Exploring Science Frontiers at Petascale

Oak Ridge National Laboratory

Petascale Computing

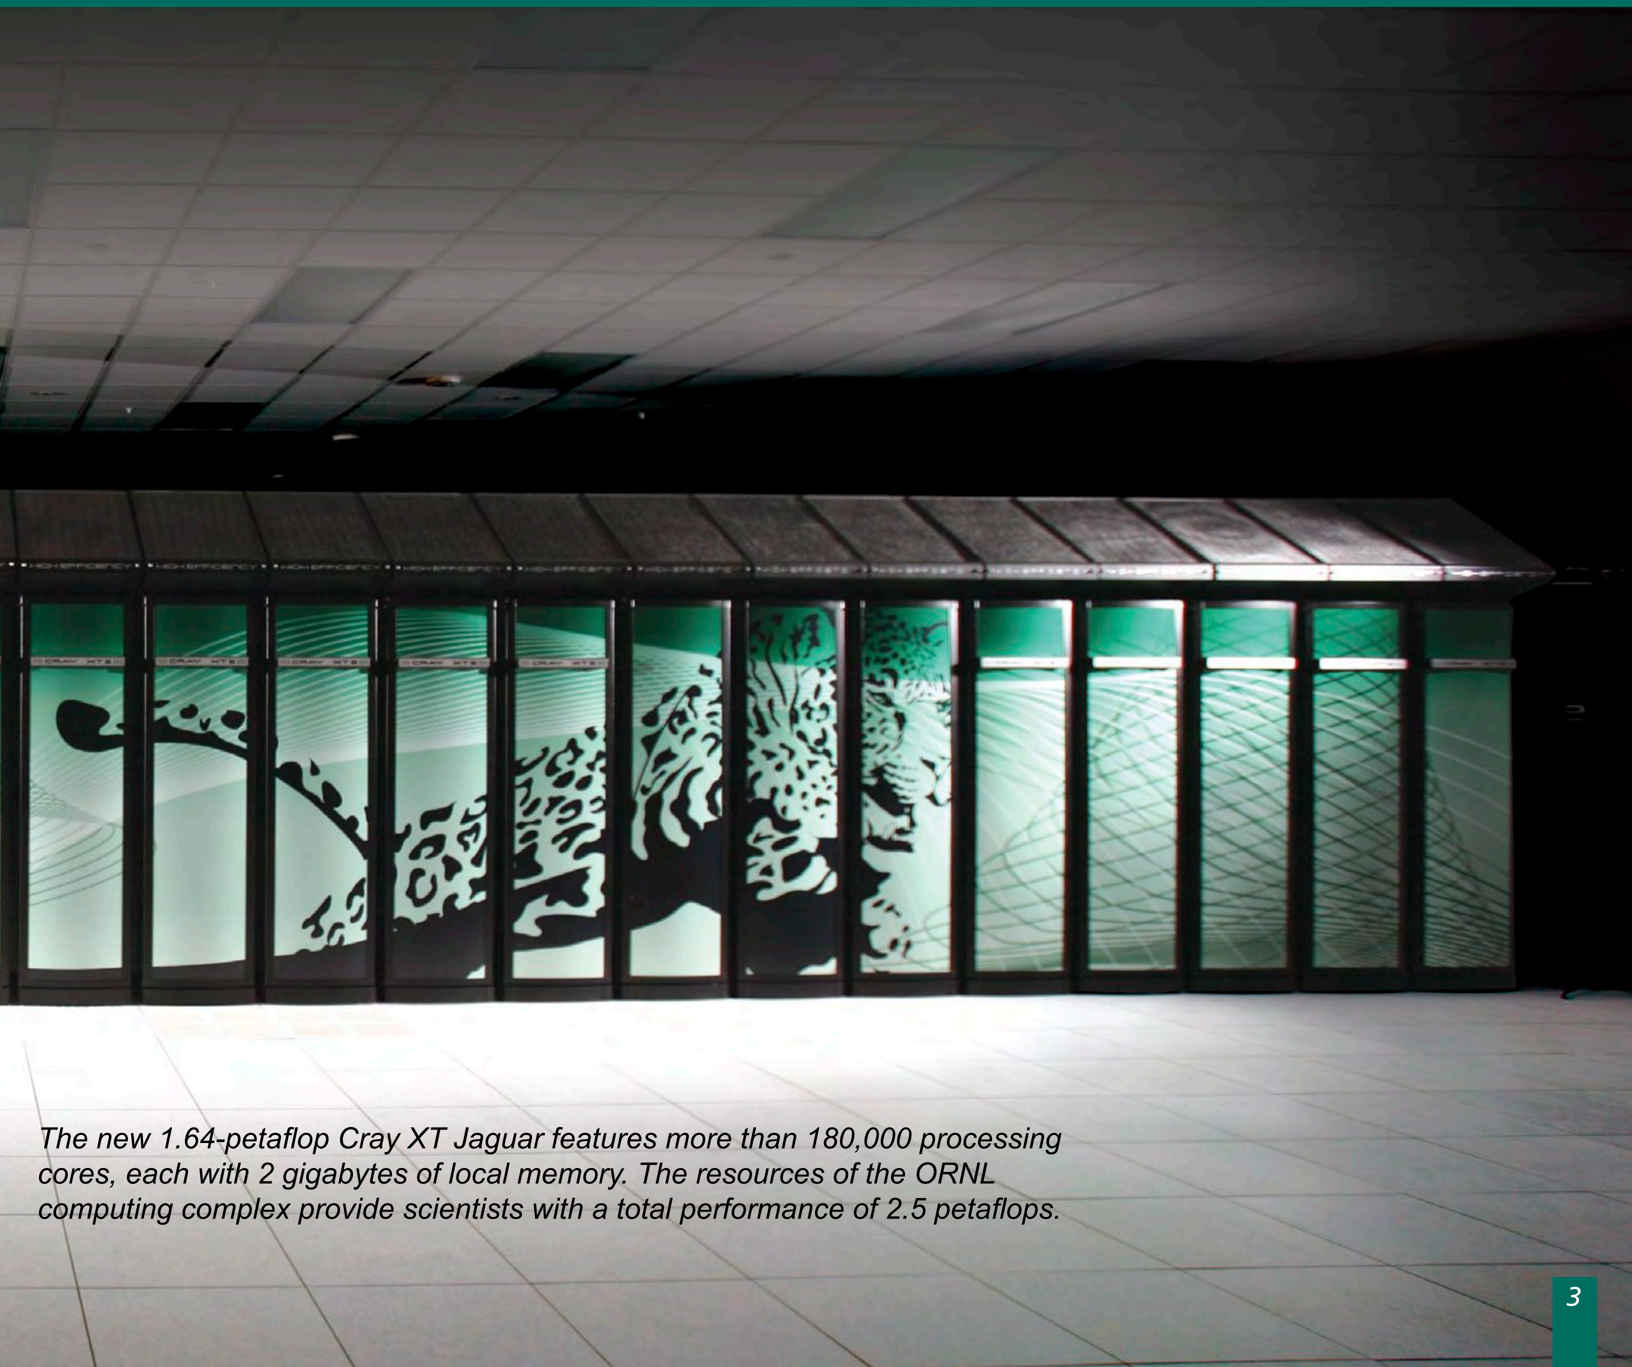

The new 1.64-petaflop Cray XT Jaguar features more than 180,000 processing cores, each with 2 gigabytes of local memory. The resources of the ORNL computing complex provide scientists with a total performance of 2.5 petaflops.

ORNL: World's MOST POWERFUL Computing COMPLEX

Petascale computing is here! The scientific computing complex at Oak Ridge National Laboratory (ORNL) will soon offer two supercomputers with a combined performance of more than 2.5 petaflops, or quadrillion floating point operations (calculations) per second. The petascale machines make ORNL the world's most powerful computing complex for open science and enable researchers to explore solutions to science's most intractable problems, including how to meet our energy needs while lessening the planetary impacts of energy use.

"This is the best place in the world to do computational science," said Arthur Bland, director of ORNL's Leadership Computing Facility (LCF) project to upgrade and install supercomputers at the National Center for Computational Sciences (NCCS). The LCF was established in 2004 to strengthen American competiveness in high-performance computing (HPC) after Japan deployed its Earth Simulator, an NEC machine that could model planetary processes with an unprecedented speed of 36 teraflops (trillion calculations per second).

If mathematics is the language of science, computation is its workhorse. Models and simulations that run on the LCF resources let scientists explore biology, chemistry, and physics in ways unimaginable when science had just two pillars—theory and experiment. Computing will continue to serve as science's third pillar into the foreseeable future.

Under the leadership of Thomas Zacharia, associate laboratory director for computing and computational sciences, ORNL's HPC facilities provide scientists with virtual laboratories unmatched by any other computing facility in the world dedicated to open, or unclassified, research. These computational tools provide deep insight into complex challenges, including the design of future car batteries, the operation of a nuclear fusion reactor running at 100 million degrees Celsius, and the elucidation of cellular ion channels, which are important in health and disease and also may be engineered to catalyze biofuels, produce unique compounds, and detoxify industrial wastes.

The most powerful supercomputer in ORNL's petascale computing complex is Jaguar, a 1.64-petaflop Cray XT system fielded by the NCCS. This enormous new system possesses more than 180,000 AMD Opteron processing cores, each with 2 gigabytes of local memory. Also on the ORNL campus, at the National Institute for Computational Sciences (NICS), is Kraken, a Cray XT4 system that is being upgraded


The New York Times
SATURDAY, APRIL 20, 2002
Japanese Computer Is World's Fastest, As U.S. Falls Back
By JOHN MARKOFF
Published: April 20, 2002


CNN
Wednesday, May 12, 2004
U.S. Moves to Build Top Supercomputer

In 2002 Japan launches the world's fastest supercomputer. America responds by establishing the Leadership Computing Facility at ORNL in 2004.

Distinguished guests at the National Center for Computational Sciences

Raymond Orbach, DOE under secretary for science; Thomas Zacharia, ORNL associate laboratory director, computing and computational sciences; David Kiefer, Cray senior vice president for customer development

Bart Gordon, U.S. representative and chairman, House Committee on Science and Technology

Al Gore, former vice president

Thom Mason, ORNL director

Marsha Blackburn, U.S. representative

John Petersen, president, University of Tennessee; Phil Bredesen, Tennessee governor

Lamar Alexander, U.S. senator; Zach Wamp, U.S. representative; Jeff Wadsworth, Battelle executive vice president, global laboratory operations and former ORNL director

Bob Corker, U.S. senator

Michael Strayer, associate director, DOE Office of Advanced Scientific Computing Research

Howard Baker, former U.S. senator

Samuel Bodman, U.S. secretary of energy

HPC and Industry Team Up for Innovation

In today's highly interconnected global economy, market leaders know that HPC is a critical ingredient in the recipe for competitive success. Modeling, simulation, and large-scale data analysis using this tool accelerate innovation while lowering its risk, resulting in reduced costs, faster time to market, and increased revenue.

ORNL's HPC Industrial Partnerships Program helps companies better exploit the competitive benefits of this powerful tool. We have unsurpassed expertise in helping our partners scale their current problems, explore new and competitively important issues beyond their current in-house capabilities, investigate and test new problem-solving approaches, gain unprecedented insight into more advanced HPC systems and software, and get a head start in preparing for such systems.

HPC is a game-changing technology, and companies that have integrated it into their research and production processes are realizing transformational results. ORNL's program is helping them make the next great leap ahead into a competitive future.

Since its inception in 1991, the computing complex at ORNL has increased the computational power available to researchers a millionfold. A thousandfold of that increase has happened in the last 5 years alone.

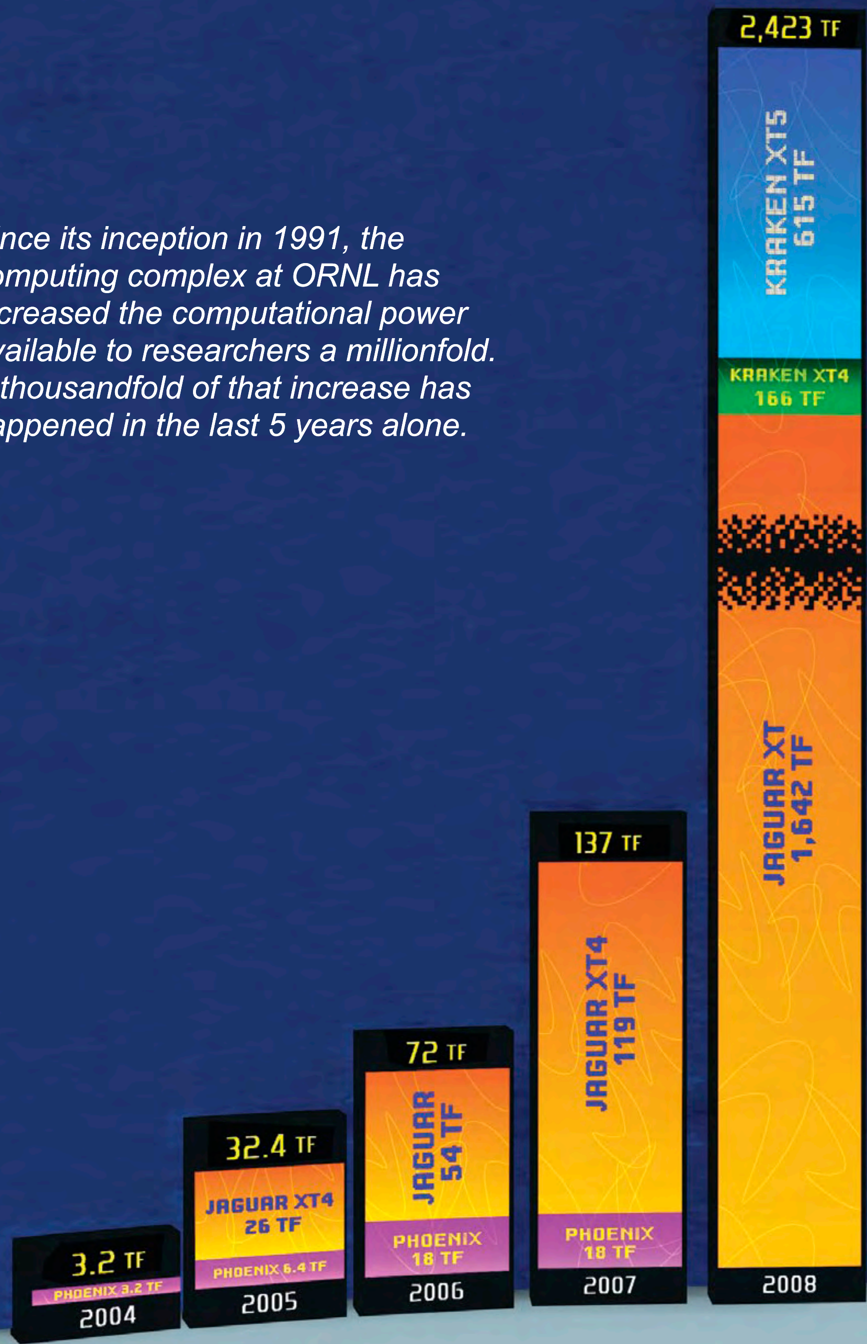

to a similarly speedy 963-teraflop XT5. Kraken is funded by the National Science Foundation's (NSF's) Office of Cyberinfrastructure through a grant to the University of Tennessee (UT).

The NCCS, sponsored by the Department of Energy (DOE) Office of Science, manages the Jaguar system for use by leading computational researchers from universities, industry, and national laboratories. NICS, sponsored by the NSF and managed by UT and ORNL, manages Kraken for use by computational researchers in the academic and industrial communities.

The combined power of Jaguar and Kraken represents a leap forward in computational science capabilities. These supercomputers will play a major role in extending the frontiers of human knowledge by helping researchers reveal the future of regional climates, develop ways to tap new energy sources, and delve into the nature of matter and the origins of life.

The Path to Petascale

Computing at ORNL has come a long way since the Center for Computational Sciences (CCS) was created at the laboratory in 1992. A year earlier ORNL had joined with three other national laboratories and seven universities to submit a Partnership in Computational Science (PICS) proposal that DOE's Office of Science fund a high-performance computer to be located at the proposed ORNL CCS. The new center, with a supercomputer operating at 35 gigaflops, or billion calculations per second, was only the beginning. Since 1991 the computational power available at ORNL has increased a millionfold, making it possible for researchers to tackle grand challenges in science, develop software and other supercomputing tools, and mentor a new generation of scientists in HPC.

That same timespan has seen the partnership between ORNL and UT in computational sciences flourish. In 1991, the same year as the PICS proposal, the two institutions collaborated in establishing the Joint Institute for Computational Sciences. In 2007 this long association made for a natural partnership as the two institutions proposed a new NSF computer center for the academic community, locating a petascale computer system at ORNL. The availability and proximity of a world-class facility with the dedicated space, power and cooling infrastructures, support systems, and in-house expertise to accommodate two petascale computers are great advantages to ORNL's NSF partner.

The support systems required for effective operation of such powerful computers are at a scale unimaginable

A Revolution Is Under Way . . .

"From every corner of science, a revolution is under way because of the growing amount of data being generated and the rapid increase in scientific understanding resulting from applying advanced computational science tools to these data." —2002, **Thomas Zacharia**, associate laboratory director for computing and computational sciences

"University, laboratory, and industrial researchers using a broad array of disciplinary perspectives are making use of the leadership computing resources to generate remarkable consequences for American competitiveness." —2007, **Raymond L. Orbach**, DOE under secretary for science

"The leadership-class computing capability . . . at ORNL will enable researchers to probe the deepest secrets of nature—and facilitate the technical, economic, and social benefits such understanding will yield. It is no exaggeration to say that this machine will give both the U.S. scientific community and industrial sector a significant competitive advantage over the rest of the world." —2004, **Spencer Abraham**, secretary of energy

"This $65 million NSF grant . . . enables an exciting new partnership between the National Science Foundation, the Department of Energy through its Oak Ridge National Laboratory, and the University of Tennessee and all of the partnering universities involved in this grant. It also makes available to these universities the considerable computational expertise and capabilities of Oak Ridge National Laboratory, the world's most powerful open scientific computing complex." —2008, **Arden L. Bement, Jr.**, director of the National Science Foundation

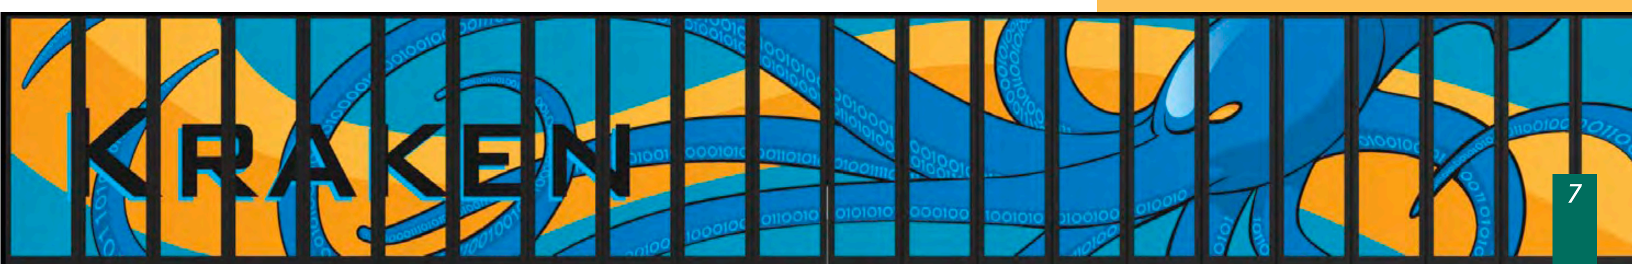

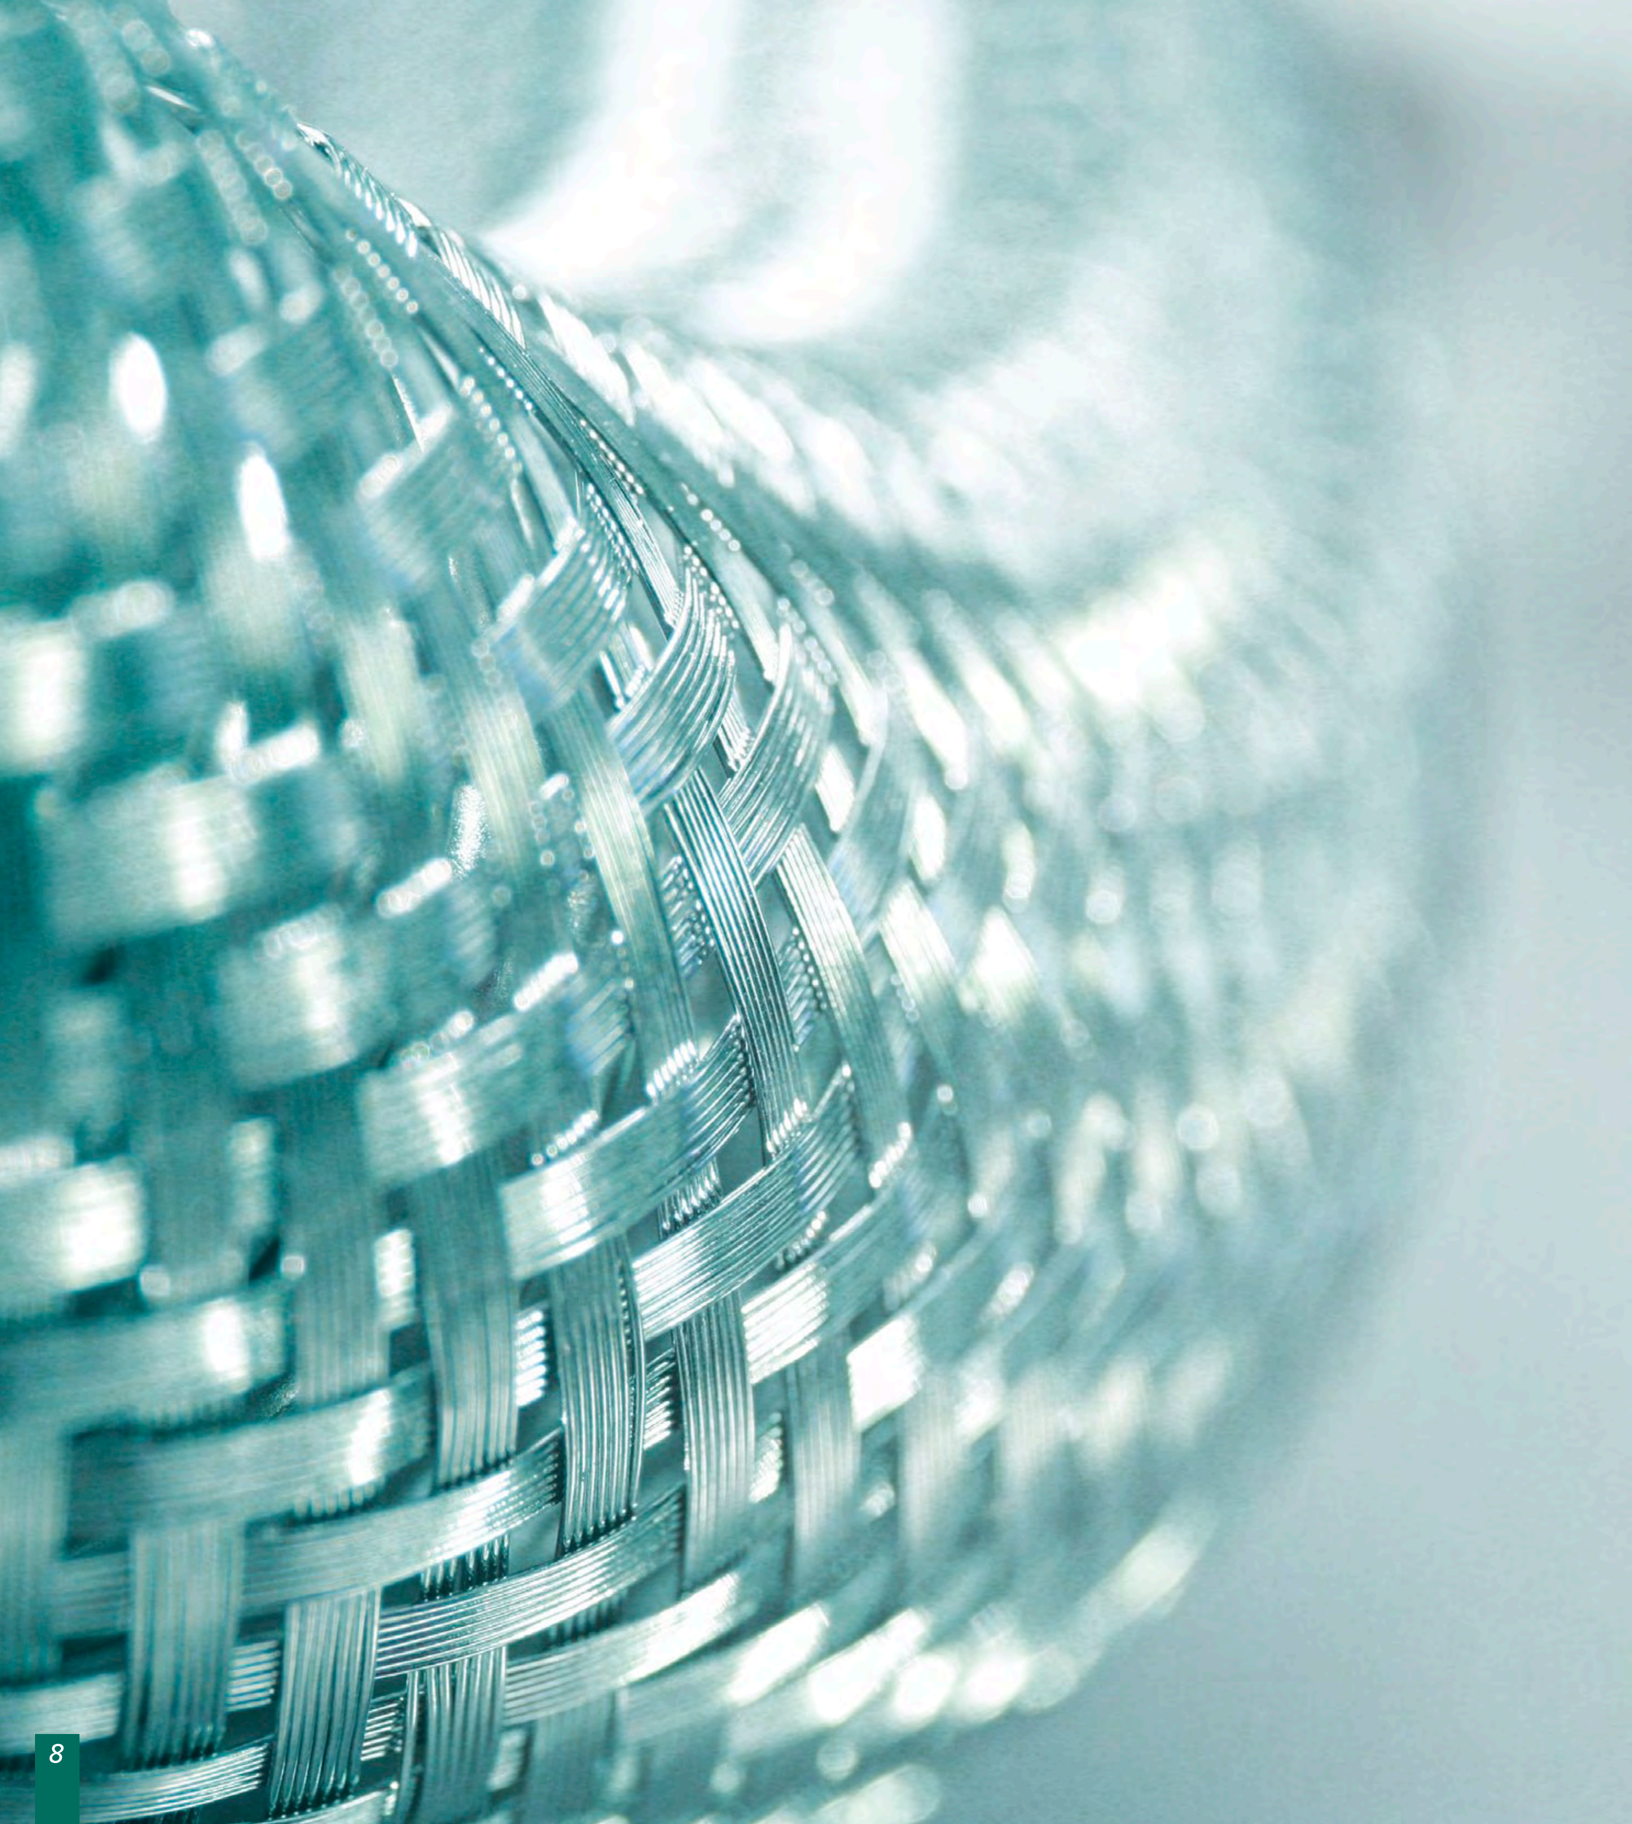

a few years ago. ORNL has established networking and data-handling resources to support the petaflop machines that include 10-gigabyte-per-second connections to the ESnet and Internet2 networks, a scalable High-Performance Storage System (HPSS) for storing simulation data, and a 10-petabyte Lustre-based shared file system (Spider) that will connect to every system in the complex. The disk subsystem can transfer data at speeds greater than 200 gigabytes per second.

As computing systems grow exponentially in power, they also grow in complexity. ORNL has amassed a pool of experienced computational scientists and other specialists to train and assist the user community in making the most of the HPC resources. These scientists are experts in both the productive, efficient use of state-of-the-art supercomputers and a range of research areas including fields such as quantum physics, astrophysics, materials science, climate, chemistry, and biology. The ORNL team ensures the smooth operation of these systems and produces tools that continue to simplify their use, allowing scientists to focus on their science instead of the mechanics of petascale computing.

Preeminence in HPC is indispensable in maintaining U.S. leadership in science and technology as well as economic competitiveness. ORNL's powerful computing complex will continue to attract world-class research collaborations and usher in the next generation of breakthroughs in climate change, materials science, energy assurance, and other global priorities. As discoveries get translated into commercial or intellectual capital, this unique scientific computing complex will continue to pay dividends to the nation and the world well into the future.

Anatomy of a Jaguar

Jaguar is the culmination of a close 4-year partnership between ORNL and Cray that has pushed computing capability relentlessly upward. The XT system grew in strength through a series of upgrades. In 2008 a 263-teraflop Cray XT4 was upgraded with the addition of a 1.4-petaflop Cray XT5. The combined system uses an InfiniBand network, the Spider file system, and approximately 182,000 processing cores to form the DOE Office of Science's 1.64-petaflop system.

Occupying 284 cabinets, Jaguar uses the latest quad-core Opteron processors from AMD and features 362 terabytes of memory and a 10-petabyte file system. It has 578 terabytes per second of memory bandwidth and unprecedented input/output (I/O) bandwidth of 284 gigabytes per second to tackle the biggest bottleneck in monster systems—moving data into and out of processors.

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 its blistering speed.

"What makes this machine unique is the balance it represents" among speed, power, and other elements essential to scientific discovery, Bland said. "Several design choices make it the best machine 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. Users have been enormously successful in using this architecture."

From a programming standpoint, the upgraded Jaguar is essentially the same as the XT4 that NCCS users have been using. A consistent programming model allows users to continue to evolve their existing codes rather than write new ones. Applications that ran on previous versions of Jaguar can be recompiled, tuned for efficiency, and then run on the new machine. "The CPU performance continues to go up, but the basic programming model of the system remains intact," said Ricky Kendall, LCF group leader for scientific computing. "This is critically important to our user community because applications typically last for 20 to 30 years."

The XT5 portion of Jaguar has a power density of more than 2,000 watts per square foot. That level of power consumption creates commensurate heat that needs to be 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 boils the refrigerant, which changes to a gaseous phase. 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. Without this extremely efficient cooling system, it would be impossible to build such a large system, said Jim Rogers, LCF director of operations. The new cooling technology also benefits the efficiency of the computer center. At the NCCS, the cooling adds only 30 percent to the power required for the computers. The average in computing centers is 80 percent.

Transforming Science

Simulations on ORNL's HPC systems consume millions of processor hours and generate an avalanche of answers. They explore complex topics including how enzymes aid biofuel production, how proteins misfold in certain diseases, and how the earth's atmosphere affects ocean circulation.

ORNL's terascale simulations have brought unprecedented scientific and economic opportunities. They allowed combustion scientists to fully resolve flame features of burning fuel—a key to designing fuel-efficient, low-emission engines. They gave physicists insight into how to use radio waves to heat and control ionized fuel in a fusion reactor. They helped researchers design materials that can recover energy escaping from vehicle tailpipes and develop advanced power-generation facilities that can trap pollutants. In addition, they allowed validation of an important model describing the behavior of high-temperature superconductors, which can transmit energy without losses.

"The Jaguar system at ORNL provides immense computing power in a balanced, stable system that is allowing scientists and engineers to tackle some of the world's most challenging problems," said meteorology professor Kelvin Droegemeier of the University of Oklahoma. "In my own work, we're making historical weather forecasts—an ensemble of ten runs every 5 days for the past 20 years, or more than 14,500 runs—to calibrate today's most advanced prediction system. The resulting data set will be of tremendous value to the research community and would not have been possible without Jaguar."

With the arrival of petascale computing, accelerated scientific productivity is certain. The importance of ORNL's computational contributions to breakthrough science was noted in a 2008 report from the DOE Office of Science, America's largest funder of basic physical science programs at universities and government laboratories. The report said six of the top ten recent significant advances in computational science used ORNL supercomputers to gain unprecedented insight into supernovas, combustion, fusion, superconductivity, dark matter, and mathematics.

Moreover, ORNL systems provided much of the simulation data used in the Fourth Assessment Report of the United Nations Intergovernmental Panel on Climate Change (IPCC), which concluded planetary warming during the twentieth century was probably due to human activity. The IPCC shared the 2007 Nobel Peace Prize with former Vice President Al Gore.

Computing Complex Facts

• Jaguar and Kraken • 420 Cray XT cabinets • 2.5 peak petaflops • 66,427 quad-core AMD Opteron processors • 265,708 processor cores • 2.3 GHz • 39,387 nodes • 479 terabytes of memory • 148,518 DDR2-800 DIMMs • 6,329 miles of interconnect cables • 480-volt power • 1,353 gallons of R134a refrigerant • 6,800 gallons per minute of 42°F chilled water • 12.7 megawatts of power • 1,278,000 cubic feet per minute of cooling air • 21,030 disks • 13.6 petabytes of disk capacity • 6 miles of InfiniBand cables • 192 Dell I/O servers • 332 gigabyte-per-second I/O bandwidth • 786-terabyte-per-second global interconnect bandwidth • National Center for Computational Sciences • National Institute for Computational Sciences • Oak Ridge National Laboratory • World’s Most Powerful Computing Complex•

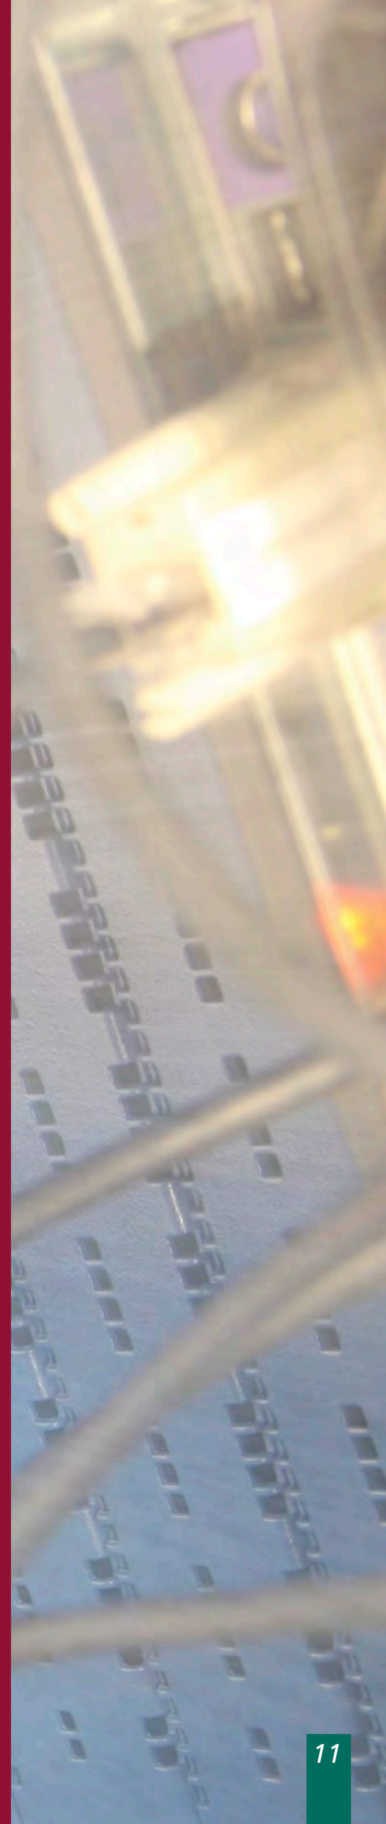

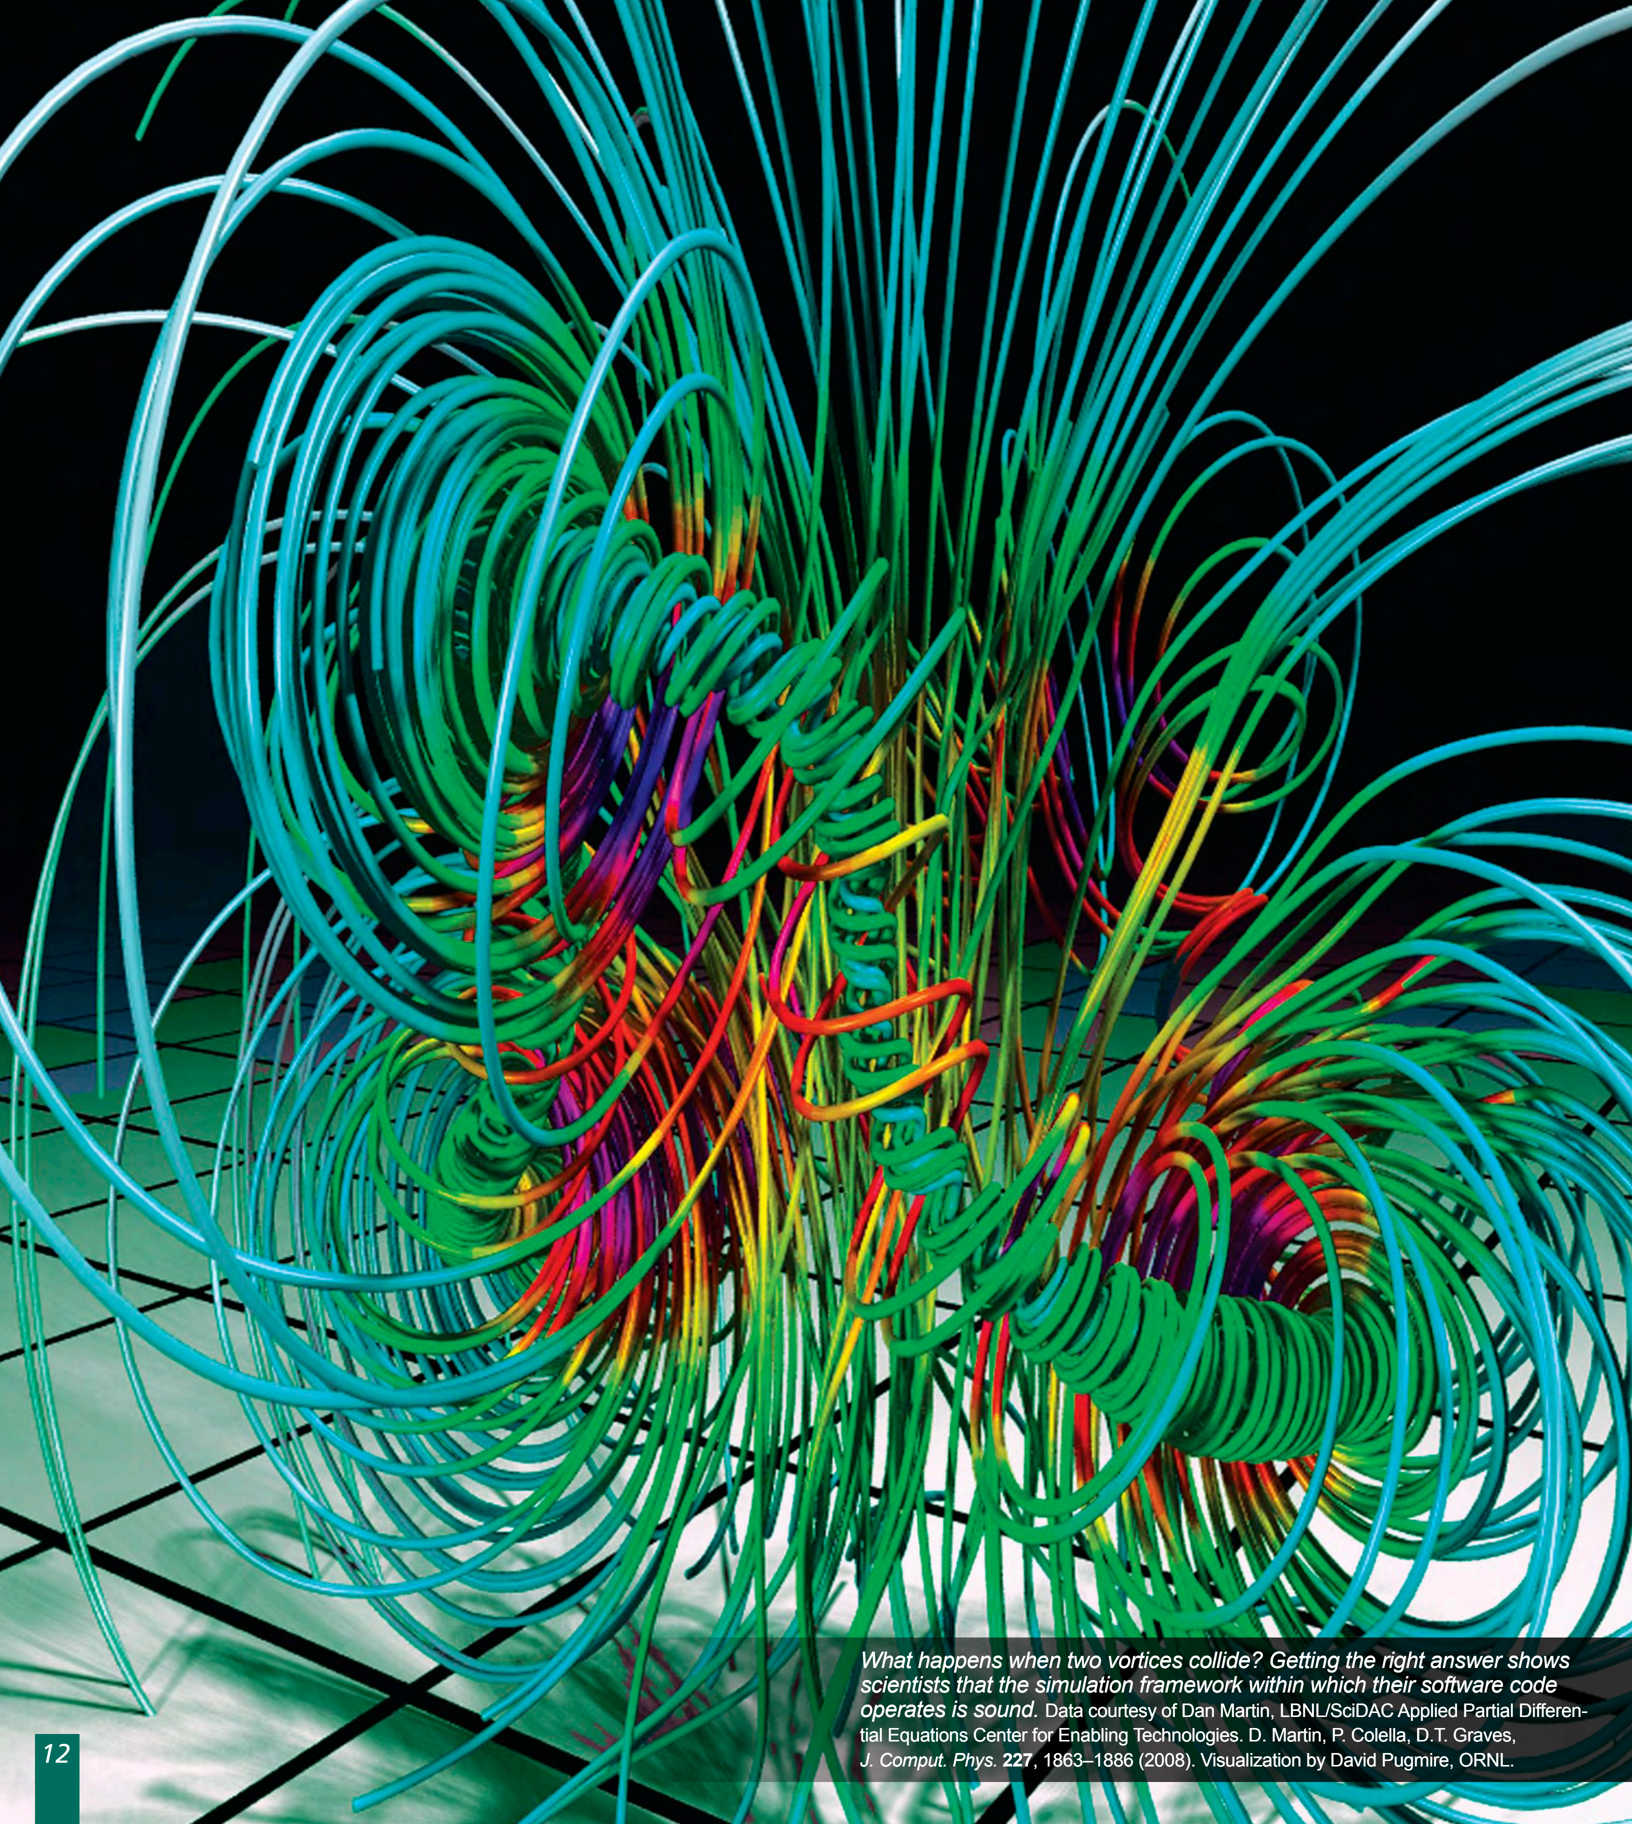

What happens when two vortices collide? Getting the right answer shows scientists that the simulation framework within which their software code operates is sound. Data courtesy of Dan Martin, LBNL/SciDAC Applied Partial Differential Equations Center for Enabling Technologies. D. Martin, P. Colella, D.T. Graves, *J. Comput. Phys.* **227**, 1863–1886 (2008). Visualization by David Pugmire, ORNL.

A Petascale Challenge: STAYING AFLOAT IN DATA FLOODS

A petascale computer produces data on a scale never seen before—hundreds of terabytes for a single run. High I/O rates as data moves in and out of processors are a good thing. But managing the flood of data and making sense of the story it tells present challenges on the same gargantuan scale. Accelerating and simplifying those tasks are crucial to turning data into discoveries.

"The first step is getting the applications ready, which includes getting both the algorithms and I/O ready," said Scott Klasky, leader of the End-to-End task at the NCCS. At I/O rates acceptable in the past, writing and reading hundreds of trillions of data points could take days. To avoid that bottleneck, NCCS experts work to accelerate data management and analysis by orders of magnitude.

Klasky and researchers at Georgia Tech and the Scientific Discovery through Advanced Computing (SciDAC) Scientific Data Management (SDM) Center have developed a high-performance I/O system called ADIOS (Adaptable I/O System). In the summer of 2008, ADIOS aided in achieving a fusion simulation that used more processing hours than any other in the world; the gyrokinetic toroidal code (GTC) used 93 percent of the NCCS's XT4 processors. ADIOS reduced I/O time for GTC tenfold.

With the help of the Georgia Tech collaborators, the NCCS is using ADIOS to support both synchronous and asynchronous data transport. In asynchronous I/O, data input and output overlap with its processing, and both the hardware and the software must be designed to handle this. Scheduling of the I/O must be optimized with the internal communication in a simulation.

But scientific understanding also hinges on visualization and analysis of the data produced by running the software applications. As data sets grow in size, so must the algorithms scientists use to analyze simulated processes. "If you can't look at the data, you don't know what you've simulated," said Sean Ahern, visualization task leader at the NCCS. Among trillions of data values collected in a simulation, often only a small fraction is of interest to the user. To help users cope with the "data flood," Ahern's team and researchers at Lawrence Berkeley National Laboratory have combined the SciDAC SDM Center's FastBit technology with the already-deployed VisIt parallel visualization system. The combination allows certain data-analysis operations, such as tracking particles with specific characteristics in fusion simulations, to run up to 1,000 times faster. The faster the analysis, the faster scientists arrive at new insights.

The cluster that drives the NCCS's high-end visualization facility, which features a 30- by 8-foot, 27-projector Powerwall displaying 35 million pixels of information, was extensively upgraded to allow analysis of larger data sets at higher speeds. A new cluster for analysis and remote visualization can handle data sets 20 times larger than the previous cluster—and do so 10 times faster. A new parallel file system enables the Powerwall to display full-wall movies at 30 frames per second. These visualization and analysis capabilities recently allowed astro physicists to track materials in a supernova that change its magnetic field and to gain new knowledge about the effect of these materials on the shock front generated by the star's collapse. As scientific computing systems push new limits of speed and complexity, visualization and analysis tools that help researchers manage and understand the results are racing to keep pace.


STORAGET

Built to MOVE and STORE BIG DATA

The breakthrough science taking place at the NCCS requires a lot of behind-the-scenes support. The supercomputers regularly grab the headlines, but their work depends on a data-handling infrastructure that includes an equally capable file system, high-performance networks, and enhanced storage capacity.

File System

A Lustre-based file system dubbed Spider will replace multiple file systems now scattered on the NCCS network with a single scalable system. It will serve all NCCS platforms and connect to every internal network. Because all simulation data will reside on Spider, file transfers among computers and other systems will be unnecessary. Eliminating file transfers will improve performance, convenience, and cost. Transferring petascale data sets between Jaguar and the visualization system, for example, could take hours, tying up bandwidth on Jaguar, slowing simulations in progress, and requiring the visualization cluster to have its own file system.

"Spider will provide 10 petabytes of storage space—about 1,000 times as much data as is contained in the Library of Congress," said NCCS Director James J. Hack. "It has an aggregate bandwidth of more than 200 gigabytes per second, which means every second it can transfer the amount of data in 50 full-length, DVD-quality movies."

Networking

Networking capability at the NCCS is being expanded in parallel with its computing capability to ensure accurate, high-speed data transfer. High-throughput networks among its systems and upgraded connections to ESnet and Internet2 have been installed to speed data transfers between the NCCS and other institutions. Speedy transfers are especially important to the many remote users of the NCCS because they facilitate movement of simulation data from the NCCS to the users' computers.

Centerwide installation of an InfiniBand network added bandwidth to facilitate movement of large data sets from the supercomputers to other platforms. The InfiniBand network SION (or scalable I/O network) connects all major NCCS systems, including Spider, analysis and visualization platforms, and the HPSS archival data storage system. More than 3,000 InfiniBand ports and greater than 3 miles of optical cable were deployed to provide high-performance I/O.

Storage

The HPSS, NCCS's archival data storage facility, has been significantly upgraded to ensure high-speed, reliable storage and retrieval of terascale to petascale data sets, which contain trillions to quadrillions of bytes. Two Sun Storage Tek SL8500 robotic libraries containing 48 tape drives were added to increase archival storage space. HPSS currently stores more than 3 petabytes of data, and between 4 and 40 terabytes are added daily. The amount stored has been doubling every year, and the addition of two petascale systems is expected to escalate that rate.

The improvements the NCCS is implementing will enable the massive data movement and storage required for the next generation of simulations.

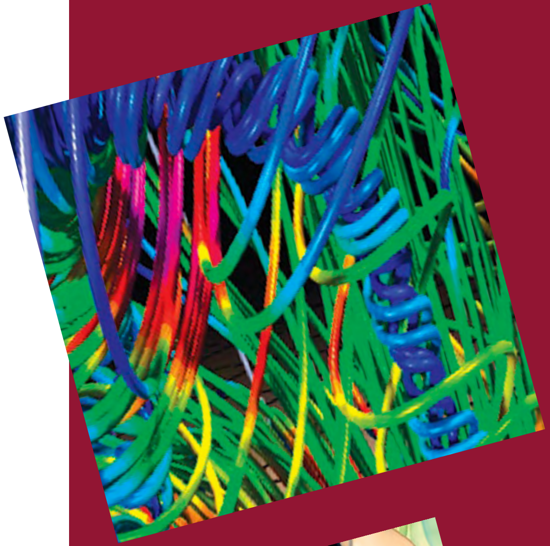

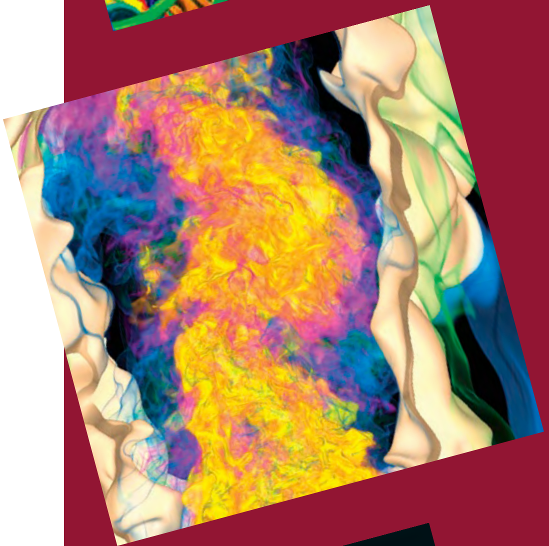

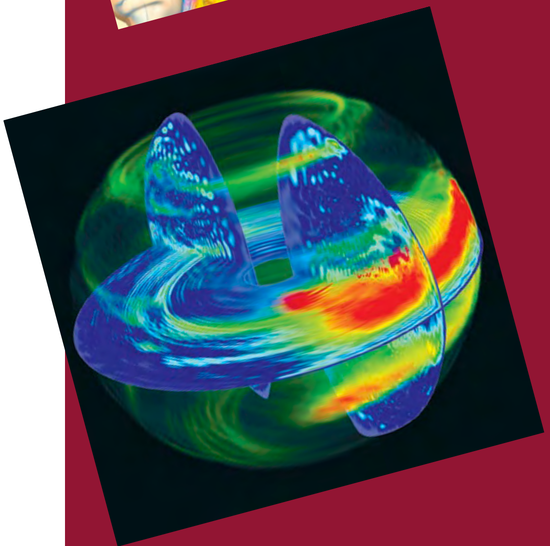


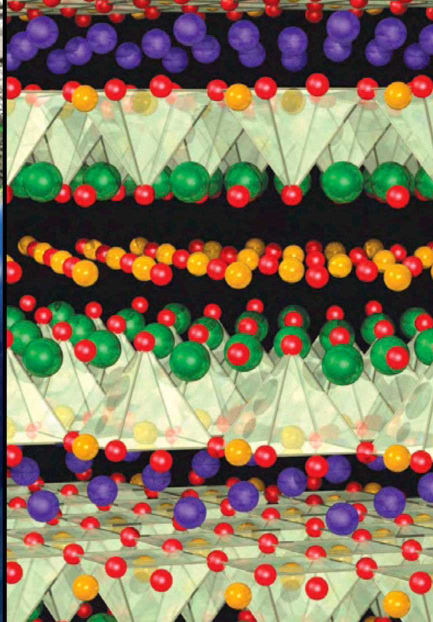
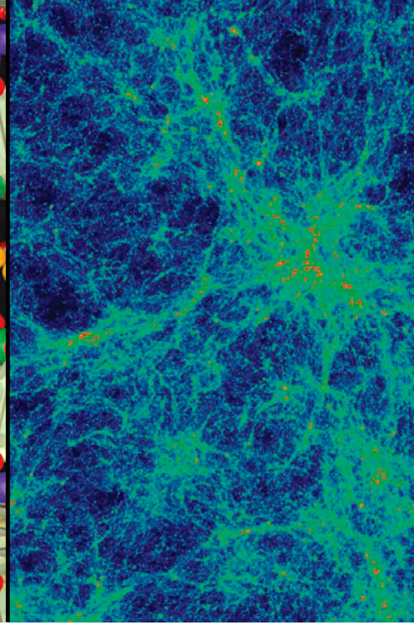
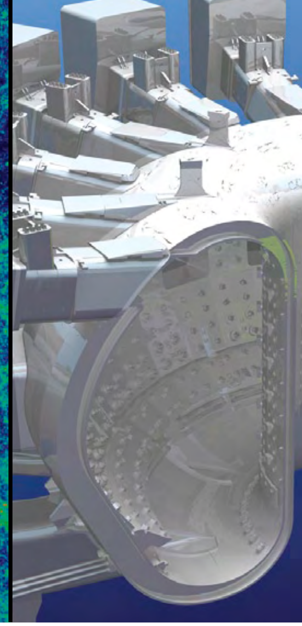

SCIENCE Drives the Need

From probing the potential of new energy sources to dissecting the dynamics of climate change to manipulating protein functions, terascale systems have been an indispensable tool in scientific investigation and problem solving. The capability offered by petascale machines to expand on these advances and address some of humankind's most pressing problems is unprecedented. With two petascale systems coming online during 2009, ORNL provides the scientific community with the most powerful tools on the planet for addressing some of the world's toughest challenges.

Energy Assurance

Petascale leadership systems will arm scientists with better data to aggressively pursue renewable energy sources and more efficiently and safely exploit conventional energy options. With the capacity to simulate systems of millions of atoms, biologists should be able to determine how the enzyme cellulase breaks down cellulose in plant fibers into sugars and use this knowledge to design more efficient enzymes for ethanol production. Simulations are aiding the design of advanced coal plants that tap energy potential while trapping pollutants and greenhouse gases. They are also helping engineers design fuel-efficient combustion devices for vehicles and power-generation equipment. The capability to couple wall, edge, and core physics in nuclear fusion reactors into one integrated ITER simulation tool will aid eventual development of a commercially viable fusion reactor. In fission energy, petascale computers will run the first coupled, geometrically faithful, and physics-inclusive simulations of an entire nuclear reactor core and provide insight into processes on diverse time and length scales that are important in recycling spent nuclear fuel. Petascale leadership computing platforms will help move energy-assurance research from simplified, single-physics studies to explorations of more realistic systems, an important next step toward predictability.

Climate

The potential of petascale simulations to clarify the evolution of the climate system is difficult to overstate. Nearly every aspect of climate simulation stands to benefit from the upcoming petascale era, promising increased knowledge of human impact on the planet and improved stewardship of Earth. Petascale simulations will enable climate scientists to incorporate increasingly sophisticated capabilities in global models, which include atmosphere, oceans, land, sea ice, and other parameters. These extensions to earlier generations of similar models will help scientists better understand the flow of carbon in the climate system and quantify the ways changes in

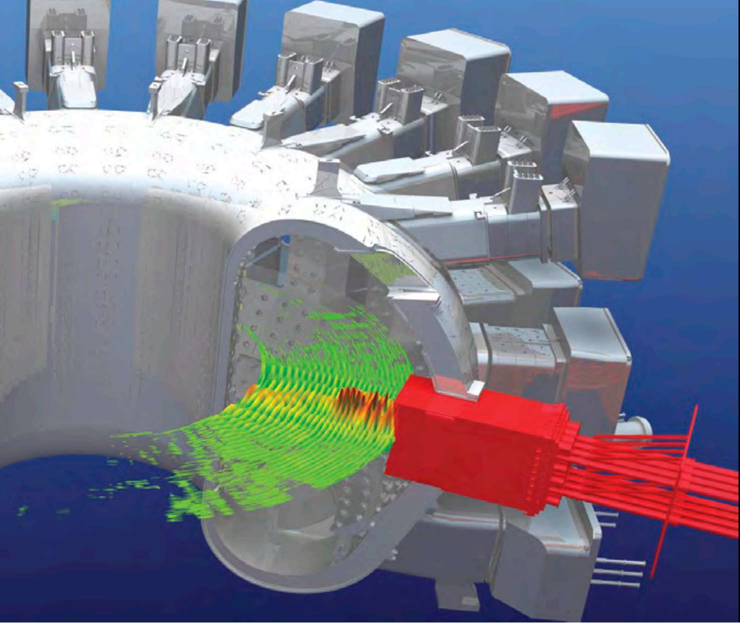
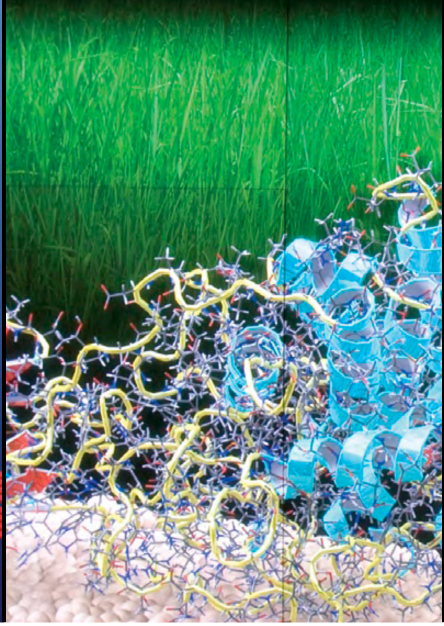
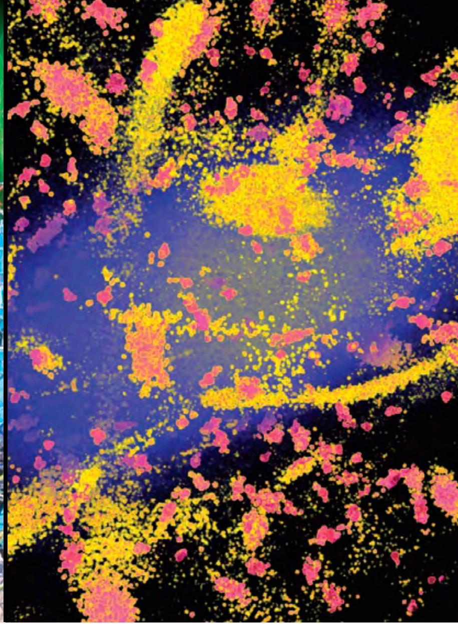

for PETASCALE Computing

atmospheric carbon feed back into other physical processes, such as the global water cycle. Policymakers are asking questions about climate change that require improved simulation fidelity, which cannot be achieved without the much higher spatial resolution in climate models that petascale computing enables. Many questions from stakeholders and resource managers focus on improved predictions over several decades.

Materials

In materials science, innovations made possible by petascale computing promise to bolster American competitiveness in multiple technological sectors. Researchers use detailed atomic simulations to investigate the behavior of materials through which current flows without resistance, and ensuing breakthroughs in superconductors could have a revolutionary effect on a range of energy issues, notably power transmission. Simulations will improve understanding of colossally magnetoresistive oxides and magnetic semiconductors for use in electronics. They will provide information needed to develop switching mechanisms in magnetic nanoparticles that will allow increased data storage on smaller devices. Additionally, they will enable design of more efficient electrical storage systems to advance the development of batteries and capacitors for plug-in hybrid cars.

Biology

Biologists will use petaflop computers for detailed studies showing how proteins carry out crucial tasks. Simulations of larger structures at longer timescales and finer resolution will allow exploration of protein structure and behavior. The aim is to devise ways to make proteins function more efficiently and with less disruption, whether to prevent disease or exploit biological processes for production of biofuels or environmental cleanup. In the case of malignant proteins—such as those that cause neurodegenerative diseases like Alzheimer's—the purpose is to aid in designing pharmaceuticals that stop the growth of abnormal protein structures in the nervous system.

Fundamental Science

Petascale computing will provide the power to reveal fundamental properties of our universe from subatomic to galactic scales. Simulations of core-collapse supernovas will illuminate the explosion mechanism and important observables, such as nucleosynthesis, gravitational waves, and neutrino signatures. Scientists will gain first-principles insights into the fundamental nature of nuclear reactions, permitting a predictive capability for nuclear properties. Data from petascale calculations will enable physicists to compute the strong-force interaction between quarks and gluons so precisely that their knowledge will no longer be limited by theoretical uncertainties. The properties of dark energy will finally come to light as petascale simulation guides the Joint Dark Energy Mission's instrumentation, observation strategy, and data analysis.

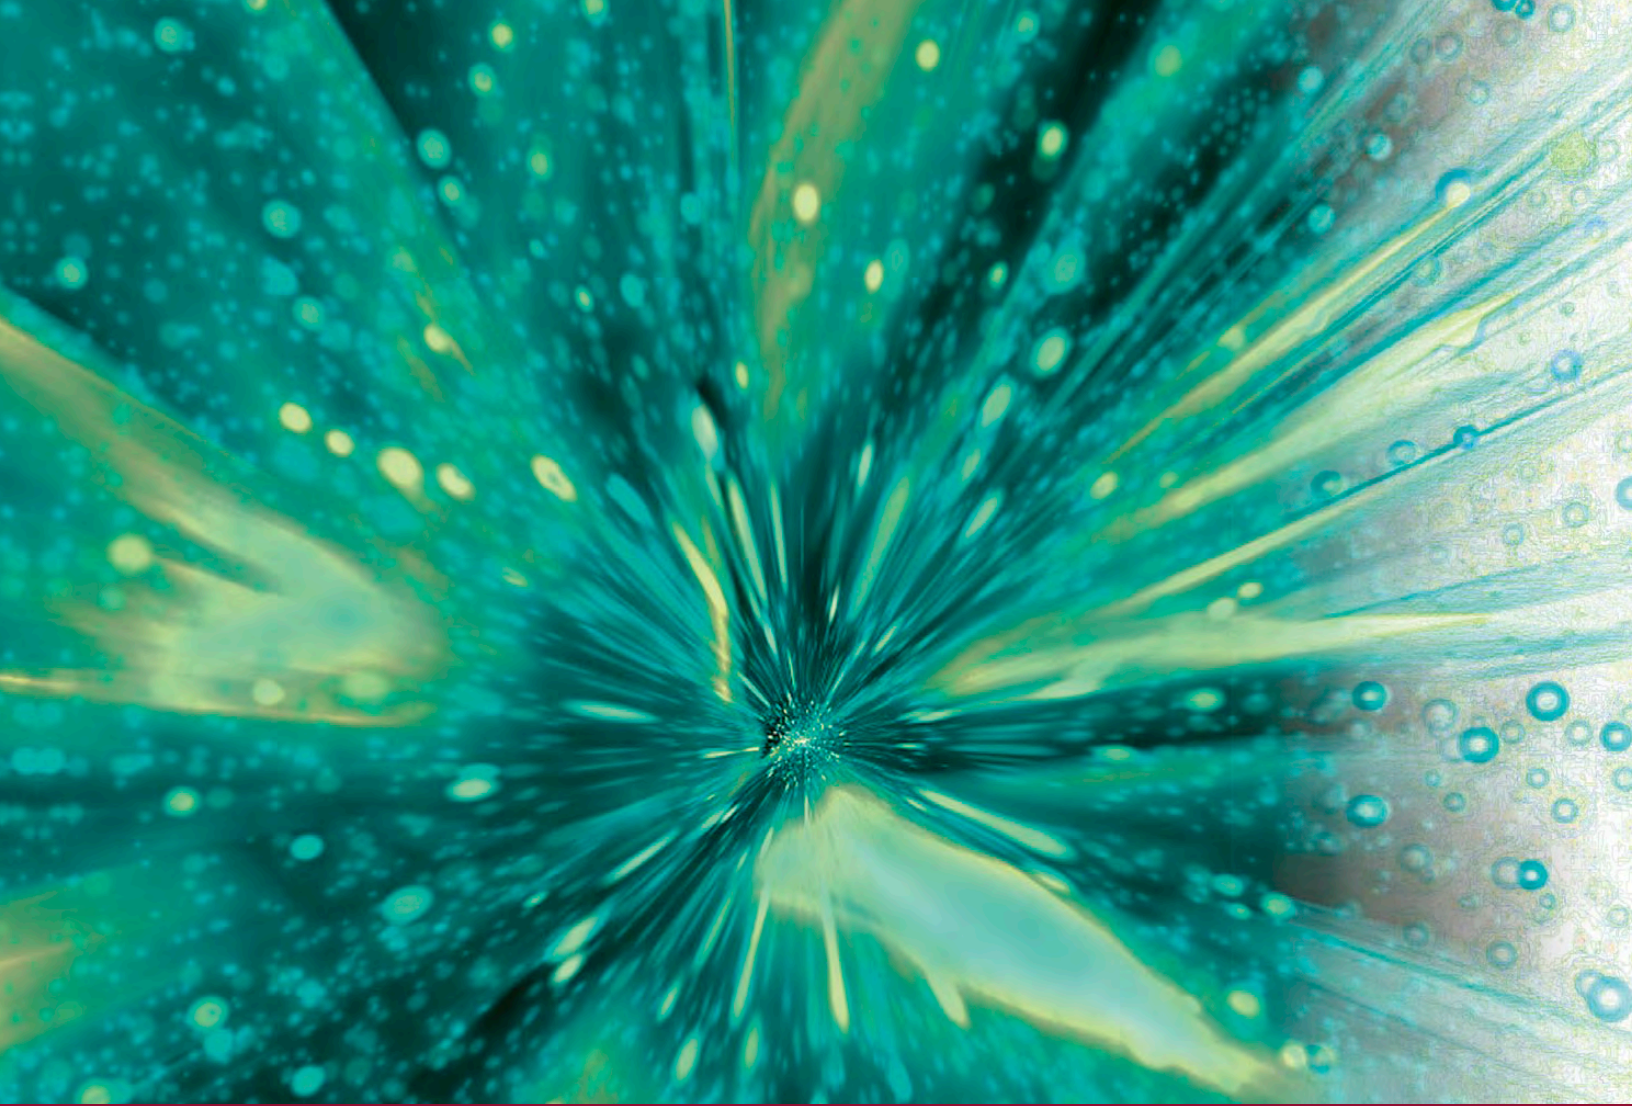

Changing the Way Science Is Done . . .

"[On Jaguar,] we got 100-year runs in 3 days. This was a significant upgrade of how we do science with this model. Forty years per day was out of our dreams." —2007, **Peter Gent**, National Center for Atmospheric Research, chairman of the Community Climate Science Model Scientific Steering Committee

"Simulation is changing the way in which science is done. While not replacing experimentation and observation, it is playing an essential role in the design of new experiments and in optimizing the insight achieved with existing ones." —2008, **Stan Woosley**, University of California–Santa Cruz

"Advanced computations at the petascale and beyond in tandem with experiment and theory are essential for acquiring the scientific understanding needed to develop whole-device integrated predictive models with high physics fidelity for fusion-energy science." —2008, **William Tang**, Princeton University

"With growing concern about environmental changes, there is a scientific and societal need to better understand climate change. The new petascale supercomputers allow scientists to examine with much higher resolution the regional and local changes in our climate system." —2008, **Warren Washington**, National Center for Atmospheric Research

"We will be able to perform simulations on length and timescales sufficient to enable fundamental biological processes to be investigated, such as protein folding, ligand binding, and the structure of plant cell walls. These simulations will, in turn, allow us to design new drugs and orient research into efficient production of cellulosic ethanol." —2008, **Jeremy Smith**, ORNL and UT

The Road Ahead: *PUSHING BEYOND PETAFLOPS*

The scientific computing community is celebrating the opportunities petascale computers present to accelerate scientific discovery. Meanwhile, for ORNL and its partners, the focus is turning to the next generation of systems—planning for the march toward tens, hundreds, and eventually thousands of petaflops of computing power that will be required to meet the voracious demands of computational scientists.

Scientists will use ORNL's petascale computing complex to explore the most important scientific challenges of our time: protecting the planet and ensuring adequate energy supplies. Climate scientists will use the systems to give planners and leaders the tools to anticipate the changing future. Petascale simulations will clarify the role the oceans play in regulating the carbon cycle and map the complex interactions of factors as diverse as plant life, land-use patterns, the atmosphere, and ice sheets.

On the energy front, petascale systems will enable fusion researchers to conduct more detailed simulations of plasma turbulence, which threatens to cool reactors and thwart the fusing of fuel to produce clean, abundant energy. Chemists and materials scientists will use the results of atomic-scale simulations to design improved catalysts that will make fuel cells practical, develop materials that convert waste heat directly into electricity, and revolutionize energy storage technologies. Biologists will gain insight into efficiently converting cellulose into ethanol. Nuclear engineers will use computation to design safer, more efficient, and cost-effective fission reactors.

ORNL is planning for tens, hundreds, and eventually thousands of petaflops

As scientists conduct their petaflop simulations, system designers will be working to provide them with even more powerful tools. To design post-petaflop systems, ORNL is working in collaboration with DOE and the Defense Advanced Research Projects Agency on a program called High-Performance Computing Systems. Cray and IBM have been selected to work toward building machines capable of more than 20 petaflops. ORNL will work closely with both companies to help them understand the strengths and weaknesses of their designs and the needs of scientific applications. By 2011 to 2012, ORNL plans to install a 25-petaflop machine built by the vendor whose design is selected. Beyond that, the goal is to install a 100- to 250-petaflop machine in the 2015 timeframe and an exaflop machine by 2018.

To achieve 100 petaflops, computers must reach beyond existing microprocessor technology. Those computers are likely to be hybrid systems incorporating application accelerators such as IBM's Cell Broadband Engine and graphics processors. More computation per watt and more efficient approaches to managing heat loads will be needed to keep the power requirements manageable for ever-larger machines.

The largest challenge in moving toward exaflop computing will be how to scale operating systems and applications from 100,000 processors or so to perhaps hundreds of millions of execution threads. With so many components operating at once, there will be constant faults in the system; applications will have to be sufficiently robust to run through the faults. Strategies must be found for addressing the "memory wall"—the failure of memory bandwidth to keep up with processor speeds.

Much active research is in progress at ORNL and other institutions to address these and other issues to clear the path toward the next great leap in computing power.

Office of Science
U.S. DEPARTMENT OF ENERGY

U.S. DEPARTMENT OF
ENERGY