Oak Ridge National Laboratory

REVIEW

• MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY •

Vol. 42 • No. 1 • 2009
www.ornl.gov/ORNLReview

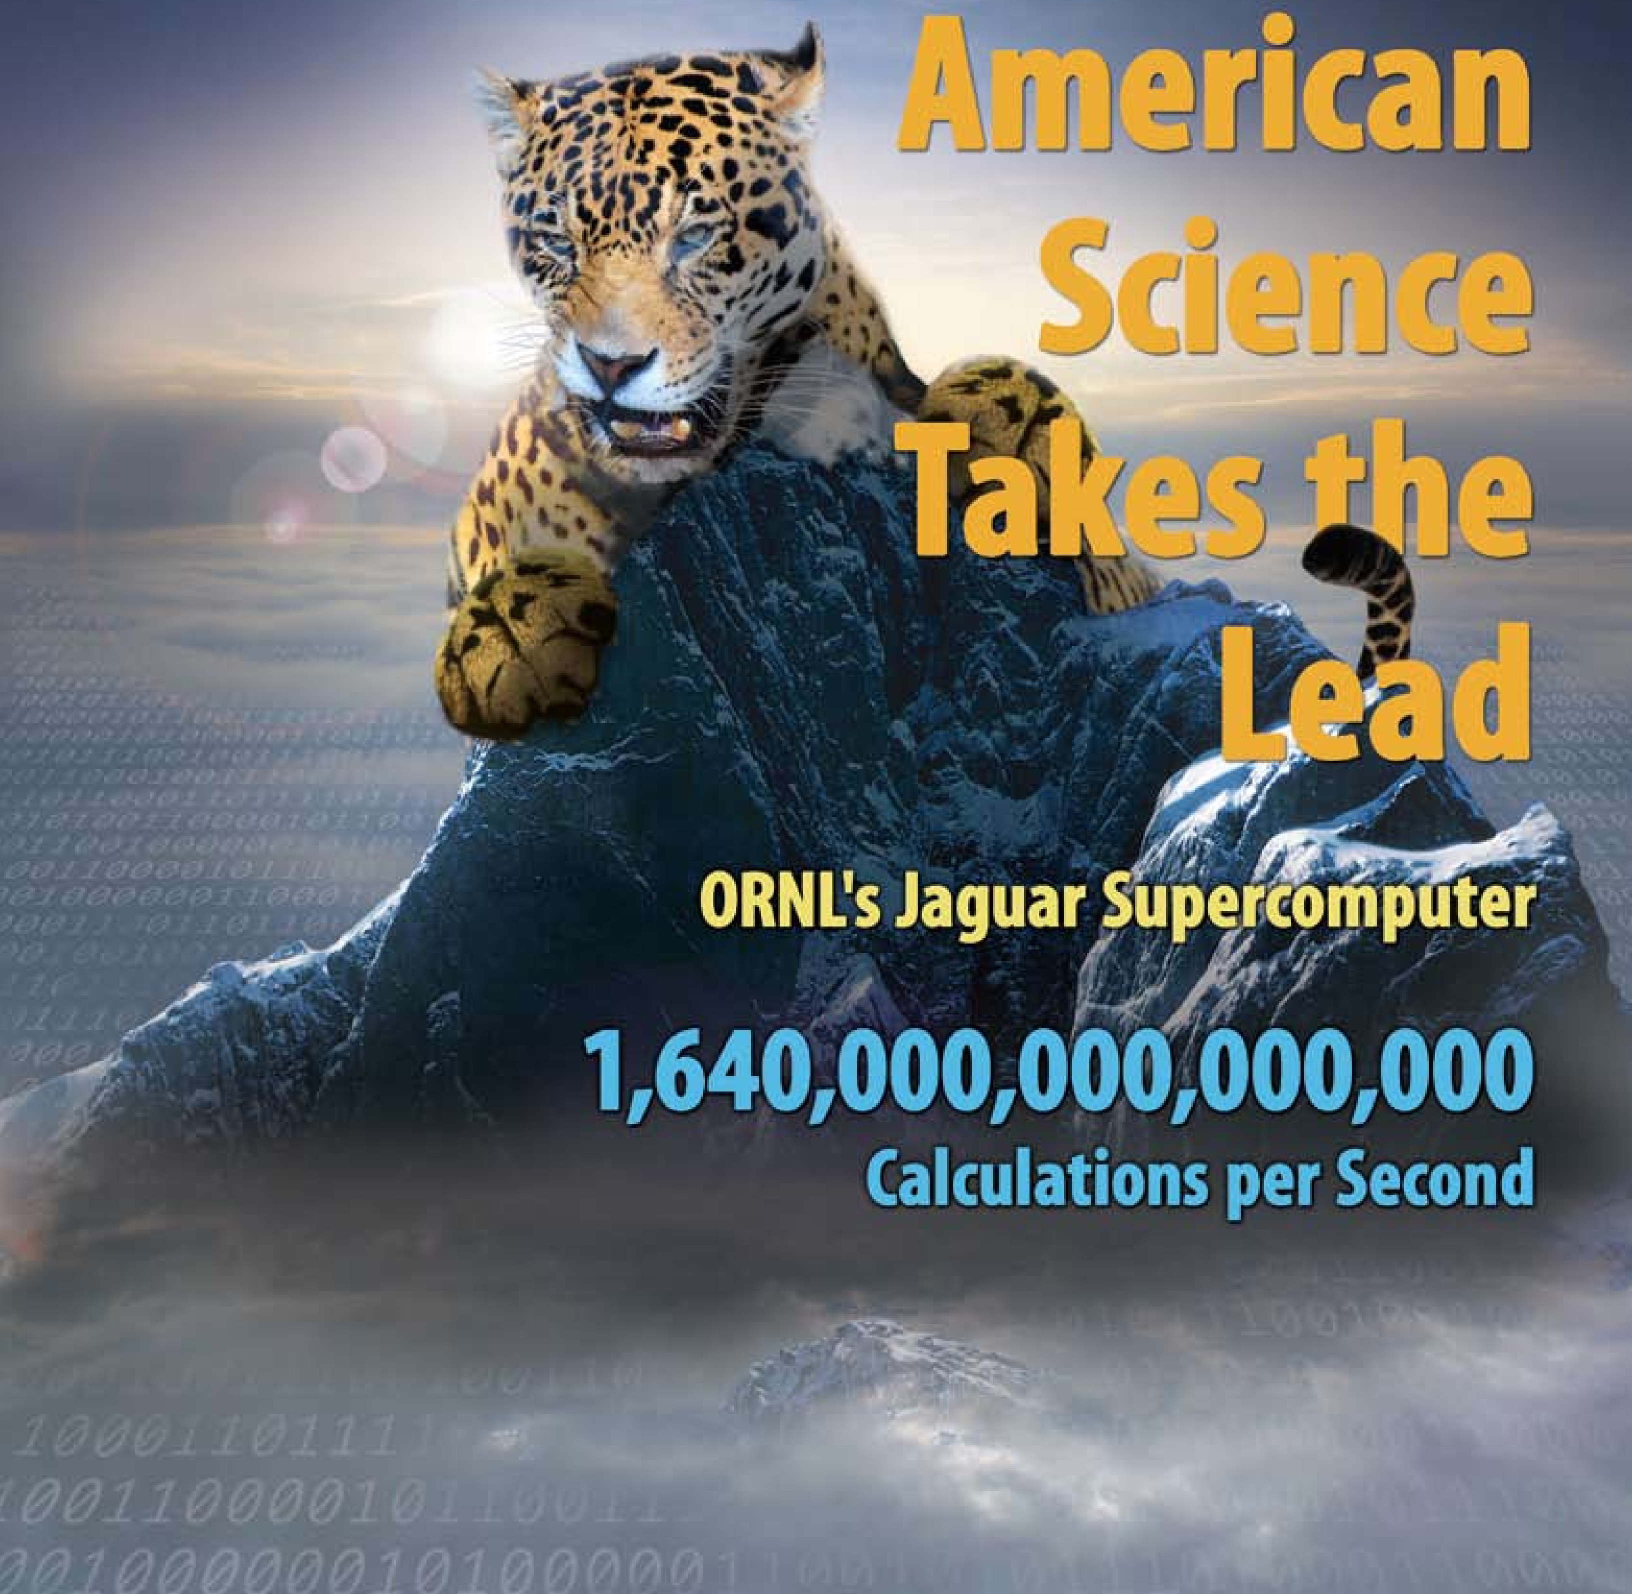

American Science Takes the Lead

ORNL's Jaguar Supercomputer

1,640,000,000,000,000 Calculations per Second

OAK RIDGE NATIONAL LABORATORY
REVIEW
• MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY •
Vol. 42, No. 1, 2009

contents

editorial

features

a closer view

research horizons

awards

news & notes

- Jaguar's Coming Out Party
- The European Spallation Source
- Clean Energy Summit

A Vision REALIZED

At approximately 3:00 a.m. on the morning of November 4, 1943, a messenger drove as fast as his car could manage through the muddy streets of Oak Ridge. Sloshing across the front yard, he knocked loudly on the door of the little two-bedroom pre-fabricated house. The short, bald man he roused from bed was Enrico Fermi, working under an assumed name on a project that would literally reshape the world's scientific and political landscape. Across town at the experimental Graphite Reactor, Fermi was leading the top-secret effort to develop the world's first sustained nuclear reaction. The goal was a revolutionary new source of power that, if used wisely, could be of enormous benefit to humankind.

The news Fermi received was worth the intrusion. Located in the complex that later became Oak Ridge National Laboratory, the Graphite Reactor had gone "critical," signaling the birth of the nuclear age. The extraordinary accomplishments by Fermi and his team were the result of an equally extraordinary vision, conceived by the scientific community, then given expression and resources by the United States government.

Almost exactly sixty-five years later, on November 18, 2008, another milestone was reached in Oak Ridge, the result of a similar vision and with, some believe, the ability to transform scientific inquiry as dramatically as its nuclear predecessor. Researchers at ORNL's National Center for Computational Sciences, located less than a half-mile from the historic Graphite Reactor, announced the successful deployment of a "petaflop" supercomputer, fittingly named the Jaguar, that had smashed through the mind-boggling threshold of 1,000 trillion calculations per second. The implications for American scientific leadership were profound. Not only had the world's center for high-performance computing returned to the United States. Of even greater significance was the ability of Jaguar to use modeling and simulation to address scientific challenges that were thought by many to be intractable only a few years earlier.

This issue of the *ORNL Review* is dedicated to the vision, articulated by the scientific community and championed by the Department of Energy, to build a Leadership-Class Facility dedicated to restoring American leadership in high-performance computing. The vision was daring: to build within four years a machine for open science 25 times more powerful than the world's largest, located at the time in Japan. How the Leadership-Class Computing Facility came to Oak Ridge, and how ORNL staff assembled the resources needed to deliver such an ambitious project with expanded scope on time and on budget—is a story with achievements as worthy as those of their predecessors more than six decades earlier.

With technological changes that occur faster than in most other scientific fields, the relevance of petaflop computing is sometimes lost in a culture more fascinated by rankings of the largest machines than by their potential for discovery. As ORNL plans to expand the Laboratory's computational capacity to 2.5 petaflops in 2009, one senses that in time Jaguar's notoriety as the world's most powerful supercomputer will be surpassed by the delivery of transformational solutions to the challenges of sustainable energy, climate change and human health.

As with the discovery that November night in 1943, the solutions made possible by Jaguar in the years ahead will be traced to a shared vision, born among the scientists who articulated its potential, led by the government that believed it possible, and delivered by the men and women at Oak Ridge bold enough to do what had never been done.

Thomas Zacharia

Thomas Zacharia

Associate Laboratory Director

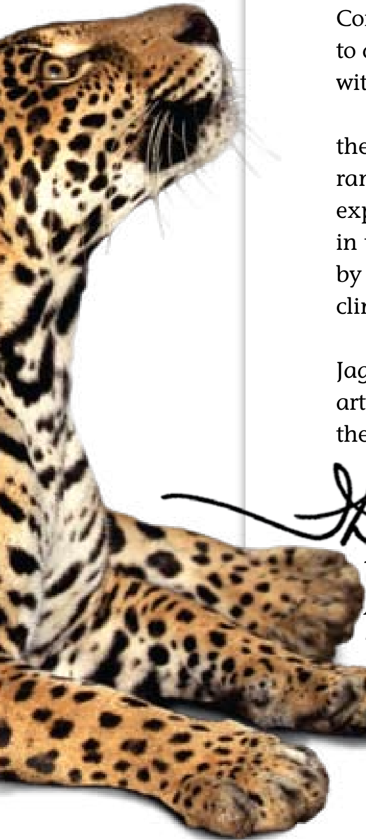

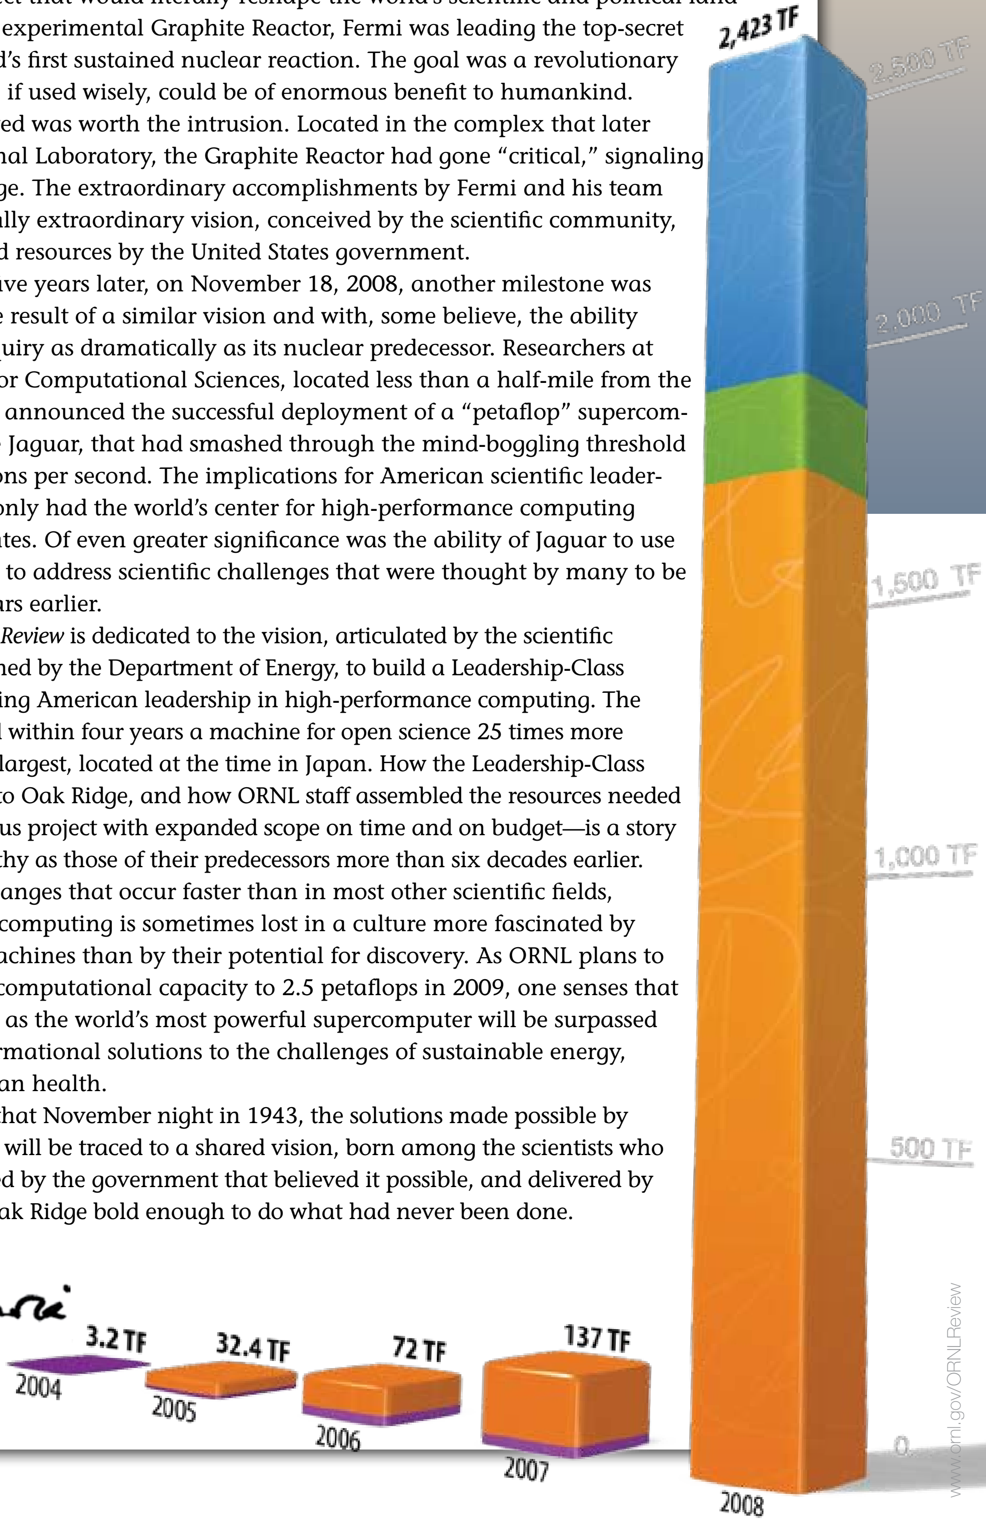

News & Notes

Jaguar's Coming Out Party

ORNL's new Jaguar supercomputer received a variety of honors at the SC08 conference in Austin, Texas—the premier international conference for high-performance computing. This recognition represented the culmination of a three-year effort to create the world's most powerful computing system.

Like its namesake, the Jaguar supercomputer's forte is speed. When the rankings of the 500 fastest supercomputers were announced at the conference, Jaguar garnered honors as the fastest supercomputer in the world for open science, with a top speed of 1.64 petaflops (quadrillion calculations per second). Only the second system to ever break the petaflop barrier, Jaguar was just fractionally slower than the first petaflop system, Los Alamos National Laboratory's Roadrunner. The Roadrunner is a classified machine used primarily to perform calculations that help to ensure the reliability of the nation's nuclear weapons stockpile.

Also announced at SC08 was the 2008 Association for Computing Machinery's Gordon Bell Prize—claimed by an ORNL team headed by Thomas Schulthess. The award recognized the feat of attaining the fastest performance ever in a scientific supercomputing application. The team, consisting of Schulthess, Thomas Maier, Michael Summers and Gonzalo Alvarez, achieved an application speed of 1.352 petaflops on Jaguar while running a simulation of superconductors.

Finally, Jaguar won awards in three of the four categories of the High-Performance Computing Challenge competition. The Oak Ridge machine took two first place awards for speed in solving a set of linear algebra equations and for speed in fetching and storing data across a network connection. Jaguar also won third place for speed in running the Global-Fast Fourier Transformation, a common scientific algorithm.

ORNL's exhibit at SC08 employed a wealth of data visualizations to highlight the power and versatility of the Jaguar supercomputer.

The European Spallation Source

Only two years have passed since ORNL's Spallation Neutron Source began to generate the most intense pulsed neutron beams in the world, but the race to build a successor to the SNS—or at least a worthy competitor—has already begun.

In Europe, interest is growing in building a "next-generation" spallation source. Recently, delegations from all three of the groups competing for the privilege of hosting the ESS—Bilbao, Hungary and Scandanavia—have visited the ORNL on fact-finding missions. The visitors are not only interested in the scientific programs at the SNS, but also in the impact the multi-billion dollar facility has had on the region, including the local economy and schools.

ORNL Director, Thom Mason, who headed up the SNS project during most of its planning and construction, recalls that "When we were building SNS, we benefitted a lot from the interactions we had with other operating facilities around the world—in Europe primarily, and in Japan. So now that the Europeans are contemplating building a source, they're trying to draw on our experience in the same way we drew on theirs."

Mason stresses that there's more involved in decisions about whether or where to build a facility than purely scientific considerations. "There's an element that's related to scientific collaboration and solving technical problems," says Mason, "but building a facility of that scale is also a big deal for the region economically, so the visits to some extent have been oriented more around the economic development impact of a major R&D facility"

The ESS website suggests that, because any European spallation source is not likely to come online before 2017, the European scientific community should plan on building a facility that's five times as capable as the SNS in many areas. So, will that sort of capability make the SNS obsolete in less than ten years? Not likely, according to Mason: "The European design goal is a five megawatt facility. SNS is currently 1.4 megawatts, but we're working on a power upgrade to increase that. So we're a bit of a moving target. The European spallation source certainly won't be five times as powerful as the SNS by the time it's finished."

Mason also explained that the proposed European spallation source is a different type of neutron source than the SNS. "The proposed European machine is a long-pulse source and the SNS is short-pulse," Mason said, "so it has different characteristics. Getting to five megawatts, for example, is easier with a long-pulse source. The second target station planned for the SNS would be a long-pulse target. I'm confident that by the time the ESS is built the SNS will have moved ahead, so we're going to stay competitive under any scenario."

Because the construction of a world-class neutron source is such a massive undertaking, both the European spallation source and the SNS are designed to be upgraded as technology improves. "Eventually we will hit some sort of limit," says Mason, "but accelerators tend to have capabilities that go beyond what was originally thought possible—just because the technology improves."

The SNS has a 20-year plan for upgrades that includes doubling the beam power, doubling the number of instruments with a second target station, and developing a long-pulse source. "That doesn't mean we won't be able to upgrade it after that," says Mason, "it's just that, at that point in time, the technology is sufficiently far out that it's hard to know what the path will be. That'll keep us busy for 10-20 years. After that it will depend on how the technology evolves."

Clean Energy Summit

Clean Energy—including its generation, conservation, and how to use it to generate jobs—was the focus among the high-profile business and research leaders at the Governor's Summit on Clean Energy held in Knoxville, Tennessee. Moderators of the summit's sessions included ORNL Director Thom Mason; Associate Laboratory Director for Energy, Dana Christensen; and Tom Rodgers, Director of Industrial and Economic Development Partnerships.

The summit's keynote speaker was Stephan Jacoby, CEO of Volkswagen Group of America. Often noted for its "green" manufacturing practices, VW's recent decision to site a $1 billion plant in Chattanooga was linked in part to the availability of ORNL's nearby research capabilities.

Mason summed up the importance of the summit for the laboratory. "The interesting thing for ORNL is that while we're rich in energy science and technology resources, in order for the laboratory to put its R&D into practice, we need partners to get things deployed—initially on a demonstration scale and eventually as a self-sustaining exercise."

Mason also noted that ORNL maintains close working relationships with a number of regional partners, including the Tennessee Valley Authority, the city of Knoxville, and Nissan USA. "These three partners are right in our neighborhood," he said. "Although we're a national lab with a national scope, it's always easier to get things going if you have proximity and familiarity."

Christensen, who also sits on the Governor's Clean Energy Task Force, indicated that the Southeast is already becoming an incubator for next-generation energy technology with demonstrations involving essentially every energy production technology and with an emerging emphasis on energy efficiency, especially in personal residences. "The Governor is looking to capitalize on this strength to position Tennessee as a national leader in many of these technology areas, such as solar photovoltaics."

A risk enables ORNL to break the petaflop barrier

THE Gamble That Paid

The announcement shook the scientific world. With no forewarning, the Japanese in 2003 unveiled a new generation of supercomputer that would literally transform the capabilities and horizons of scientific research. Called the Earth Simulator, the Japanese machine could perform at an unheard of speed of 36 "teraflops," or 36 trillion floating operations per second. Dedicated primarily to climate research, the Earth Simulator was not only the world's most powerful computer—its extraordinary performance exceeded the combined capabilities of the 20 largest American machines.

In the United States, the implications of the Earth Simulator were clear. If America yielded leadership in high-performance computing, the likely result would be a similar loss of status in the international scientific community, followed by an inevitable decline of American economic competitiveness.

In a response reminiscent of the American reaction to the Russian launch of the Sputnik satellite in 1957, the Department of Energy announced plans to build what the agency called a "Leadership-Class Computing Facility." DOE's Office of Science envisioned a facility that possessed, in addition to a collection of intellectual talent, power and cooling infrastructure that surpassed by a degree of magnitude any of the existing computer centers currently housed in American universities or national laboratories. Indeed, the Leadership Class Facility would provide a machine

not just larger than the Earth Simulator, but eventually capable of attaining the unimaginable speed of a “petaflop,” or 1,000 trillion calculations per second.

The goal of this computational “arms race” was not just the bragging rights that come with the world’s largest computer. Of far greater significance was the potential of high-performance computing to make possible transformational discoveries in some of science’s most important and most daunting challenges. With computational tools that could barely be conceived only a decade earlier, scientists predicted that the modeling and simulation made possible with a petaflop machine would place in an entirely new context discussions about the ability to understand long-term scientific challenges, such as fusion energy and climate change. The ability to reduce the time needed to process and analyze massive volumes of data—in some cases from months to days—held out the potential for comparable reductions in the time required to move new technologies from the laboratory to the marketplace.

The high-performance computing sector is one where technological breakthroughs often are measured in months rather than years. Despite the Department of Energy’s desire to move rapidly, building a Leadership-Class Computing Facility faced an immediate dilemma. While a number of institutions possessed the intellectual talent necessary to design the software and hardware architecture for a petaflop machine, virtually none of these computing centers had readily available either the building or the infrastructure needed to support a supercomputer that demanded power and cooling at levels five hundred to one thousand times greater than most existing machines.

One exception was at Oak Ridge National Laboratory. Using a creative new scheme that utilized private funds and land deeded from the Department of Energy, UT-Battelle, the laboratory’s managing contractor, had just completed construction of a $73 million, 350,000 square foot facility that included a full acre of America’s most modern computational space. An adequate and reliable source of power, a critical element that was increasingly unavailable in some states, was provided by a new substation built by the Tennessee Valley Authority on the ORNL campus.

The design for ORNL’s computational center also included networking and data-handling resources to support a future petaflop machine. The new facility boasted 10-gigabyte-per-second connections to the ESnet and Internet2 networks, and a scalable high-performance storage system for storing simulation data. The disk subsystem could transfer data at speeds greater than 200 gigabytes per second.

An enormous risk

From one perspective, the decision by ORNL officials to build the nation’s largest computer center was an enormous risk. The facility, named the National Center for Computational Sciences, was constructed with essentially no large-scale program to operate, leaving open the possibility that Oak Ridge could have been home to the world’s most modern roller skating rink. In this instance, the new Oak Ridge facility represented the confluence of foresight, boldness, and luck. More than a year before the announcement of the Japanese Earth Simulator, ORNL officials privately concluded that it was only a matter of time before the U.S. government made major investments in building a high-performance supercomputer. While others may have shared this prediction, **the Oak Ridge team was certainly among the first to grasp the implications of such a machine on the need for a supporting infrastructure that dwarfed previous computer centers in size and cost.**

The prediction proved accurate. When in 2003 the Department of Energy invited competitive bids for a $500 million, four-year project to build a petaflop machine, the Oak Ridge proposal contained a critical feature: ORNL would have a new facility ready on day one, with state-of-the-art connectivity, reliable power and space for future expansion. Unknown to many was the fact that, in anticipation of a new program, ORNL had quietly been hiring some 200 computational scientists from around the world with expertise in quantum physics, astrophysics, materials science, climate, chemistry, and biology.

The gamble paid off. In May 2004, after a spirited competition among America's leading computational programs, Oak Ridge was selected as the site for the Leadership-Class Computing Facility. DOE's charge was straightforward. Models and simulations on the supercomputer would offer scientists a "third pillar of science," a transformational addition to the historic pillars of theory and experiment. By December 2008, equipped with a new generation of software and operated with high standards of efficiency, a petaflop machine would enable researchers to explore biology, chemistry, and physics in ways previously unimaginable. DOE, quite simply, expected ORNL to provide scientists with virtual laboratories unmatched by any other computing facility in the world.

In a class by itself

Beginning with a 26-teraflop system in 2005, Oak Ridge embarked upon a three-year series of aggressive upgrades designed to build the world's most powerful computing system. The existing Cray XT was upgraded to 119 teraflops in 2006 and to 263 teraflops in 2007.

Four years later, the dream of building a petaflop machine and restoring U.S. leadership in high-performance computing is a reality. On November 14, 2008, Oak Ridge officials announced the successful testing of the new Cray XT, called Jaguar, with a peak performance of 1.64 petaflops. Smashing through the petaflop barrier, the Jaguar incorporated 1.382-petaflop XT5 and 263-teraflop XT4 systems. With approximately 182,000 AMD Opteron processing cores, the new 1.64-petaflop system is more than 60 times larger than its original predecessor.

Aided by modern facilities, Jaguar is also the culmination of a close partnership between ORNL and Cray, dedicated to pushing computing capability relentlessly upward through a series of upgrades. The most recent upgrade occurred in 2008, when a 263-teraflop Cray XT4 was linked to a 1.4-petaflop Cray XT5. The combined system uses an InfiniBand network, a 10-petabyte file system, and approximately 182,000 processing cores to form Oak Ridge's current 1.64-petaflop system. Occupying 284 cabinets, Jaguar uses the latest generation of quad-core Opteron processors from AMD and features 362 terabytes of memory. The machine has 578 terabytes per second of memory bandwidth and unprecedented input/output bandwidth of 284 gigabytes per second to tackle the biggest bottleneck in supercomputing systems—moving data into and out of processors.

Keeping the machines from melting through the floor is no small task. The XT5 portion of Jaguar has a power density of more than 2,000 watts per square foot, creating commensurate heat that must be constantly dissipated. To cool the system, Cray worked with its partner, Liebert, to develop ECOphlex, a technology that pipes a liquid refrigerant through an evaporator on the top and bottom of each cabinet. Fans flush heat into the evaporator, where it vaporizes the refrigerant. The vaporization process absorbs the heat. The coolant is then condensed back to the liquid phase in a chilled-water heat exchange system, transferring the heat to chilled water. This extremely efficient cooling process is a critical element in making possible the design of increasingly powerful supercomputers. The new cooling technology also benefits the computer center's efficiency. While cooling often adds some 80 percent to the volume of power required at computing centers, at Oak Ridge the new cooling process adds only 30 percent.

It's all about the science

As the world's first petaflop system available for open research, Jaguar is already in high demand by scientists who are honing their codes to take advantage of the machine's blistering speed. Jaguar represents a unique balance among speed, power, and other elements essential to scientific discovery. Several design elements make

Jaguar the machine of choice for computational sciences—more memory than any other machine, more powerful processors, more I/O bandwidth, and the high-speed SeaStar network developed specifically for very-high-performance computing.

Researchers thus far have been enormously successful in utilizing Jaguar's architecture. From a programming standpoint, the upgraded Jaguar is essentially the same as the XT4 that scientists in Oak Ridge have been using for three years. A consistent programming model enables users to continue to evolve existing codes rather than develop new ones. Applications that ran on previous versions of Jaguar can be recompiled, tuned for efficiency, and then run on the new machine. As the CPU performance continues to grow, the system's basic programming model remains intact. For users, such continuity is critically important for applications that typically last for 20 to 30 years.

Speed and efficiency aside, Jaguar's ultimate value will be measured by the science the machine can deliver. The Department of Energy has dedicated Jaguar, unlike most supercomputers, to addressing a relatively small of number of "grand scientific challenges" too large and complex for most

“We now for the first time have the tools to address some of science’s most intimidating questions"

To keep Jaguar cool, Cray worked with its partner, Liebert, to develop ECOphlex, a technology that pipes a liquid refrigerant through an evaporator on the top and bottom of each cabinet.

existing systems. A single project on Jaguar might consume millions of processor hours and generate an avalanche of data. Proposed projects are peer-reviewed and funded by DOE's Innovative and Novel Computational Impact on Theory and Experiment program.

Early results are encouraging. A report released in October 2008 by the DOE Office of Advanced Scientific Computing Research showcased ten scientific computing milestones, including five projects conducted at Oak Ridge National Laboratory. Among the highlighted ORNL-based research was one of the largest simulations ever produced of plasma confinement in a fusion reactor, which could potentially pave the way for carbon-free sustainable energy production. Jaguar also performed a billion-particle simulation of the dark matter halo of the Milky Way galaxy, in which researchers performed the largest simulation to date of the dark matter cloud holding our galaxy together. The report noted ORNL scientists who completed combustion simulations that dissected how flames stabilize, extinguish, and reunite, showing the path to cleaner, more efficient diesel-engine designs.

Kraken Rising

Hot on the heels of Jaguar's success, ORNL is partnering with the University of Tennessee on the assembly and start-up of Kraken, another massive Cray XT5 system. Initially, the new system will have more than 600 teraflops of processing power, but upgrades scheduled for later this year will enable Kraken to become the second ORNL-based system to crack the petaflop barrier.

When fully deployed, Kraken—named after the giant, multitentacled, sea creature of Norse legend—will use its 100,000 processor cores to grapple with scientific questions that have eluded the grasp of previous computing systems.

While the Kraken and Jaguar are nearly identical in terms of hardware, they will serve the research needs of two different sets of users. Kraken, sponsored by the National Science Foundation, will be available primarily to university-based users around the country. Jaguar, on the other hand, is available to a smaller group of users concentrating on high-impact projects of national importance—as determined by the U.S. Department of Energy. The users on Jaguar are chosen annually through a peer-reviewed competition from proposals submitted from industry, academia and government agencies.

"Kraken's research agenda is determined by the National Science Foundation, the University of Tennessee and the greater academic community," said Buddy Bland, project director for ORNL's Leadership Computing Facility. "The NSF's mission enables more curiosity-driven research, providing a scientist with an interesting concept time on Kraken for investigation. Jaguar, on the other hand, is focused directly on solving critical national problems within the Department of Energy's mission."

One expects that the experience gained in bringing Jaguar on line will enable Kraken to rise from the depths somewhat faster than its predecessor. When that occurs, the ORNL-UT research alliance may boast the two fastest supercomputers in the world dedicated to open scientific research.

The University of Tennessee has long been accustomed to national rankings in athletics; now the university's Kraken ranks as the largest academic computer in the world. As with Jaguar, the implications for the university's research program are just beginning to unfold.

ORNL's Associate Laboratory Director for Computational Sciences, Thomas Zacharia, believes the research community is just beginning to understand Jaguar's potential. "We now for the first time have the tools to address some of science's most intimidating questions: How does the earth's atmosphere affect ocean circulation? How do enzymes aid biofuels production? How do proteins misfold in certain diseases?" Zacharia says. The answers, he contends, will open up dramatic opportunities, not just for science, but for American economic growth. Already, he notes, leading companies such as Boeing and General Motors have used Jaguar's simulations to improve materials for their products.

To Zacharia and his colleagues at ORNL, the end is by no means in sight. Indeed, his plans for 2009 call for Oak Ridge to offer two supercomputers with a combined performance of more than 2.5 petaflops. The goal is not considered a fantasy at Oak Ridge, where since 1991 computational power has increased a millionfold. From their perspective, the risk is worth the opportunity.

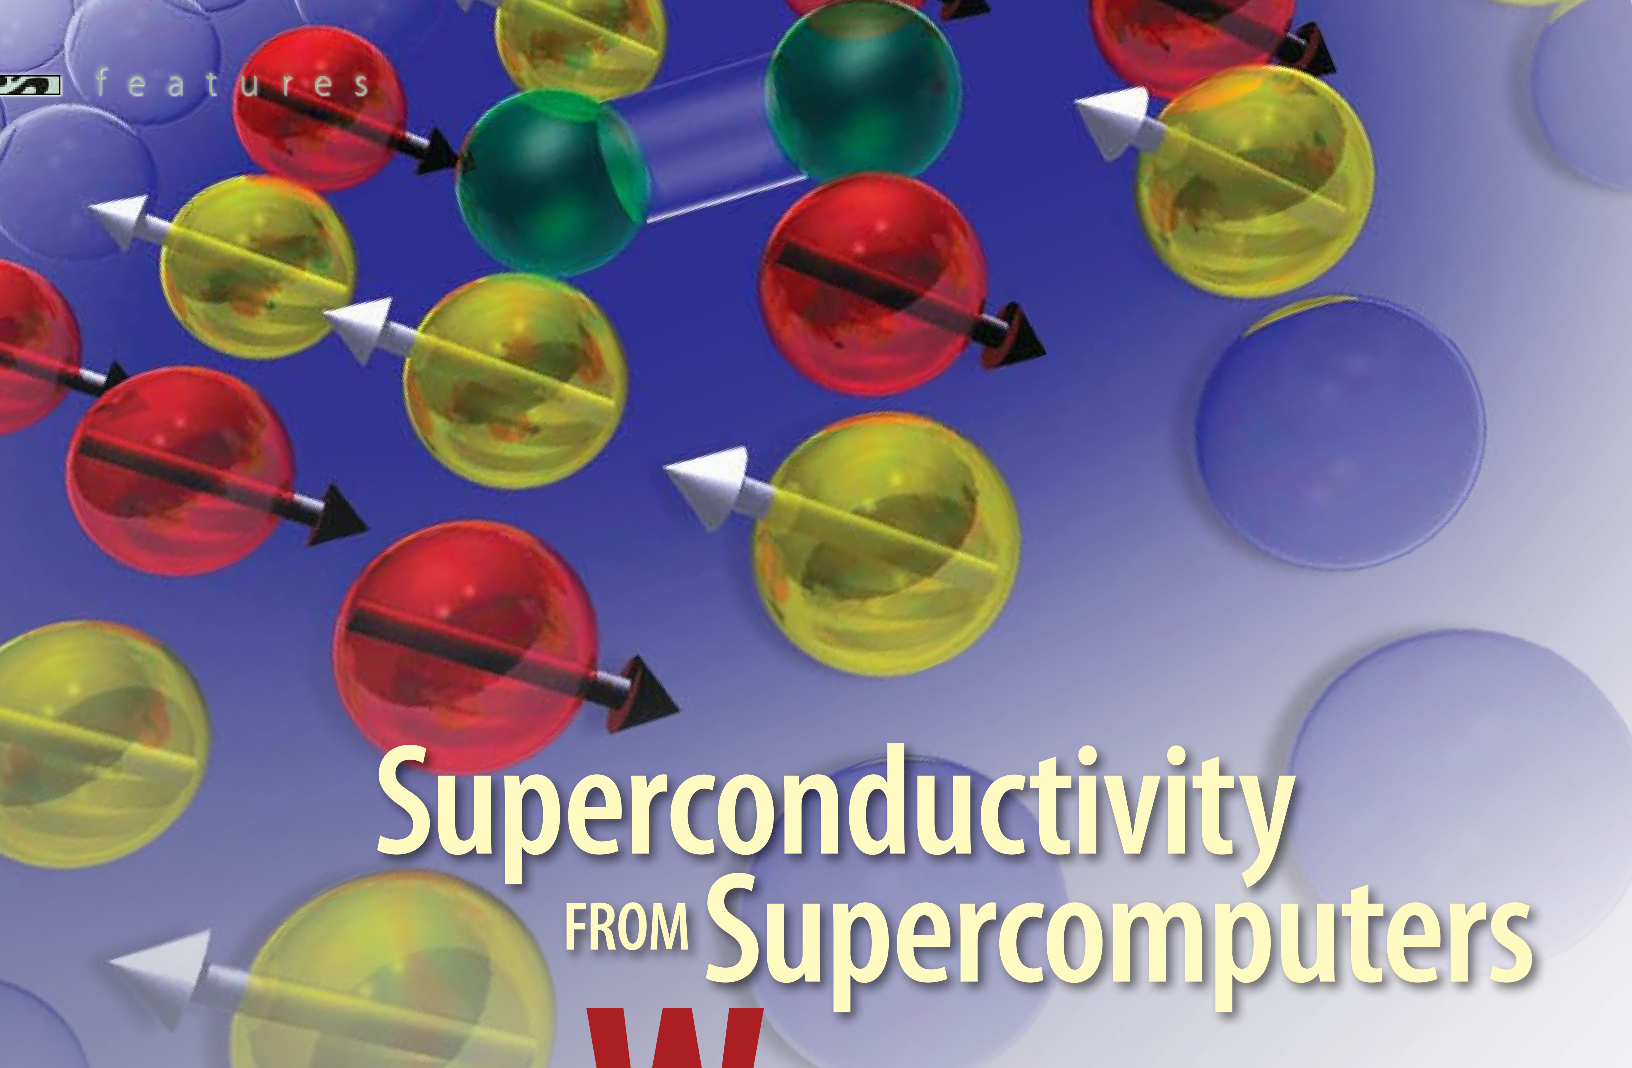

Superconductivity FROM Supercomputers

New supercomputers at ORNL are unlocking the mysteries of superconducting materials

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

While physicists rarely have the image of the 1960s counterculture, the first American Physical Society session on high-temperature superconductors, held in 1987 at New York City's Hilton Hotel, has been termed by some the "Woodstock of Physics." Excited by the discovery of high-temperature superconductors (HTSCs) the previous year, more than 2,000 physicists packed a hotel meeting room and spilled out into the corridors, fighting for the privilege of hearing the technical papers first-hand. The session lasted all night. The conference is perhaps the only meeting of physicists ever characterized as "a riot."

Twenty years after "Woodstock," the potential of high temperature superconductors both tantalizes and taunts the scientific community. Not unlike Woodstock, the jubilation of 1987 eventually encountered reality. Fulfilling the promise has proved to be much harder than predicted by many of those starry-eyed physicists two decades ago. HTSCs are being used in some applications, and the technology is advancing, but superconductivity is far from being a part of daily life.

New tools

The technological breakthroughs needed to accelerate the pace of superconductivity research may now be possible with a new generation of supercomputers. In calculations conducted during 2006, a team of scientists used the unparalleled computing resources at Oak Ridge National Laboratory to identify a mechanism that provides clues to the mystery of how HTSC materials work. The discovery may eventually help to realize the potential of these highly energy-efficient materials.

HTSCs are materials that conduct electricity without resistance at temperatures as high as 150 K (roughly –190°F). While extremely cold, the temperature is relatively balmy compared with the level at which conventional superconductors must operate—near absolute zero, or below –400°F. Because the higher temperatures are more manageable, HTSCs have much more technological potential than conventional superconductors.

A theoretical understanding of why HTSCs lose their resistance to electricity remains elusive. "There has been a huge amount of theoretical and experimental work on high-temperature superconducting systems, but no complete understanding," says ORNL's Thomas Maier. "One especially would like to understand what causes the pairing interaction and why these systems become superconducting."

The question is a very important one, particularly in a world confronting rising temperatures and looming energy shortages: Understanding why HTSCs behave as they do would open the door to the development of new materials operating at higher temperatures, perhaps even at room temperature. These new materials would pave the way for breakthrough applications that could save enormous amounts of energy. Researchers envision power cables that transmit electricity with little or no losses, mass production of consumer electric vehicles, super-efficient high-speed trains and leaps in energy efficiency for a variety of electrical machinery.

No resistance

Maier is part of an Oak Ridge team working to develop a theoretical description of HTSCs. One mystery they are trying to solve is why the electrons in some materials bond to form pairs called "Cooper pairs" that settle into a state in which they conduct electricity without resistance. "The ultimate reason a material becomes superconducting is that the electrons join into Cooper pairs," Maier explained. "If a current is driven through the system, the superfluid phase formed by the Cooper pairs does not 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 identifying the force. 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, spin fluctuation mediates the pairing that leads to superconductivity," says Maier.

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

…calculations conducted at Oak Ridge are the first at a scale adequate to analyze a microscopic model for the source of the pairing interaction…

Doug Scalapino of the University of California–Santa Barbara was among 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 Oak Ridge project, said the new findings help relate the Hubbard model to lab experiments on cuprates, particularly neutron scattering and angle-resolved photoemission scattering. The Oak Ridge supercomputers revealed ways to address the pairing mechanism by using data obtained from actual materials by these techniques.

Researchers do not know how to measure the pairing mechanism directly in materials, Scalapino explains, but the Hubbard model used to reveal the pairing mechanism can be used to obtain neutron and angle-resolved photoemission scattering 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." Researchers then can use that formula to do the same thing for the actual material—use neutron scattering and angle-resolved photoemission data obtained directly from materials to calculate the pairing interaction in actual 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, says Maier. In 2007, his team added inhomogeneities to their model to see how they affected the pairing interaction and the transition temperature.

The calculations for the inhomogeneities probably were conducted on ORNL's Cray XT4 Jaguar supercomputer because, whereas the previous simulations required a smaller number of faster processors, the new problem required a large number of processors.

"The interest ultimately is in how to increase the transition temperature to room temperature," Maier says. "If we can understand why a material is superconducting at 150 K, then we can ask what we must do to raise that temperature. There is no guarantee once we understand the process that we can raise the transition temperature, but it's a big step."

Had Maier's colleagues at the 1987 "Woodstock" conference had access to modern supercomputers, one can imagine that the path to understanding the mysteries of high-temperature superconductors would have been an easier one.

Going Underground

Jaguar makes possible the largest groundwater simulations

As the American scientific community develops a long-term plan to address climate change, not every response focuses on new energy sources. Even as we develop promising new carbon-free technologies for solar power, biofuels and nuclear energy, the economy faces the prospect of being tethered for some time to the old energy sources—primarily fossil fuels such as coal and oil.

Climate change represents a collision of economic and environmental concerns. Coal is very abundant and critical to much of America's economic base. With coal power, however, comes a variety of serious environmental problems, including the discharge of carbon dioxide, or CO_2, into the air by coal-fired power plants. Carbon dioxide is the most worrisome of the greenhouse gases. According to the Intergovernmental Panel on Climate Change, levels of CO_2 in the atmosphere are 35% higher than they were before the Industrial Revolution and are in fact higher than at any time in the last 650,000 years. Climate scientists believe it is no coincidence that the planet is experiencing a string of the warmest years since the taking of measurements began more than 150 years ago.

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

A team of researchers led by Peter Lichtner of Los Alamos National Laboratory is using Oak Ridge National Laboratory's Jaguar supercomputer to simulate this process, known as carbon sequestration, and is searching for ways to maximize its benefits and avoid its potential drawbacks. Indeed, many view carbon sequestration as an absolutely critical component of worldwide efforts to lower substantially the volume of carbon emissions. Using Jaguar, the team has been able to conduct the largest groundwater simulations to date, pursuing the research with an application known as PFLOTRAN.

The process being simulated by Lichtner's team involves taking CO_2 that has been separated from a power plant's emissions and injecting it nearby into a deep saline aquifer 1 to 2 kilometers below the surface. If all goes according to plan, the CO_2 would disperse under a layer of impermeable rock with the opportunity to dissolve into the surrounding brine.

When pumped into the ground, the CO_2 would be in a state known as a supercritical phase, which is present when the gas is kept above 50°C (120°F) and more than one hundred times atmospheric pressure. The plan assumes that the CO_2 would be kept in the supercritical phase by the heat and pressure naturally present deep underground. According to Lichtner, CO_2 in this phase is in some ways like a liquid and in some ways like a gas. The primary benefit is the avoidance of the rapid expansion that would accompany changes between the two phases.

Lichtner's team is investigating a process known as "fingering," that speeds the rate at which the CO_2 dissolves. Fingering grows out of the fact that while CO_2 in the supercritical phase is lighter than the surrounding brine, the brine in which CO_2 has been dissolved is actually heavier than unsaturated brine. The result is a convection current, with "fingers" of the heavier, saturated brine sinking. This fingering in turn increases the surface area between the CO_2 and the brine and speeds the dissolution of the supercritical CO^2 into the brine.

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

"The problem is that we are injecting huge amounts of CO_2 by volume," Lichtner explained. "If we were injecting it into a deep saline aquifer, for example, we would initially have to displace the brine that was present. The question then would be, 'Where does that go?' If we inadvertently pushed the brine up into the overlying aquifers we might contaminate, say, the drinking water for the whole Chicago metropolitan area. The dissipation of CO_2 is literally a race against time."

Other hazards must be better understood before large volumes of CO_2 can be pumped underground. If the CO_2 were to rise to the surface, another substantial hazard would be created. The process of dissolving CO_2 into groundwater is, in fact, known as

Carbon dioxide dissolving over time in a deep saline aquifer.

carbonation. The unintended rise of CO_2 to the surface could rapidly turn the groundwater into seltzer water.

"There are natural occurrences of CO_2 shooting out of the ground," Lichtner noted. "As the pressure and temperature are artificially lowered, the bubble of injected CO_2 starts approaching the surface, with a change of phase from supercritical to liquid to gas. The result could suddenly occupy a much larger volume, forming a geyser like those in Yellowstone National Park.

"So long as the supercritical phase exists, the possibility that the CO_2 could escape through fractures, abandoned boreholes, or boreholes that leak presents a hazard to people living in the vicinity. Therefore, understanding the rate of dissipation is important to knowing how rapidly we can move from the supercritical phase."

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

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

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

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

The team looks forward to using ORNL's new Jaguar supercomputer, capable of speeds greater than 1,000 trillion calculations a second, or a petaflop. Jaguar will make possible the simulations needed to address questions on the scale of the Illinois Basin posed by carbon sequestration.

Meanwhile, researchers realize that the scale and complexity of the climate change challenge, equaled only by the scale of the stakes involved, will require scientific tools such as Jaguar that are equal to the task.—*Leo Williams*

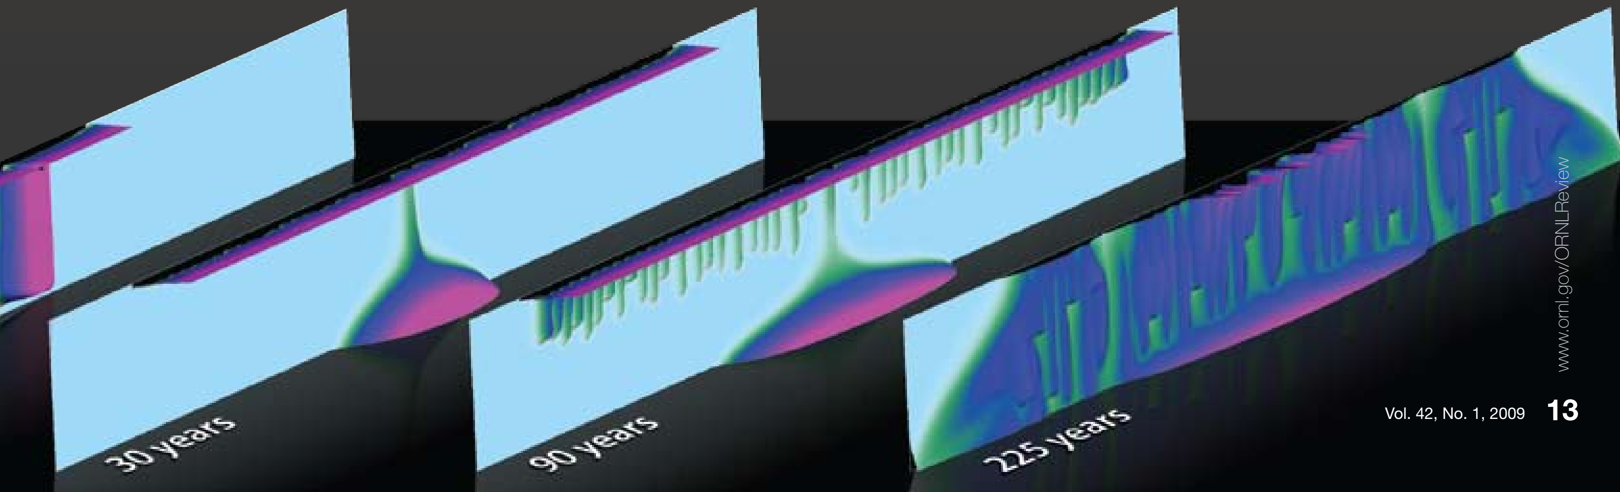

Hotter Than The Sun

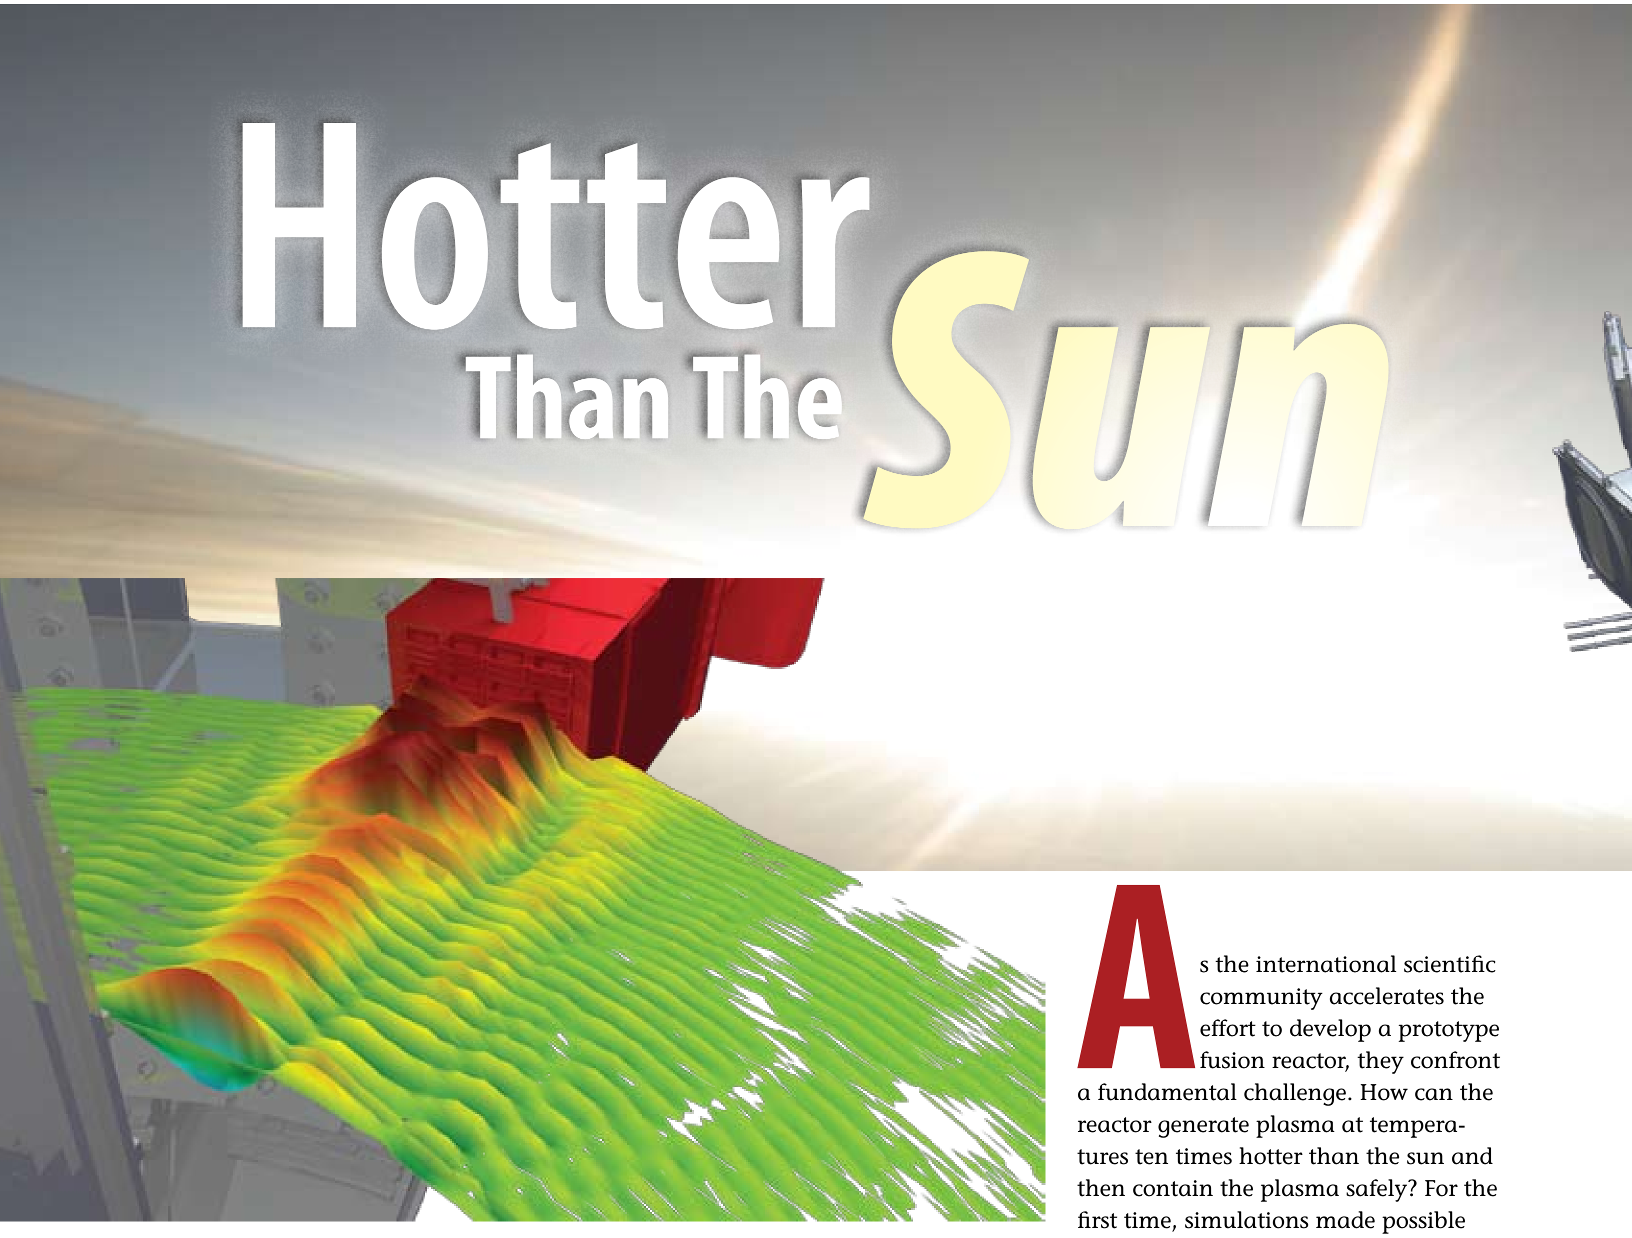

Progress continues on the quest for a lasting source of clean energy

As the international scientific community accelerates the effort to develop a prototype fusion reactor, they confront a fundamental challenge. How can the reactor generate plasma at temperatures ten times hotter than the sun and then contain the plasma safely? For the first time, simulations made possible with extraordinary computing power may provide a portion of the answer.

In 2007 a fusion research team led by Fred Jaeger and Lee Berry of Oak Ridge National Laboratory achieved a performance of more than 87 trillion calculations per second, or teraflops, on the Cray XT4 Jaguar supercomputer at the National Center for Computational Sciences. The simulation provided insight into how best to heat an experimental reactor scheduled to begin operating in 2016 in Cadarache,

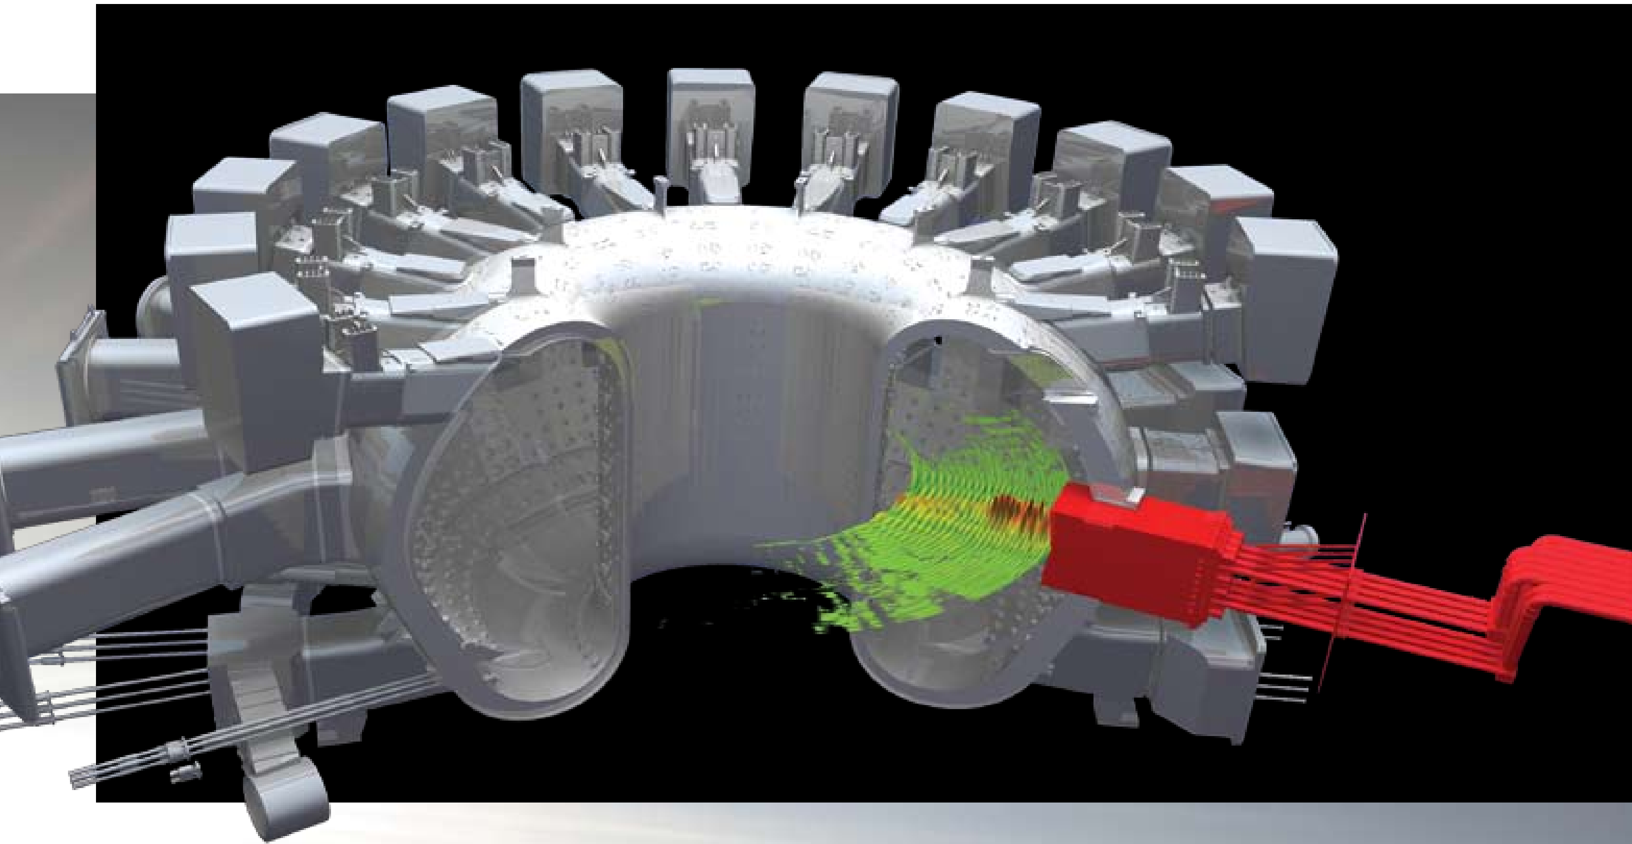

The radio frequency antenna (red) launches three-dimensional wave fields into the ITER plasma. The waves heat deuterium and tritium fuel to fusion temperatures about ten times hotter than the surface of the sun. Image credit: Sean Ahern/U.S. ITER Project Office.

France. The ambitious project, named ITER, is a coalition comprising the United States, the European Union, Russia, India, South Korea, China and Japan formed to provide the collective funding and scientific expertise needed to develop commercial fusion power plants.

ITER will use antennas to launch radio waves carrying 20 megawatts of power into the reactor, the equivalent of a million compact fluorescent light bulbs. The waves will heat the deuterium and tritium fuel to fusion temperatures—or more than 400 million degrees Fahrenheit. The deuterium and tritium form plasma, a state of matter created when gases become so hot that electrons get energized and fly off their atoms. As conceived, the radio waves would drive currents that help confine the plasma. Jaeger's simulations will contribute to understanding how to make the most of the wave power in both heating and controlling the plasma.

"The 2007 run was the first two-dimensional simulation of mode conversion in ITER," said Jaeger, who used the simulation to explore the conversion of fast electromagnetic waves to slow electrostatic waves. Before the run, mode conversion in ITER was simulated in only one dimension, although scientists could simulate mode conversion in two dimensions for smaller tokamaks. "We need to know which types of waves are present because different waves can interact differently with the plasma."

Jaeger's team uses AORSA, a software code that solves Maxwell's equations for the electromagnetic wave fields in the plasma. The AORSA team is part of a Scientific Discovery through Advanced Computing project known as the SciDAC Center for Simulation of Wave-Plasma Interactions. The team includes plasma scientists, computer scientists, and applied mathematicians from ORNL, the Massachusetts Institute of Technology, Princeton Plasma Physics Laboratory, General Atomics, CompX, Inc., Tech-X Corporation and Lodestar Research Corporation.

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

Upon analyzing the energy distributions of the very-high-energy ions created when radio waves heat the plasma, the scientists found that, in some cases, the ions increased the fusion reaction rate. They learned the optimal frequency for driving current in the ITER plasma and identified the heat-loss channels that limit the current. Comparing the ITER model with current tokamaks, they also found stronger central focusing of radio waves in ITER.

These findings at ORNL, made possible by the new capabilities of supercomputing, bode well for the goal of keeping plasma hot enough for fusion, and for the world's quest for a lasting source of clean energy.
—Dawn Levy and Leo Williams

THE PERFORMANCE POLICE

Software engineers monitor problems that threaten climate simulations

The climate debate has shifted. The large majority of scientists accept data that show the world is getting warmer as the result of human activity. The discussion now focuses on a variety of questions about how climate change will proceed and how the process can be slowed and mitigated. Should forests and food croplands be converted to produce plants for biofuels? What technologies best capture and store carbon? How intense will hurricanes and heat waves be? Will the release of methane trapped in permafrost accelerate climate change?

Answers to these and other questions depend increasingly on sophisticated climate simulations, a digitized world that mirrors our past and probes the future. Unlike the conventional laboratories of

Recent summers have seen, for the first time, open water in the Northwest Passage.

test tubes and microscopes, this new world would not function without software applications such as the Community Climate System Model (CCSM), a megamodel coupling four independent models whose codes describe Earth's atmosphere, oceans, lands, and sea ice.

Such simulation tools are now commonplace on ORNL's supercomputers. Indeed, much of the U.S. climate community conducts simulations in Oak Ridge, including the Department of Energy, the National Center for Atmospheric Research, National Oceanic and Atmospheric Administration and the National Aeronautics and Space Administration.

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

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

One validation of the climate simulations lies in the volume of publications. In the months after simulation data were posted, scientists produced approximately 300 journal articles. The IPCC cited the studies in the Fourth Assessment Report, which concluded planetary warming during the twentieth century was likely the result of human activities. In 2007 the IPCC shared the Nobel Peace Prize with Al Gore.

Weather versus climate

Drake frames the question that drives the project. "How can we simulate climate 100 years from now if we do not know what the weather will be 100 days from now? Climate is statistical, or average, weather," Drake explains. "Climate data can tell us if hurricanes will be more likely, less likely, or stronger, but the data cannot tell us when they will occur."

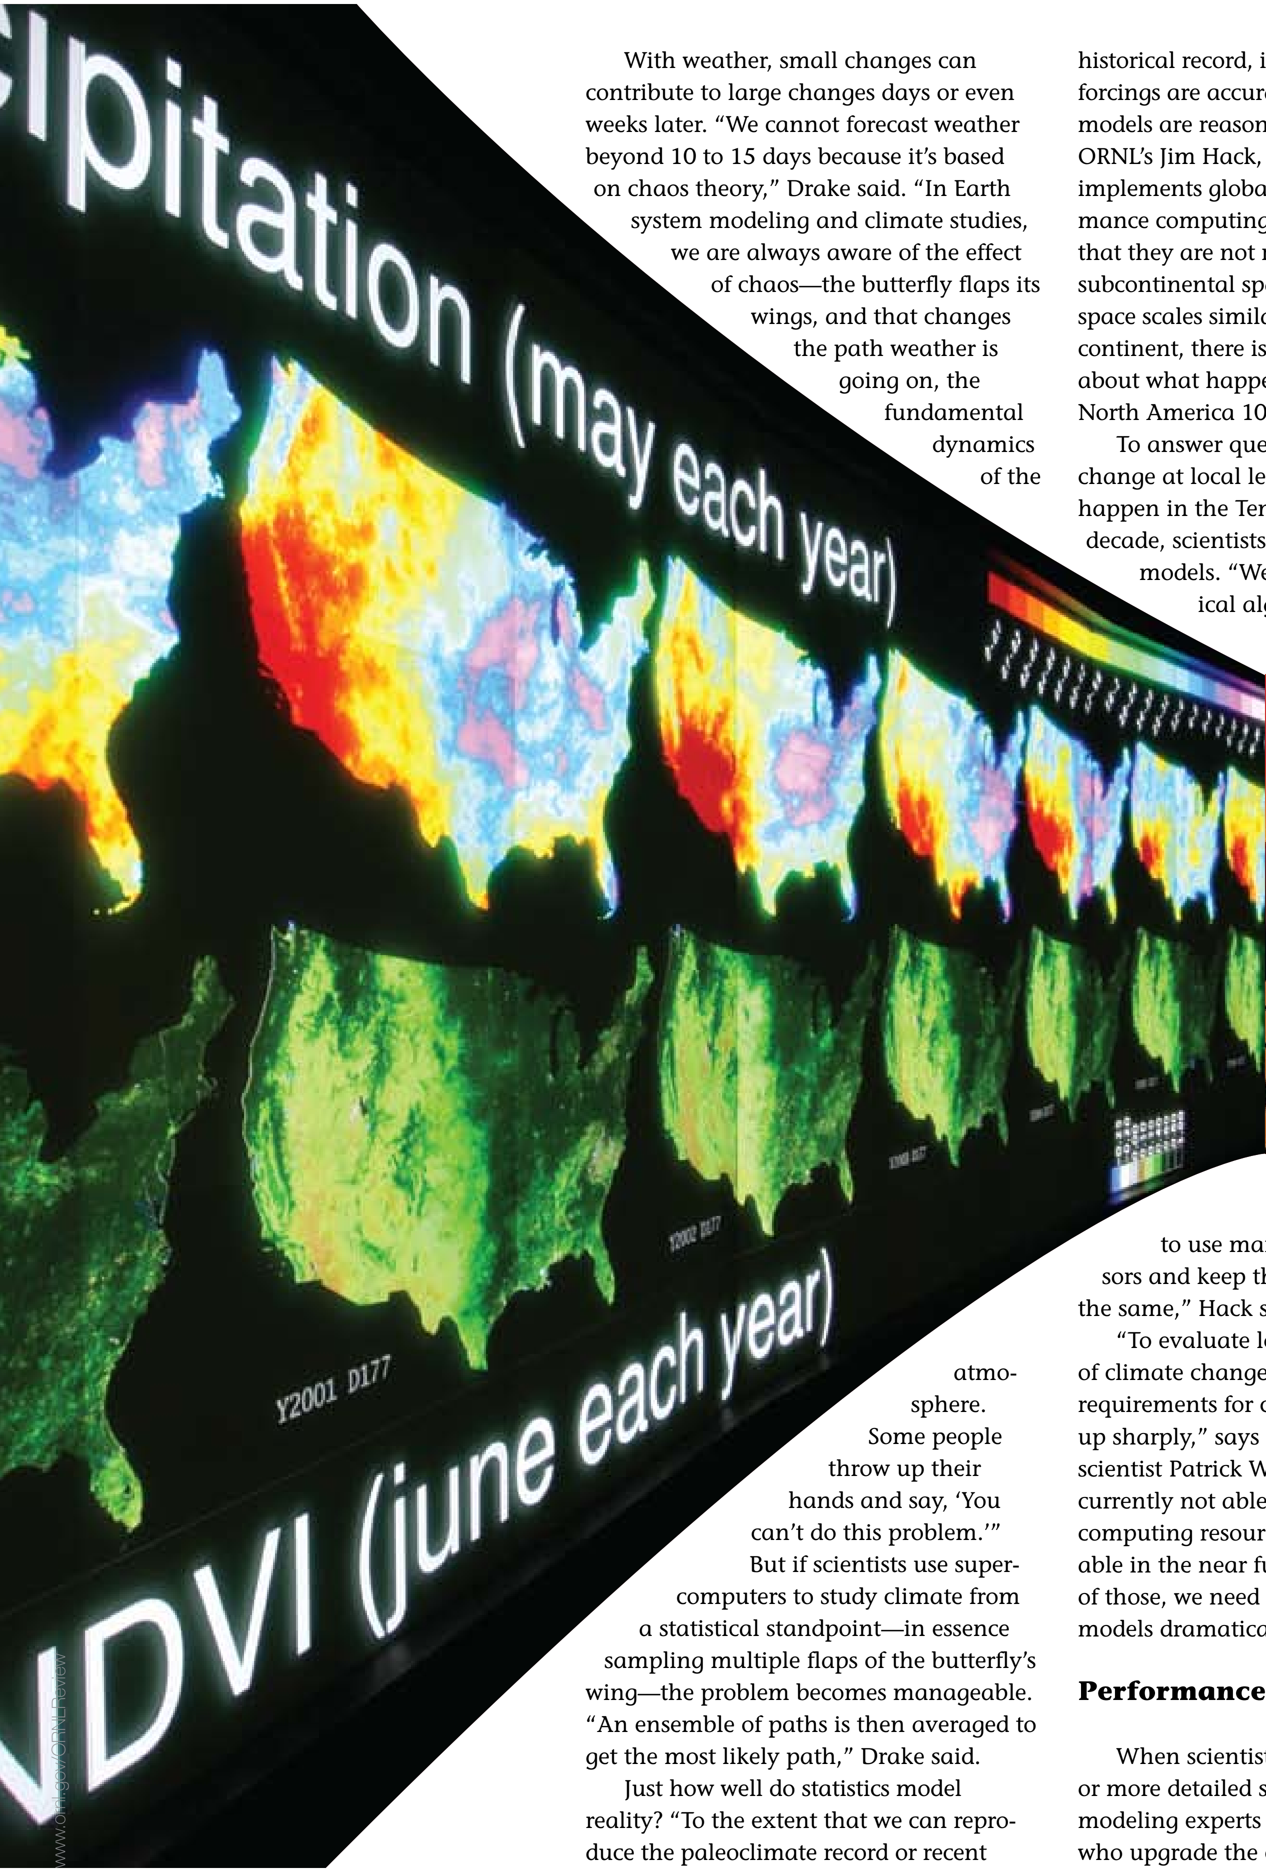

With weather, small changes can contribute to large changes days or even weeks later. "We cannot forecast weather beyond 10 to 15 days because it's based on chaos theory," Drake said. "In Earth system modeling and climate studies, we are always aware of the effect of chaos—the butterfly flaps its wings, and that changes the path weather is going on, the fundamental dynamics of the atmosphere. Some people throw up their hands and say, 'You can't do this problem.'" But if scientists use supercomputers to study climate from a statistical standpoint—in essence sampling multiple flaps of the butterfly's wing—the problem becomes manageable. "An ensemble of paths is then averaged to get the most likely path," Drake said.

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

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

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

Performance police

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

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

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

On Cray and IBM systems, the group has improved performance through both algorithmic and implementation efforts. Recent work improved performance 2.5-fold on benchmark problems on ORNL's Cray XT4 Jaguar. "With the improvements to the scalability of the CCSM software by Pat and his colleagues, along with the dramatic growth in the performance of Jaguar, the CCSM developers are seriously considering model resolutions and advanced physical processes that were not on the table before," said Trey White, who as ORNL's liaison to the CCSM project helps the scientists maximize the machines' capabilities.

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

Worley's team is currently working with a large multilab project to extend the physical climate model by including chemical and ecological processes. The computer allocations are provided through the Climate Science Computational End Station, an Innovative and Novel Computational Impact on Theory and Experiment program award led by NCAR's Warren Washington on Jaguar at Oak Ridge.

"For the Department of Energy, which is very concerned with the carbon cycle and with the impact of climate change on ecology and ecosystem services, this kind of Earth system model is invaluable," Drake said. "We are doing everything we can to get there as quickly as possible."

—Dawn Levy

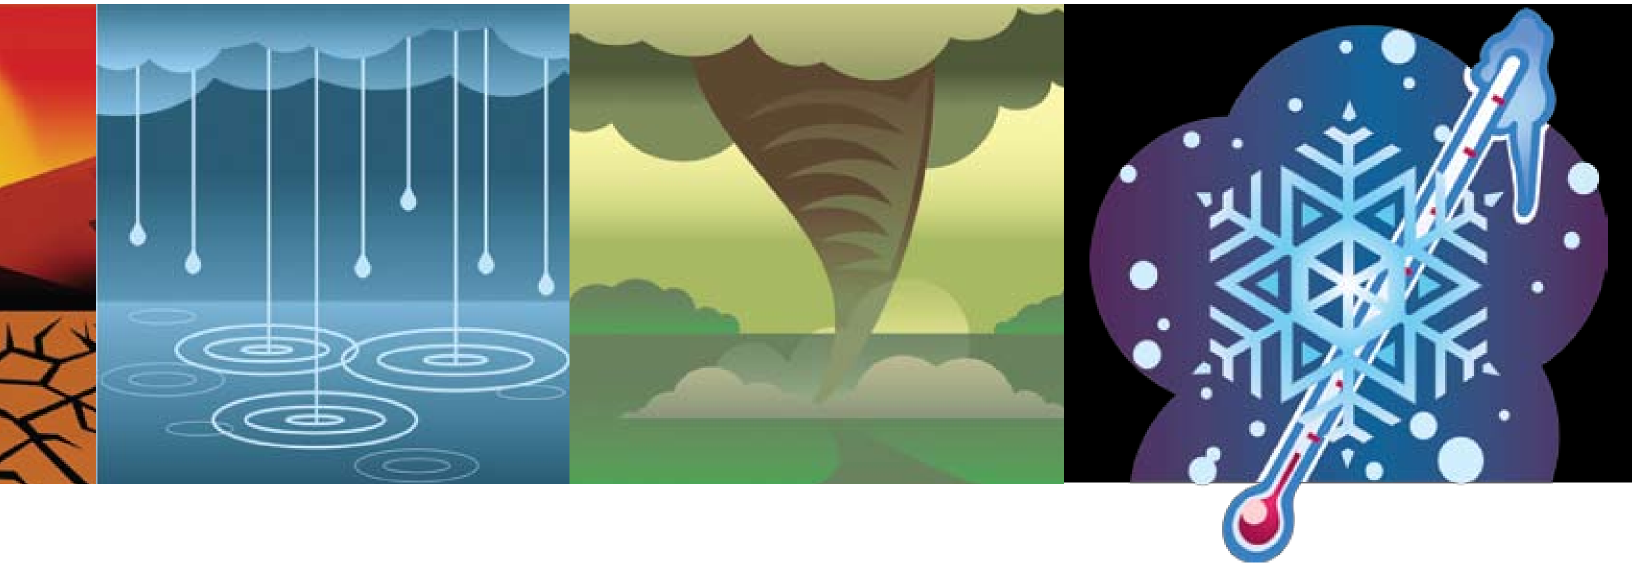

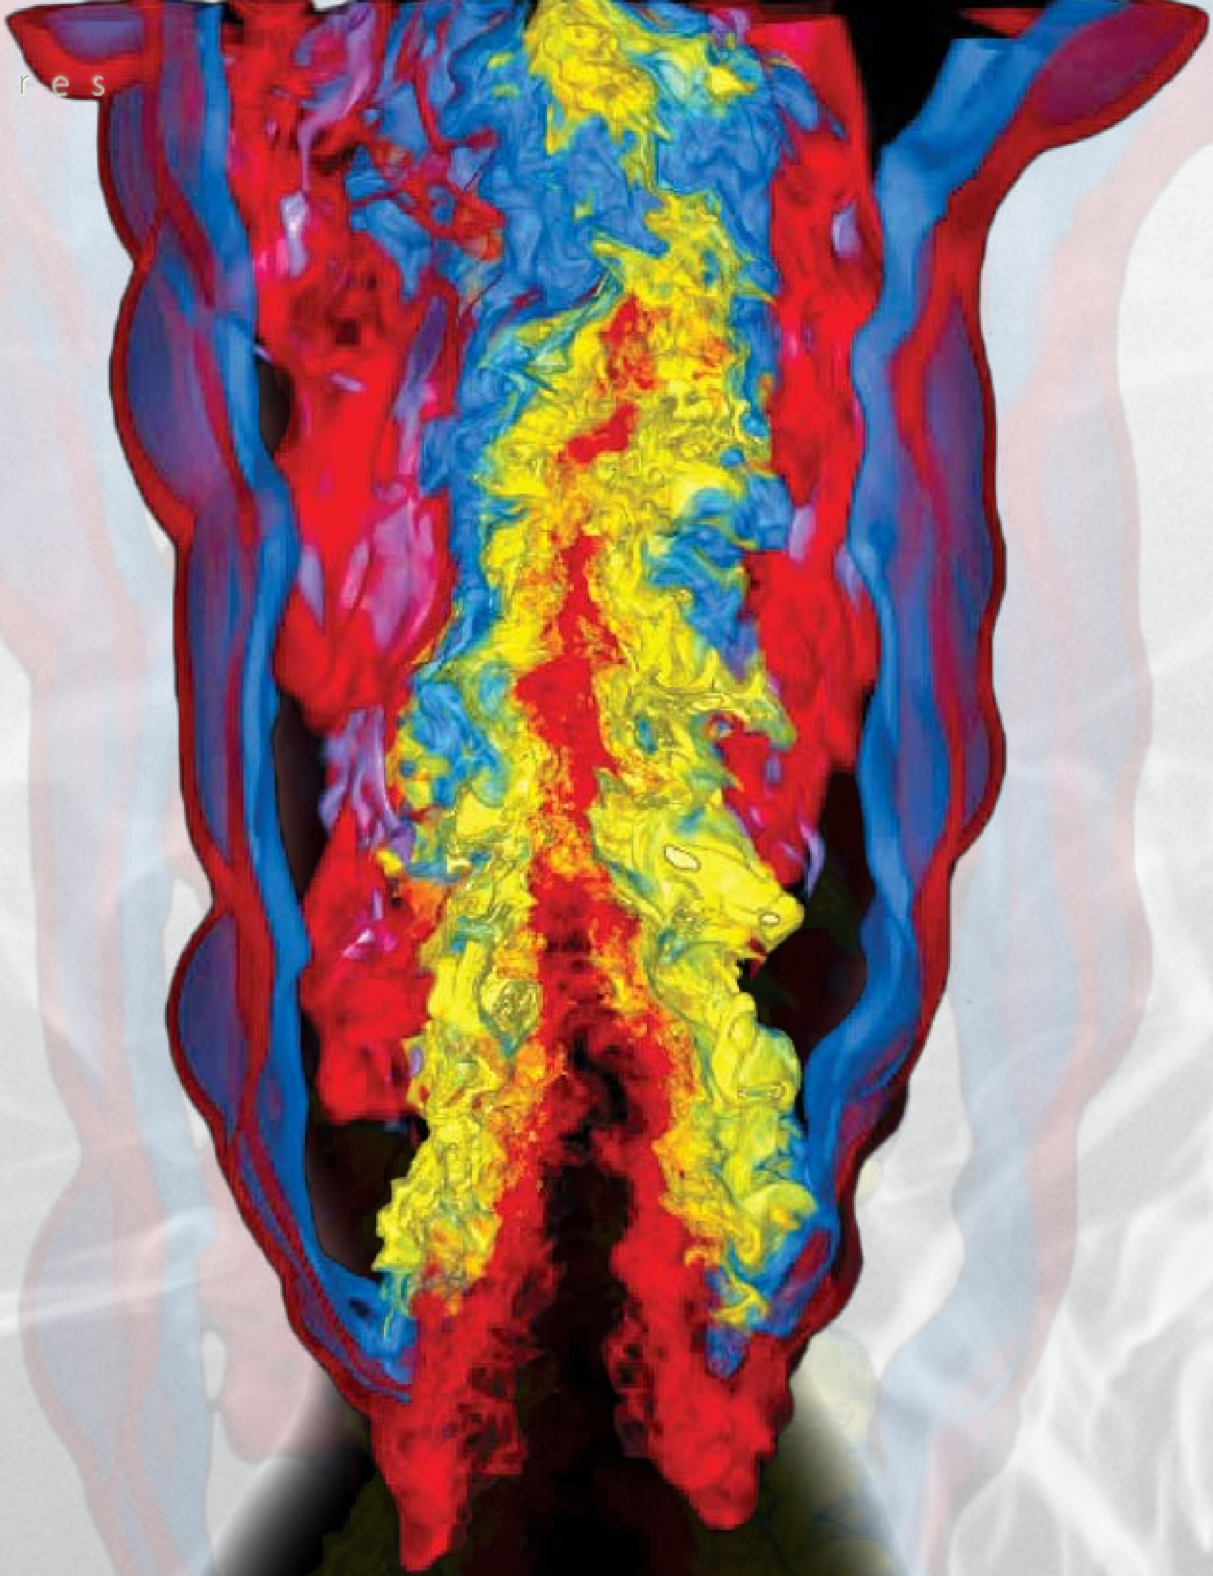

DOUBLING FUEL EFFICIENCY

Low-temperature combustion technology may save energy and cut emissions

The model shows feature detail on all size scales—the biggest, the smallest, and everything in between—of a turbulent fuel jet igniting in a hot co-flowing airstream

To flip an old saying, where there's fire, there's smoke. Perhaps no one knows that better than Jacqueline Chen, a mechanical engineer at Sandia National Laboratories who employs one of the world's fastest supercomputers to model combustion.

With Sandia colleague mechanical engineer Chun Sang Yoo and computational scientist Ramanan Sankaran of Oak Ridge National Laboratory, Chen used the Cray XT4 Jaguar supercomputer at Oak Ridge to generate 120 terabytes (trillion bytes) of data about flames similar to those occurring during ignition and stabilization of diesel-engine jets. The data equaled more than five times the printed contents of the U.S. Library of Congress. In their simulation, the researchers burned one of the simplest hydrocarbon fuels—ethylene molecules. Whereas diesel fuels are more complex, ethylene is a commonly used fuel in laboratory experiments and a fuel fragment resulting from high-temperature initiation reactions for autoignition of *n*-heptane, a diesel surrogate.

Engineers are using Chen's data library to develop predictive models that will be used in the future to optimize designs for diesel engines and industrial boilers with reduced emissions and increased efficiency. Since diesel fuel powers most semi trucks, delivery vehicles, buses, trains, boats and farm, construction and military vehicles and equipment in the United States, development of advanced diesel technology is a leading near-term option by which the country could reduce fuel consumption and greenhouse gas emissions.

"If low-temperature compression ignition systems employing dilute fuel mixtures make their way into next-generation autos, fuel efficiency could increase by as much as 25 to 50 percent," Chen says. The new technology would also make it possible to meet future low-emission vehicle standards with almost undetectable emissions of nitrogen oxide, a major contributor to smog, she adds.

Chen, Yoo and Sankaran created the first three-dimensional simulation that fully resolves flame features, such as chemical composition, temperature profile and flow characteristics. The model shows feature detail on all size scales—the biggest, the smallest and everything in between—of a turbulent fuel jet igniting in a hot coflowing airstream.

The researchers modeled in unprecedented detail what happens in the so-called "lifted flames" relevant to industrial boilers and diesel engines. Unlike spark-plug ignition systems in automobiles powered by gasoline in which the fuel and oxidizer (air) are premixed, diesel-injection systems have the diesel fuel entering the engine full of hot air via jet nozzles. Turbulence mixes the fuel and air. Pistons subject the air/fuel mixture to pressure, and the mixture heats further, spurring a chemical pathway that sharply increases the concentration of a highly reactive chemical, hydroperoxyl radical, Chen says. The hydroperoxyl radical produces heat that spurs the production of other radicals that eventually leads to thermal runaway. At about 2,330°F, the fuel/air mixture autoignites, or bursts into flame, as a result of the rapid, heat-producing oxidation of its own constituents, regardless of heat from external sources. This process creates a lifted flame. The temperature peaks in excess of 3,140°F.

"Autoignition is helpful because it stabilizes the flame," Chen says. "A hot, vitiated coflow supports its existence."

Before this work by Chen and her colleagues, scientists had modeled only large eddies, or turbulent curlicues, in a burning fuel. They had not simulated the full range of scales down to the smallest eddies, which dictate the viscosity of the system and dissipate heat. Equipped with greater computational power, researchers can now resolve the nitty-gritty of the small eddies responsible for flame extinction and reignition in canonical flows with a moderate Reynolds number, which indicates the range of scales in a system.

The combustion research at ORNL is supported by the Department of Energy through the agency's Innovative and Novel Computational Impact on Theory and Experiment program. The program seeks to address scientific "grand challenges" by granting large allocations of supercomputing time to approximately 25 peer-reviewed projects each year, including proprietary projects in partnership with private industry. The program is unique in the ability to provide scientists with the computational tools required to address problems too large and complex for most research institutions.—*Dawn Levy*

Image: Volume rendering of a lifted autoigniting hydrogen/air jet flame with hydroperoxy radical (ignition marker, red and yellow) and hydroxyl radical (flame marker, blue).

INVISIBLE MEANS OF SUPPORT

Astrophysicists simulate the dark matter that cradles a galaxy

In the early 1930s, the eminent Swiss astronomer Fritz Zwicky noticed something very odd as he was looking to the skies: Galaxies seemed to move around each other too fast.

Zwicky was scrutinizing a group of eight galaxies orbiting one another more than 350 million light years away in the Coma Galaxy Cluster. Drawing from early work by Issac Newton and Albert Einstein, he understood the balance of forces necessary to keep the galaxies in this dance. Like a yo-yo swung by a child, they need both the centrifugal force pushing them outward and the string—in this case gravity—pulling them back in. Too much force inward and the system collapses; too much outward and the galaxies fly apart.

From a yo-yo to a galaxy, every object in the universe with mass exerts a gravitational pull on other objects. To Zwicky, the galaxy cluster he was observing appeared to have too little mass, and therefore too little gravity, to keep the galaxies from flying off into space. He and his colleagues theorized that these and all galaxies must be dominated by matter invisible to the eye. They called it dark matter.

Equipped with an understanding of how gravity works and extensive observations of planets, stars and galaxies, scientists have in fact concluded that less than one-fifth of the matter in the universe is visible. The remainder, dark matter, has no interaction with regular matter except through the force of gravity. Nevertheless, so much dark matter exists in the universe that its gravitational force controls the lives of stars and galaxies.

What, then, would dark matter look like if we could see it? A team led by astrophysicist Piero Madau of the University of California–Santa Cruz has taken a substantial step toward answering this question. Using the power of Oak Ridge National Laboratory's Jaguar supercomputer, Madau's team has run the largest simulation ever of dark matter evolving over billions of years to envelop a galaxy such as our own Milky Way. The envelope is known as a dark matter halo.

Madau and his collaborators—including Juerg Diemand and Marcel Zemp, both of UCSC, and Michael Kuhlen of the Institute for Advanced Study in Princeton, New Jersey—reviewed the simulation and their findings in the journal *Nature*. The simulation followed a galaxy worth of dark matter through nearly the entire history of the universe, dividing the dark matter into more than a billion separate parcels. The effort was staggering and involved tracking over 13 billion years the evolution of 9,000 trillion trillion trillion tons of invisible materials spread across 176 trillion trillion trillion square miles. Each parcel of dark matter was 4,000 times as massive as the sun.

Hypothetical particles with real gravity

Scientists are still trying to determine exactly what dark matter is. Candidates include hypothetical particles such as the neutralino, the sterile neutrino, the axion or some other weakly interacting massive particle. Fortunately, researchers do not need to fully understand dark matter in order to simulate it. All they need to know is that dark matter interacts with other matter only through gravity and is cold, meaning the matter is made up of particles that were moving slowly when galaxies and clusters began to form. Using initial conditions provided by observations of the cosmic microwave background, Madau and his team were able to simulate dark matter through a computer application called PKDGRAV2, developed by a group of numerical astrophysicists at the University of Zurich, who ignored visible matter and focused entirely on the gravitational interaction among a billion dark matter particles. The project had a major allocation of supercomputer time through the Depart-

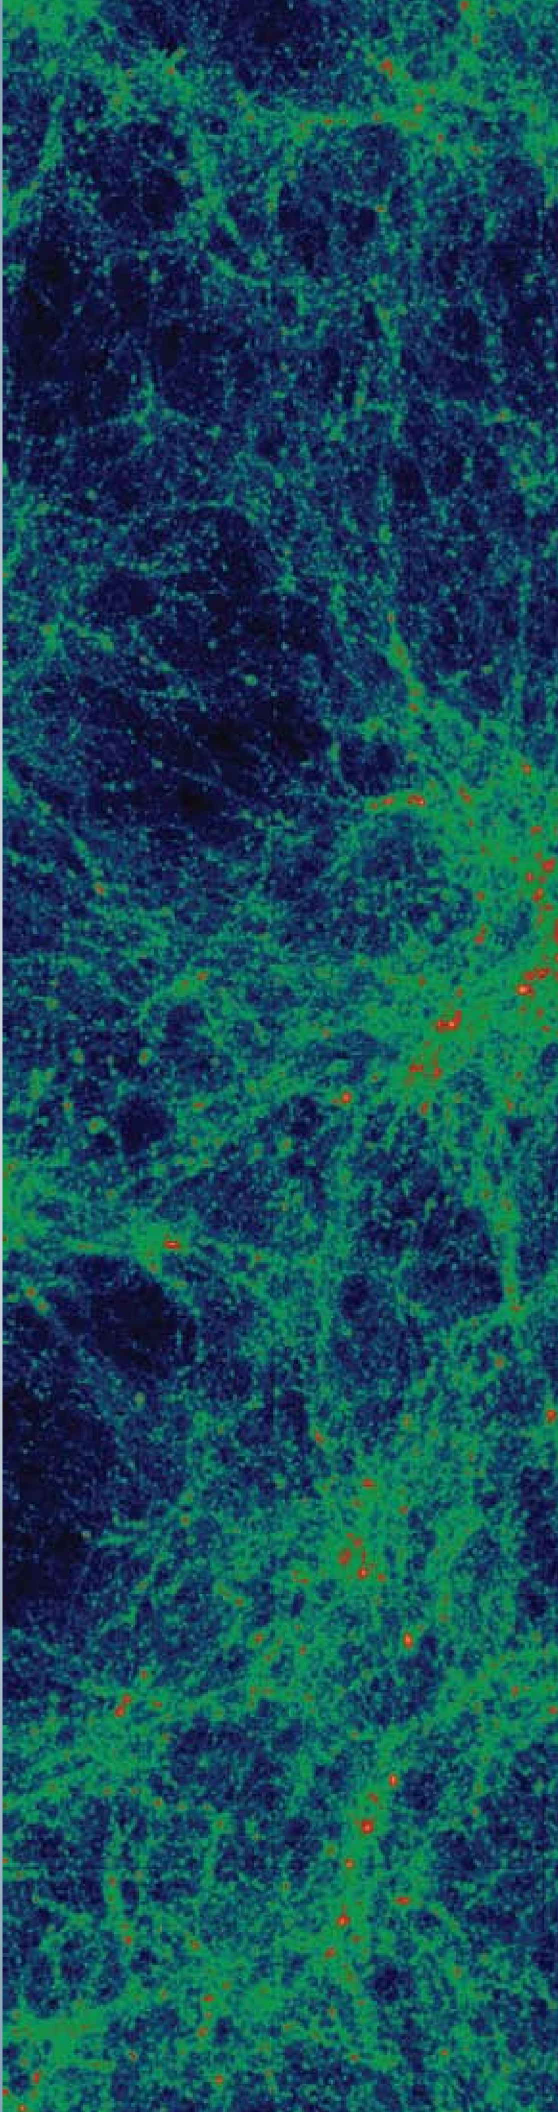

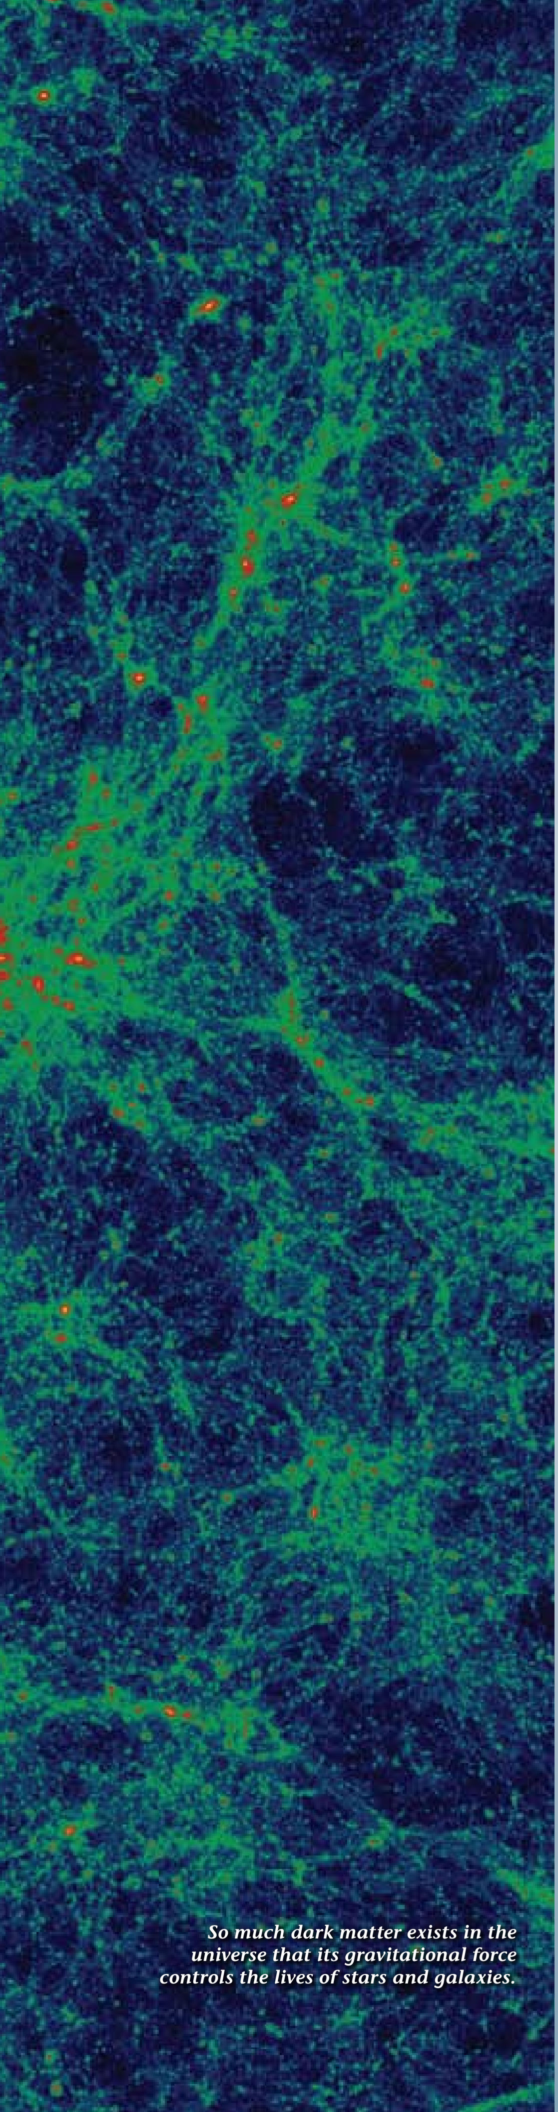

So much dark matter exists in the universe that its gravitational force controls the lives of stars and galaxies.

ment of Energy's Innovative and Novel Computational Impact on Theory and Experiment program. The simulation used about 1 million processor hours on the Jaguar system, located at ORNL's National Center for Computational Sciences.

"The computer was basically just computing gravity," Madau explained. "We have to compute the gravitational force among 1 billion particles, and to do that is very tricky. We are following the orbits of these particles in a gravitational potential that is varying all the time. The code allows us to compute with very high precision the gravitational force due to the particles that are next to us and with increasingly less precision the gravitational force due to the particles that are very far away because the gravity becomes weaker and weaker with distance."

Dark matter is not evenly spread, although researchers speculate it was nearly homogeneously distributed immediately after the Big Bang. Over time, however, gravity pulled the matter together, first into tiny "clumps" having more or less the mass of Earth. Over billions of years these clumps were drawn together, a process that continued until they combined to form halos of dark matter massive enough to host galaxies.

One lingering question was whether the smaller clumps would remain identifiable or would smooth out within the larger galactic halos. The answer required a state-of-the-art supercomputer such as Jaguar, which at the time of the simulations in November 2007 was capable of nearly 120 trillion calculations a second. Because earlier simulations did not have the resolution to resolve any unevenness, the results appeared to show the dark matter smoothing out, especially in the galaxy's dense inner reaches. Madau's billion-cell simulation, however, provided enough resolution to verify that the earliest forms of dark matter do indeed survive and retain their identity, even in the very inner regions, where our solar system is located.

"We expected a hierarchy of structure in cold dark matter," Madau explained. "What we did not know is what sort of structure would survive the assembly because as these subclumps come together they are subject to tidal forces and can be stripped and destroyed. Their existence in the field had been predicted. The issue was whether they would survive as assembled together to bigger and bigger structures."

"What we find," he continued, "is the survival fraction is quite high."

Madau's team will be able to verify its simulation results using the National Aeronautics and Space Administration's Gamma-Ray Large Area Space Telescope. Launched on June 11, 2008, The telescope will scan the heavens to study some of the universe's most extreme and puzzling phenomena: gamma-ray bursts, neutron stars, supernovas and dark matter, just to name a few. While dark matter particles cannot themselves be detected (direct detection of dark matter is being pursued by large underground detectors), researchers believe that dark matter particles and antiparticles may be annihilated when they bump into each other, producing gamma rays that can be observed from space. The clumps of dark matter predicted by Madau's team should bring more particles together and thereby produce an increased level of gamma rays.

A second verification comes from an effect known as gravitational lensing, in which the gravity exerted by a galaxy along the line of sight bends the light traveling from faraway quasars in the background. If the dark matter halos of galaxies are as clumpy as this simulation suggests, the light from a distant quasar should be broken up, like a light shining through frosted glass.

"We already have some data there," Madau noted, "which seems to imply that the inner regions of galaxies are rather clumpy. The flux ratios of multiply imaged quasars are not as you would predict with a smooth intervening lens potential. Instead of a smooth lens, there is substructure that appears to be affecting the lensing process. Our simulation seems to produce the right amount of lumpiness."

Madau's simulations in less than two years have reshaped the discussion about how our universe is held together. As researchers have access to increasingly powerful supercomputers, their findings could enable them to join their predecessors Newton and Einstein in unlocking the door to some of humankind's most fundamental questions.—*Leo Williams*

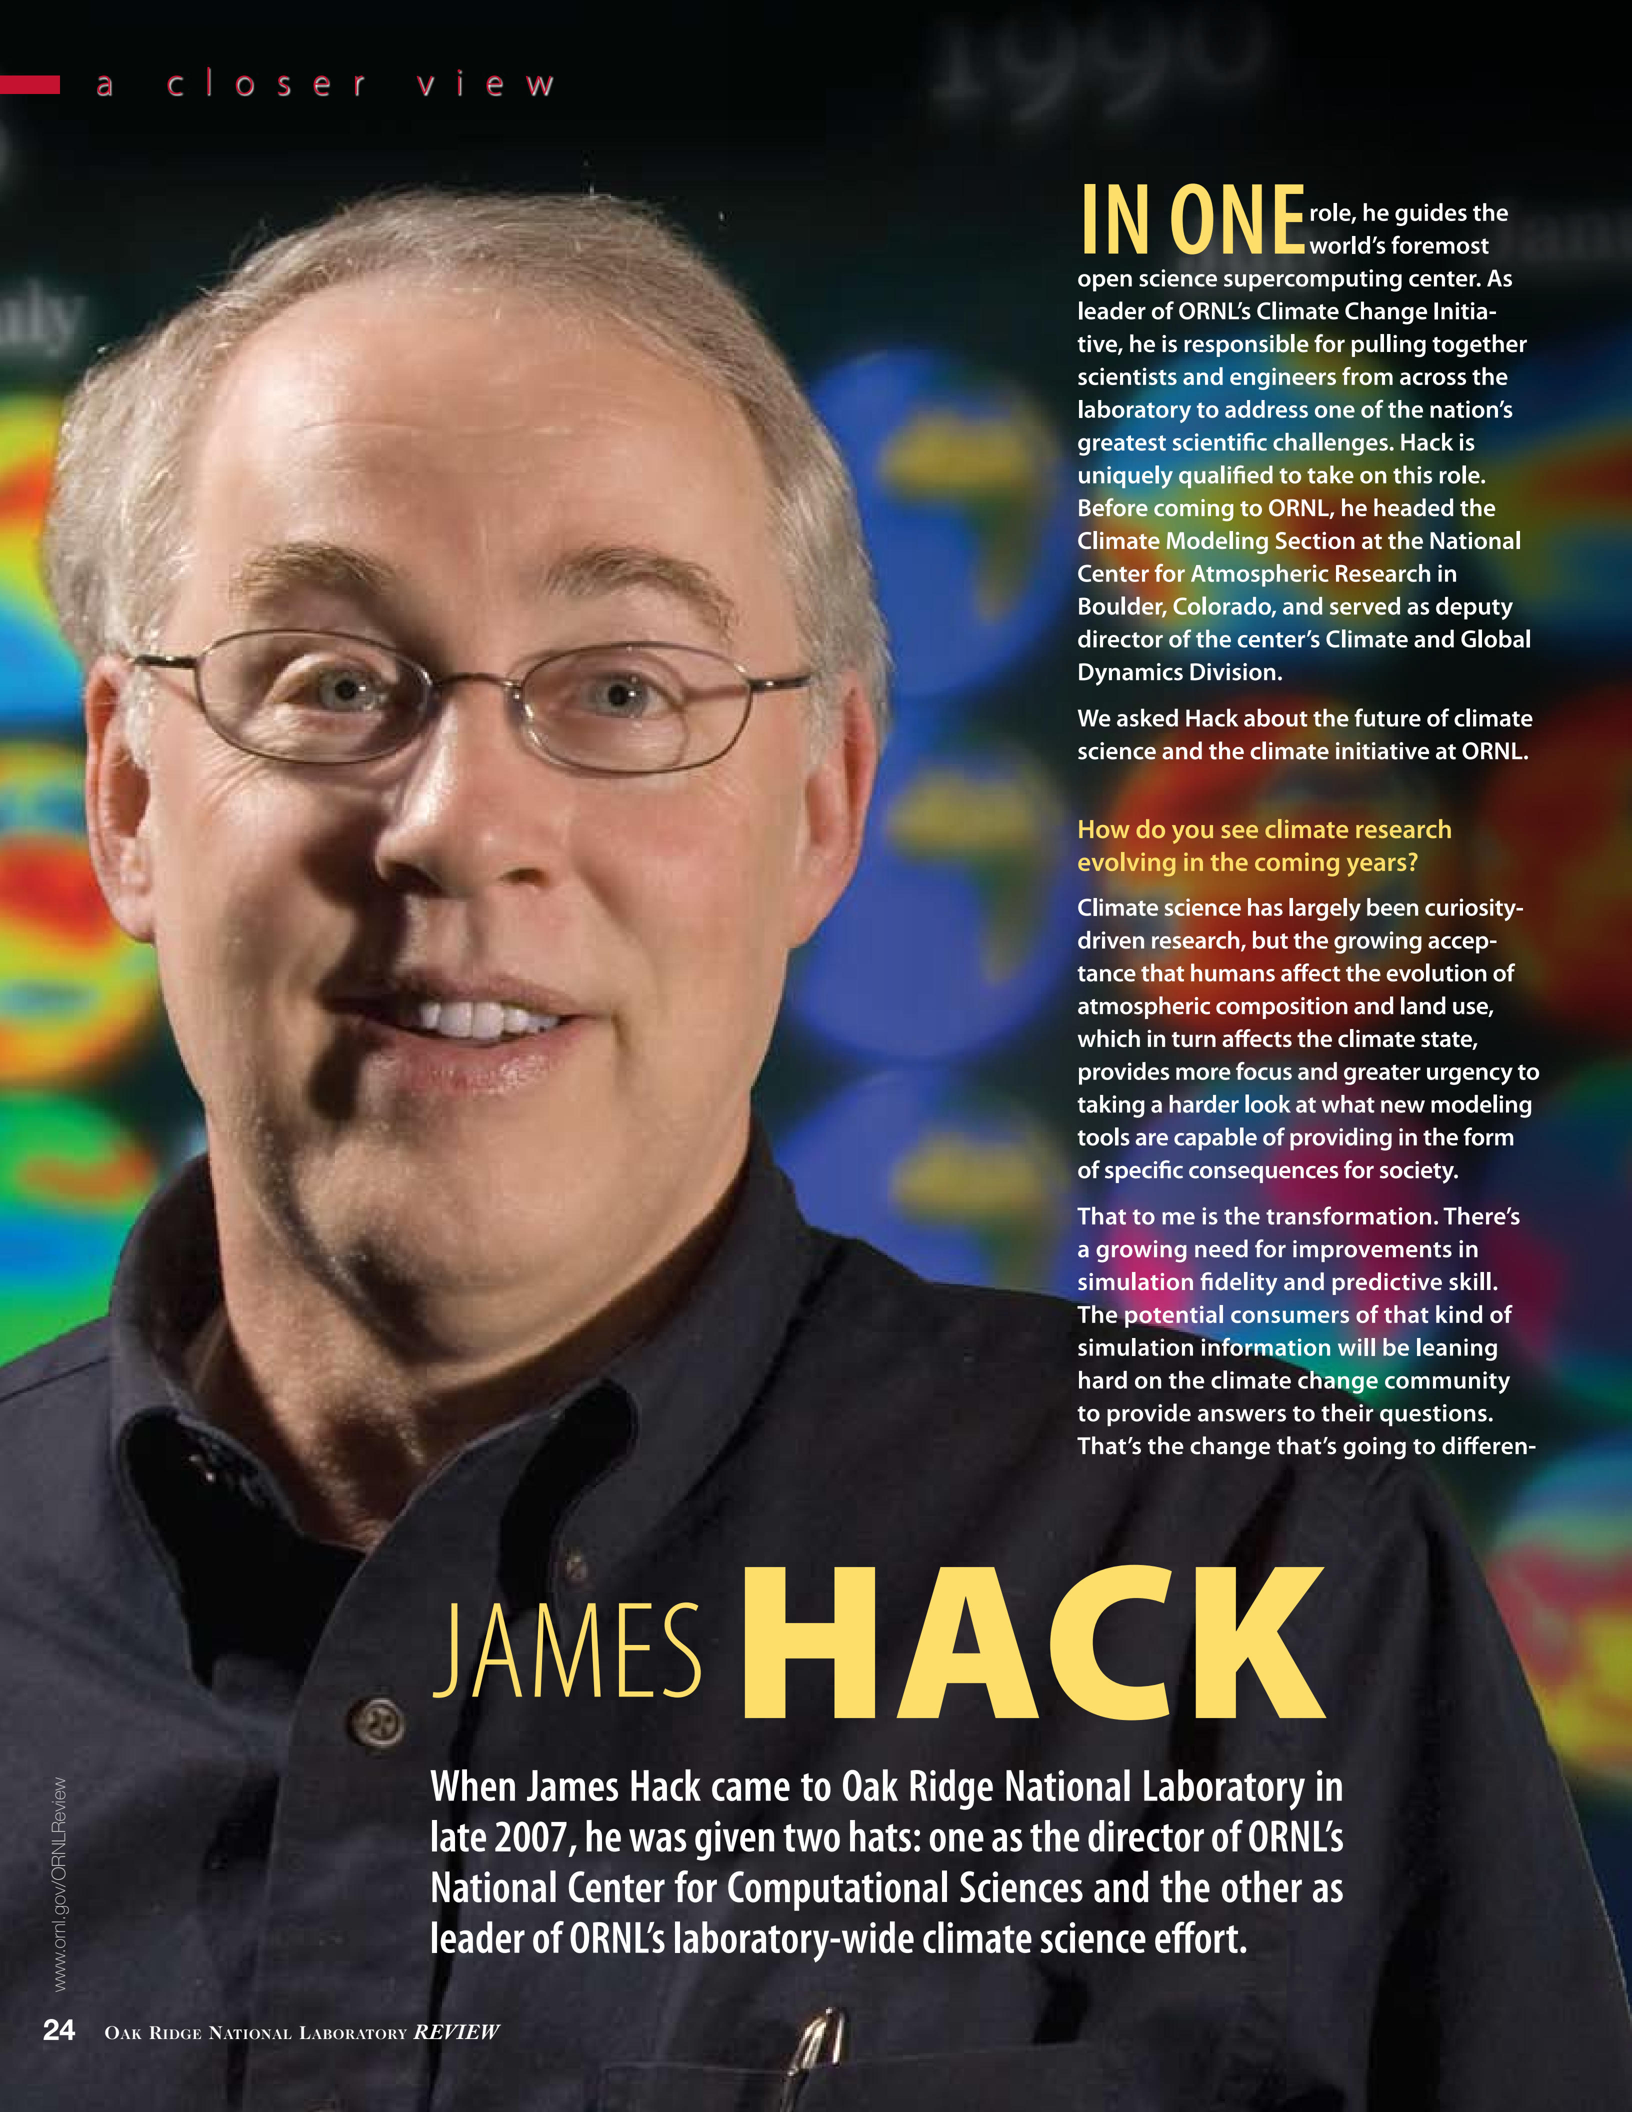

JAMES **HACK**

When James Hack came to Oak Ridge National Laboratory in late 2007, he was given two hats: one as the director of ORNL's National Center for Computational Sciences and the other as leader of ORNL's laboratory-wide climate science effort.

IN ONE role, he guides the world's foremost open science supercomputing center. As leader of ORNL's Climate Change Initiative, he is responsible for pulling together scientists and engineers from across the laboratory to address one of the nation's greatest scientific challenges. Hack is uniquely qualified to take on this role. Before coming to ORNL, he headed the Climate Modeling Section at the National Center for Atmospheric Research in Boulder, Colorado, and served as deputy director of the center's Climate and Global Dynamics Division.

We asked Hack about the future of climate science and the climate initiative at ORNL.

How do you see climate research evolving in the coming years?

Climate science has largely been curiosity-driven research, but the growing acceptance that humans affect the evolution of atmospheric composition and land use, which in turn affects the climate state, provides more focus and greater urgency to taking a harder look at what new modeling tools are capable of providing in the form of specific consequences for society.

That to me is the transformation. There's a growing need for improvements in simulation fidelity and predictive skill. The potential consumers of that kind of simulation information will be leaning hard on the climate change community to provide answers to their questions. That's the change that's going to differentiate the next 10 years of climate change science from the previous 30.

Give us an example of this information.

We know from observations over the last 50 years that the snowpack in the Pacific Northwest has been decreasing. At the same time, temperature in the same region has been increasing. If that trend continues, it raises lots of concerns for water resource managers who have counted on storing their water in the form of snow until a certain time of year when it starts melting.

If precipitation never comes down as snow or if it starts melting sooner than we need it, we may not able to meet water demands. This is a good example of an infrastructure that's vulnerable to specific changes in a region's climate state. Many of the solutions to this problem may also bring with them other environmental consequences.

How accurate is climate prediction?

We think we might currently have sufficient skill to project climate change on regional scales about the size of the Southeast, Pacific Northwest, Rocky Mountain West or Farm Belt. As a scientific community we need to demonstrate the potential and quantify the uncertainties. Although thus far we haven't done a very good job with this challenge, climate researchers are starting to realize that we have an opportunity to take a step back and ask, "What can we do on regional scales and timescales that we think are predictable?"

For example, there's a belief that climate statistics have some predictive skill on decadal timescales. The driver for that is going to reside in the ocean, where motion scales have a very, very long time frame. There is a belief in the scientific community that the ocean's behavior can be predicted several decades into the future.

If we can solve the ocean part of the problem, given the fact that 70 percent of the planet is covered with water, we have a very strong constraint on the other parts of the system. The question then becomes, "Will the other component models follow?"

How do you convince critics that you're getting it right?

We develop numerical experiments to assess whether the global model can produce useful information on the timescales and space scales of most importance to resource managers and planners. They may want to know where the temperature's headed locally, how the hydrological cycle is likely to behave, or how extreme events will change. Do the models provide us with the kind of predictive skill we need, and if not, how can they be improved?

What is the role of computing in this effort?

Computing is a big part of the effort. To fully evaluate the skill in our modeling tools, we need very large computer systems—petascale machines. Assimilating data streams that will be used in the evaluation of modeling frameworks requires very large computer and data systems.

Clearly, a significant computational piece is modeling—building models that have all the components they need to accurately predict the evolution of the earth's climate system. That's computationally very intensive. Incorporating the complexities of the carbon cycle in these models, using the expertise of ORNL's Environmental Sciences Division, contributes to the computational demands. And then mining the data to deal with questions of human impacts and climate extremes is also very computationally intensive.

So, computation in fact ties the whole effort together. It cuts across all the various climate science applications. There are certain areas of science where you need a virtual laboratory to explore the "what-if" experiments, which is what computation provides for the climate problem.

You are leading a new multidisciplinary effort at ORNL focused on climate science. What is the goal?

ORNL has identified climate change as an opportunity that could very effectively exploit existing competencies, particularly high-performance computing and ORNL's long history in contributing to fundamental knowledge about carbon science and in global modeling. The lab also has expertise in evaluating impacts on societal infrastructure. Take rising sea levels. Most of the folks living around the world live close to coastlines, so if the sea level rises even a meter, it has a huge societal impact. The people who are displaced must go somewhere else, maybe moving into areas that were previously used for agriculture. That displaces agricultural activities. ORNL has a very strong geographic information systems group that can contribute to quantification of these scenarios.

We are looking at how we can bring these various competencies together to provide a capability that's unique among the laboratories. The goal is to provide stakeholders, resource managers and others with information they need to deal with the consequences of climate change.

What will ORNL's climate initiative look like?

We're trying to engage people from across the laboratory to stretch the kind of work they're doing in such a way that it requires partnerships with other ORNL staff. So far, many of the more promising proposals include collaborations that cut across the directorates of Biological and Environmental Sciences and Computing and Computational Sciences.

As the initiative matures, I hope we'll begin to incorporate more people in the energy arena, another strong part of the ORNL scientific program. These things could include ways to link climate change and the hard questions we're facing in energy production, like bioenergy and renewable energy technologies, as well as energy consumption. Dealing directly with climate mitigation questions, such as strategies for the sequestration of carbon, is an opportunity for this initiative.

From an energy production point of view, planning has a multidecadal timeframe. Anyone planning investments in the energy infrastructure needs to understand what role the environment might play. That's the goal—to be able to say 20 years from now, "Here's what we anticipate will happen with environmental change on a regional scale."

—Leo Williams *(Originally published in HPCwire.)*

COSMIC CAULDRONS

Researchers challenge conventional wisdom concerning how heavy elements are formed in exploding stars

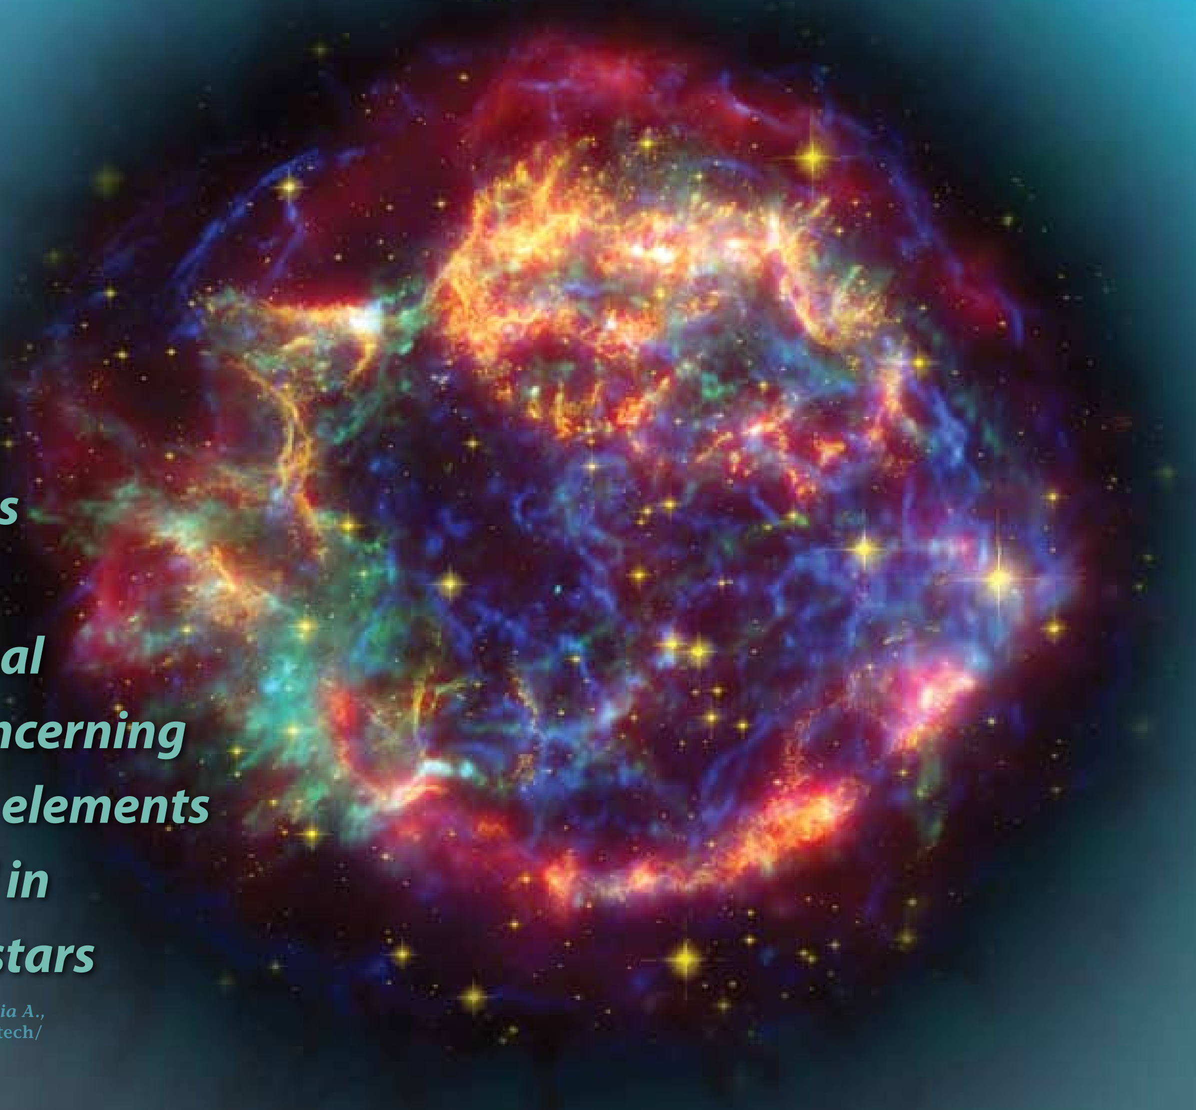

Supernova remnant *Cassiopeia A.*, Image Credit: NASA/JPL-Caltech/STScI/CXC/SAO

Tin prevents steel from corroding in cans of food and provides the ringing sound in church bells. As recently discovered at Oak Ridge National Laboratory, tin also plays a pivotal role in the creation of hundreds of heavy elements in catastrophic stellar explosions. The finding debunks a long-held belief that the formation of any one of these elements has little effect on the creation of others.

To understand how heavy elements evolved billions of years ago into the ancient stardust that created and sustains life on Earth, ORNL theoretical astrophysicist W. Raphael Hix and his collaborators build sophisticated computer models that track how heavy subatomic nuclei, such as tin, antimony and lead, become heavier by capturing neutrons. The models need, as input, the probabilities that thousands of different types of heavy nuclei will capture free neutrons under the conditions in which stars become supernovae. The conditions are indeed extreme—temperatures of a few billion degrees, with relatively few heavy nuclei in a bath of a billion trillion free neutrons per cubic centimeter.

In this hot stellar gas, photons—energetic particles of invisible and visible light—not only carry the temperature of the gas but also "boil off" neutrons (photodisintegration) that have just been captured by heavy nuclei. When the neutron capture rates and the boiling-off rates balance out, the system is in equilibrium. Calculations of the number of nuclei of any one isotope at this moment are relatively straightforward, relying on properties such as nuclear masses rather than the rate at which fast heavy nuclei capture free neutrons.

As the system cools, however, this balance is broken. Scientists previously thought that some neutron capture rates might have an effect—but only to influence a small number of nuclei. As hundreds of heavy nuclei are formed in the explosion, too many nuclei and not enough time (less than seconds) are available for a global impact to occur. This is one reason that theoretical estimates were thought sufficient to use the neutron capture rates as input to the models.

Necessity was another reason that estimates were used. The nuclei responsible for heavy element formation are radioactive isotopes that do not occur naturally on Earth. Therefore, measuring their neutron capture rates before they spontaneously change into other nuclei is a daunting task.

"ORNL has the only facility in the world that can produce and accelerate a beam of radioactive tin nuclei to a high enough energy that enables an indirect experimental determination of the rate of the neutron capture reaction," says Michael S. Smith, experimental physicist at ORNL's Holifield Radioactive Ion Beam Facility. "We have conducted such experiments on five different nuclei, with a goal of providing an empirical foundation for the simulations.

"To improve the explosion models, we hope to determine experimentally neutron capture rates of select isotopes with our accelerator," he adds. "Raph and his collaborators will put our new rates in the model, run a new simulation of an exploding star and get new predictions of the rapid creation of heavy isotopes." The researchers did not, however, expect any single new rate to change the predictions dramatically.

One of the five ORNL experiments was led by Ray Kozub of Tennessee Technological University. Kozub accelerated tin-130 nuclei and aimed the beam at a polyethylene target containing deuterium—heavy hydrogen with one proton and one neutron. As the beam interacts with the target, many neutrons leave the deuterium nuclei in the target and attach to passerby tin-130 nuclei in a so-called "transfer reaction," while the protons from the broken deuterium nuclei are ejected from the target. Sophisticated detector systems measure the products of the reactions—protons and heavy tin-131 nuclei.

"This indirect approach is one of the few ways by which we can determine how many times per second a neutron attaches itself to a tin nucleus in an exploding star," Smith says. "This capability is one of the major motivations for building facilities that produce radioactive beams."

The view that measuring certain individual neutron capture rates may actually be important evolved during simulations performed by Josh Beun, then a Ph.D. graduate student at North Carolina State University. In collaboration with his adviser Gail McLaughlin, Hix and Rebecca Surman of Union College in upstate New York, Beun was simulating the rapid neutron capture process, or the r-process, which produces about half of the elements present in stars and on Earth. These nuclei captured neutrons in milliseconds to seconds. "Astrophysicists are not sure where the *r*-process happens," Hix says, "but our latest finding should help point us in the right direction."

When Beun changed the neutron capture rates for four of the nuclei studied in ORNL experiments, he found the expected result. "These increased capture rates would make a little difference but only locally—that is, only for nuclei with similar masses," says Hix. However, he adds, "For tin-130 we saw a large effect that really puzzled us. We needed more than two years to figure out how tin-130 can make a huge impact on the abundances of other isotopes.

"What we found is that a particular isotope's neutron capture rate can matter on a global scale, affecting the rate of production of 100 heavier isotopes. As the gas cooled, tin-130 nuclei more likely held onto captured neutrons, making them unavailable for other captures."

Tin-130 is not the only nucleus that has this special role. "We found by happenstance that a fraction of these neutron capture reactions actually matter early enough to make a global change," says Hix.

The team found that tin-130 has three important characteristics that make it special: First, tin-130 is abundant at the point when the *r*-process is fervently operating because of the presence of at least 100 free neutrons for every heavy nucleus. Second, tin-130 has a slow radioactive decay rate, enabling the isotope to hold its neutrons for a longer time.

A third characteristic is that the mass difference between tin-130 and tin-131 is considerably larger than differences found in other neighboring pairs of nuclei. As the temperature in the explosion drops and fewer photons are present, neutron capture rates matter for a longer time.

"The earlier a nucleus breaks out of the balance between neutron capture reactions and boiling-off reactions, the better the chance of effecting global change," Hix says. "The heavier nuclei have more time to form." The breaking of this balance depends on the mass difference between the nucleus and its nearest neighbors.

The group is currently surveying thousands of nuclei to see which others satisfy these criteria and therefore can have a global impact on element formation in stellar explosions. Experimental nuclear physicists are very excited about this change in the conventional wisdom, primarily because it suggests that measurements of other neutron capture rates must be carried out to predict accurately heavy element formation in supernovae. Perhaps more significant, the finding also provides an additional reason for the construction of a $600 million, next-generation facility in the United States, called the Facility for Rare Isotope Beams, that would produce intense beams of radioactive nuclei.—*Carolyn Krause*

...and the WINNERS are...

Accomplishments of Distinction at Oak Ridge National Laboratory

The annual UT-Battelle Awards Night celebration was highlighted by the presentation of the UT-Battelle Director's Awards for Outstanding Individual Accomplishment. The **Director's Award in Science and Technology** went to **David Singh**, honored for his development of theoretical approaches and their application to key problems in the fields of condensed matter physics, materials science and related areas. The **Director's Award in Laboratory Operations** went to **Michael J. Pierce** for his efforts to ensure the availability of ORNL's nuclear facilities in support of the laboratory's nuclear research and development mission. **Mark Dobbs** garnered the **Director's Award in Community Service** for his long history of traveling with organizations locally and worldwide to provide humanitarian services.

Thomas Wilbanks has been appointed a member of the Committee on America's Climate Choices by the **National Academy of Sciences**.

Laboratory Director **Thom Mason** has received **McMaster University's Distinguished Alumni Award for the Sciences, 2008,** for his "high level of distinction and achievement through scholarship, research, teaching, and creative contributions to the sciences and service to society."

Will Minter was presented with the **Black Engineer of the Year Award** by **Lockheed Martin Corporation**, **the Council of the Engineering Deans of the Historically Black Colleges and Universities** and ***US Black Engineer & Information Technology Magazine*** for promotion of higher education.

Nermin Uckan has been appointed a member of the **Fusion Energy Sciences Advisory Committee** of the **U.S. Department of Energy**.

Jun Qu has received the **John G. Bollinger Outstanding Young Manufacturing Engineer Award** from the **Society of Manufacturing Engineers**.

ORNL received the **High Performance Computing Challenge Award** in the Linpack Test category and the **High Performance Computing Challenge Runner up Award** in the Global Fast Fourier Transform category in a competition sponsored by the **Defense Advanced Research Projects Agency, the National Science Foundation** and the **U.S. Department of Energy**.

Jeff Smith has received the **Oak Ridge "Muddy Boot" Award** from the **East Tennessee Economic Council**.

Louis Mansur has received a **Lifetime Achievement Award** from the **American Nuclear Society's Materials Science and Technology Division**.

Rebecca Efroymson and **Paul Hanson** have been elected **fellows** of the **American Association for the Advancement of Science**.

Pencheng Dai, **Chong Long Fu**, **Amit Goyal**, **Soren Sorenson**, **Charles Vane** and **Andrey Zheludev** have been elected **fellows** of the **American Physical Society**.

Mike Miller has been elected a **fellow** of the **Minerals, Metals and Materials Society** for his leadership in the development and utilization of atom probe tomography for characterization of materials.

David Singh

next issue

ORNL Tackles
10 Big Problems
for Science

online extras

www.ornl.gov/ORNLReview

on the web

Reference desk:

- Read journal papers on research described in this issue.

OAK RIDGE NATIONAL LABORATORY REVIEW
P.O. Box 2008, Oak Ridge, Tennessee 37831-6266

ORNL *Review*
Editor and writer—Jim Pearce
Writers—Dawn Levy, Leo Williams
Designer—LeJean Hardin
Illustrator—Andy Sproles
Photographers—Charles Brooks, Curtis Boles,
Larry Hamill, Jason Richards
Web developer—Dennis Hovey

Editorial office telephone: (865) 574-7183
Editorial office FAX: (865) 574-9958
Electronic mail: ornlreview@ornl.gov
Web addresses: www.ornl.gov
www.ornl.gov/ORNLReview/

Oak Ridge National Laboratory is managed by
UT-Battelle, LLC, for the U.S. Department of
Energy under contract DE-AC05-00OR22725

ISSN 0048-1262

UT-BATTELLE
OAK RIDGE
National Laboratory
U.S. DEPARTMENT OF
ENERGY