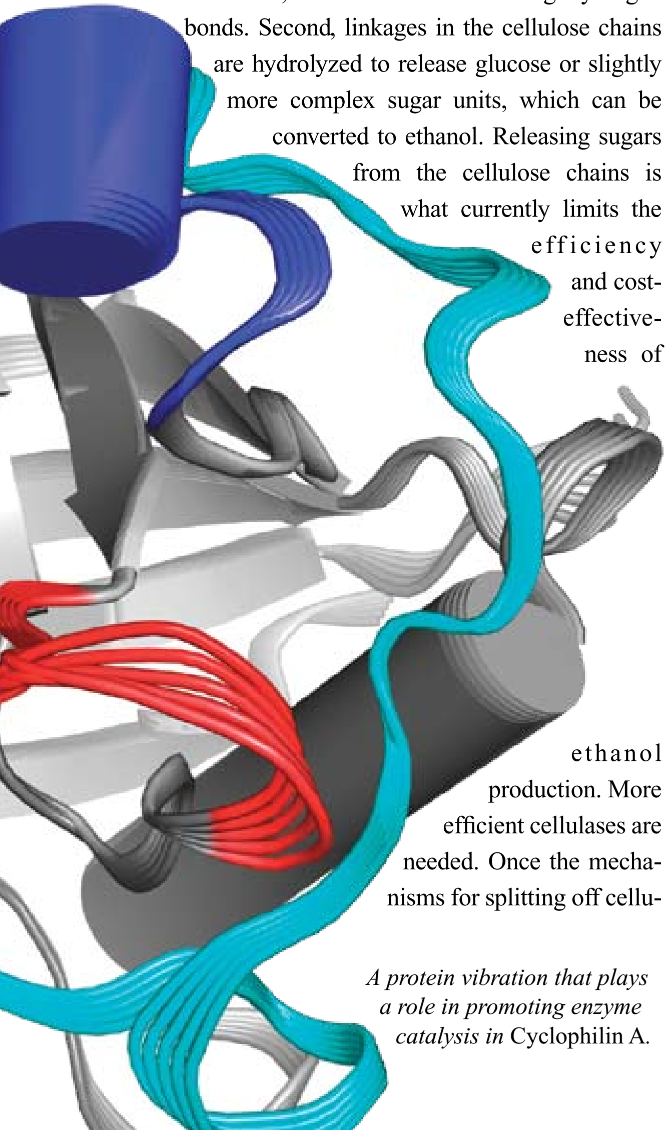

A protein vibration that plays a role in promoting enzyme catalysis in Cyclophilin A.

Another enzyme being investigated is *Cyclophilin A* (CypA), a protein of medical interest because it plays a role in HIV-1 infection. HIV-1 cannot infect a cell in the absence of CypA because, scientists believe, the viral particle cannot assume its proper shape without the enzyme. CypA also influences protein folding (the process by which a chain of amino acids folds into a three-dimensional protein structure). Correct folding is essential for proteins to function properly, and misfolded proteins cause disorders such as Alzheimer’s disease and brain-wasting illnesses. The NCCS simulations have revealed a network of protein vibrations that plays a role in making energy available for CypA to function, Agarwal said. Understanding the mechanism of CypA functioning and its role in the HIV-1 life cycle could guide researchers in finding a way to slow the spread of the AIDS virus.

It is difficult and expensive to investigate biomolecules through experimental techniques. The effort required for understanding their functions at the cellular level is very high because every aspect of a molecule requires a separate experiment, Agarwal noted. “The power of computational techniques such as the ones we are using is that they provide a very good collective set of information on a variety of time and length scales. So we get a very complete picture from computational studies.”

Agarwal’s project is using the AMBER code package, including its PMEMD simulations engine, and LAMMPS. The 2006 runs used from 500 to 2,000 processors of the Cray XT3. “These simulations on NCCS resources allowed us to investigate multiple proteins from different species simultaneously to provide new insights,” he said.

Large-scale simulation is providing insights into the functioning of protein complexes such as F1-ATP.

In 2007 the team will continue to investigate the nature of protein vibrations and how they promote catalysis, Agarwal said. They will also continue to explore the conservation of vibration patterns as organisms evolve into more complex forms. The results of those simulations will help researchers understand what aspects of the shapes of proteins can and cannot be changed—knowledge that has implications for engineering proteins to operate more efficiently.

The 2007 simulations will use 4,000 to 8,000 processors of the Jaguar XT4 (or up to 70 percent of its capacity), Agarwal said. The team is working on scaling its codes so they can use as many as 20,000 to 40,000 processors once computers of that capacity come online. They are collaborating with ORNL's Future Technologies Group on that task. They are also considering other codes such as NAMD for their suitability to run on a high number of processors.

As the NCCS prepares to provide computing power at the petascale and beyond, Agarwal's team looks forward to employing the expanding resources to gain a more accurate, complete picture of the micromachinery of living things.

The Little Engines that Can

Harvard researcher Martin Karplus and his team used groundbreaking simulations on Jaguar to reveal the essential role of two molecular motors, F1ATP-synthase and DNA polymerase, that power everyday biological processes.

Karplus was the principal investigator for the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program's "Molecular Dynamics Simulation of Molecular Motors" project.

F1ATP-synthase produces ATP, or adenosine triphosphate, from adenosine diphosphate and inorganic phosphate. ATP-synthase contains a rotary motor with a rotation generator that provides the energy required for ATP synthesis.

The ATP molecule carries out reactions necessary for cellular functioning and supplies energy for our every action. The average person consumes at least 40 kilograms of ATP every day at rest, more if active. Typically, a human will use his/her body weight of ATP over the course of a day.

The ATP-synthase rotary motor, believed to be the smallest in existence, turns food into usable energy. It resembles automobile engines in many ways, but its performance leaves them in the dust: its shaft rotates 1,000 times in a second, producing 3 molecules of ATP each time, for a total of 3,000 molecules per second.

The second motor, DNA polymerase, copies nucleic acid to make new cells or reproduces DNA. Basically, the motor uses one strand of DNA as a template and makes a copy. DNA polymerase conducts almost error-free duplication (1 error in 100,000 copies).

The polymerase motor essentially sits on a nucleic acid to be copied and jumps from step to step, adding one base (letter of the genetic code: C, G, A, and T) at a time. The movement from base to base, known as translocation, is at the heart of Karplus's research, in which his team tries to elucidate the relationship between energy and translocation.

Karplus's research has a variety of potential applications. For example, understanding the high-fidelity mechanism in polymerase might make it possible to eliminate mistakes in duplication. (Mistakes lead to disease or irregularities; too many, and the organism will die.) Inhibiting the ATP-synthase motor could lead to techniques to prevent cancers from spreading. ATP-synthase is also being explored for uses in making nanodevices, and the sheer efficiency of these miniature machines makes them of interest to all branches of science.

A white dwarf star is destroyed, leaving a radioactive cloud of nickel and other elements that glows as bright as a galaxy for weeks.

WHEN YOU DETONATE A STAR...

In 1998 two teams of astronomers—one based at Lawrence Berkeley National Laboratory and the other at Mount Stromlo Observatory in Australia—decided to see if the universe is still expanding, and how fast.

If what we saw of the universe were all there is, the primary force acting on it would be the gravitational pull of stars, galaxies, and other matter, which should be slowing it down. Instead, the two teams discovered not only that the universe is expanding, but that the expansion is accelerating.

Their conclusion, dubbed the "Breakthrough of the Year for 1998" by *Science* magazine, has profound implications. If the expansion of the universe is accelerating, something else must be at work. Therefore, the teams' observations provide compelling evidence for a newly discovered force—dubbed "dark energy" by physicists—that pervades empty space.

The American and Australian teams made their startling discovery by observing the universe's biggest thermonuclear explosions, stars that become as bright as galaxies for a few weeks as they blow apart. Known as Type Ia supernovas, these explosions serve a key role in cosmology by revealing both how far away the stars are and how fast they are receding.

Astronomers can calculate the distance to a Type Ia supernova because it produces a characteristic light curve, with the absolute brightness of the explosion closely linked to the time it take to brighten and fade. Given information from expansion of the nearby universe, they also have an idea of how big the universe should be. In short, the most distant Type Ia supernovas are farther away than expected, making the universe bigger than predicted and leading to the conclusion that its expansion is accelerating rather than slowing.

So supernovas are being used to answer some profound questions. But to fully understand what they are telling us, scientists need to resolve a few critical questions about the supernovas themselves. The most important could be those that focus on the relationship between the supernova's light curve and its brightness. Simply put, does this relationship hold for all Type Ia supernovas? Astrophysicists discovered the relationship by observing nearby supernovas whose distances can be verified in other ways. They assume that the relationship holds for more distant explosions as well, an assumption that allows them to gauge the distance to events billions of light-years away. They would love to see proof that their assumption is valid.

To get that proof, they will need a better understanding of how Type Ia supernovas blow up. "We would

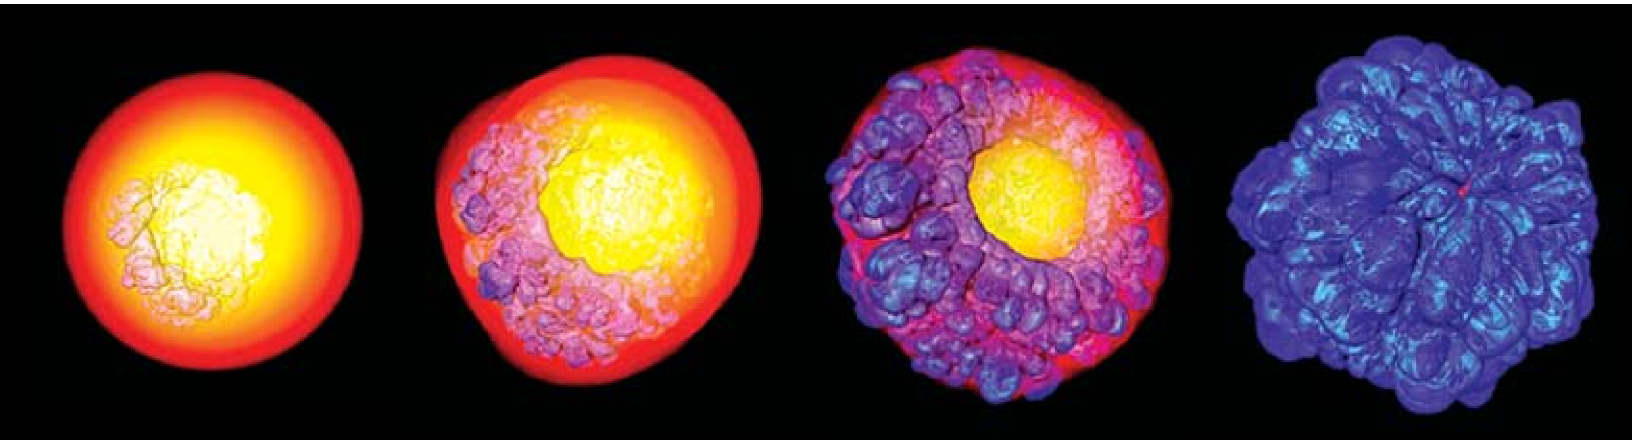

Using Jaguar, astrophysicists are able to examine realistic 3-D models of the ignition of Type Ia supernovas.

like to understand the physics underneath," explained Stan Woosley, astrophysicist at the University of California–Santa Cruz. "Type Ia supernovas are still one of the most important pieces of evidence for an accelerating expansion of the universe."

A team led by Woosley is using the enormous computing power of the NCCS and its Cray XT4 Jaguar supercomputer to provide explanations by simulating the evolution of Type Ia supernovas.

Scientists believe a Type Ia supernova begins with a dead star known as a white dwarf. Having burned through its lighter elements and not being massive enough to burn through its heavier ones, a white dwarf is essentially a massive chunk of carbon and oxygen floating through space. To become a supernova, it must be closely teamed with another, less evolved star—a red giant—which rains hydrogen onto the white dwarf.

In this scenario the white dwarf becomes increasingly more massive as it takes on this new material, and its core becomes hotter as the dwarf's increasing gravity adds pressure. It goes through a centuries-long period of convection, in which hot material at the core rises and sinks like water boiling in a pot. As the white dwarf approaches a critical mass—nearly half again the mass of the sun—it goes through a runaway fusion reaction that eventually blows the star apart.

This runaway reaction also produces nearly a sun's mass of iron and radioactive nickel. It is the decay of this nickel, with a half-life of about 6 days, that makes the supernova billions of times as bright as the sun.

Woosley's team is simulating each aspect of the supernova: the long period of churning convection that leads to the supernova's first flames, the expansion of these flames to create an explosion, and the radiation at the end of the event that gives the supernova its brightness.

The job is exceptionally demanding, with scales ranging from the microscopic to the stellar. The simulations must include the initial flames, which have a thickness in the neighborhood of four one-hundred-thousandths of an inch, and the entire white dwarf, which is roughly the size of the earth. Without the leadership computing resources of the NCCS, calculations of such size and scale would be impossible.

As it moves forward, the team will be addressing many areas about Type Ia supernovas that are not well understood. Does the flame start at the center or off center? Is there just one flame or many? Does the explosion exceed the speed of sound and thereby become a detonation? And why are these events so consistent?

As it moves to answer these questions, Woosley's team will provide a key piece in the puzzle of where our universe is headed.

The Spin Zone

In 1967 a Cambridge University graduate student poring over data from a newly constructed radio telescope noticed something very odd—a radio signal blinking regularly from a far corner of the sky.

The signal came from a pulsar, the spinning remnant of a great stellar explosion. Pulsars are the enormously massive—smaller than the moon, heavier than the sun—spinning leftovers from core-collapse supernovas, the cataclysms that provided many of the elements on earth and made our own lives possible. A pulsar appears to blink because radiation shoots out of its magnetic poles, which, as with the earth, can be tilted a little from its axis of spin. As a result, a pulsar behaves like a stellar lighthouse, pointing at an observer once with each rotation.

Nearly 4 decades later, a team of scientists using ORNL supercomputers has discovered the first plausible explanation for a pulsar's spin that fits the observations made by astronomers. The discovery, from a team led by Anthony Mezzacappa of ORNL and John Blondin of North Carolina State University, was that the spin of a pulsar comes not from the spin of the original star but from the shock wave created when the star's massive iron core collapses.

Instabilities in that shock wave lead to two rotating flows—one in one direction directly below the shock wave and another, inner flow that travels in the opposite direction and spins up the core (see image at right). The discovery, featured in the journal *Nature*, comes from the team's use of three-dimensional simulations and the enormous computing power that made those simulations possible.

The team eagerly anticipates further advances in high-performance computing that will be coming to ORNL, as more powerful systems will be needed to describe the core-collapse supernovas with sufficient realism and detail.

"In a nutshell, this rapid advance in supercomputing technology will give us the tools to understand these events and their role in our universe," Mezzacappa explained. "This is a very, very exciting and very satisfying thing."

Visualization by Kwan-Liu Ma, UC-Davis

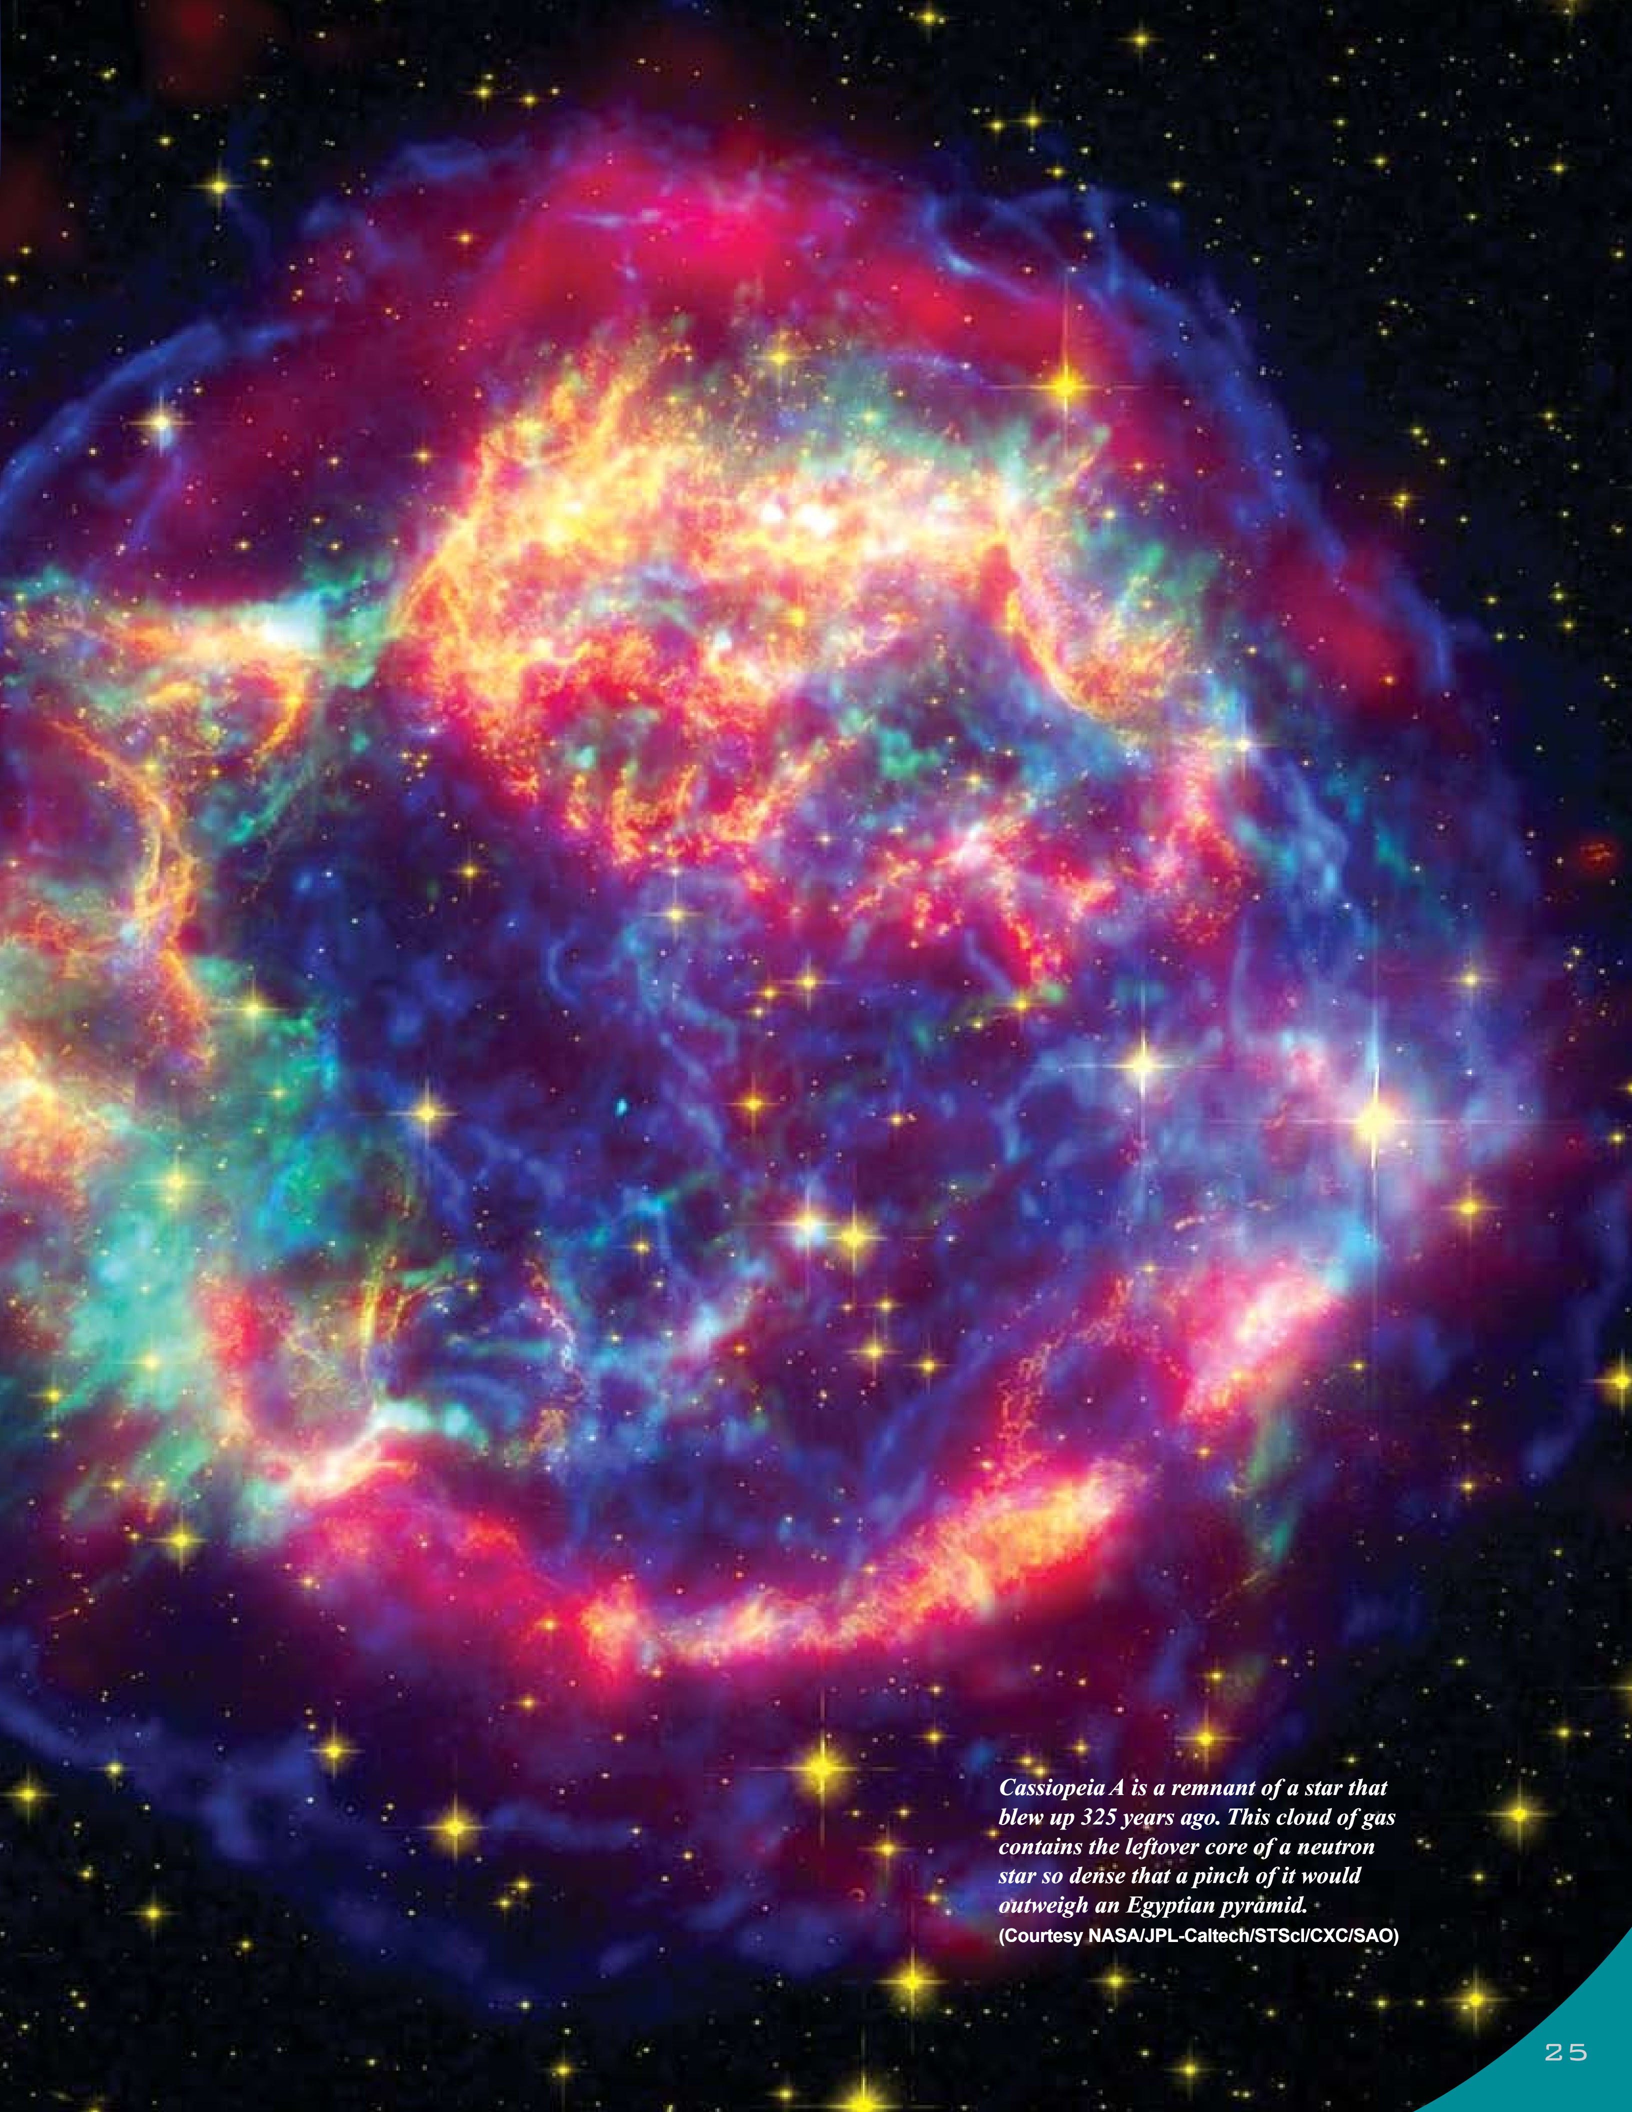

Cassiopeia A is a remnant of a star that blew up 325 years ago. This cloud of gas contains the leftover core of a neutron star so dense that a pinch of it would outweigh an Egyptian pyramid.

(Courtesy NASA/JPL-Caltech/STScI/CXC/SAO)

INDUSTRY USERS JOIN NCCS COMMUNITY

The NCCS opened its resources to industry in 2006, allowing researchers at private companies to join in the breakthroughs being achieved by university and national laboratory scientists.

The Boeing Company, General Atomics, and Dreamworks Animation joined university researchers and became NCCS users through the INCITE program. The decision to bring private business into the program grew out of a recommendation from the Council on Competitiveness, which recognized that private researchers were using high-performance computing to conduct important research vital to U.S. technical leadership and were prepared to make use of the center's leadership facilities.

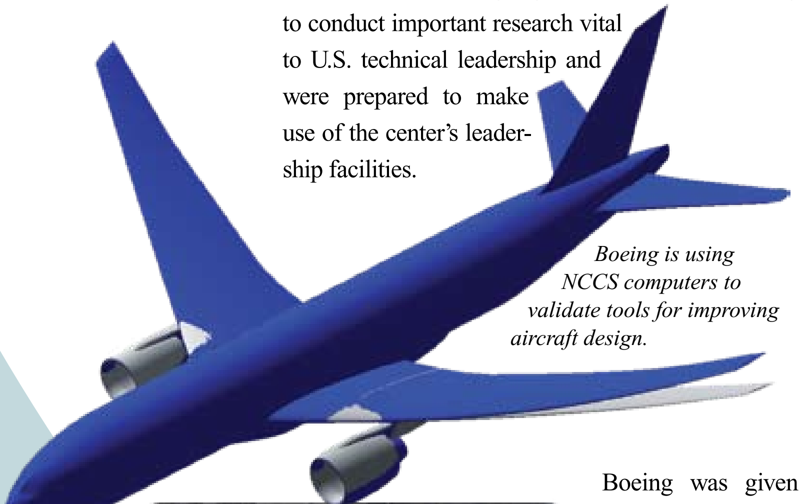

Boeing is using NCCS computers to validate tools for improving aircraft design.

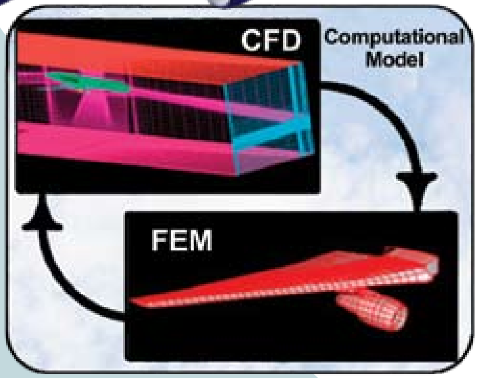

Images used with permission from Boeing

Boeing was given 200,000 processor hours on the Cray X1E Phoenix supercomputer to validate the use of powerful computational fluid dynamics tools in designing future aircraft. The Boeing simulations were not used directly in airplane production. Rather, they demonstrate the applicability and predictive accuracy of the computational tools.

One key experiment in this project simulated the effect on a wing when a flap is suddenly deployed. In the past such analysis had not included both the structural response of the wing and the airflow around it. Boeing's simulations required the coupling of two complex codes: one simulating aerodynamics and the other simulating the structural response of the wing.

General Atomics was given 400,000 processor hours on Phoenix to simulate turbulence in the plasma used to power fusion reactors such as the proposed ITER reactor. Turbulence is among the most vexing problems faced by fusion researchers, and the General Atomics team focused on the interaction of two types: one caused by temperature differences among ions and another, much smaller-scale turbulence caused by temperature differences among electrons.

Dreamworks was given 950,000 processors hours on the center's Cray XT3 Jaguar supercomputer to expand the limits of real-time ray tracing animation, in which light rays are simulated interacting with objects. The project promises to benefit fields such as medical imaging, defense simulations, and visualizations in general as well as the production of animated feature films.

All told, the NCCS hosted five INCITE projects in 2006. Besides the teams from private industry, the center welcomed users from Harvard University and the California Institute of Technology (Caltech) into the program.

This year was the NCCS's first in the INCITE program. In 2007 more than 80 percent of the NCCS leadership systems will be dedicated to INCITE projects. These projects and others promise to help American industry maintain its competitive position in the world.

OPERATIONS PERSPECTIVES

NCCS INFRASTRUCTURE

Continuing upgrades in computing capacity in 2006, and plans for even larger projects to come, drove modifications to the NCCS infrastructure that supports the leadership computers and the research that depends on them.

Fresh off the task of installing the mammoth Cray XT3 Jaguar in 2005, ORNL and NCCS staff turned to putting another, larger Cray in place in 2006 and laying the groundwork for further expansion through 2008. Space, power, cooling, and operational infrastructure were enhanced to prepare for the installation of 68 cabinets for a new Cray XT4 system on the second floor of the Computational Sciences Building (CSB), the relocation of the 56 cabinets housing the original Jaguar (the Cray XT3) to the second floor, and another planned upgrade in 2007.

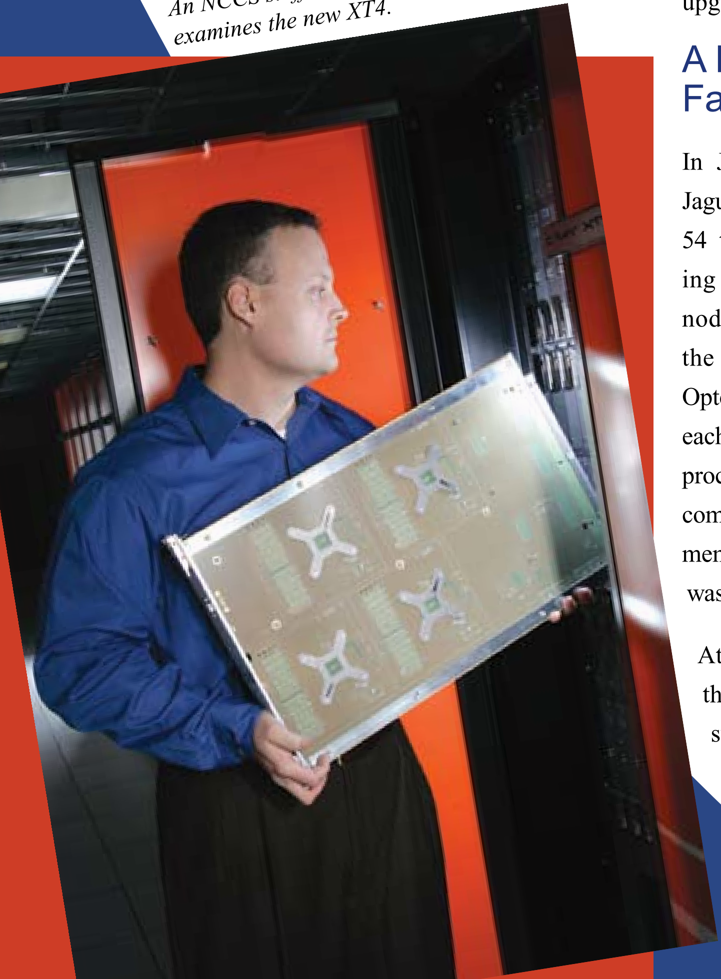
An NCCS staff member examines the new XT4.

A Bigger, Faster Jaguar

In July the 25-teraflop Jaguar was upgraded to 54 teraflops by removing its 1,303 compute nodes and replacing the four single-core Opteron processors on each node with dual-core processors, doubling the computer's capacity. The memory on each node was also doubled.

At the same time, the bandwidth of the system was doubled by upgrading to 23-foot cables, two new power supplies were added to each cabinet, and the diagnostic controls in each cabinet were updated to support the dual-core chips and their increased power consumption.

Then in November 2006 the NCCS took delivery of a new 68-cabinet Cray XT4 with a peak processing speed of 65 teraflops. Each of its cabinets is configured with 24 compute nodes, and each node contains four dual-core 2.6-gigahertz, 64-bit AMD Opteron processors and 4 gigabytes of memory. In all, the system has a total of 6,528 processors and 26 terabytes of memory.

The XT4 was installed on the second floor of the CSB. Once it was configured, accepted, and prepared for users, applications running on the XT3 on the first floor were moved to the new system. Installing the XT4 and bringing it online was only a first step in the upgrade process, however. Once the XT4 was up and running, the XT3 on the first floor was disassembled, moved to the second floor, and integrated into the new XT4 to form an even higher-capability system. The combined XT3/XT4 system, also called Jaguar, delivers a peak processing speed of 119 teraflops and 46 terabytes of aggregate memory.

The new Jaguar boasts a total of 11,708 processors, 11,508 of which are compute processors and the rest service processors. Each node is connected to a Cray Seastar router through Hypertransport, and the Seastars are interconnected in a three-dimensional-torus topology. The interconnect is designed for high bandwidth, low latency, and extreme scalability. Users can scale applications


The Cray XT4 Jaguar computer.

from 200 to more than 23,000 cores without sacrificing performance.

The operating system is UNICOS/lc, a combination of Linux on the service nodes and the Catamount microkernel on the computing nodes. Catamount is designed to minimize system overhead, thus allowing scalable low-latency global communication. Approximately 600 terabytes of scratch disk space are available in the system for use by the Lustre file system.

The next upgrade will take Jaguar to 250 teraflops in late 2007. Each processing node will be pulled from the racks, and the dual-core processors will be replaced with quad-core elements.

Other Primary Systems

The 18.5-teraflop Phoenix Cray X1E, another NCCS primary system, stayed put on the first floor of the CSB. Phoenix has 1,024 multistreaming vector processors and 2 terabytes of globally addressable memory. Its vector processors, highly scalable system architecture, high memory bandwidth, and memory system that allows each node direct access to memory on other nodes enable Phoenix to perform at higher speeds than many larger machines. It remains the computer of choice for applications that need its vector architecture.

Also on the first floor is Ram, a 256-processor SGI Altix with 2 terabytes of shared memory and a peak performance of 1.5 teraflops. Ram is available to users of the primary systems for pre- and postprocessing tasks. Its large shared memory facilitates the analysis of very large data sets.

Networking and Storage

Faster computers demand faster networks to efficiently deliver data.

The Phoenix Cray X1E computer.

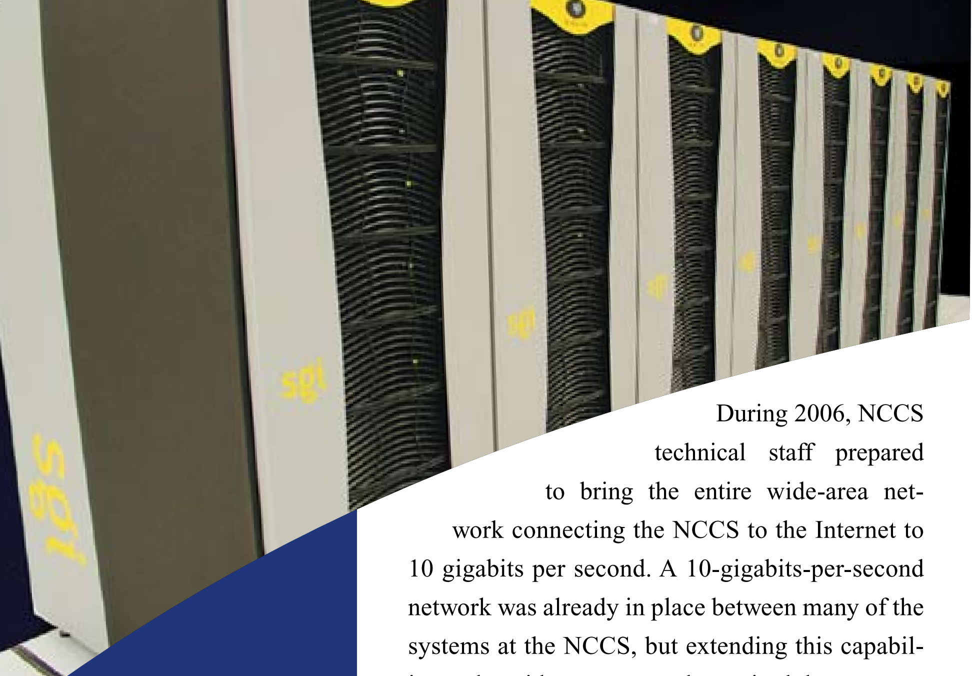
The SGI Ram computer.

During 2006, NCCS technical staff prepared to bring the entire wide-area network connecting the NCCS to the Internet to 10 gigabits per second. A 10-gigabits-per-second network was already in place between many of the systems at the NCCS, but extending this capability to the wide-area network required that new cybersecurity methods and procedures be adopted and reviewed prior to deployment.

With the faster network in place, system users can move multiterabyte data sets to their home institutions or to the larger research community at several hundred megabytes per second. The rollout is scheduled for April 2007.

Implementation of a Lustre file system, called Spider, at the NCCS was close to completion in late 2006. Spider will provide a file system that is directly accessible on the various NCCS systems, enabling users of NCCS computers to store data sets from their simulations and immediately perform data analysis and visualization using other NCCS resources. NCCS staff worked with Cluster File Systems, Inc., the maker of Lustre, to commission new features to support the NCCS configuration. Spider will allow users to connect to Jaguar through router nodes that connect to Jaguar's compute nodes so users can see the file system directly. Spider will initially be supported by 21 server systems and provide 85 terabytes of disk space. It is expected to be ready for users by spring of 2007. Eventually, it will connect to all of the leadership computing machines. Implementation will start with periphery systems such as the visualization systems and move up to the scale of Jaguar.

The upgrades to the networks and file systems were accompanied by the development of tools for effectively using file systems and efficiently moving large data sets between parallel file systems. To reap the maximum benefit from Spider, it will be important to move data rapidly from the local file systems of the computers to the Lustre system. In addition to an emphasis on data management tools, 2007 will also see increased attention to developing user tools such as compilers, debuggers, and performance-analysis tools that can scale to fit the needs of petascale computers and the applications than run on them.

The High-Performance Storage Sytstem (HPSS), which allows NCCS users to store and rapidly access their simulation data, continued to evolve this year. In mid-2006, version 6.2 of the software that operates HPSS was deployed. Planning also began for the release of version 7.1 in 2007 or early 2008. The HPSS enhancements include greater speed and more efficient handling of smaller data files.

Plans for 2007 include moving some of the devices

The NCCS 1st-floor computer room.

The High-Performance Storage System.

making up HPSS, particularly the StorageTek mass storage silos, to an area of ORNL's 4500 North building that housed earlier parallel processors before the CSB was built.

Things to Come

The next giant step for the NCCS will be the installation of a petaflop computer, scheduled for late 2008. In preparation, four new power transformers will be purchased in 2007 for installation in 2008. The new transformers will upgrade the power at the CSB from the present 7.3 megawatts to 15 megawatts. A new emergency generator and an additional 1-megawatt uninterruptible power source will also be installed. Additional piping will also be needed to supply adequate chilled water to cool the petaflop machine.

The petaflop system will be installed in the first-floor computer room at the NCCS. Additional space on the second floor will be reconfigured to accommodate up to 200 cabinets of disk storage for the petaflop machine and other potential projects.

By the time a petascale computer is in operation, the Spider file system is expected to be operating at 200 gigabytes per second. An InfiniBand® network will be used to connect Spider to the computing platforms, and teams will be working to create the InfiniBand fabric between the systems, scale it up, and ensure that the software can effectively move data between systems.

In addition to the petaflop computer planned for 2008 (commissioned by the DOE Office of Science), the University of Tennessee (UT), ORNL, and several other partners have submitted a proposal to the National Science Foundation to develop another petaflop computer, which would also be located at the NCCS. If that bid is successful, another four transformers will be needed to ramp up the power at the CSB to 35 megawatts to drive two petaflop machines.


TOUCH
GLASS
DO N
TOUC
GLAS

EXTREME MAKEOVER

The Jaguar XT4 installation.

Adding a room for a new world-class supercomputer isn't easy. Creating a new computer room for the combined XT3/XT4 system required total renovation of about 8,000 square feet of space of the second floor of the CSB. Four labs were relocated, walls were removed, computer equipment was relocated, and a 2-foot raised floor was taken up and replaced by a 3-foot raised floor.

To supply chilled water to heating, ventilation, and air-conditioning units that will remove the heat from the giant machines, construction crews installed two 16-inch water lines to the second floor from the Central Energy Plant behind the CSB and put a chilled-water loop under the new raised floor. The expanded cooling system for the new Jaguar includes 32 new 30-ton cooling units.

The XT3/XT4 will be air cooled. The cooling units will pump air chilled to 51 degrees Fahrenheit into the area under the raised floor, which will serve as a giant air plenum. Each of the 124 cabinets will have a 4,000-cubic-feet-per-minute variable-speed fan to constantly pull chilled air into the cabinet from the under-floor area. The cold inlet air will blow through the computer cabinet (a journey of about 6 feet) and exit at an average temperature of 125 degrees Fahrenheit.

Crews began work in early 2006 on a power upgrade for the CSB that was finished in May. A 3.3-megawatt transformer and a 4,000-amp switchboard were installed. Eleven 300 kilovolt-amp power distribution units were installed, eight new and three moved from the first floor computer room. That upgrade will provide adequate power through the upgrades planned for 2007.

To supply additional power for growth at the NCCS and other areas of ORNL, the Tennessee Valley Authority, the lab's power supplier, installed a new power substation in 2006. The $19 million facility, which came online in the fall, has the potential to supply 140 megawatts of electricity, compared with 40 megawatts for the 60-year-old substation it replaced. The new substation has advanced equipment and safety features, including controls that tie into the ORNL computer network to enable remote monitoring of the equipment. A fiber optic arc-detection system will automatically shut off power through the substation if a circuit shorts. The new substation also has room to add more capacity as the lab grows.

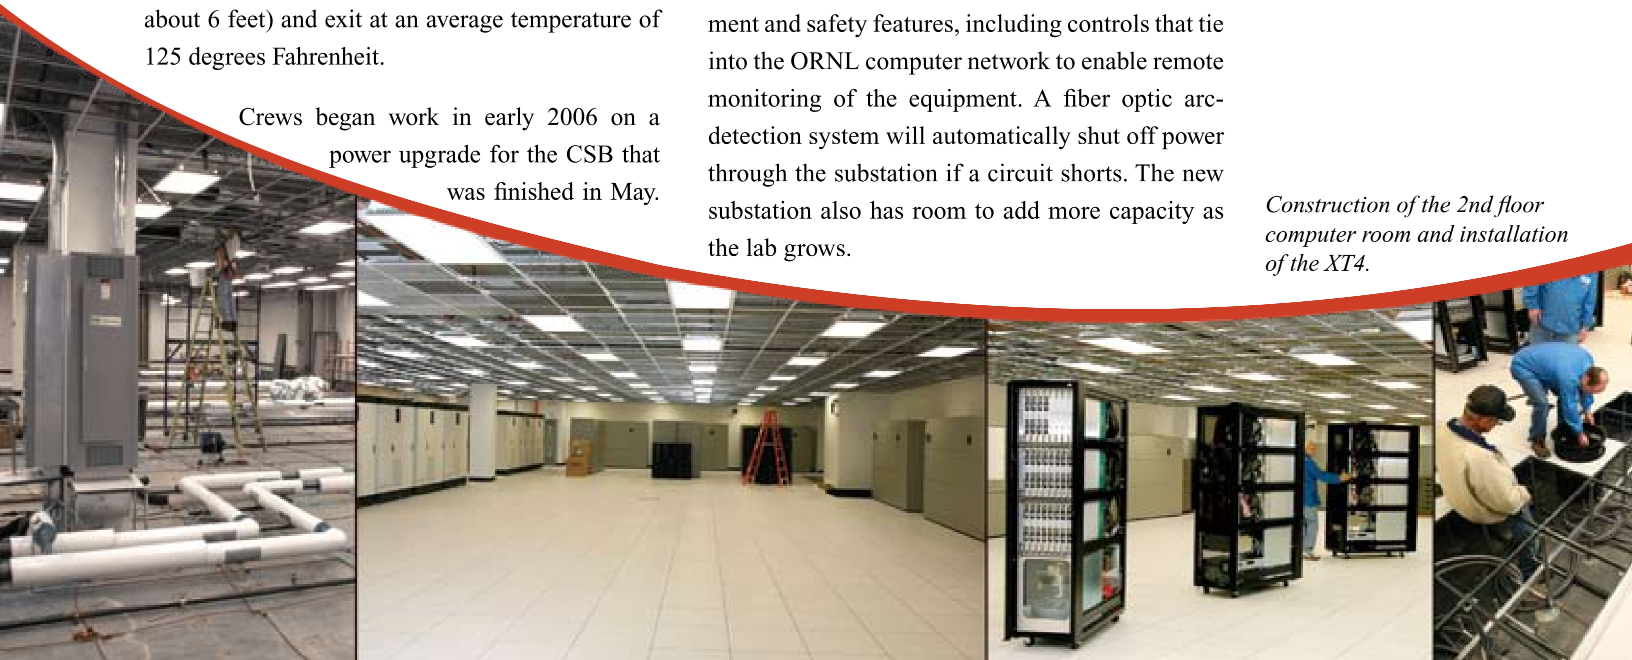

Construction of the 2nd floor computer room and installation of the XT4.

VISUALIZATION & COLLABORATION

As scientific simulations scale up exponentially, so does the pressure on processes and tools for making use of the data. Developing tools that enable researchers to efficiently manage, analyze, and visualize simulation results was a primary focus for NCCS staff in 2006.

A data management process dubbed "end-to-end" was developed to free NCCS users to produce data and interpret it, rather than puzzle over how to move, store, and manage it, said Scott Klasky, leader of the end-to-end task. Many routine data management tasks have been either automated or taken on by NCCS staff so that scientists can concentrate on their research.

One of the goals of end-to-end is to ensure that calculations produce all the information the researchers need. One way of doing that is to generate more metadata, or self-describing information, along with the scientific data. For example, a researcher might produce scores of binary files containing thousands of data points from calculations of different variables. If the researcher doesn't keep detailed notes, it could be difficult later to match data sets with the calculations that produced them. To avoid such situations, NCCS staff can help researchers set up calculations so that they generate metadata to provide a context and description in every data file stored. Because generating metadata uses computer time, it is important to make the process as efficient as possible.

Researchers confer on the end-to-end process.

The workflow process has been automated to automatically archive data on tape as simulations run to avoid filling up the disk space on the supercomputers. The provenance information for the data—what is being archived, where the data are, which run they are from—is also preserved to provide a context for all data files.

Network improvements during the past year made it easier for off-site NCCS users to keep tabs on their simulations. Researchers can watch their work in progress via a Web browser anywhere they have Internet access. Faster input/output allows quick delivery and constant updates. The capability to monitor simulations in progress using the Internet enables users to catch mistakes and correct them early in a run before valuable processing time has been wasted.

Faster networks also make it possible to move simulation results from the NCCS supercomputers to computers at a user's own institution for processing. The increased portability of data allows users to get a head start on reviewing their results and react to interesting trends they observe as they analyze them.

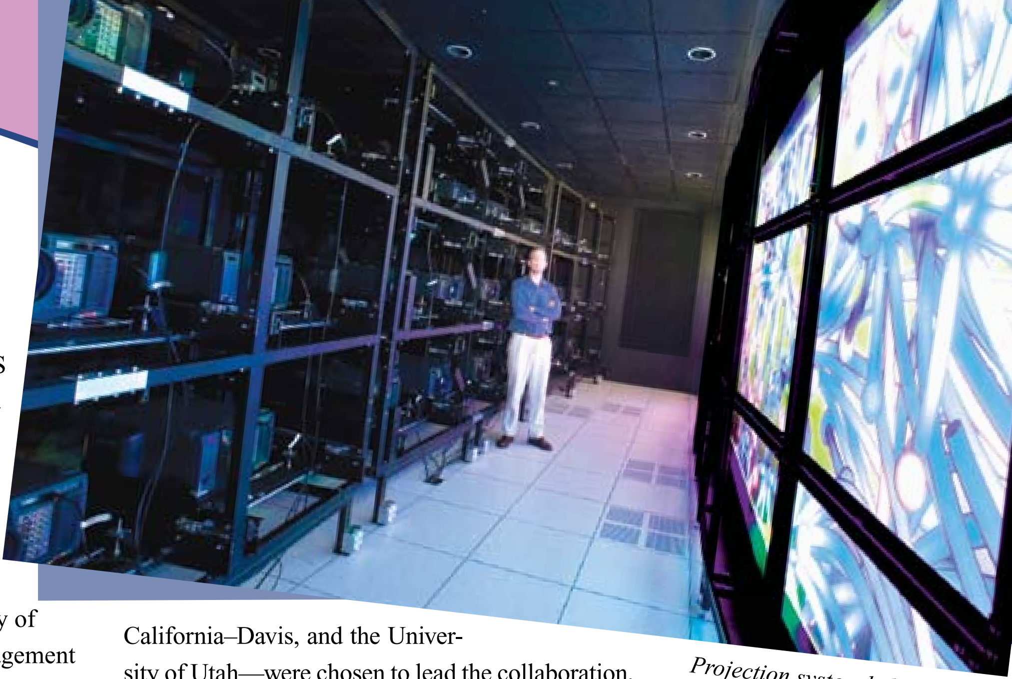

Projection system behind the EVEREST PowerWall.

As resources become available, the NCCS intends to be able to save all the metadata from simulations into databases so users can search and query them quickly. Another goal is to establish metadata catalogs that track where all the data from a run are stored so researchers can quickly locate pieces they need. The increasing complexity of high-performance computing and data management is making it impossible for researchers to do all their own record keeping for their projects, Klasky said. “We have to make it easier for the users so they can get back to doing their science rather than becoming experts in 10,000 different data management technologies.”

The same philosophy is at work in the data analysis and visualization process, a closely related step in the work flow. Scientists can’t just look at a printout of tens of billions of data points from terascale calculations and see the patterns in them, said visualization team leader Sean Ahern. Helping them extract the patterns hidden in those data, and providing the tools to do so, is the job of the analysis and visualization team at the NCCS.

The team’s biggest accomplishment for 2006 was a winning bid by the NCCS and its partners to establish a Scientific Discovery through Advanced Computing (SciDAC) Visualization and Analysis Center for Enabling Technologies, Ahern said. The NCCS and four partners—national labs Lawrence Livermore and Lawrence Berkeley, the University of California–Davis, and the University of Utah—were chosen to lead the collaboration, which will develop and improve large-scale data analysis tools to support SciDAC researchers.

Another accomplishment was embedding analysis and visualization specialists with some of the research teams using NCCS resources. These liaisons participate in meetings with the scientists and help them understand, at the early stages, the data coming out of their simulations. Analysis is becoming part of the research process from the start, rather than something done after simulations are finished. The liaisons are part of the climate and astrophysics teams and have worked with biology, combustion, fusion, and materials teams. Ahern plans to expand the practice as resources become available.

Plans for 2007 include deploying a new analysis and visualization computer cluster that can act either as a single large system or as several smaller systems. It will equip the NCCS to analyze much larger data sets. The new cluster will be coupled with the new NCCS Lustre file system, Spider. Spider will take data directly from Jaguar, and the cluster will take data directly from Spider for analysis and visualization.

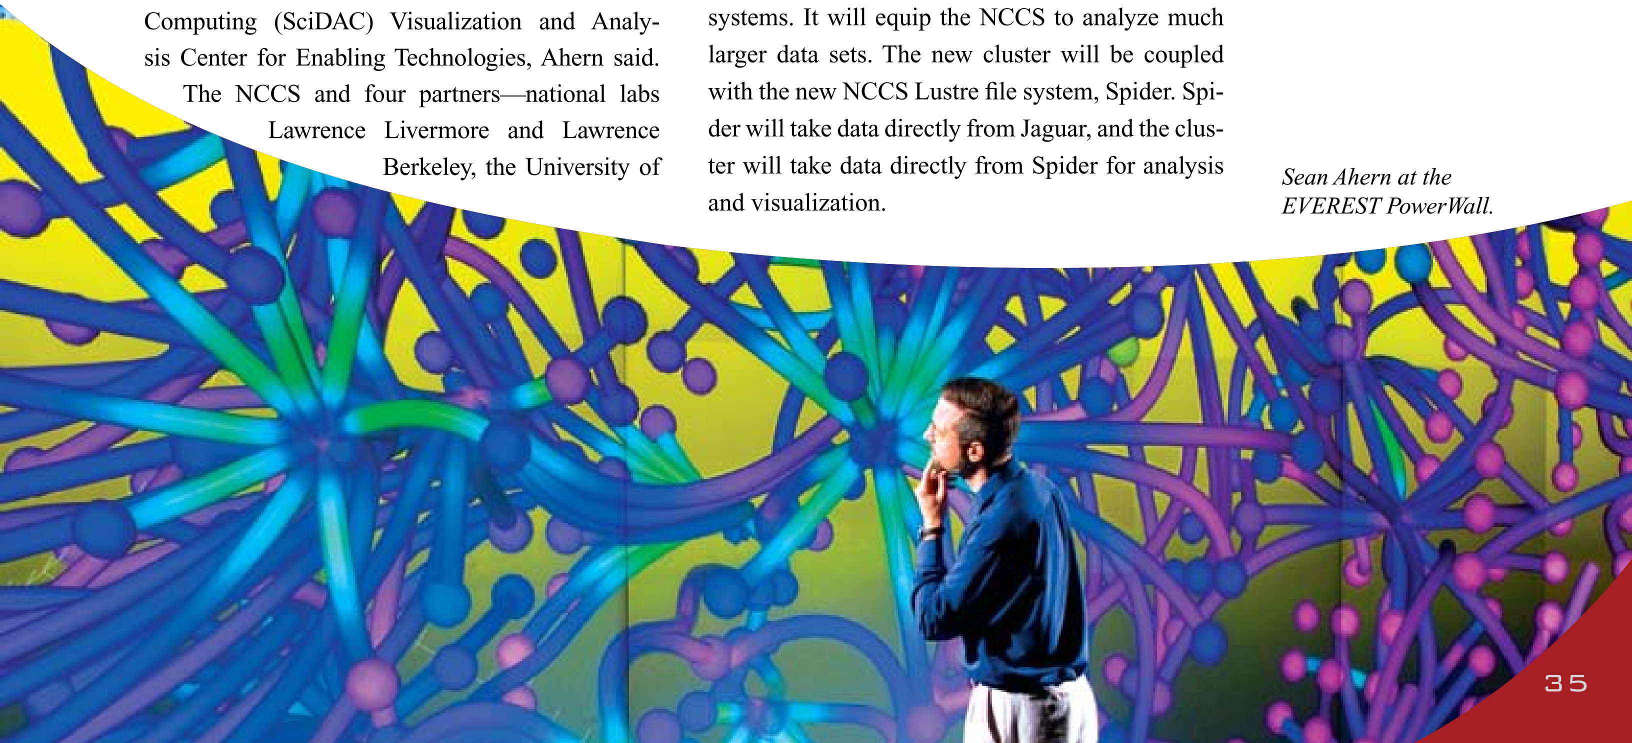

Sean Ahern at the EVEREST PowerWall.

NCCS STAFF

Each of the four NCCS groups—Scientific Computing, High-Performance Computing Operations, User Assistance and Outreach, and Technology Integration—added staff in 2006 as the center continued to grow. 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. Fortunately, the NCCS's staff has the talent to bring the center into the era of petascale computing.

Scientific Computing Group

The SCG comprises research scientists, visualization specialists, and work-flow experts who work directly with users, allowing them to take full advantage of the center's unique, next-generation systems to advance the boundaries of knowledge.

At its core is a group of liaisons to the research projects. The liaisons are experts in their fields—chemistry, physics, astrophysics, mathematics, numerical analysis, computer science—but they are also experts in designing code and optimizing it for the NCCS systems. Some applications must be reworked, an especially big challenge for code that dates back years or even decades, while others must be written from scratch. SCG liaisons work with the research teams at a variety of levels, from full-fledged partners in developing code to more informal troubleshooting consultants.

A new set of challenges arises once the programs are running: researchers must analyze, organize, and transfer an enormous amount of data. The SCG's end-to-end task works to streamline the work flow for system users so that their time is not eaten up by slow and repetitive chores. The goal is to automate tasks such as monitoring the system, analyzing results, and moving the data to their final destination. Ultimately, the center is working toward a dashboard application to give researchers all the information they need day to day on a single screen.

Once users have completed their runs, specialists in the SCG's visualization team help them make sense of the sometimes overwhelming amount of information they generate. This mission is exemplified by the EVEREST PowerWall at the NCCS. EVEREST, which stands for "Exploratory Visualization Environment for Research in Science and Technology," is a 30-foot-wide, 8-foot-tall tiled

The Scientific Computing Group.

display that can show 35 million pixels of information. It provides a unique opportunity to view more data and to view them in larger groups.

Members of the visualization team use both open-source and commercial software, choosing packages that are both flexible and able to accommodate large amounts of data. The group has made its presence felt most in areas such as materials science, astrophysics, and climate and fusion science.

High-Performance Computing Operations Group

The HPC Ops Group is responsible for ensuring that the NCCS's unique leadership systems are up and running, allowing users to perform the groundbreaking research for which the center was established.

HPC Ops is far more than a typical information technology organization. It manages unique, state-of-the-art supercomputers, and it installs and runs systems that make use of bleeding-edge technologies to handle the stunning speed of the NCCS's leadership computers and the enormous amounts of data they produce.

To accomplish this daunting task, HPC Ops staff take on a variety of roles: system administration, configuration management, cybersecurity, and around-the-clock operations support. The group works with the center's infrastructure systems as well as its supercomputers.

It's a big job. The center's premier Jaguar Cray XT4 supercomputer, for instance, comprises 124 separate cabinets with more than 11,000 dual-core AMD processors. Its manufacturer, Cray Inc., cannot keep a comparable system on hand to perform the rigorous testing of its power and cooling supply that would go into a commodity desktop or laptop computer before it is sold.

As a result, HPC Ops staff must perform thorough testing each time they install or upgrade a new system. They must also keep tabs on all the changes that are made to a system throughout its lifetime. This job, known as configuration management, involves tracking changes made to a system and allows the group to rebuild the system if the need arises.

The group is also responsible for making certain the systems are running as smoothly as possible. This task involves monitoring them around the clock and anticipating problems before they arise. By using diagnostic tools, HPC Ops is able to make certain there is ample disk space and identify hardware components that might be near failure.

The group is also charged with keeping the systems secure. It monitors activity that is out of the ordinary—such as port scans, a technique used by hackers to compromise a system by surveying it for open access points—and it is responsible for ensuring that NCCS systems conform to ORNL cybersecurity policy.

The High-Performance Computing Group.

The Technology Integration Group.

Technology Integration Group

TechInt staff work behind the scenes to develop the infrastructure that supports the NCCS systems, stay on the leading edge of technology advances, and provide systems-level expertise. They research and evaluate emerging technologies in areas such as networks, file systems, and archival storage and integrate technologies and tools into the infrastructure as they are adopted.

Putting a faster network and a new file system into place was the focus of much of the TechInt group's effort during 2006. They prepared to expand the NCCS wide-area network to 10 gigabits per second to allow rapid movement of data sets among locations. Much of the preparation involved restructuring the network to accommodate a new cybersecurity model, integrating technical plans and security plans, and establishing required security measures. The new network is expected to be in operation in early 2007.

The rollout of a Lustre file system, dubbed Spider, is also expected in early 2007. Spider will provide a file system that is directly accessible by NCCS systems. This will enable users to store data sets from their simulations and immediately perform data analysis and visualization with other NCCS resources. TechInt staff worked with Cluster File Systems, Inc., the makers of Lustre, to commission new features to support the NCCS configuration. The Spider system will evolve and continue to add capability to keep pace with the NCCS computer systems.

Members of TechInt also worked with other developers of the HPSS to develop and deploy a new release of the software during 2006. This task included deploying the new software on the production HPSS system at the NCCS. Work also began on a successor version of HPSS expected to be released in early 2008.

The development of tools to fit the specialized needs of the NCCS is a constant task for TechInt staff. As NCCS systems grow to tens of thousands of processors, it is increasingly difficult for conventional approaches to scale to accomplish the required tasks—for example, debugging an application that might contain 80,000 tasks. During 2007, the TechInt team will specifically focus on developing tools (e.g., compilers, debuggers, and performance analysis) that enable users to take full advantage of the expanded capabilities of the leadership-class systems.

User Assistance and Outreach Group

UAO staff are the most familiar representatives of the NCCS to computer users and the public. They

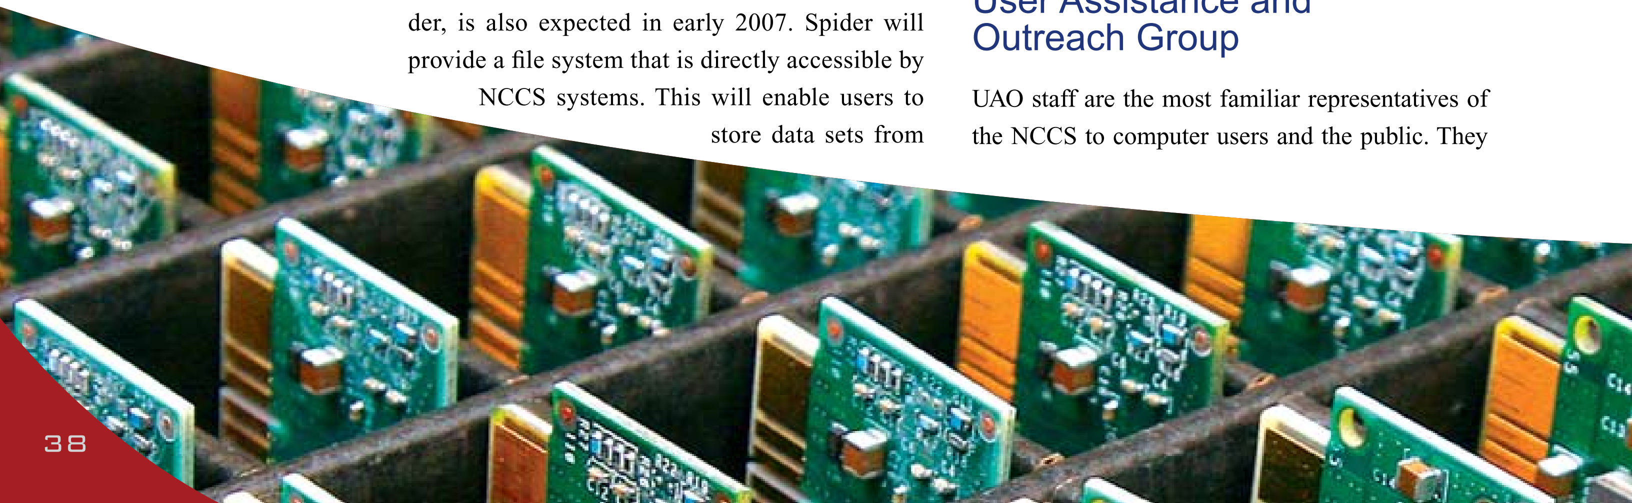

manage the liaison of users with the NCCS, mediate between the outside world and the NCCS, and acquaint the public with the work being conducted.

The key UAO function is to be user advocates. UAO staff help users with organizational tasks such as understanding the user agreement, setting up accounts, and monitoring the status of their allocations as well as technical issues like debugging and compiling codes. The group both insulates users from the mechanics of NCCS processes and represents their interests in those processes. Its staff attend the working meetings of the other NCCS groups, sharing information and suggestions from users and participating in decisions as the users' voice.

During 2006, UAO set up several avenues of communication with users—e-mail distribution, the NCCS Web site, monthly teleconferences, a user survey, an annual user meeting, and periodic workshops.

It also is in charge of communicating outward. To reach the community of potential users, it publishes research highlights and *HPCwire* articles. And it oversees the production of printed and electronic media aimed at the larger scientific community and the public—fact sheets, annual reports, special publications, Web sites, and displays. UAO compiles and provides all NCCS metrics (except system availability data) and responds to requests for NCCS statistics.

A series of Leadership Computing Facility seminars was instituted in 2006 that brings scientists using the NCCS supercomputers to the ORNL campus for open presentations on their work. Plans are to expand the seminars to once monthly during 2007.

The User Assistance and Outreach Group.

Outreach to educational institutions is a continuing priority. An informal "speakers bureau" sprang up to talk to groups at high schools and colleges; it will be formalized and expanded in 2007. Another plan is to expand collaboration with smaller colleges in the area, such as Maryville College and local community colleges, and perhaps establish seminars and internships with the smaller schools.

Another growth area for 2007 will be the Director's Discretion projects—smaller projects chosen by the NCCS, rather than through INCITE, for which 20 percent of the computer time is reserved. The process for allocating time to those projects will be formalized, and UAO will be more directly engaged with them.

LOOKING AHEAD

PATH TO PETASCALE

WHAT IF ...

- we understood our own impact on the earth's climate so completely that we could head off global warming,
- we understood exactly how stars exploded to provide us with the building blocks of life,
- we were able to exploit an unlimited source of energy free of pollution and greenhouse gases,
- we could economically produce wire and other materials that had no electrical resistance?

The NCCS was established with one primary goal: to expand the boundaries of science with the world's fastest supercomputers. The center has achieved breakthroughs in a variety of fields with its terascale systems, high-performance computers capable of well over a trillion calculations per second. As it looks to 2008 and beyond, the NCCS will continue its leadership with petascale systems, machines that will operate at 1,000 trillion calculations per second and beyond.

Following are just a few examples of the achievements that can be expected from these extraordinary systems.

Our Changing Climate

Climate scientists are faced with the challenge of accurately anticipating how the world in which we live is changing and how we are influencing that change. Current terascale supercomputers allow climatologists to effectively simulate the ocean, land, and atmosphere to show how the systems influence one another. In addition, they are able to predict the effect a given level of greenhouse gases such as carbon dioxide will have on global warming.

However, this is only the beginning. With petascale computers, scientists will accurately simulate the exchange of carbon between the ocean, the land, plants and animals, and the atmosphere. Researchers will be able to predict how climate change will lead to melting ice sheets, a phenomenon that will help determine future ocean levels. And they will be able to focus on a specific geographic region and predict changes in temperature, rainfall, vegetation, and land use.

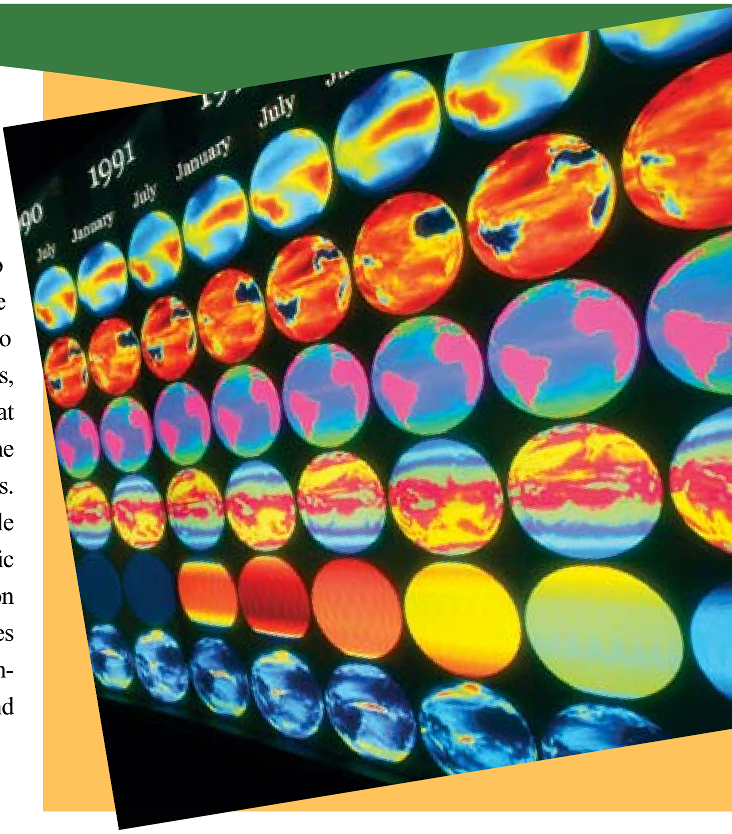

The Sky Above Us

Astrophysicists are striving to understand the process that makes a dying star explode. Core-collapse supernovas mark the end of stars more than eight times the mass of the sun, providing the universe with many elements and making life itself possible. Without them, we would not be here to study the heavens.

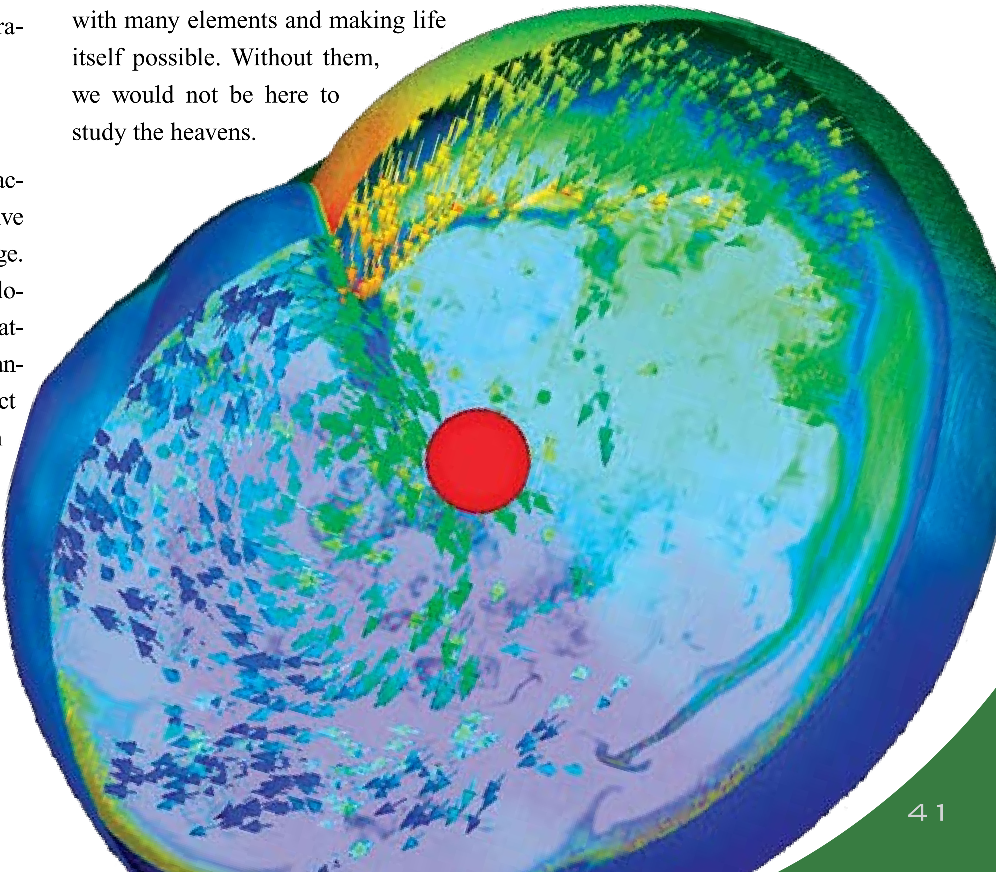

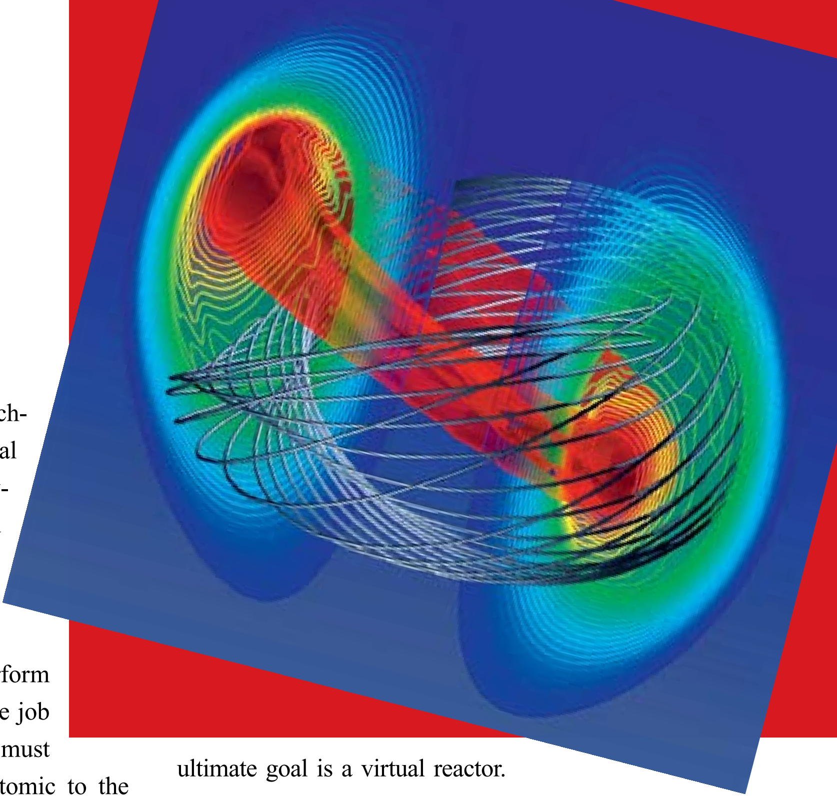

Using terascale supercomputers, researchers are able to perform two-dimensional simulations that approximate the behavior of key but elusive particles called neutrinos. To get the information they need, however, these researchers will need petascale systems to perform more complex simulations—and perform them in three dimensions. To make the job even more difficult, the simulations must handle scales ranging from the subatomic to the stellar, encompassing an area significantly larger than the earth's orbit around the sun.

The Future of Energy

Researchers around the world are striving to harness the process that fuels the sun and stars. Fusion reactors will use fuel that is abundantly available from ocean water, generate power without emitting pollutants and greenhouse gases, and produce far less waste—and waste that is far less dangerous—than current nuclear power plants.

In the meantime, however, researchers must figure out how to control an ionized gas far hotter than the sun using magnetic fields and radio waves. Terascale systems allow them to simulate only one source of instability at a time in this immensely complex system, and then only for brief periods. Petascale systems will allow them to move toward an integrated simulation that incorporates all the competing forces. The ultimate goal is a virtual reactor. Such a tool will put us one step closer to a source of power we can use without worrying about its impact on future generations.

Materials of the Future

Researchers are developing new materials that conduct electricity with no resistance—and no loss. Known as superconductors, these materials promise to revolutionize such diverse areas as electrical transmission, transportation, medicine, and computer design. However, there is a catch: to exhibit superconductivity, the materials must be very cold. Even so-called high-temperature superconductors must be cooled to −190 degrees Fahrenheit or lower.

By understanding the process of superconductivity, researchers hope to ultimately develop materials that display superconductivity without being cooled. Terascale supercomputers have allowed them to validate a model that explains the behavior of high-temperature superconductors, opening a new pathway for scientific advancement. More complex simulations will allow for more detailed chemistry and describe the behavior of different types of electrons. At petascale and beyond, researchers will be able to begin the search for materials with properties designed to meet the challenge.

2007 INCITE PROGRAM

ORNL's NCCS is providing leadership computing to 29 programs in 2007 under DOE's INCITE program.

Leading researchers from government, industry, and the academic world will use more than 75 million processor hours on the center's Cray leadership computers to advance the boundaries of knowledge in areas ranging from profoundly small nanostructures to profoundly large stellar systems. All told, the center's Cray XT4 (Jaguar) and Cray X1E (Phoenix) systems will provide more than 75 percent of the computing power allocated for the INCITE program.

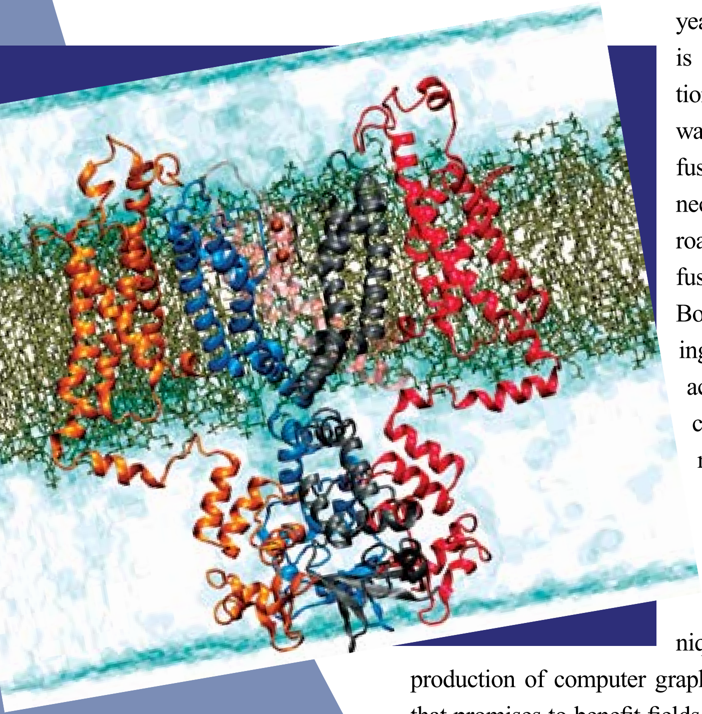

The INCITE program gives industry a unique opportunity to use the world's fastest supercomputers for solving challenging problems. General Atomics, the Boeing Company, and Dreamworks Animation will be participating in the program for a second year. General Atomics is exploring the interaction of long- and short-wavelength turbulence in fusion power reactors, a necessary step along the road to clean, abundant fusion power generation. Boeing is demonstrating the effectiveness and accuracy of powerful computational fluid dynamics simulation tools in designing the next generation of aircraft. And Dreamworks is developing techniques that improve the production of computer graphics imagery, a project that promises to benefit fields such as medical imaging, defense simulations, and visualizations in general as well as the production of animated feature films.

In addition, Corning, Inc., will be a new user to the NCCS supercomputers. The company will perform the first-ever calculation of realistic transition ranges for silica and silicate glasses. The company's project promises substantial scientific and practical benefits, as silicates constitute the vast majority of both traditional and high-tech glass products.

In other projects, climate scientists are developing definitive climate models, examining turbulent transport in the earth's oceans, and studying carbon dioxide sensitivity and abrupt climate change. Fusion scientists are working to understand the process that powers the sun to aid in the development of clean, abundant power. Biologists are unlocking the secrets of membrane proteins and biomolecular structures. Chemists are developing approaches to designing chemical catalysts. Nanoscientists are simulating structures at the molecular and atomic levels. And astrophysicists are performing simulations on supernovas, the Milky Way's dark matter halo, and black holes in binary star systems.

Most of this work, more than 70 million processor hours, will be performed on the Jaguar system. Jaguar is the most powerful supercomputer in the United States for open scientific computing. With a peak performance of 119 teraflops and 46 terabytes of memory, it is one of the ten most powerful systems in the world according to the TOP500 list of the world's most powerful computers.

Another 5 million-plus hours have been allocated to Phoenix, which is currently the most powerful unclassified vector computer in the United States. Phoenix has 1,024 multistreaming vector processors, each having 2 megabytes of cache and a peak computational rate of 18 gigaflops. The system delivers a peak performance of 18.5 teraflops. Phoenix uses custom-designed vector processors to get high performance for scientific codes.

2007 INCITE PROJECT LIST

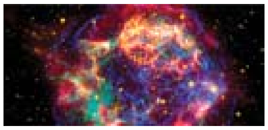

Multi-Dimensional Simulations of Core-Collapse Supernovae

Anthony Mezzacappa (Oak Ridge National Laboratory)

Jaguar - 7,000,000 processor hours and Phoenix - 300,000 processor hours

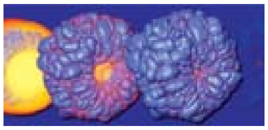

First Principles Models of Type Ia Supernovae

Stan Woosley (University of California, Santa Cruz)

Jaguar - 4,000,000 processor hours

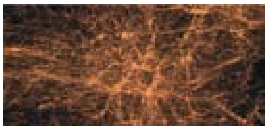

Via Lactea: A Billion Particle Simulation of the Milky Way's Dark Matter Halo

Piero Madau (University of California, Santa Cruz)

Jaguar - 1,500,000 processor hours

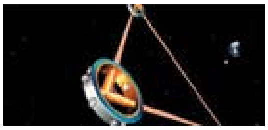

Numerical Relativity Simulations of Binary Black Holes and Gravitational Radiation

Joan Centrella (NASA/Goddard Space Flight Center)

Jaguar - 500,000 processor hours

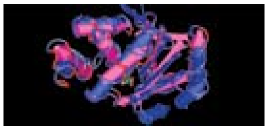

Next Generation Simulations in Biology: Investigating Biomolecular Structure, Dynamics and Function Through Multi-Scale Modeling

Pratul Agarwal (Oak Ridge National Laboratory)

Jaguar - 1,000,000 processor hours

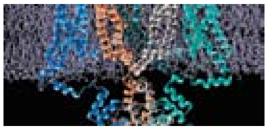

Gating Mechanism of Membrane Proteins

Benoit Roux (Argonne National Laboratory and University of Chicago)

Jaguar - 4,000,000 processor hours

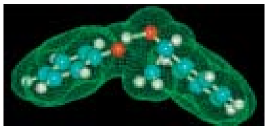

An Integrated Approach to the Rational Design of Chemical Catalysts

Robert Harrison (Oak Ridge National Laboratory)

Jaguar - 3,000,000 processor hours and Phoenix - 300,000 processor hours

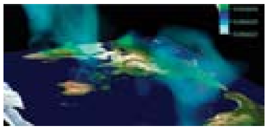

Climate-Science Computational End Station Development and Grand Challenge Team

Warren Washington (National Center for Atmospheric Research)

Jaguar - 4,000,000 processor hours and Phoenix - 1,500,000 processor hours

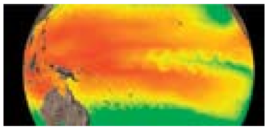

Eulerian and Lagrangian Studies of Turbulent Transport in the Global Ocean

Synte Peacock (University of Chicago)

Jaguar - 3,000,000 processor hours

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

Zhengyu Liu (University of Wisconsin - Madison)

Phoenix - 420,000 processor hours

2007 INCITE PROJECT LIST *(CONTINUED)*

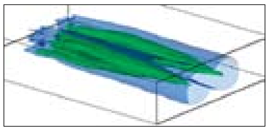

Modeling Reactive Flows in Porous Media

Peter Lichtner (Los Alamos National Laboratory)
Jaguar - 1,000,000 processor hours

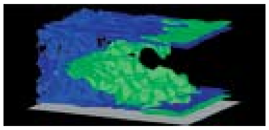

High-Fidelity Numerical Simulations of Turbulent Combustion - Fundamental Science Towards Predictive Models

Jackie Chen (Sandia National Laboratories)
Jaguar - 6,000,000 processor hours and Phoenix - 50,000 processor hours

Performance Evaluation and Analysis Consortium End Station

Patrick Worley (Oak Ridge National Laboratory)
Jaguar - 1,000,000 processor hours

Real-Time Ray-Tracing

Evan Smyth (Dreamworks Animation)
Jaguar - 900,000 processor hours

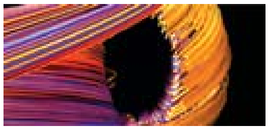

Gyrokinetic Plasma Simulation

W. W. Lee (Princeton Plasma Physics Laboratory)
Jaguar - 6,000,000 processor hours and Phoenix - 75,000 processor hours

Simulation of Wave-Plasma Interaction and Extended MHD in Fusion Systems

Don Batchelor (Oak Ridge National Laboratory)
Jaguar - 2,000,000 processor hours

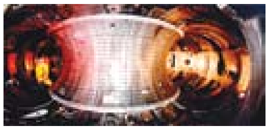

Interaction of ITG/TEM and ETG Gyrokinetic Turbulence

Ronald Waltz (General Atomics)
Phoenix - 500,000 processor hours

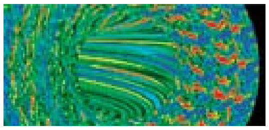

Gyrokinetic Steady State Transport Simulations

Jeff Candy (General Atomics)
Jaguar - 1,000,000 processor hours

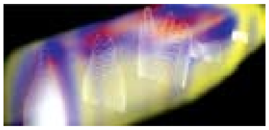

High Power Electromagnetic Wave Heating in the ITER Burning Plasma

Fred Jaeger (Oak Ridge National Laboratory)
Jaguar - 500,000 processor hours

2007 INCITE PROJECT LIST *(CONTINUED)*

Computational Design of the Low-loss Accelerating Cavity for the ILC

Kwok Ko (Stanford Linear Accelerator Center)

Phoenix - 400,000 processor hours

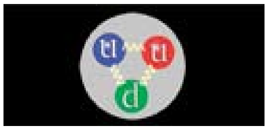

Lattice QCD for Hadronic and Nuclear Physics

Robert Edwards (Thomas Jefferson National Accelerator Facility)

Jaguar - 10,000,000 processor hours

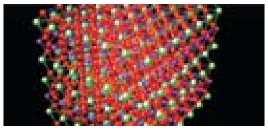

Predictive Simulations in Strongly Correlated Electron Systems and Functional Nanostructures

Thomas Schulthess (Oak Ridge National Laboratory)

Jaguar - 7,000,000 processor hours and Phoenix - 500,000 processor hours

Development and Correlations of Large Scale Computational Tools for Flight Vehicles

Moeljo Hong (The Boeing Company)

Phoenix - 200,000 processor hours

Ab Initio Modeling of the Glass Transition

John Mauro (Corning Incorporated)

Phoenix - 100,000 processor hours

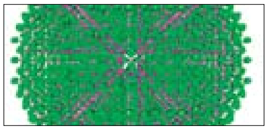

Linear Scale Electronic Structure Calculations for Nanostructures

Lin-Wang Wang (Lawrence Berkeley National Laboratory)

Jaguar - 1,500,000 processor hours

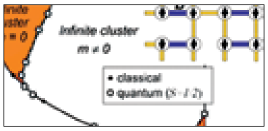

Bose-Einstein Condensation vs. Quantum Localization in Quantum Magnets

Tommaso Roscilde (Max-Planck Gesellschaft)

Jaguar - 800,000 processor hours

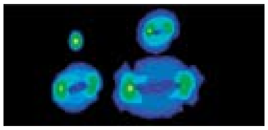

Computational Atomic and Molecular Physics for Advances in Astrophysics, Chemical Sciences, and Fusion Energy Sciences

Michael Pindzola (Auburn University)

Phoenix - 750,000 processor hours

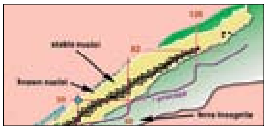

Ab-Initio Nuclear Structure Computations

David J. Dean (Oak Ridge National Laboratory)

Jaguar - 5,000,000 processor hours

APPENDICES

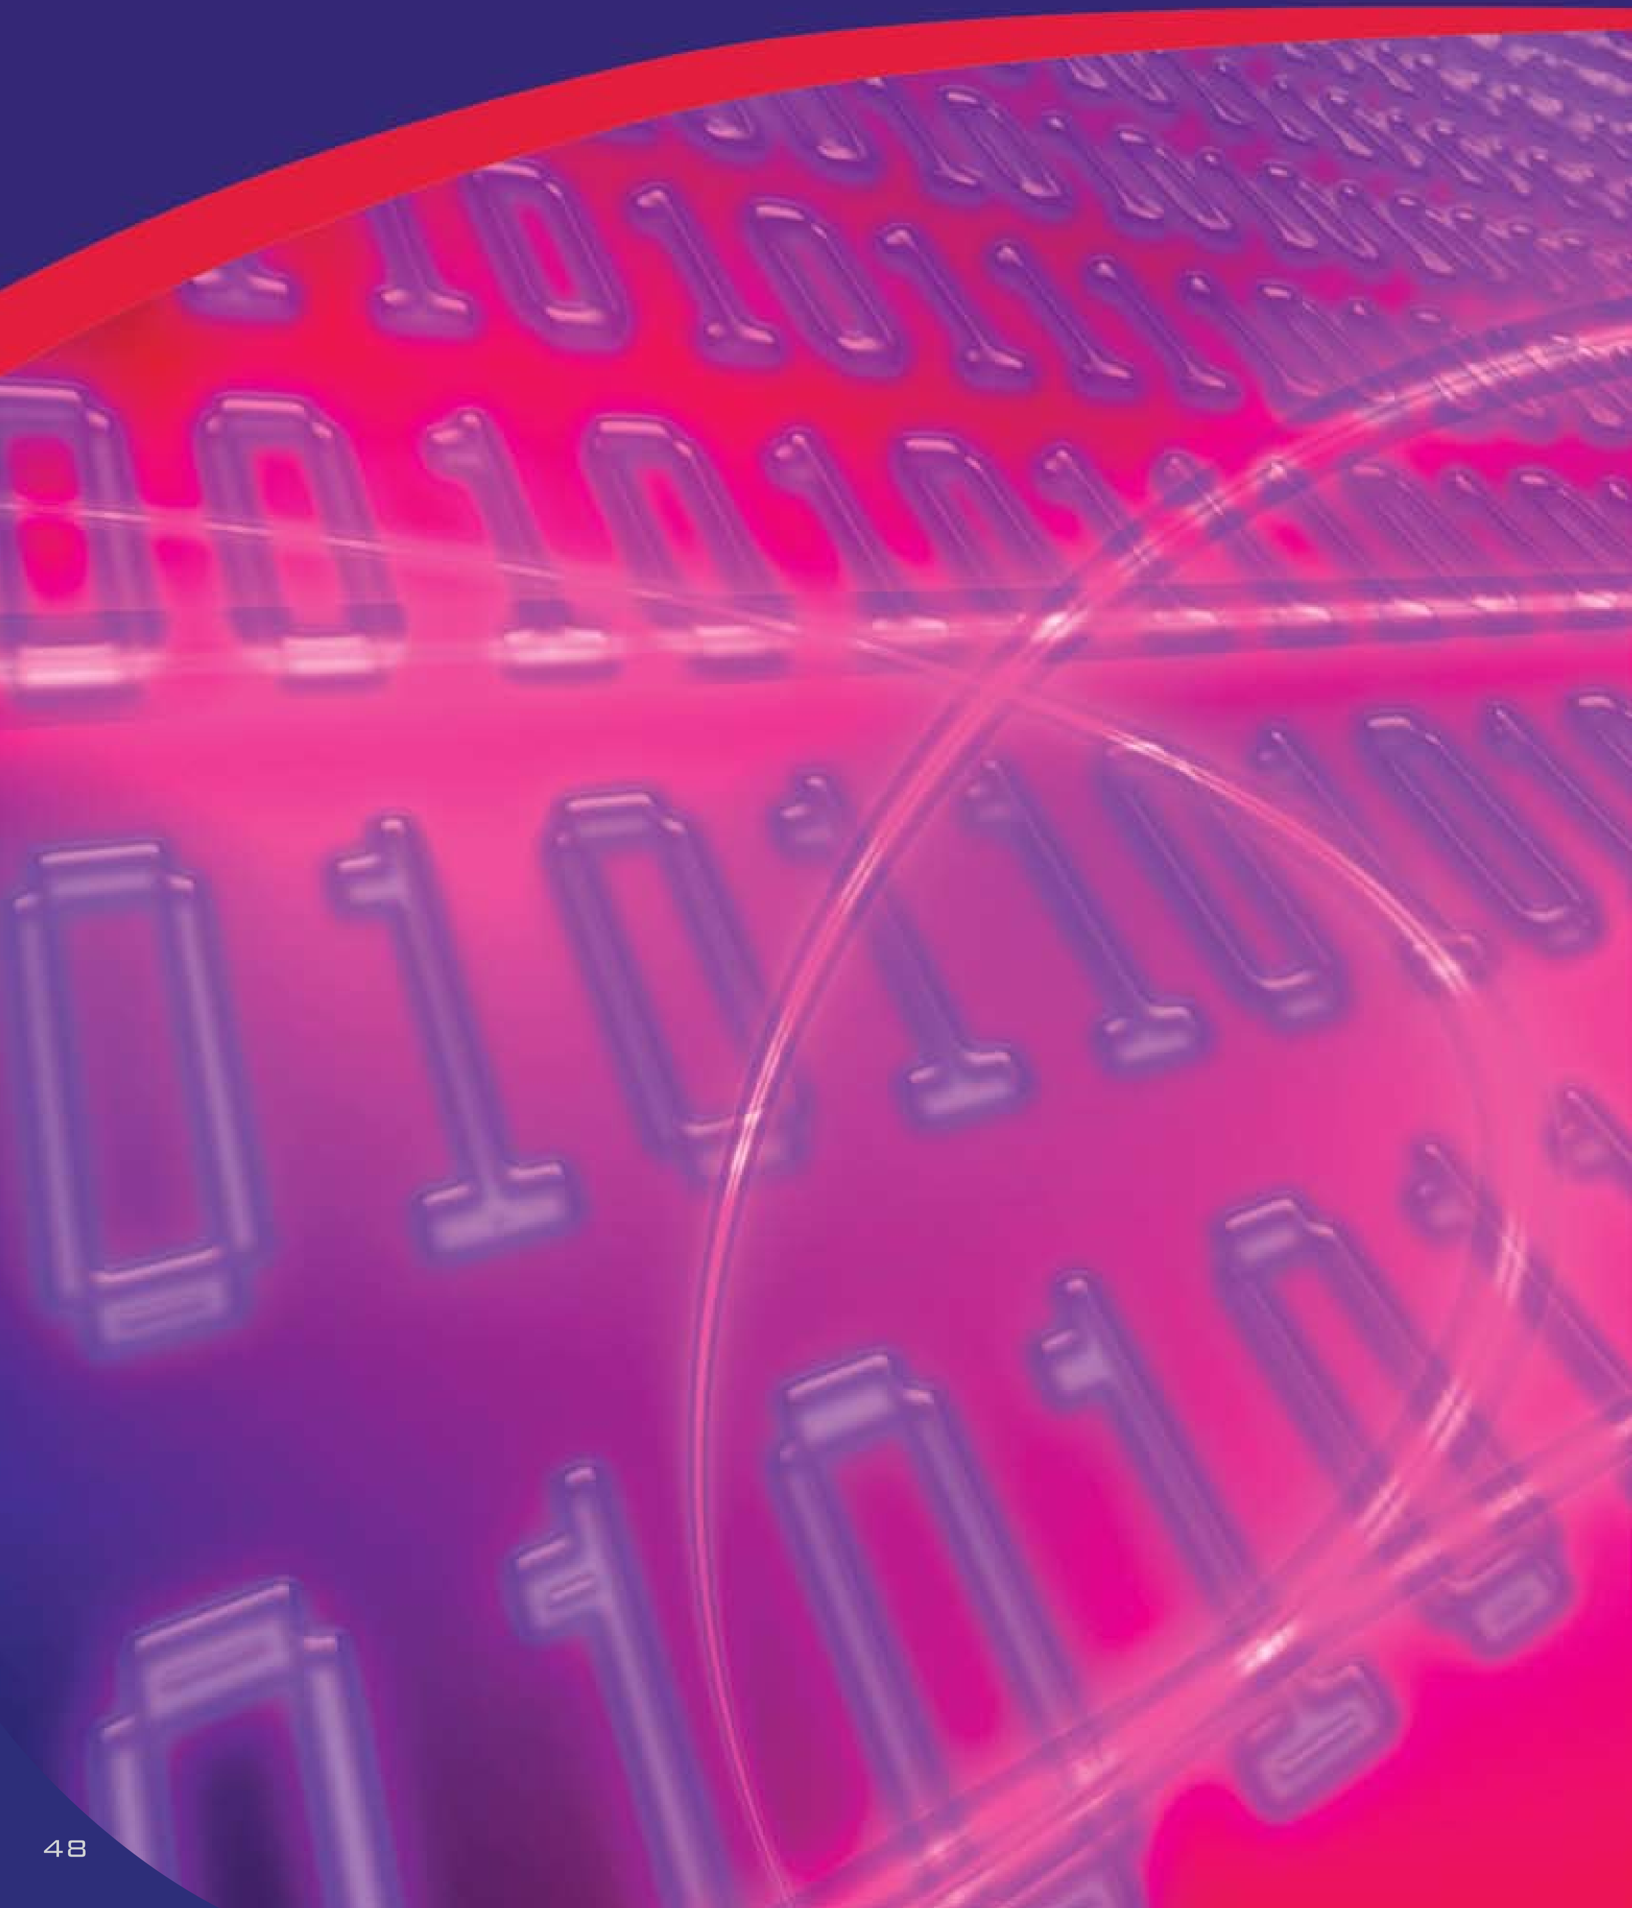

NCCS PARTNERSHIPS

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.

Joint Institute for Computational Sciences

JICS provides a framework for ORNL to work with the University of Tennessee (UT) and its other university partners to realize three strategic goals:

1. Create new modeling and simulation capabilities for computers operating at the terascale and beyond and use these capabilities to solve the science and engineering challenges that are the most important to DOE and the United States.
2. Train scientists and engineers to model and simulate systems on supercomputers and educate a new generation of researchers expert in applying computational simulation in research and education.
3. Create a leading cyber infrastructure for science and engineering in the southeastern United States.

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. The institute has also pursued an outreach program to historically black colleges and universities for more than a decade, beginning with a MetaCenter Regional Alliance grant awarded in 1995 and funded by the National Science Foundation.

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. One suite has a raised floor, electrical power, and cooling equipment so it can accommodate a cluster of computers. The suites can be made available to commercial firms for proprietary research in collaboration with ORNL researchers or used for basic research and nonproprietary collaborations among researchers from ORNL and its university partners.

The JICS staff includes university-ORNL joint appointees, JICS research affiliates, postdoctoral fellows, graduate students, administrative personnel, and support staff to assist researchers and students in using the NCCS supercomputers.

The joint appointees work as faculty members in university departments and staff

The interior atrium of the JICS facility.

Researchers and university partners share the state-funded JICS facility on the ORNL campus.

scientists in ORNL research groups. Electronic collaboratory tools allow them to continue their research when they return to their home institutions and to televise their classes from ORNL to their home universities. Distance education tools also enable graduate students doing research at JICS to take classes from their home institutions. These communication capabilities are supported by the 10 gigabit-per-second network constructed to connect ORNL with UT and the core universities.

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. The committee reports to the ORNL Associate Laboratory Director for Computing and Computational Sciences. Responsibilities of the committee include providing

- advice on priorities and strategies to effectively execute the mission of the NCCS;
- scientific advice—for example, which domains might be ready to achieve “breakthrough science,” and potential scientific directions for the NCCS user program.

In addition, the Advisory Committee should advocate and promote effective communication between ORNL leadership, DOE, other federal agencies, and the user community to help facilitate mutual understanding in support of achieving maximum impact by NCCS users.

Operations Council

The NCCS is one of 18 major user facilities at ORNL. The mission of the Operations (Ops) 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. It is composed of representatives of each of the operational elements needed to provide the underlying NCCS infrastructure: HPC Ops, TechInt, UAO, Scientific Computing, Networking, Visualization, and Facility Management. Representatives for the Human Resources; Recruiting; Cyber and Information Security; Quality Assurance; Environment, Safety, and Health; Finance; and Procurement organizations meet monthly with the Ops Council.

The Research Alliance in Math and Science

The Research Alliance in Math and Science (RAMS) program provides opportunities for

RAMS students present posters of their summer research projects.

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 increase the representation of individuals from minority groups among holders of advanced degrees in science, mathematics, engineering, and technology. RAMS is based on the idea 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 a summary paper discussing the research project, give an oral presentation of research results, and present the project in poster sessions and meetings at ORNL and national conferences. Academic credit for the internship can be arranged through the college or university.

Students from 13 predominantly minority-serving colleges and universities participated in RAMS in 2006. Schools represented were Alabama A&M, Dillard, Fisk, Florida A&M, Jackson State, Knoxville College, North Carolina Central, Savannah College, Savannah State, UT–Knoxville, Virginia Tech, Winston-Salem State, and Wofford College. Research topics ranged from computational biology to computational materials science to sensor networks to high-end visualization of scientific research and population studies.

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 the Mathematical, Information, and Computational Science Division of DOE's Office of Advanced Scientific Computing Research.

Tours of the NCCS computer room were given to summer RAMS students.

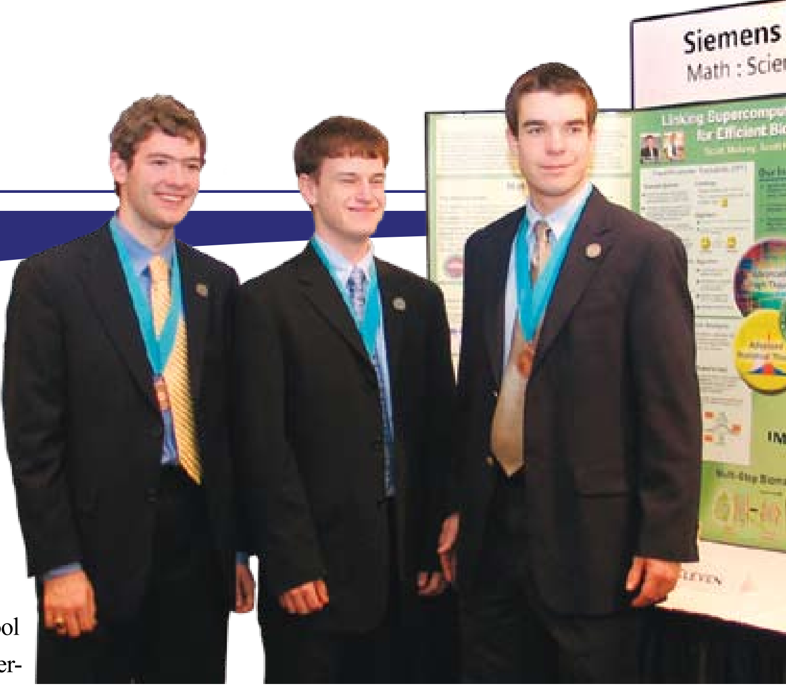

NCCS OUTREACH

STUDENTS LAND NATIONAL HONOR

Three Oak Ridge High School seniors used NCCS supercomputers in 2006 to improve the process for producing biofuel. As a consequence, they also won a national math and science competition and pulled in a $100,000 scholarship.

The students— Steven Arcangeli, Scott Molony, and Scott Horton— took the Grand Prize Scholarship at the 2006–2007 Siemens Competition in Math, Science and Technology. The three worked intensively with senior ORNL researchers Tatiana Karpinets, Hoony Park, Chris Symons, and Nagiza Samatova to achieve this recognition.

The students worked full time at the lab through the summer of 2006 and about 1 day a week in the following months to learn the fundamentals of graph theory, statistical theory, systems biology, bioinformatics, artificial intelligence, and programming in C and C++.

Their research results are preliminary but promising. By looking at a reasonably well-studied trait (aerobic versus anaerobic growth) and using only a small number of microbes, the students identified a target list of important genes. Many were known, thus confirming the validity of their idea. Others provide fodder for the formulation of future hypotheses and their experimental validation.

Top: Oak Ridge Operations Manager Gerald Boyd, left, and Assoc. Lab. Director Thomas Zacharia watch as Deputy Secretary of Energy Clay Sell signs the Cray.

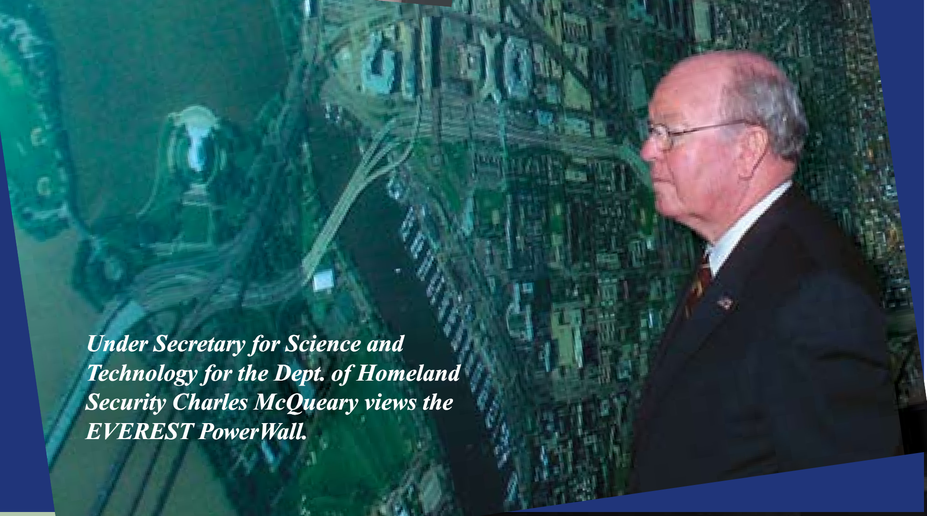

Under Secretary for Science and Technology for the Dept. of Homeland Security Charles McQueary views the EVEREST PowerWall.

TOURS

Leadership computing was a big draw for visitors to ORNL in 2006. More than 200 tours of NCCS computing and visualization facilities were conducted during the course of the year. Among the more prominent dignitaries to tour the facilities were Congressman Zach Wamp;

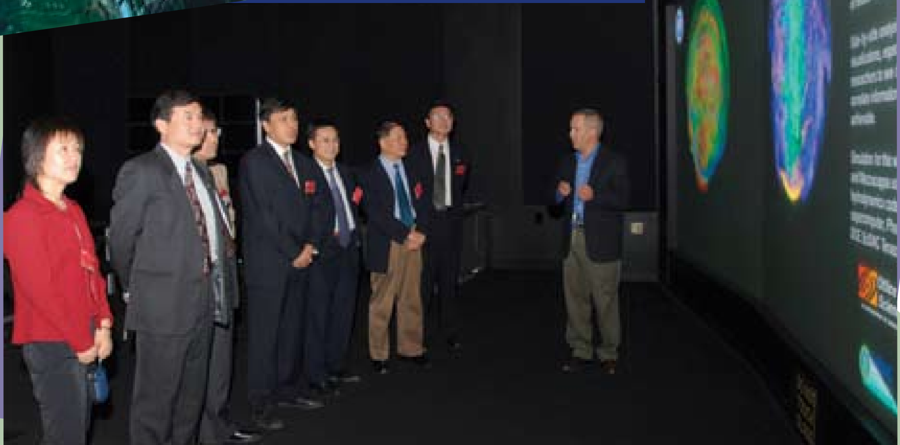

Ross Toedte of NCCS shows the PowerWall to visitors from the Chinese Academy of Science.

Oak Ridge High School students Scott Molony (left), Scott Horton, and Steven Arcangeli won the Grand Prize Scholarship at the 2006–2007 Siemens Competition in Math, Science and Technology.

Deputy Secretary of Energy Clay Sell; Hector Ruiz, chief executive officer of chip maker Advanced Micro Devices; a delegation from the Chinese Academy of Sciences; and Phillip Fulmer, head football coach at UT.

Warren Washington, National Center for Atmospheric Research, left, and Steven Ashby of the Center for Applied Scientific Computing, center, were given a tour of the NCCS by John Drake and Thomas Zacharia, right.

MEETINGS AND WORKSHOPS

2006 NCCS Users Meeting

February 14–16

Computational scientists gathered at ORNL for a 3-day meeting in mid-February to discuss ways to make the best use of the NCCS's Cray XT3 Jaguar and Cray X1E Phoenix supercomputers.

The NCCS Users Meeting gave researchers with allocations on the systems an opportunity to meet with their peers and become better acquainted with the nation's most powerful open science computing resources. The meeting featured presentations and tutorials on getting the most out of the two systems. It also featured presentations focusing on visualization tools and services and on data analysis and storage solutions.

A wide range of NCCS users shared the cutting-edge science they were achieving with the systems. Presenters came from diverse areas of research, including astrophysics, climate science, biology, fusion science, materials science, and computer science.

Attendees at the 2006 Users Meeting met in the state-of-the-art JICS conference room.

Quantitative Quantum Chemistry

March 17–20

Computational scientists gathered in Santa Fe, New Mexico, to discuss the profound advances that have been achieved in quantum chemistry and the challenges that lie ahead.

Attendees at the Visualization with VisIt workshop.

SCG lead Ricky Kendall was among the session chairs at this conference, which provided an opportunity for researchers to discuss their current work as well as theory and methods used in the field and their application to chemical systems.

Cray User Group 2006

May 8–11

NCCS staff gave a wide variety of presentations at this meeting in Lugano, Switzerland, touching on both the center's Cray leadership systems and the research being performed on them.

The theme for the meeting was "Scaling to New Heights," as technical experts and other Cray users from around the world gathered for tutorials, meetings focusing on shared interests, and technical presentations.

NCCS presentations during the meeting touched on subjects such as the center and its plans, challenges of the Cray XT3 and X1E systems at ORNL, architectures and development platforms used on the systems, and specific areas of research.

Visualization with VisIt

May 23

Researchers running simulations on a modern supercomputer are left with a daunting challenge: how can they make sense of the mountain of data they have created?

They were given answers in an all-day workshop, entitled "Visualization with VisIt," held at ORNL. Visualization experts were present for this hands-on event, in which researchers learned the basics of VisIt, a powerful visualization and data analysis tool developed by experts at DOE.

ORNL was a perfect location for the workshop; Sean Ahern of the NCCS SCG and Jeremy Meredith from ORNL's Computer Science and Mathematics Division were among the application's developers.

In the workshop, instructors demonstrated the use of VisIt to access data and create simple visualizations. They also delved into more sophisticated uses of the application to explore data with VisIt's interactive tools and to create graphics and movies.

SciDAC 2006

June 25–29

The NCCS hosts a number of research collaborations funded through the SciDAC program. Researchers affiliated with the program gathered in Denver to showcase successes and discuss research in various areas, from the universe's biggest events to its smallest.

NCCS users highlighted their journeys toward petascale computing through simulations of fusion, chemistry, combustion, and astrophysics.

Code Camp

August 9–11

NCCS users got personalized help on their applications at a 3-day workshop held at ORNL.

The Code Camp was an opportunity for researchers to work directly with NCCS experts to improve the scaling and stability of their codes and troubleshoot any issues they might have had on Jaguar or Phoenix. Members of the SCG and UAO groups were on hand to provide personalized help.

SC06

November 11–17

The NCCS was once again a major player at the country's premier supercomputing conference, touting the scientific breakthroughs being made on the center's state-of-the-art platforms and previewing the path to petascale science.

SC06, the international conference for high-performance computing, networking, storage, and

Both photos: The ORNL NCCS booth at SC06.

Computational Biology Workshop attendees.

analysis, was held in Tampa, Florida. The NCCS unveiled a new booth for the conference, featuring all-electronic content focused on advanced networking, leadership computing and applications, and the many scientific applications that are enabled through high-performance computing at the center. Interactive kiosks and a large PowerWall allowed visitors to learn more about current activities in areas such as supernova evolution, fusion plasma simulation, and atmospheric carbon dioxide patterning.

NCCS presentations throughout the conference focused on the work of leading computational scientists and researchers from around the country, representing academia, national laboratories, and industry.

Computational Biology Workshop

December 11–12

Researchers and computational scientists from across the country gathered at ORNL for a 2-day meeting, “Future of BioMolecular Simulations: From Ab Initio to Nano-molecular Machines,” to discuss biological computer simulation in the age of petascale computing.

The NCCS-sponsored workshop allowed members of the biological sciences community—computational biochemists, biophysicists, and developers of biomolecular simulation codes, among others—to discuss the successes and challenges of a field working to take advantage of increasingly powerful computer systems. Attendees came from Harvard, Stanford, and New York universities; the San Diego Supercomputing Center; and a variety of other institutions.

Workshop attendees agreed that as supercomputers continue to grow in power, the codes used in biological simulation must advance as well. The workshop enabled researchers to identify and discuss the issues they will be facing, including a variety of software and development challenges.

NCCS Seminar Series

The NCCS Seminar Series got started on October 3 with a talk from University of California–Irvine physicist, NCCS user, and American Physical Society fellow Zhihong Lin.

Lin spoke on the topic “Gyrokinetic Particle Simulation of Turbulent Transport in Fusion Plasmas.” He is a member of the fusion research team headed by W. W. Lee of Princeton Plasma Physics Laboratory, which has large INCITE allocations on both Jaguar and Phoenix.

He discussed performance of the Gyrokinetic Toroidal Code, which simulates turbulence in the multinational ITER fusion reactor, and opportunities and challenges for fusion researchers in using future petascale supercomputers.

The seminar series brings speakers to ORNL to share their research and expertise and interact with NCCS staff. It also provides an opportunity for collaboration building and helps to communicate NCCS objectives and accomplishments throughout ORNL.

STATISTICS

The charts on this and the following page present a statistical overview of the NCCS in 2006: 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.

The charts on this page lay out NCCS research by program office within DOE's Office of Science and by discipline. Reflecting DOE's core mission, the Basic Energy Sciences program used the most processor hours on the center's Cray XT3 Jaguar system, followed by the Fusion Energy Sciences program. The Biological and Environmental Research program made the most use of the center's Cray X1E supercomputer, followed by Basic Energy Sciences. Among disciplines, fusion scientists took the greatest advantage of Jaguar, followed by materials scientists and combustion scientists. Climate scientists made the most use of Phoenix, followed by fusion scientists and chemists.

NCCS USAGE BY PROGRAM

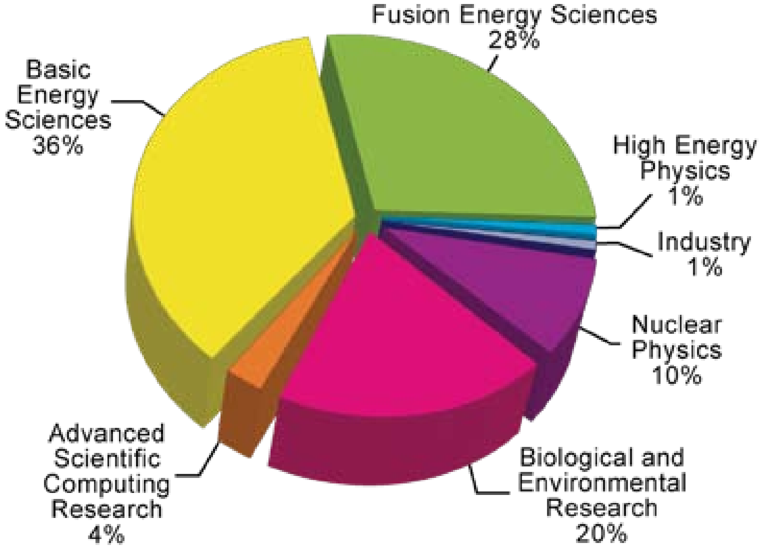

NCCS USAGE BY DISCIPLINE

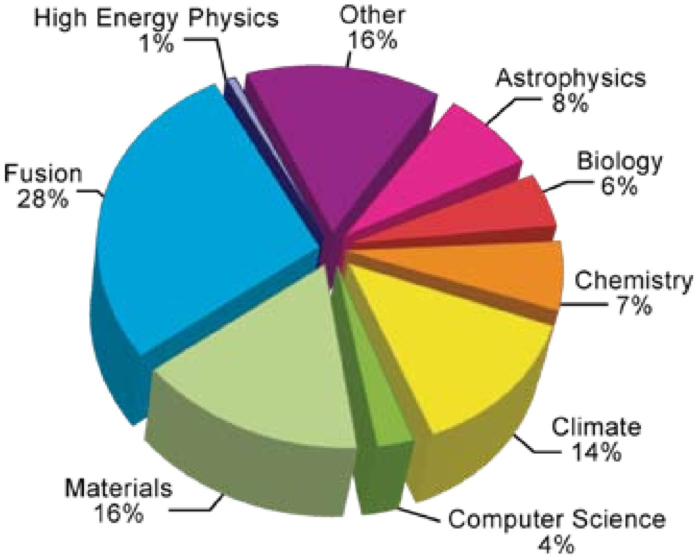

Active Users by Sponsor

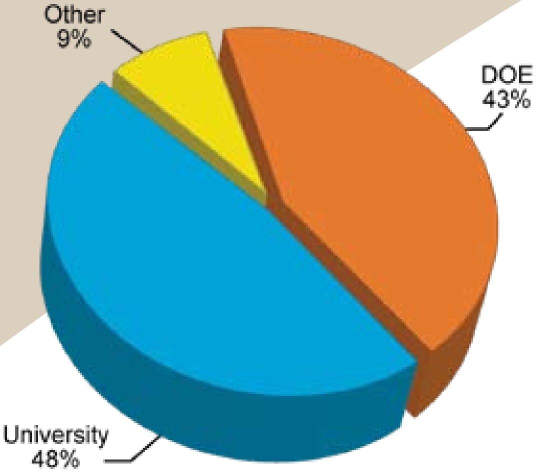

The pie charts on this page show where NCCS projects are sponsored as well as the disciplines that supply the center's active users. The chart at the top illustrates why university researchers are indispensable to the center, as universities provide nearly half of its active users. As the chart on the bottom indicates, the greatest number of active users, more than a quarter of the total, comes from computer science. Fusion science, climate science, and chemistry each also provide more than 10 percent of the center's active users.

Active Users by Discipline

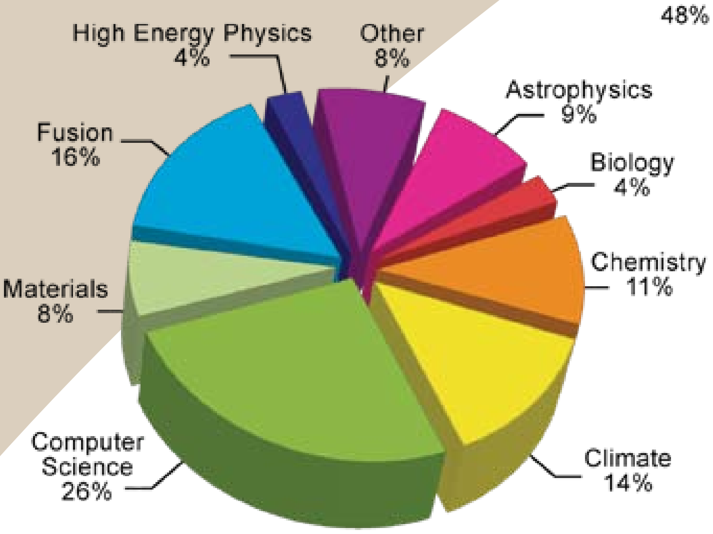

The usage chart for the High Performance Storage System shows the trends in storage of data from 1998 to 2006. It indicates slow but steady growth until about 2002, when the amount of data stored began to climb sharply. The chart shows an increase of over 400 terabytes in data archived from 2005 to 2006.

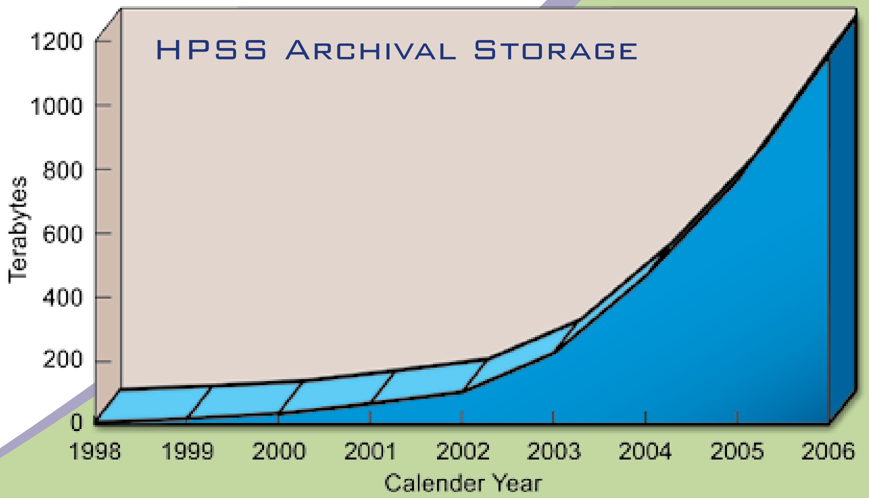

SELECTED USER PUBLICATIONS

Using the resources of the NCCS, researchers continue to produce numerous scientific breakthroughs. Listed below are a small sampling of the more than 100 publications from 2006, grouped by related discipline, that highlight a portion of the work being achieved through the combination of talented researchers, leadership-class machines, and the dedicated staff of the NCCS. For the complete list, please see www.nccs.gov.

ASTROPHYSICS

Bruenn, S. W., et al. 2006. ”Modeling core collapse supernovae in 2 and 3 dimensions with spectral neutrino transport.” *Journal of Physics Conference Series* **46**, 393–402 (Sept.).

Dessart, L., et al. 2006. “Multi-dimensional radiation/hydrodynamic simulations of protoneutron star convection.” *Astrophysical Journal* **645**, no.1, 534–550 (July).

Ropke, F. K., et al. 2006. ”Multi-spot ignition in type Ia supernova models.” *Astronomy & Astrophysics* **448**, no. 1, 1–14 (Mar.).

BIOLOGY

Luo, G. B., et al. 2006. “Dynamic distance disorder in proteins is caused by trapping.” *Journal of Physical Chemistry B* **110**, no. 19, 9363–9367 (May).

CHEMISTRY

Jiang, D. E., B. G. Sumpter, and S. Dai. 2006. “Structure and bonding between an aryl group and metal surfaces.” *Journal of the American Chemical Society* **128**, no. 18, 6030–6031 (May 10).

CLIMATE

Dickinson, R. E., et al. 2006. “The Community Land Model and its climate statistics as a component of the Community Climate System Model.” *Journal of Climate* **19**, no. 11, 2144–2161 (Jan.).

Hack, J. J., et al. 2006. “CCSM CAM3 climate simulation sensitivity to changes in horizontal resolution.” *Journal of Climate* **19**, no. 11, 2267–2289 (June).

Friedlingstein, P., et al. 2006. “Climate-carbon cycle feedback analysis, results from the C4MIP model intercomparison.” *Journal of Climate* **19**, no. 14, 3337–3353 (July)

COMBUSTION

Sankaran, R., et al. 2006. “Direct numerical simulation of turbulent lean premixed combustion.” *Journal of Physics Conference Series* **46**, 38–42 (Sept.).

FUSION

Candy, J., et al. 2006. “Relevance of the parallel nonlinearity in gyrokinetic simulations of tokamak plasmas.” *Physics of Plasmas* **13**, no. 7, article no. 074501 (July).

Baek, Hoyul, et al. 2006. “Neoclassical polarization drift of collisionless single ions in a sheared radial electric field in a tokamak magnetic geometry.” *Physics of Plasmas* **13**, article no. 012503 (January).

Waltz, R. E., et al. 2006. “ Gyrokinetic simulations of off-axis minimum-q profile corrugations.” *Physics of Plasmas* **13**, no. 5, article no. 052301 (May).

HIGH ENERGY PHYSICS

Adamson, P., et al. 2006. “First observations of separated atmospheric nu(mu) and (nu)over-bar(mu) events in the MINOS detector.” *Physical Review D* **73**, no. 7 (April).

MATERIALS

Maier, T. A., M. Jarrell, and D. J. Scalapino. 2006. “Pairing interaction in the two-dimensional Hubbard model studied with a dynamic cluster quantum Monte Carlo approximation.” *Physical Review B* **74**, no. 9, article no. 094513 (Sept.).


OAK RIDGE
CENTER FOR COMPUTATIONAL SCIENCES
5600

ORNL 2007-G00545/jcp

DEPARTMENT OF ENERGY
UNITED STATES OF AMERICA