Oak Ridge National Laboratory
Vol. 37 • No. 2 • 2004
www.ornl.gov/ORNLReview

REVIEW

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

America's New National Leadership Computing Facility

Oak Ridge National Laboratory
REVIEW
Vol. 37 • No. 2 • 2004

Contents

Cover: Thomas Zacharia, ORNL Associate Lab Director for Computing and Computational Sciences, leads the effort to build the world's fastest supercomputer starting with the Cray X1. *Photo by Larry Hamill*

Regaining America's Leadership in

High-Performance Computing

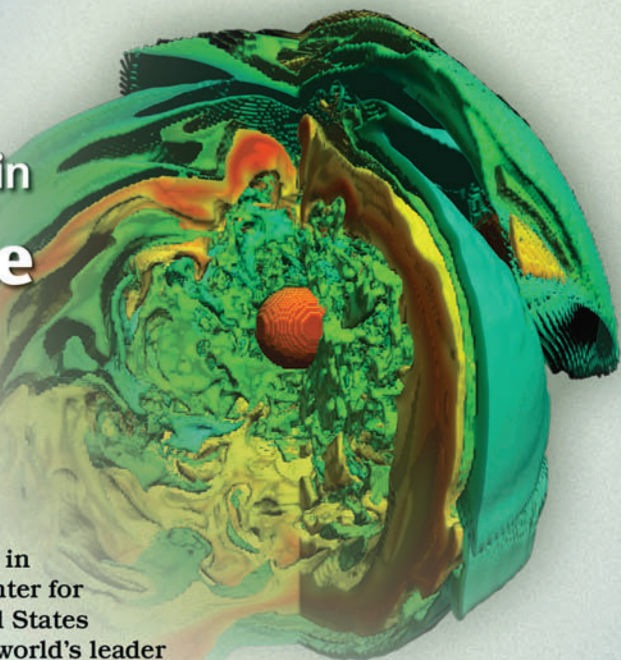

On May 12, 2004, Energy Secretary Spencer Abraham announced a new initiative that promises to reshape the next generation of scientific discovery in America. Addressing an audience at the national Center for Competitiveness, the Secretary vowed that the United States will regain from Japan our historical position as the world's leader in high-performance computing. We will do so by investing in a National Leadership Computing Facility, to be housed at the Department of Energy's Center for Computational Sciences at Oak Ridge National Laboratory. ORNL and its principal partner national laboratories, Argonne and Pacific Northwest, have been tasked by DOE to build a machine with a sustained capacity of 100 trillion calculations per second. The National Leadership Computing Facility will be a new, open, unclassified national resource that will enable breakthrough discoveries in biology, fusion energy, climate prediction, nanoscience, and other fields that will fundamentally change both science and its impact on society.

In addition to national laboratories, research institutions, and academia, ORNL's partnership to support the Department of Energy's effort includes a significant role by the private sector. Three industry-leading vendors—Cray, IBM, and Silicon Graphics—will be engaged as we explore different architectures and experimental approaches on a single infrastructure at ORNL.

Teaming with the Department of Energy, our collective vision is in some respects a simple one. We believe that meeting America's great scientific challenges will require the confluence of various disciplines. We also believe that in each instance, discovery will rest upon the foundation of computational science. Within three years we could be witness to achievements in areas such as climate prediction and protein folding that will literally alter the future of humankind.

This issue of the *ORNL Review* is dedicated to an enormous risk taken by ORNL in the pursuit of that vision. The risk included a decision by UT-Battelle to build a $72 million, privately financed structure to house the Center for Computational Sciences on land transferred from the Department of Energy. The Center is a dramatic statement about ORNL's vision of scientific discovery. Both the vision and the risk were embraced by the state of Tennessee, which provided $10 million for the construction of a second new facility, the Joint Institute for Computational Sciences, to be shared by ORNL and the University of Tennessee on the ORNL campus. Together, these two new facilities represent one of the world's foremost computational programs, with the vision, the personnel, and the infrastructure needed to provide the foundation for the National Leadership Computing Facility.

The following pages provide an overview of the unique resources that the ORNL Center for Computational Sciences will bring to deliver leadership computing for the Department of Energy. We hope the reader will sense our excitement as we anticipate the challenges, opportunities, and discoveries that await us on the horizon.

Thomas Zacharia
Associate Laboratory Director
Computing and Computational Sciences

Leadership-Class Computing for Science

DOE's Center for Computational Sciences at ORNL leads a partnership that will provide America with the world's fastest scientific computing facility.

Energy Secretary Spencer Abraham adds his signature to the first Cray X1 cabinet at CCS along with that of others including President George W. Bush.

Thanks in part to a $300 million modernization program, Oak Ridge National Laboratory has the largest facility for computational science research supported by the Department of Energy's Office of Advanced Scientific Computing Research. DOE's Center for Computational Sciences (CCS) at ORNL is paving the way for groundbreaking scientific and technical accomplishments in the 21st century.

Because ORNL has the people, the computational resources, and the infrastructure, DOE selected CCS to lead a partnership with a goal of creating the world's most powerful supercomputer by 2007. CCS will host the National Leadership Computing Facility (NLCF), and ORNL will execute a plan that will pool the partnership's computational resources to achieve a sustained capacity of 100 trillion calculations per second (teraflops or TF). The NLCF partnership's plan is to surpass the world's current fastest supercomputer, Japan's 40-TF Earth Simulator, by 2005.

The current ORNL Cray X1 computer will be boosted to 20 TF. A 20-TF Cray Red Storm-based machine will be added in 2005 when a key member of the partnership, DOE's Argonne National Laboratory, will install a 5-teraflop IBM Blue Gene computer. In 2006 the100-TF Cray X2 supercomputer will be installed at CCS, and the National Aeronautics and Space Administration and industrial partners will help boost the NLCF's peak computing capacity to 250 TF by 2007.

ORNL has gathered some of the finest experimental researchers and theorists, as well as some of the best analytical and computational tools, in the world. The Laboratory now boasts the world's highest-resolution electron microscope. With the completion in 2006 of the Spallation Neutron Source, Oak Ridge will be the world leader in neutron science. Supporting these and other assets, CCS supercomputers perform 13.4 TF and will quadruple that computational speed in the near future.

ORNL's new computational facilities house not only unclassified leadership-class supercomputers but also state-of-the-art data storage. The 150,000 sq. ft computational sciences building has space to spare for adding numerous cabinets of processors to greatly increase the computational speed and scientific modeling capability of CCS supercomputers.

The CCS also serves as an evaluation center. ORNL scientists evaluate different supercomputer architectures to determine which science codes work best on the flagship supercomputer (Cray X1) and which ones work better on IBM or SGI Altix supercomputers. These experts advise vendors on how to design next-generation supercomputers to improve scientific productivity. Researchers develop software tools that enable CCS supercomputers to run science codes more efficiently.

ORNL's computational facilities are bolstered by state-of-the-art connectivity, with a strong research capability for building even better and faster networks to connect CCS supercomputers with national networks and with links to Atlanta, Memphis, Chicago, the Research Triangle, and other sites. These networks will enable industrial firms to collaborate more efficiently with ORNL researchers on projects of interest to industry. Also at CCS, first-class visualization expertise and equipment help researchers obtain insights from their calculation results and communicate their significance.

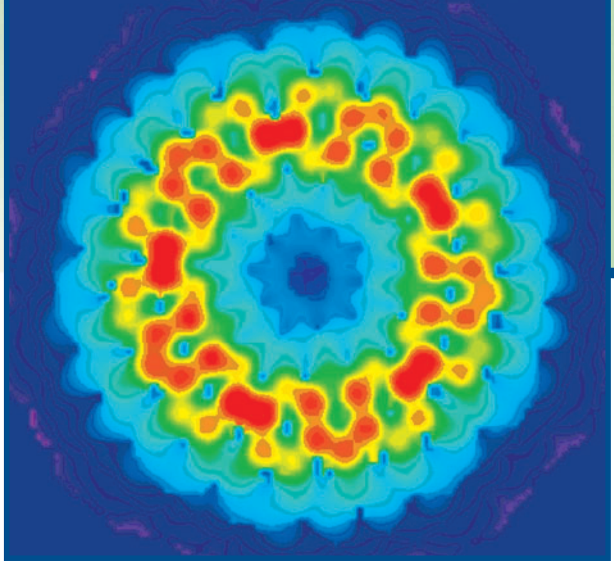

Computer model of the charge density on an equatorial plane of a carbon nanotorus.

Attracting the Best and the Brightest

Many of the world's best computational researchers are coming to Oak Ridge. The promise of new facilities and new research opportunities is attracting bright young computational scientists, such as our new Alston S. Householder Fellow, Jennifer Ryan, and our new Eugene P. Wigner Fellow, Thomas Maier. At the University of Tennessee–ORNL Joint Institute for Computational Sciences (JICS), housed in a new building on the modern ORNL campus, some 30 joint faculty appointments will be available to attract the best and brightest university instructors who conduct research in the computational sciences. More than 40 science and technology staff members, 56 postdoctoral researchers, 46 graduate students, and a dozen joint appointments were recently attracted to CCS.

Combining experiment, theory, and computational simulation, ORNL researchers in collaboration with peers from partner universities and other organizations are conducting outstanding science. Collaborating groups are making discoveries in biology, climate, materials, fusion, nanoscience, and astrophysics. These collaborations extend to dozens of outstanding American universities. For example, ten universities are involved in the TeraScale Supernova Initiative, which is based at ORNL.

ORNL and UT researchers together have delivered science of distinction. Several have received R&D 100 Awards, the Discover Award, the Presidential Early Career Award for Scientists and Engineers, the Gordon Bell Award, and the IEEE Sidney Fernbach Award. One Fernbach Award winner also was elected a member of the National Academy of Engineering.

ORNL researchers have published several papers in *Nature* and *Science* based on work using CCS supercomputers. For example, Ed Uberbacher, researcher with the Computational Biology Institute at CCS, was one of more than 100 contributors to the landmark paper on sequencing the human genome published in the February 15, 2001 issue of *Nature*. His development of the GRAIL gene-finding program used on CCS supercomputers was listed in *Science* magazine's time line depicting the history of the Human Genome Project.

A recent ORNL research effort involving a CCS supercomputer was published in *Science*. Marco Buongiorno Nardelli, assistant professor of physics at North Carolina State University with a joint faculty appointment at ORNL, created a detailed model of the Schottky barrier. Experimentalists Rodney McKee of ORNL and Fred Walker of UT tested the model's predictions. They found that barrier height on semiconductor chips represents more of an opportunity than a problem, opening the way for smaller, faster, and smarter computers.

CCS Partnerships

To serve DOE, the Laboratory's primary customer, ORNL has built enduring partnerships on an international and national scale. ORNL has become, for example, a key resource for the Intergovernmental Panel on Climate Change. Other interagency partnerships include the National Science Foundation, National Aeronautics and Space Administration, National Security Agency, National Nuclear Security Administration, and Department of Homeland Security.

Working closely with the University of Tennessee, ORNL has developed a variety of academic and industrial partnerships through the Joint Institute for Computational Sciences. The institute is the Laboratory's gateway to computational scientists from academia nationwide that have partnerships with UT-Battelle, which manages ORNL for DOE. Academic collaborations also include the Research Alliance for Minorities.

ORNL also has strong industrial partnerships with supercomputer vendors—Cray Inc., Silicon Graphics Inc. (SGI), and International Business Machines (IBM). The 6.4 TF Cray X1 (Phoenix) at CCS is the largest of its type in the world. Add that total to 4.5 TF of the IBM Power4 (Cheetah), 1.5 TF of the SGI Ram, and 1 TF of the IBM Power3 (Eagle), CCS has a collective computing capability of 13.4 teraflops. The Cray X1 has passed a milestone acceptance test and is undergoing evaluation on a suite of scientific computer programs including global climate modeling, high-temperature superconductivity, astrophysics, and fusion energy.

Identifying the Grand Challenges

CCS has four institutes: biology, climate, materials, and fusion. The institutes provide an intellectual home and computational infrastructure for community building. The institutes share several common goals: extended simulations in areas of science important to DOE; repositories of community codes optimized for high-end computing; a testbed for evaluations of new computer hardware and application of innovative software engineering techniques; interactions with CCS's future technologies group to push hardware beyond original vendor design specifications to achieve science missions; workshops to enhance researcher skills and train the next generation of modelers; increased interactions between research scientists and computer scientists and mathematicians; and collaborations to interpret and improve simulation results and to strengthen links between predictive modeling and experimental research.

Most important, CCS institute researchers will interact with the community of scientists in their respective fields to identify the unclassified "grand challenge" problems that can be solved only by CCS supercomputers. The institutes are a key to ensuring that CCS has the synergy of skillful research partnerships and world-class computational technology to meet the challenges of solving national scientific problems.

BIOLOGY: Uncovering Secrets of Living Cells

4

The Computational Biology Institute at CCS will develop software tools to enable understanding of the molecular interactions of protein networks in bacteria and in mice.

Probing microbes to determine what they are made of and what drives them requires more than mass spectrometers, microarrays, and microscopes. Computational models run on supercomputers have been key contributors to our growing understanding of these single-cell organisms.

How Microbes Help DOE

The Department of Energy seeks to understand the diverse range of biochemical pathways that enable microbes to survive under extreme conditions—high temperature, high radiation, and high concentrations of toxic chemicals. DOE is interested in harnessing the genes of microbes whose talents could help DOE meet its missions in environmental bioremediation, climate change, and energy production.

For example, the bacterium *Deinococcus radiodurans* can withstand high doses of radiation because its cells efficiently repair radiation damage. These bacteria also might be able to convert radioactive uranium in storage ponds from a soluble to an insoluble form so that this toxic metal stays put in the sediments instead of dissolving in water that may flow off-site. Thus, it might be possible to harness the genes of *D. radiodurans* for remediating sites with mixed wastes—combinations of radioactive materials and toxic metals. Use of genes with the right abilities from bacteria such as *D. radiodurans* and *Shewanella oneidensi* (studied at ORNL) could potentially save DOE billions of dollars in toxic waste cleanup activities.

Genes and other DNA sequences contain instructions on how and when the cell should build proteins. Proteins form complexes, or molecular machines, that do the work of the cell.

Certain bacteria in the ocean and on land absorb carbon dioxide from the atmosphere and perform photosynthesis. Harnessing the genes from these bacteria would help DOE achieve its goal of finding ways to halt the buildup of atmospheric carbon dioxide from energy production to counter global warming.

DOE is also interested in microbes that produce clean fuels, such as methane, methanol, and hydrogen. ORNL researchers are focusing on *Rhodopseudomonas palustris*, whose genes might be harnessed to produce hydrogen for possible use in power-producing fuel cells for cars and buildings in the envisioned hydrogen economy. ORNL researchers and their colleagues are studying these microbes as part of DOE's Microbial Genome Program and Genomes to Life (GTL) Program.

Computational biologists working with supercomputers at DOE's Center for Computational Sciences (CCS) at ORNL have a long history of contributing to an understanding of microbial genes. They have identified many genes in bacterial, mouse, and human genomes and have computationally analyzed the human genome using ORNL-developed gene-finding computer programs. ORNL researchers also have written and used assembly programs and analysis tools to produce draft sequences of the 300 million DNA base pairs in chromosomes 19, 16, and 5 for DOE's Joint Genome Institute (JGI) as part of DOE's Human Genome Project. Some have analyzed 60 complete and draft microbial genomes containing 230,000 genes and used computers to keep up with JGI sequencing rates of a genome per day. Others have predicted the structures of proteins from amino-acid sequences using an ORNL-developed protein-threading computer program.

Bioinformatics specialists from ORNL and the University of Tennessee have written algorithms and developed other tools to make it easier for biologists to use computers to find genes and make sense out of the rising flood of biological data. These data are produced in studies of biochemical pathways and processes, cellular and developmental processes, tissue and organism physiology, and ecological processes and populations. Through ORNL's user-friendly Genome Channel web site, its Genomic Integrated Supercomputing Toolkit, and CCS supercomputers, the international biology community, including pharmaceutical industry researchers and academics, have easily obtained genetically meaningful interpretations of their DNA sequences and other data. ORNL's web site, especially the pages supporting the Human Genome Project, is the focus of approximately 150,000 sessions per month in the biological community.

ORNL computational scientists are now working with research partners in the GTL Program to develop high-throughput computational tools for rapidly analyzing, interpreting, and communicating the volumes of data on, for example, five novel proteins that the partners discovered dur-

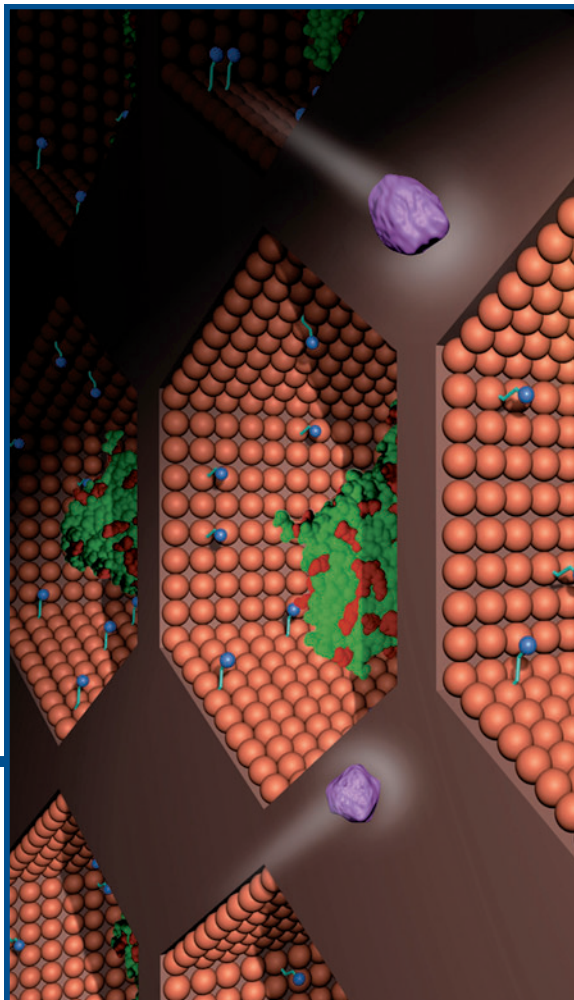

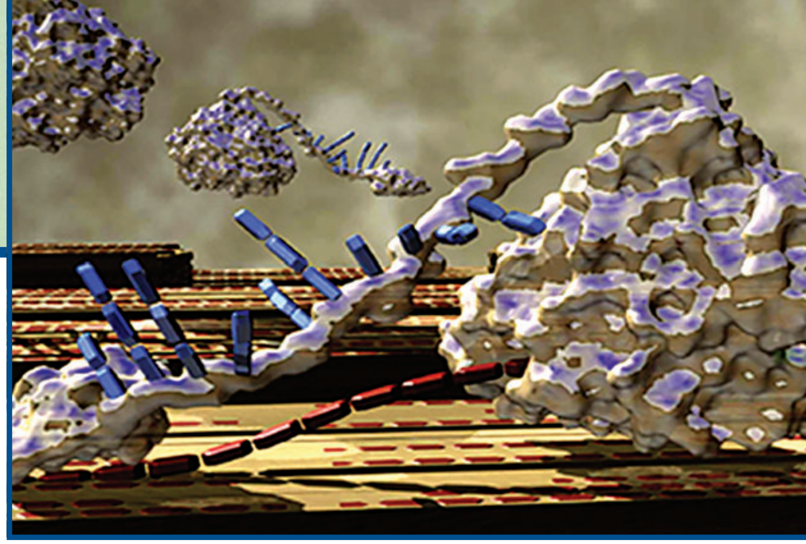

Cellulose breakdown in plant cell walls.

ing research on *R. palustris*. Analytical tools and algorithms will be needed to determine how proteins interact, stimulate chemical reactions, and move materials inside and out of cells when exposed to different conditions. Proteins turn genes on and off, regulating their activities. When a bacterial cell is moved from clean water to polluted water, proteins in cells capture environmental signals and turn on genes that make special proteins enabling the cell to adapt to a new environment. Computer models will be built to characterize this cascade of changes.

Institute Missions

The Computational Biology Institute (CBI), led by Jeff Nichols, has been formed as a multidisciplinary partnership to develop and provide innovative computational algorithms, analysis tools, and data and hardware infrastructure to enable a scientific understanding of the molecular interactions typical of networks of proteins in complex microbial and metazoan systems—primarily bacteria and mice. ORNL traditionally has conducted research on mice to determine the genetic effects of radiation and toxic chemicals on mammalian systems; radiation and toxic chemicals are by-products of weapons development and energy production, which have long been missions of the U.S. government. CBI will analyze, model, and simulate molecular interactions and networks of interactions among proteins and cells from microbes and mice.

What are the CBI focus areas? One is microbial genome analysis—determining which genes are present in each genome. Another is mass spectrometry analysis—modeling data from mass spec experiments to determine which proteins are made in the cell and when they are used. Another focus area, molecular interaction image analysis, investigates which proteins interact with each other, and when and where. CBI scientists will also use molecular machine modeling, docking, and dynamics to determine which molecular machines are made to do the work. A final focus area is molecular interaction networks modeling and simulation, which describes the web of interactions that transmit information to control the cell.

To describe these bio-molecular interactions involving networks of cells and biochemical pathways, CBI will foster interactions and networking among researchers who need access to supercomputers to better understand data produced by experimentalists at ORNL and at universities. Most of the research is sponsored by DOE and the National Institutes of Health, which provides funding for neuroscience studies by members of the Tennessee Mouse Genome Consortium with whom ORNL mouse biologists work. Some small research projects at CBI involving single principal investigators are sponsored by the National Science Foundation.

CBI comprises researchers largely from ORNL's Life Sciences, Chemical Sciences, Environmental Sciences, and Computer Science and Mathematics divisions. These researchers also collaborate with researchers from the University of Tennessee and other universities.

Computational research planned for the future will require leadership-class scientific computing. This capability should enable researchers to simultaneously track 100 moving proteins in a live microbe with the help of imaging technology and to meet a GTL goal of completely characterizing a microbe in a year.

By combining experimenters' analytical capabilities with the mathematical and simulation capabilities of CBI, the biology community will have a better understanding of the function of large macromolecular complexes, the control of gene expression, cell membrane dynamics, metabolism, and signaling and environmental responses. Single cells are very small, but the complexities of their workings and interactions demand large networks of interacting researchers using very large computers.

Courtesy of Pacific Northwest National Laboratory

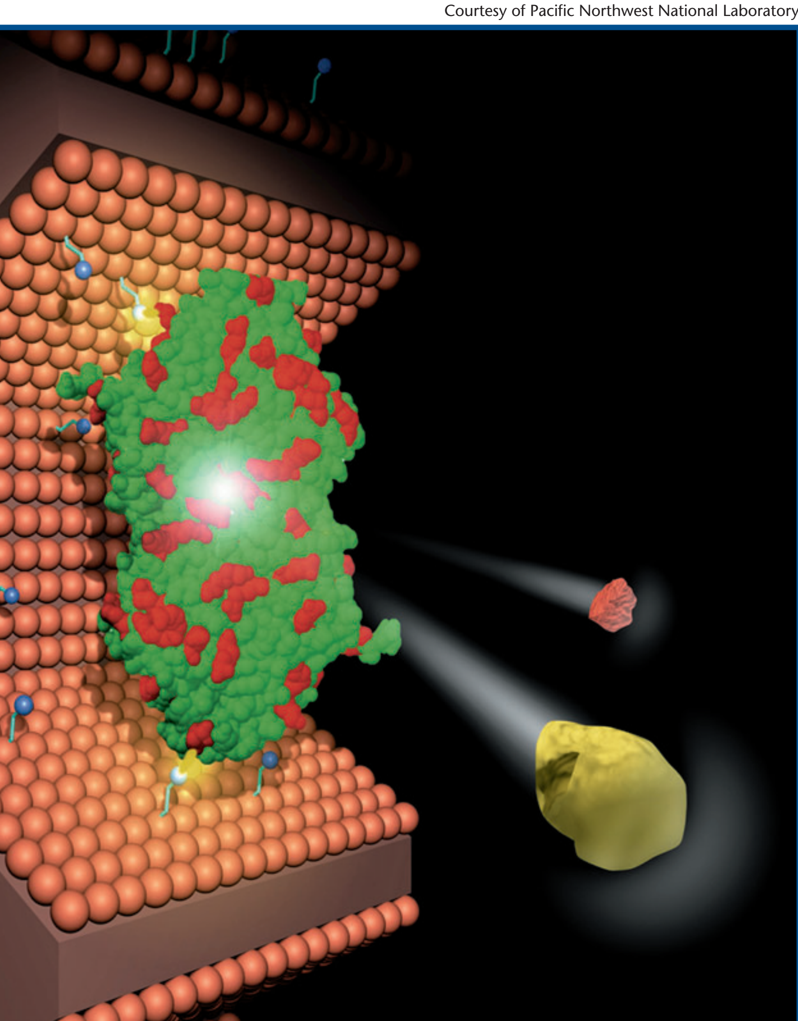

Enzyme (green) embedded in a synthetic membrane that increases the enzyme's stability and activity. The enzyme converts toxic materials (purple molecules at left) into harmless substances (yellow and red molecules at right).

CLIMATE AND CARBON RESEARCH: Glimpses of Global Warming

Advanced climate models and faster supercomputers must be developed to ensure that policymakers have better global and regional predictions of future climate and its effects.

As greenhouse gases accumulate in the atmosphere, many questions arise concerning how fast and in what ways Earth's environment will change. For example, in the United States, will increased emissions of carbon dioxide from coal combustion in the 21st century make the Southeast wetter or drier over the next 100 years? Will changes in temperature and moisture conditions make certain U.S. regions more vulnerable to insect-borne diseases? By the year 2100, will the world's glaciers be largely melted and will some low-lying coastal lands be flooded by rising sea levels?

Detailed answers to questions like these being asked by policymakers and researchers will require more sophisticated climate models, faster supercomputers to run them, and larger data storage repositories that can be networked nationwide to store and exchange large data files. The Department of Energy's climate science mission is to improve the scientific basis for assessing potential consequences of climate change on decade-to-century time scales. The goals of the Climate and Carbon Research Institute (CCRI) at DOE's Center for Computational Sciences at ORNL are focused on this endeavor. DOE's computer centers at ORNL and elsewhere are making great progress in understanding future climate—a predicted average of weather patterns—under different scenarios. Already scientists are running global warming "experiments" on supercomputers.

Because carbon is a waste product of fossil energy combustion in power plants and energy use by cars and trucks, DOE's primary interest is the effect of carbon dioxide emissions on climate warming. DOE wants researchers to represent the carbon cycle correctly in their models to provide an objective framework for investigations of interactions of processes and feedback involving the atmosphere, land, and oceans.

CCRI, led by David Erickson and John Drake, is coupling carbon and climate models with the help of researchers in ORNL's Environmental Sciences Division. Earlier under the guidance of Drake, CCS researchers modified an important climate model so it could run on massively parallel supercomputers. The model predicts interactions between the atmosphere and land and between the atmosphere and oceans. CCRI has collaborated with the National Center for Atmospheric Research (NCAR) in Boulder, Colorado, and with Los Alamos National Laboratory in modeling interactions among processes generated by and affecting the land, atmosphere, and oceans.

Climate modelers face difficult challenges. For example, changes in the chemical makeup of the atmosphere can have confounding effects. Greenhouse gases such as carbon dioxide emitted from the land to the atmosphere can absorb infrared radiation from Earth's surface and prevent the escape of heat. However, sulfate aerosols from coal-fired power plant emissions can have a cooling effect, moderating the temperature signal and altering weather patterns.

DOE, which has funded research that has led to important breakthroughs in climate modeling, is pushing researchers to make climate simulation models more comprehensive and more detailed over the next 20 years. These models will consume a record number of compute cycles using the fastest computers ever built.

Oak Ridge and Japan

Climate simulation is an international research activity of importance to policymakers in a variety of nations. The Japanese government has invested in climate science by building the world's largest supercomputer dedicated to fine-scale climate simulations.

The Japan Earth Simulator, ranked number one in supercomputer power and speed on the latest Top500 list of Jack Dongarra of the University of Tennessee, has a theoretical computing speed of 35 to 40 teraflops, or 35 to 40 trillion calculations per second. Significantly, the scientific codes used on supercomputers at DOE sites such as the CCS have made America competitive with Japan in advancing the understanding of climate. For its second priority for future facilities, DOE's Office of Science is supporting development of a supercomputer at ORNL that will surpass Japanese computational abilities in climate prediction and other areas.

U.S. computing sites also are collaborating with Japan in climate prediction. In 2001 the Intergovernmental Panel on Climate Change (IPCC), relying partly on re-

sults from global climate models run on the world's supercomputers, concluded that "there is new and stronger evidence that most of the warming observed over the last 50 years is attributable to human activities." The IPCC projected that, by the end of 2100, the global average temperature of the earth could rise by 2.7 to 10.4°F.

ORNL's CCRI has joined Japan, NCAR, and the National Energy Research Scientific Computing (NERSC) Center in California in running simulations to provide answers for IPCC's Fourth Assessment, due out in 2007. The CCRI simulations require an average of 25% of the compute hours on the IBM Power4 at CCS. The four participants are all using the same code on scenarios with different carbon concentrations in the atmosphere.

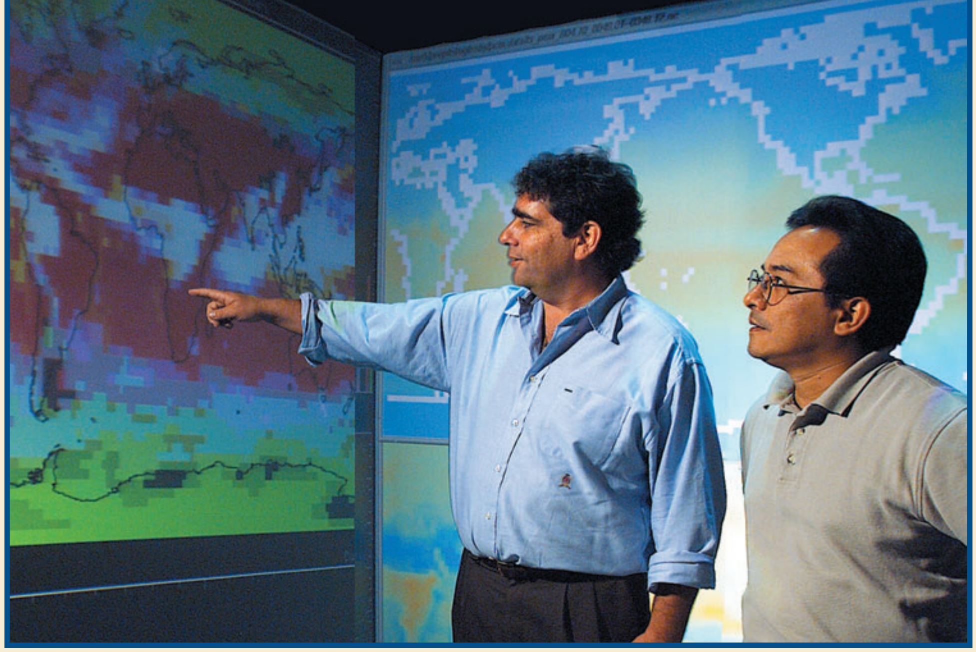

David Erickson (left, with José Hernandez) is director of CCS's Climate and Carbon Research Institute. He is a member of the scientific planning team of the Surface Ocean-Lower Atmosphere Study (SOLAS). He co-authored a chapter on the interactive effects of ozone depletion and climate change, which appears in the United Nations Environment Programme document entitled Environmental Effects of Ozone Depletion and Its Interactions with Climate Change: 2002 Assessment.

NERSC is predicting the climate and its effects through 2100 assuming that the carbon content stays the same as it is today—370 parts per million (ppm). NCAR is predicting the effects of a more optimistic scenario in which humankind finds a way to stop the buildup of atmospheric carbon when it reaches 876 ppm by 2100.

About a half dozen researchers at CCRI, aided by 20 staff members in CCS, have collaborated in building and optimizing the Community Climate System Model on the IBM Power4. The scenario simulated at ORNL assumes that humankind stabilizes atmospheric carbon concentrations at 550 ppm by 2100 and then reverses the buildup by sequestering carbon and replacing the fossil fuel economy with a hydrogen economy in which buildings and transportation vehicles are powered by fuel cells and fossil power is replaced with fission and fusion power. The IPCC runs at ORNL will generate 30 terabytes of data. Much of the data will be retained at the High Performance Storage System at CCS in Oak Ridge.

IPCC runs must be completed in 2004, so that the results can be analyzed and papers can be written, peer-reviewed, and published by 2006. The schedule will allow IPCC participants to assess the papers and prepare a report by 2007.

Drake says that the simulation at CCS for the IPCC is the largest set of coordinated runs he has ever seen for any project. Indeed, the climate simulation may consume the largest amount of computing resources ever used on a single set of codes with a single objective.

There are many questions about carbon and the climate that policymakers would like answered. For example, is the U.S. terrestrial system in the East taking up more carbon from the atmosphere than it is giving off? The forest that has grown in New England in the last century is storing carbon, but when the forest leaves fall, considerable carbon comes out of the decaying leaves and re-enters the atmosphere.

This and many other questions involving interactions and feedbacks in the climate system remain unanswered. To answer them, historical datasets must be better developed, the climate models must be improved to make more accurate predictions about regional as well as global climate, and the supercomputers must be faster to accommodate more data and model calculations. The CCRI hopes that leadership-class computing will provide the detailed answers that researchers and policymakers need.

Sulfate particles from the combustion of fossil fuels contribute to the haze in Great Smoky Mountains National Park and exert a cooling effect that is considered in climate models.

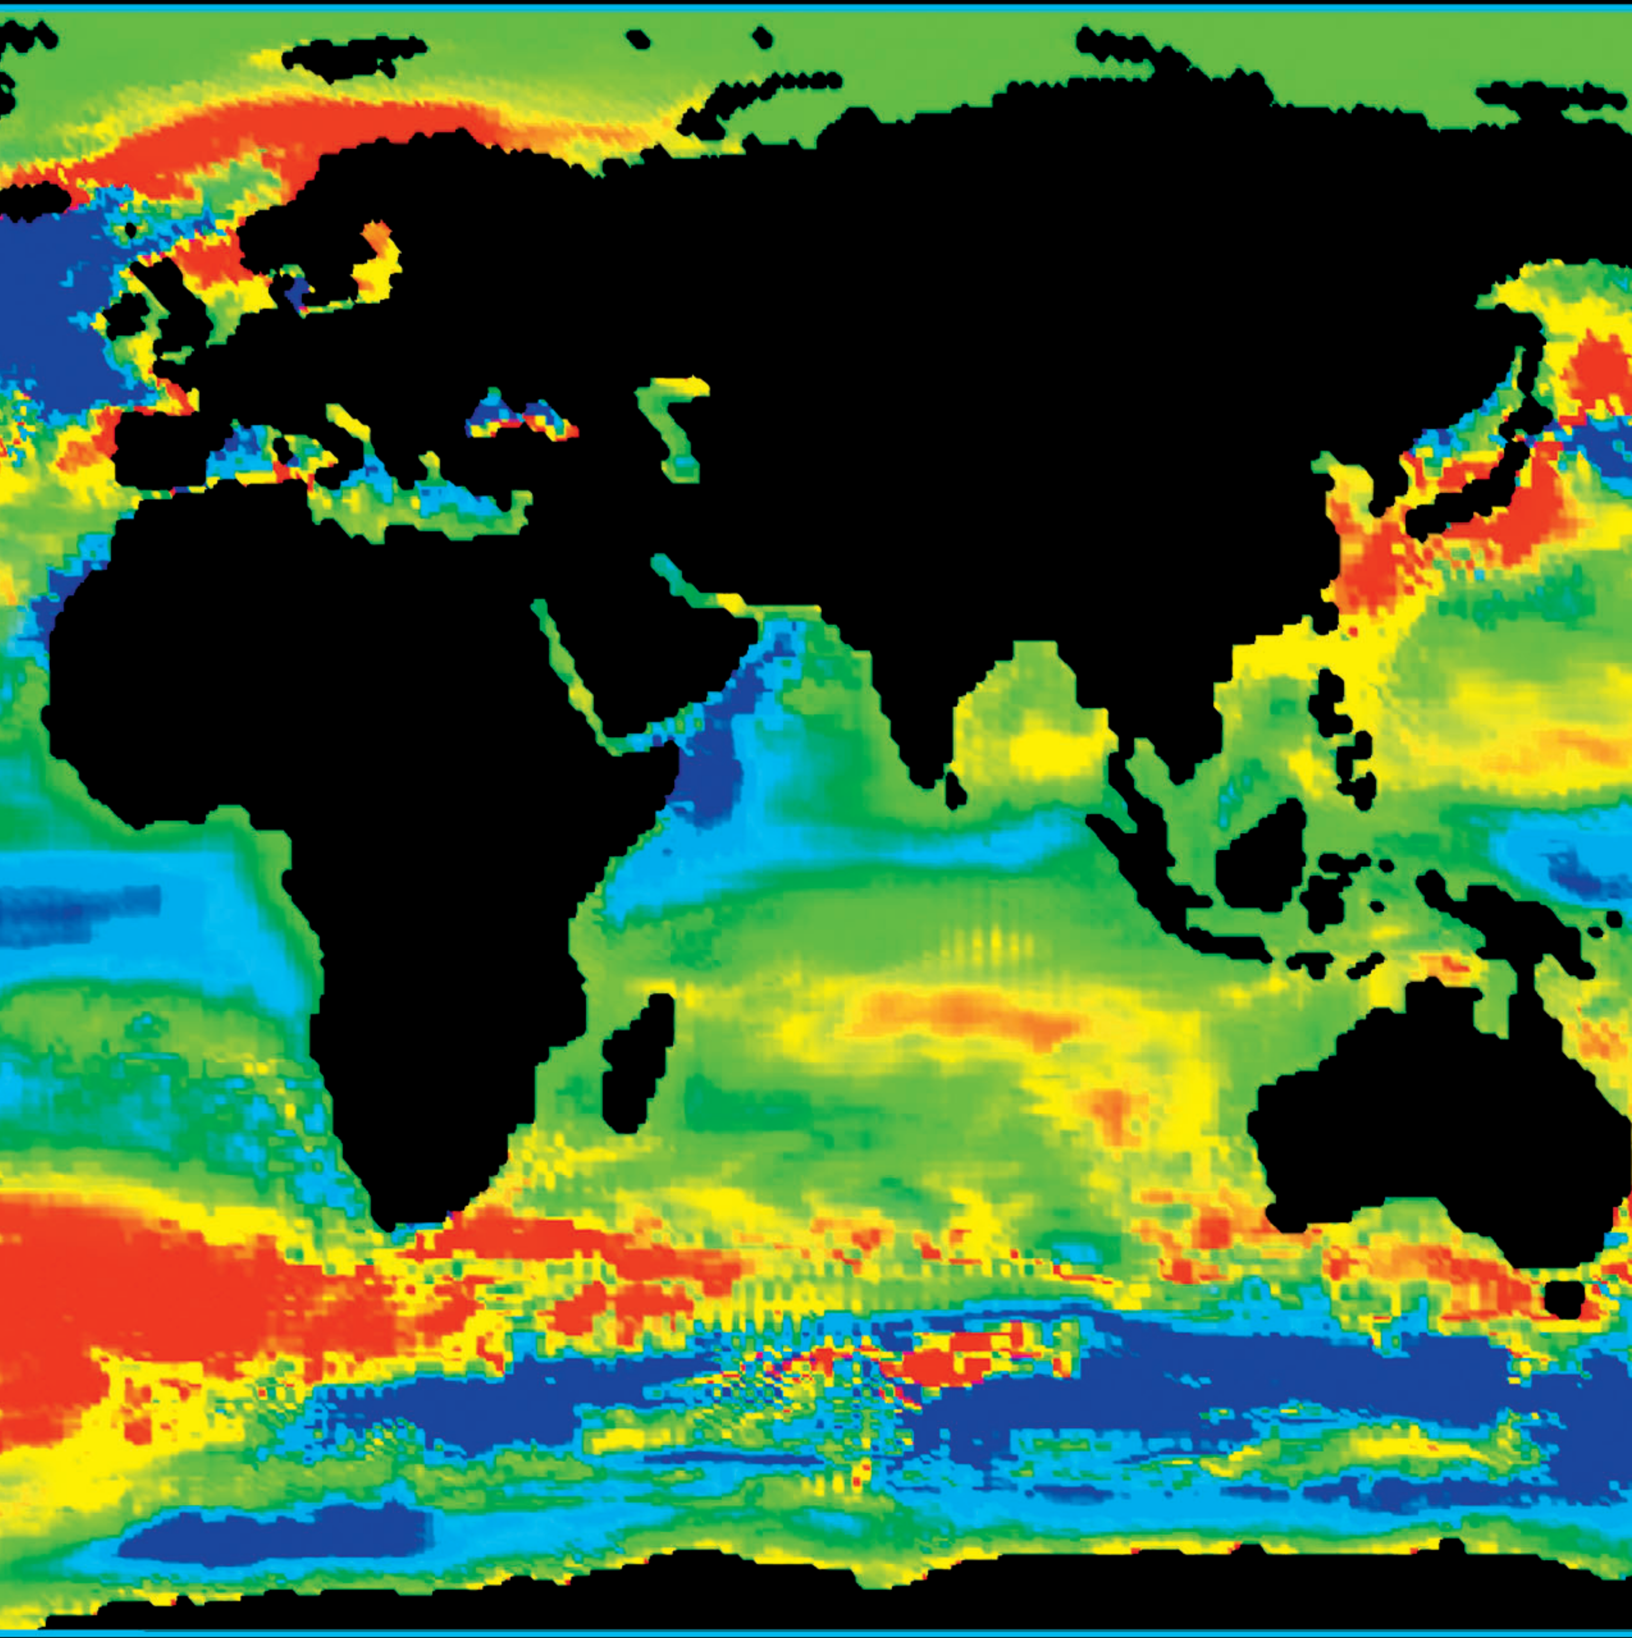

Depiction of the amount of carbon flux between the oceans and the atmosphere (peak heights) and amount of biological activity (color).

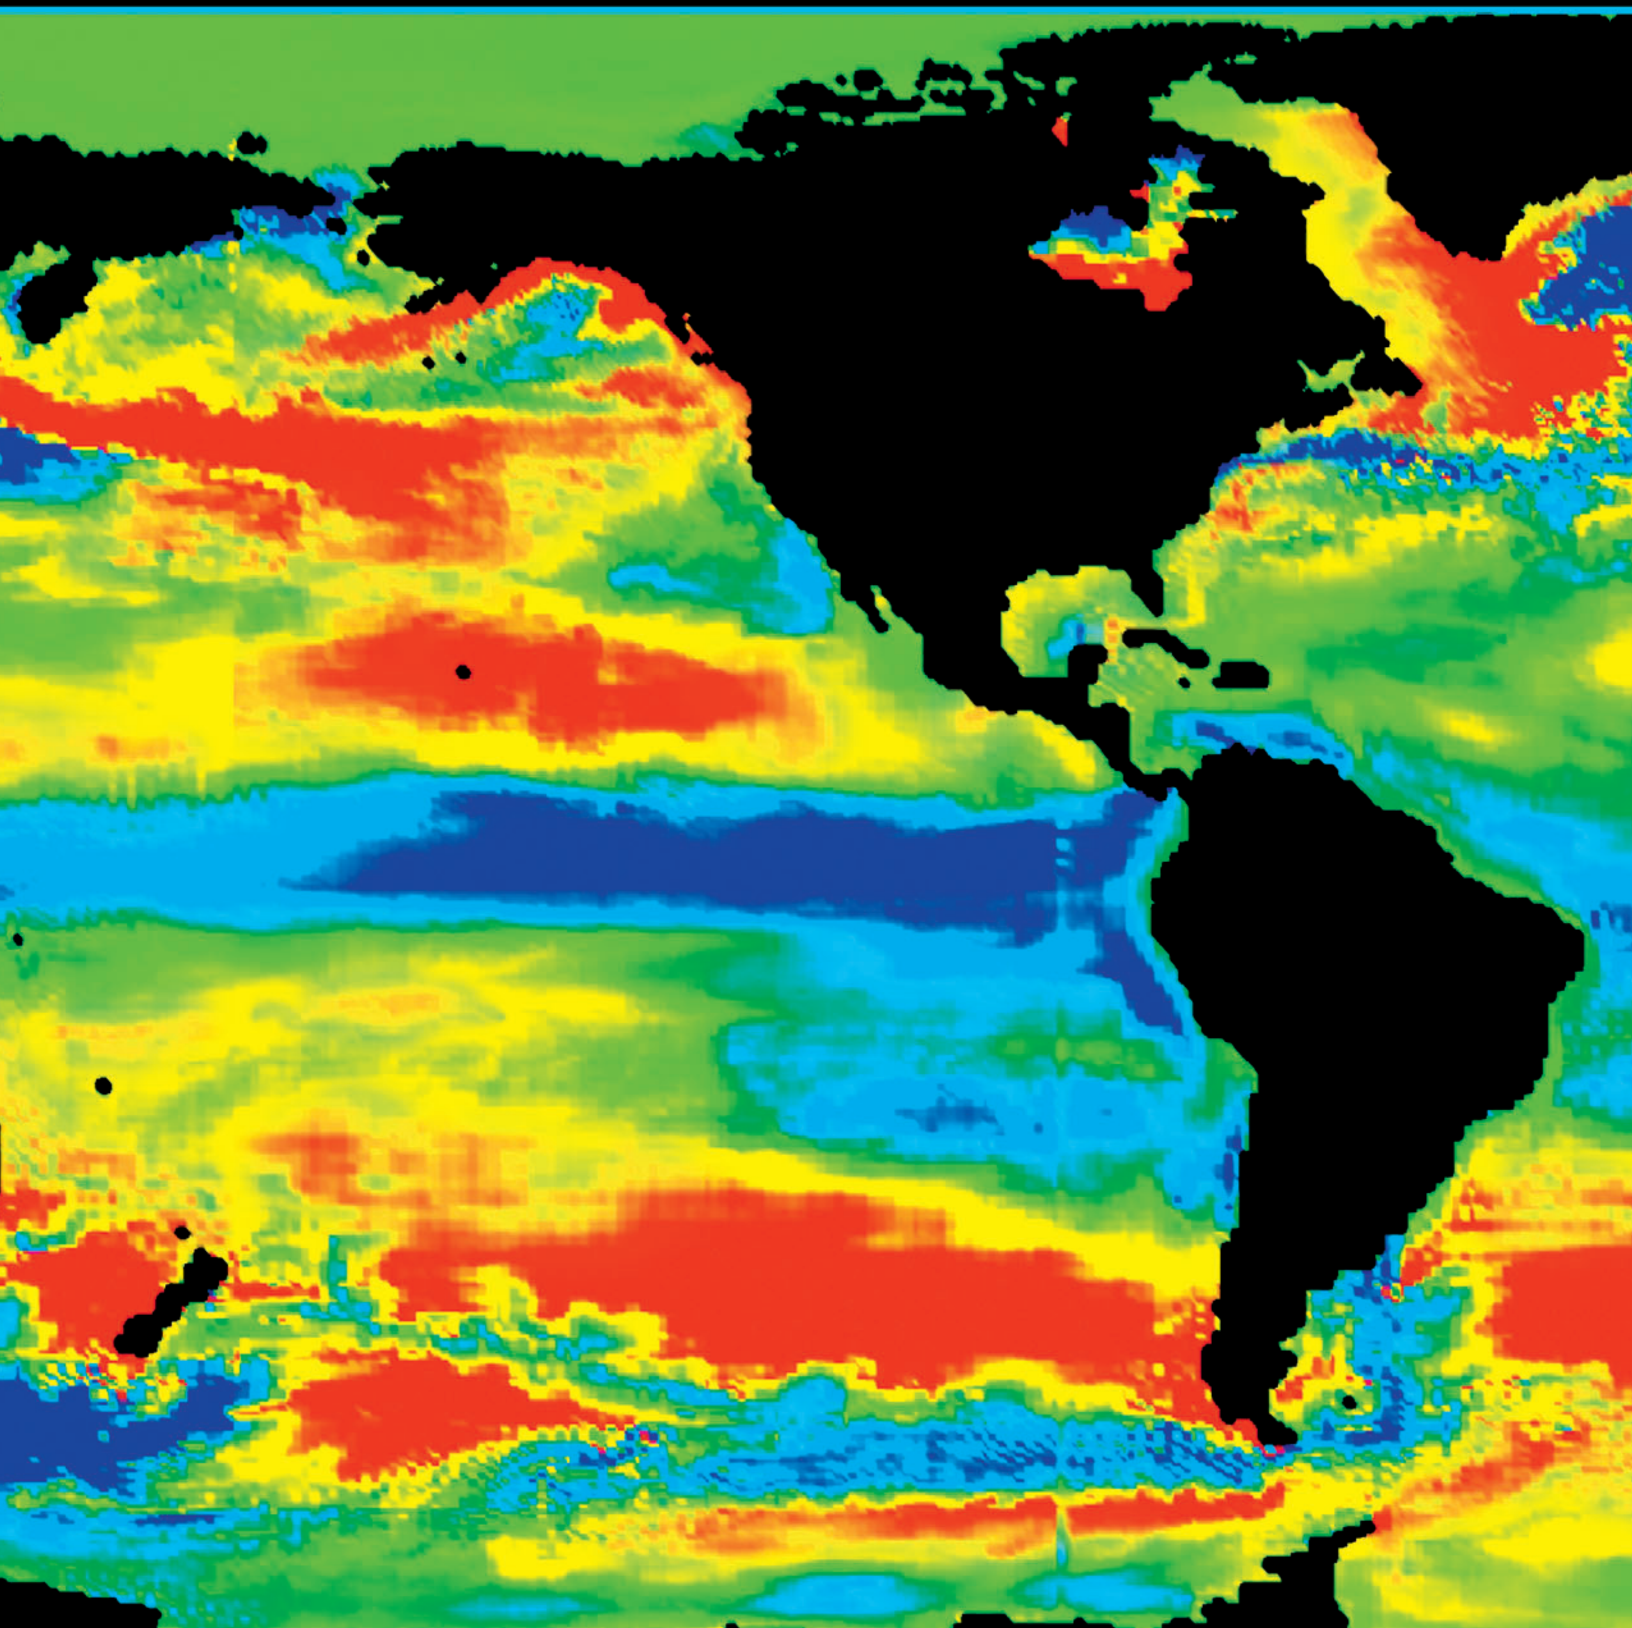

FUSION: Progress through Computation

Ultrascale scientific computing, combined with the nation's best physicists and computer scientists in research partnerships, will be needed to solve difficult problems to hasten the design of a cost-effective fusion power plant.

If we continue to burn fossil fuels for energy at the current rate, they will last only another few hundred years. In the context of civilization, the fossil fuel era is drawing to a close. In addition, it would be wise to reduce our combustion of oil, gas, and coal because the process produces pollutants that are detrimental to our health and carbon dioxide that could change our climate in undesirable ways. One possible future source of electricity for the world is fusion energy. When the fusion process that produces the energy of the sun and stars is harnessed on the earth, we will have a virtually unlimited supply of clean energy.

Achieving this goal will occur when a hot plasma gas containing heavy hydrogen atoms—deuterium obtained from the ocean and tritium bred by the fusion process itself—is heated to a high enough temperature, squeezed to a high enough density, and held in this state by magnetic fields for a long enough time to sustain fusion reactions. The heat produced from these reactions could then be used to drive a turbine and generate electricity.

Which geometry should external magnetic fields have to confine the plasma in a way that ensures sustained fusion reactions? Can radiofrequency (rf) heating—high-frequency radio waves generated by oscillators outside the doughnut-shaped plasma vessel—possess the right set of frequencies to heat and control the plasma? Heating involves transferring rf energy to the charged particles in the plasma, which in turn collide with other plasma particles, boosting the temperature of the bulk plasma. Controlling the plasma involves suppressing its instabilities so fewer particles crash into the plasma vessel wall and cool the plasma. The most efficient way to answer these questions and to identify all the parameters needed for heating and controlling fusion plasmas in an economic fusion device is through fusion simulation on supercomputers and validation of these simulations by experiments.

In the past five years, the theory group in ORNL's Fusion Energy Division in collaboration with computer scientists and applied mathematicians in ORNL's Computer Science and Mathematics Division, has developed powerful computer codes to simulate fusion plasmas in two and three dimensions and predict their responses to simulated magnetic fields and radiofrequency heating. During that time, the Department of Energy's Center for Computational Sciences at ORNL obtained powerful supercomputers, including the IBM Power4 (Cheetah), which is a primary production resource for DOE's Scientific Discovery through Advanced Computing (SciDAC) program, and the SGI Altix, which has the extra advantage of having a large memory that holds 2 trillion bytes of data.

ORNL's fusion theorists have been using these two supercomputers for their SciDAC-sponsored research and have also ported their codes to the Cray X1 system at CCS. They are glad DOE is providing funding to build a supercomputer at ORNL that will exceed the computational speed record currently held by Japan's Earth Simulator supercomputer. A supercomputer with this much capability will hasten the solution of design problems related to future fusion power devices.

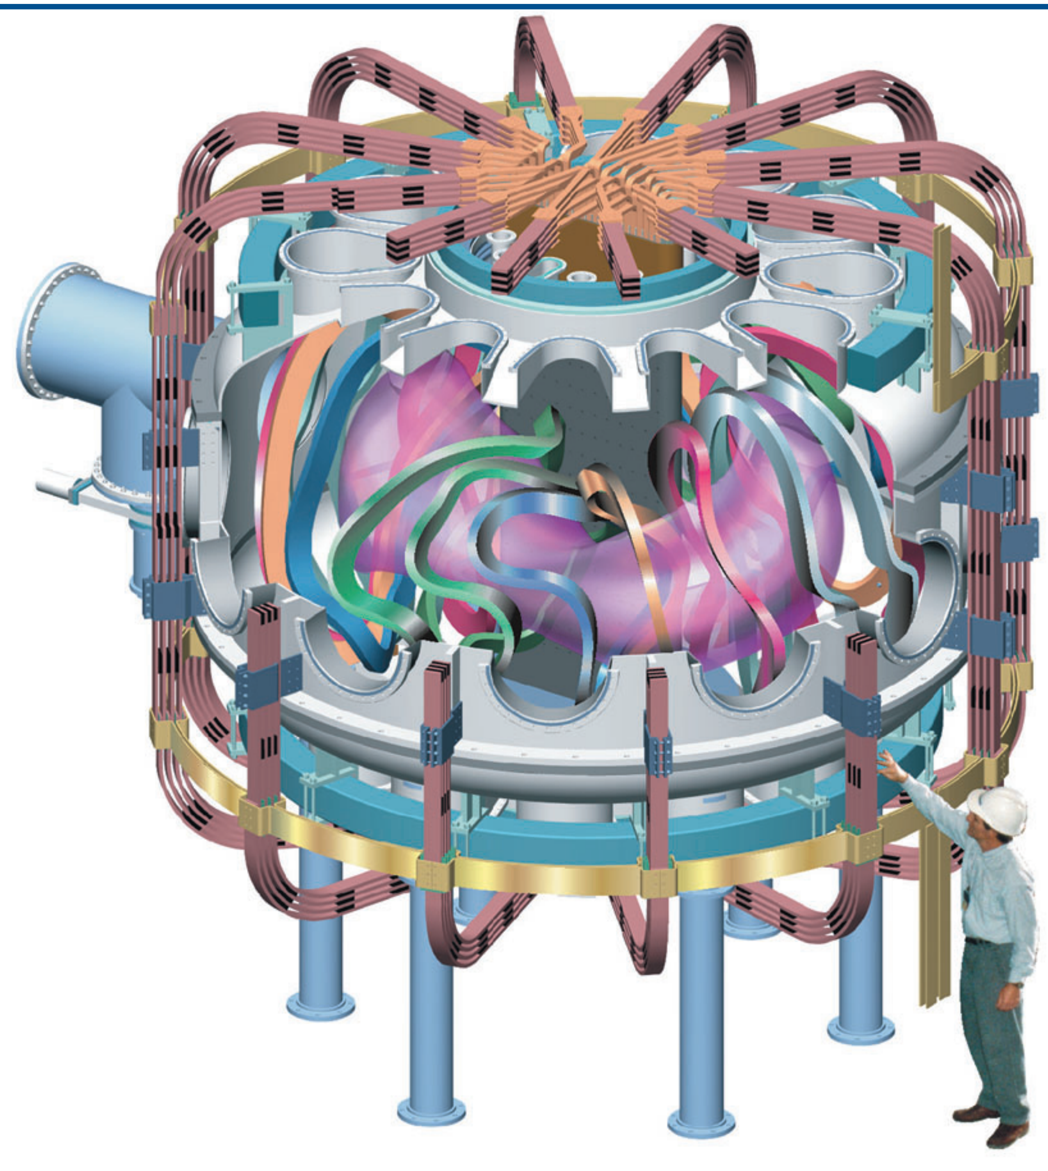

The Quasi Poloidal Stellarator was designed with the help of a computer optimization model that used an extensive suite of physics design tools. The QPS is expected to be built by 2008.

ORNL's All Orders Spectral Algorithm computational model for analyzing plasma-wave interactions is a success story in integrating physics and software.

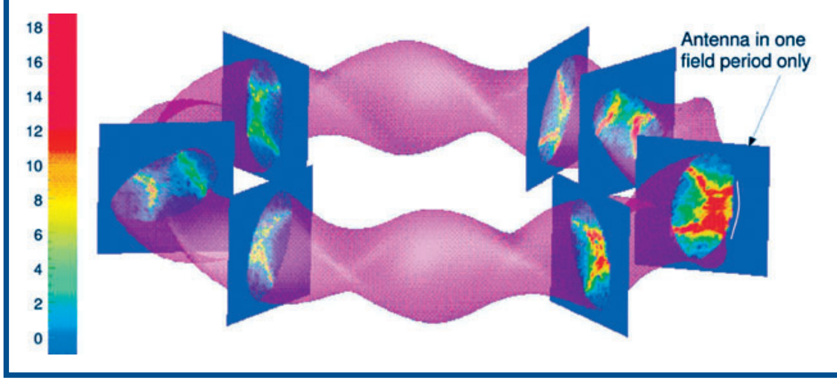

ORNL and ITER

One DOE goal is to demonstrate a sustained, self-heated (burning) fusion plasma, in which the plasma is maintained at fusion temperatures by the heat generated by the fusion reaction itself. To attain this ambitious goal, DOE recently declared the International Thermonuclear Experimental Reactor (ITER) to be its number one near-term-science facility of the future. ITER, the first facility capable of producing a sustained burning plasma, will demonstrate the scientific and technological feasibility of fusion energy. To fully exploit ITER's capabilities and to guide the design and development of subsequent fusion power devices will eventually require a predictive capability for fusion systems. To accomplish this goal, DOE will launch a major effort to advance state-of-the-art computational modeling and simulation of plasma behavior. ORNL's fusion simulation work will contribute to this effort and ORNL's involvement with ITER will allow fusion researchers to use data from the ITER project to validate simulation programs and reliably design future fusion reactors. The fusion community at ORNL in collaboration with other major U.S. research institutions, such as DOE's Princeton Plasma Physics Laboratory (PPPL), will leverage American successes in fusion simulation to regain a U.S. leadership role in this $5 billion international project.

Simulation—A Shortcut to Solutions

Fusion energy research has a long history of employing supercomputers to solve highly complex mathematical equations. Physicists know the equations, and applied mathematicians and computer scientists bring expertise in solving them. In addition to fusion specialists, ORNL is home to a team of world-class physicists, mathematicians, and computer scientists who collaborate with their peers from other national laboratories and international organizations. To solve complex scientific problems, the research community must develop new algorithms, optimize computer codes for modern computer architectures, and help vendors design and build ultrascale supercomputers.

The ultimate goals of fusion simulation are quite ambitious—to predict, reliably, the behavior of plasma discharges in a toroidal magnetic fusion device on all relevant time and space scales. This effort would bring together into one framework all the codes and models that presently constitute separate disciplines within plasma science. Scientific and numerical issues awaiting solutions are those associated with the coupling of codes with different natural time- and length-scales.

In the near term, ORNL scientists will be carrying out massively parallel computations on the IBM Power3 supercomputers at DOE's National Energy Research Scientific Computing Center (NERSC) and on the IBM Power4 and Cray X1 supercomputers at CCS. These fusion simulation efforts, including collaborative projects funded by SciDAC, focus primarily on grand challenge problems in four areas: transport processes in plasmas, wave/plasma interaction, advanced stellarator designs, and simulation of plasma evolution.

Advanced stellarator designs, for example, have the potential for leading to a fusion reactor that is smaller and more economically attractive than existing stellarators and would eliminate the potentially damaging plasma disruptions that plague conventional research tokamaks (including ITER). Stellarators employ 3-D magnetic fields as opposed to the 2-D fields used for tokamaks. This added complexity has historically limited our ability to design optimized systems. New software and faster computers have lifted this constraint. ORNL, in collaboration with PPPL, has developed a computer optimization model that employs an extensive suite of stellarator design tools. These tools enabled the design of compact stellarator facilities valued at $100 million, including the Quasi-Poloidal Stellarator that DOE plans to build at ORNL later this decade.

Simulation of plasma evolution will help validate recent evidence that non-linear processes called self-organized criticality exists for fusion plasmas. This approach could provide insight into the vulnerability of complex systems such as power grids and communication networks.

Fusion Energy Institute

In the past, tough physics problems could be attacked only piecemeal. Now, thanks to advanced computer simulations, more powerful supercomputers, and improved data storage, these once intractable problems are now within the realm of being solved systematically. To speed this process, ORNL experts in physics, applied mathematics, and computer science have formed a Fusion Energy Institute in CCS. The Institute will draw on the expertise of the world's best fusion scientists to identify grand challenges in fusion that can be solved only by using the leadership-class computing capabilities available at CCS. Because CCS will become a site for the world's fastest scientific supercomputer, it is more likely that today's intractable fusion problems will be solved, enabling a more rapid design of cost-effective fusion devices for power production.

MATERIALS: Exploring and Modeling 21st Century Materials

Computational materials scientists at ORNL are using the high-performance computing infrastructure at CCS to explore superconductivity and magnetic nanostructures.

The 1986 discovery of high-temperature superconductivity sparked the quest for room-temperature superconductors that could transmit electrical current without heat losses and without the need for an expensive coolant such as liquid helium. Room-temperature superconductors could make possible ultra-efficient power transmission lines, practical electric cars, and superconducting magnets that could bring high-speed levitated trains and smaller, more efficient, and less costly rotating machinery, appliances, particle accelerators, electric generators, and medical imaging devices.

High-temperature superconductors are being used commercially. A few urban utility companies have tripled their capacity to carry power simply by replacing existing underground cables with liquid-nitrogen-cooled superconducting cables. Cellular telephone towers have extended their reception range and call-handling ability with superconducting signal filters.

Understanding Superconductivity

No one understands why certain copper-oxide materials exhibit high-temperature superconductivity, says ORNL Corporate Fellow Malcolm Stocks, co-leader of the Materials Research Institute (MRI) at DOE's Center for Computational Sciences at ORNL. MRI co-leader Thomas Schulthess is working on this problem with Thomas Maier, Eugene P. Wigner Fellow at CCS, and solid-state physicist Mark Jarrell of the University of Cincinnati. When Maier was a postdoctoral fellow at Cincinnati he, Jarrell, and other colleagues were working on a theoretical understanding of high-temperature superconductivity. Jarrell holds that a microscopic understanding of why the current layered oxide (high-T_c) materials are superconducting will lead to the design and synthesis of new room-temperature superconductors.

Jarrell is a co-developer of a new theoretical approach, called the Dynamical Cluster Approximation (DCA). Here, a cluster of atoms, within which the complex quantum mechanical interaction between electrons are treated essentially exactly, is embedded into an effective medium that accounts for the effects of the rest of the material in a computationally feasible way without compromising the mathematical rigor of the theory.

Using a massively parallel implementation on the supercomputers at CCS and at Pittsburgh Supercomputing Center, Jarrell and Maier studied the two-dimensional (2-D) Hubbard model. This model has been widely accepted as the theoretical framework for capturing the physics underlying high-T_c superconductivity. Because structurally the high-T_c materials are a series of copper-oxide planes, with no apparent interactions between them, researchers believed they could be modeled as 2-D systems.

However, based on their work, Jarrell and Maier began to suspect that the 2-D Hubbard model may not provide a complete description of high-T_C superconductivity. When Maier came to ORNL as a Wigner Fellow, he and his CCS colleagues were able to pursue this suspicion using the power of the newly installed Cray X1, a machine ideally suited to performing the complex DCA calculations. The most recent calculations are confirming Maier's suspicions that the strictly 2-D Hubbard model is insufficient, a conclusion that runs counter to most conventional wisdom. If true it will be necessary to improve the underlying model by, for example, introducing some 3-D features such as coupling between the 2-D copper oxide planes.

The CCS team was able to carry out this work because the Cray X1 makes the DCA calculations up to 35 times faster than the computers used previously. This improved performance allows the modeling of larger DCA clusters containing up to 64 (8 × 8) sites, rather than only 4 (2 × 2)

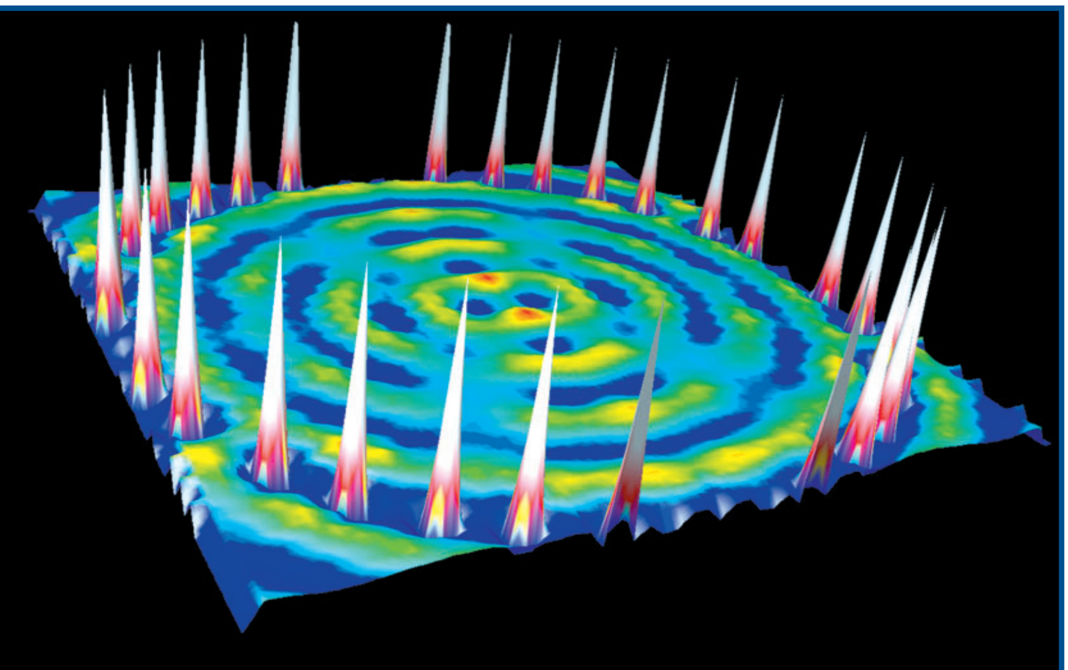

Using a CCS supercomputer, researchers calculated the magnetic structure of a quantum corral nanostructure, which consists of magnetic iron atoms deposited on a copper surface that "corral" copper electrons.

Thomas Maier, recently named a Eugene P. Wigner Fellow at ORNL, used the Cray X1 supercomputer at CCS to confirm his suspicion that a widely used computer model was insufficient to describe high-temperature superconductivity.

sites previously possible. According to Schulthess, when the Cray system has 4 to 10 times its current capability, the ORNL-Cincinnati collaboration will likely have a model that describes the physics of high-temperature superconductors, which eventually will enable computational design and prediction of new materials that are superconducting at room temperature.

Nanomagnetic Modeling

The revolution in magnetic storage that has occurred over the past 40 years has always depended on the timely introduction of new materials that allow researchers to cram increasing amounts of digital data into smaller and smaller regions. In the very near future, whole new classes of materials will be needed if the magnetic storage industry is to maintain its current rate of increase in storage density for computers and digital media. Increasingly, basic research is turning to nanoscience for solutions. Nanoscience covers materials whose size is a few thousand to a few million atoms. At the nanoscale, materials frequently exhibit new and unexpected phenomena. Understanding and harnessing these “emergent” phenomena could have untold future benefits to science and technology.

For some time Stocks, Schulthess, and others at CCS have been modeling magnetism and electrical transport in magnetic materials as part of the effort to develop new and improved magnetic materials. Now, the group has also turned to studying the new frontier of the properties of magnetic nanostructures. Recently, using the IBM supercomputers at CCS, the group performed the first fully self-consistent *ab initio* calculations of the magnetic structure of scientifically fascinating nanostructures called quantum corrals.

A quantum corral is a magnetic nanostructure in which magnetic atoms, are positioned in the shape of a stadium on a copper surface. In the mid-1990s Don Eigler of IBM used a scanning tunneling microscope (STM) not only to make rings of magnetic atoms but also to measure the electron density, the concentration of the copper surface’s electrons that are corralled by the magnetic atoms—hence, the name “quantum corral.” On a (111)-facet of copper, some of the electrons form an effective two-dimensional electron gas just above the surface. When the magnetic atoms are deposited, the 2-D electron gas is perturbed, causing standing waves that are captured by STM images.

The CCS researchers have not only calculated the oscillations in the surface charge obtained experimentally but have also predicted oscillations in the magnetic density inside the corral as well as the state of the magnetic atoms themselves. These predictions present a challenge to next-generation STMs that will measure the magnetic structure of nanostructures. This achievement was made possible by use of codes developed by the University of Tennessee’s Balazs Ujfalussy and his colleagues in Budapest and Vienna.

Materials Institute

The rapid pace at which new complicated materials must be introduced into technology, as well as the drive to miniaturize devices, increasingly dictates a more systematic approach to design and fabrication that is built on a firm theoretical underpinning. Here computation can be expected to play a central role. Because simulation can study features close to the atomic limit that are difficult to access experimentally, computational materials science can accelerate the development of new materials.

To facilitate these goals, the Materials Research Institute has been created within CCS. MRI’s missions are to build and maintain a community of leading computational materials scientists and to involve them in the development of computational methods, algorithms, and simulation software, as well as in the early evaluation of new computer hardware at CCS. MRI’s long-term goal is to enable the design of advanced materials by using leadership-class computing envisioned by DOE.

The computed molecular structure of a nanoscale assembly of six polymer nanoparticles.

ASTROPHYSICS: Simulating Supernovae

Answers to huge scientific questions may come from findings made in realistic simulations of exploding stars on the world's fastest supercomputers.

During the catastrophic death throes of massive stars, known as core-collapse supernovae, many elements were created, including those necessary for life on the earth. How and why these stars that were greater than 10 times the mass of our sun and that had evolved over millions of years died explosively in a few hours are mysteries that scientists cannot solve in laboratory experiments. However, simulations on supercomputers hold out hope of unraveling the secrets of supernovae.

The five-year TeraScale Supernova Initiative (TSI), an ORNL-centered, multidisciplinary computational initiative involving 10 universities, has already made a major discovery about core-collapse supernovae, thanks to funding from the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) program. The discovery was made by three researchers, including TSI leader Tony Mezzacappa, an ORNL computational astrophysicist. One of the supercomputers used was the IBM Power3 supercomputer at DOE's Center for Computational Sciences at ORNL.

Mezzacappa's team discovered that the shock wave—thought to stall while trying to propagate through the stellar core until re-energized by neutrino heating—can become unstable and change shape to look more like a cigar and less like a sphere than previously thought. If the shock wave generates the explosion and if this "supernova accretion shock instability" (SASI) occurs, the characteristics of a supernova explosion could be fundamentally altered. Observational evidence suggests that the light from a core-collapse supernova explosion is polarized. One explanation for the observed polarization could be that the explosion is shaped like a cigar. The SASI may also be the underlying mechanism whereby neutron stars born in supernovae are boosted to mean velocities of 450 kilometers per second.

TSI also discovered the importance of modeling both the microphysical and macrophysical aspects of core-collapse supernovae. Researchers found that the dynamics of an exploding star's gross characteristics depend on the nature of the atomic nuclei of iron and heavier elements that compose the stellar core. Theorists believe the star's explosion is caused by a shock wave that arises when the star's iron core collapses on itself, compressing its subatomic particles to the point where they repel each other and force the core to rebound. The rebounding force may be further energized by neutrino heating—the interaction of neutrinos in the star with its core. Neutrinos are nearly mass-less, radiation-like particles believed to be responsible for powering supernova explosions.

This TSI finding resulted from a merger of nuclear structure theory and astrophysics theory pertaining to core-collapse supernovae. Use of more sophisticated models of nuclei in the core to calculate the interactions of neutrinos with the stellar-core nuclei led to changes in the astrophysics models that will strongly impact scientists' understanding of the mechanisms responsible for supernova explosions and the synthesis of elements. To test new nuclear models that provide input to supernova models, neutrino-nucleus interaction experiments have been proposed for the accelerator-based Spallation Neutron Source being built at ORNL. The reason: neutrinos emerging from its target will have the same energy as supernova neutrinos.

ORNL researchers in TSI have found that one of their radiation transport codes that simulates neutrino heating and one of their hydrodynamics codes that simulates the core's turbulent fluid and rotation run 10 to 20 times faster on the CCS Cray X1 than on the IBM Power3 and Power4 supercomputers.

With leadership-class scientific computing, TSI will develop and run three-dimensional models that more realistically simulate the complex processes that drive a supernova and synthesize elements. As simulation needs escalate, huge leaps will be needed in data management, data analysis, networking capabilities, the development of software tools, and visualization. Simulation results will provide the theoretical foundation for understanding physics data from premier experimental and observational facilities, as well as insights into combustion, climate, fusion energy, and nuclear medicine. For scientists, simulating the stars will have benefits on the earth.

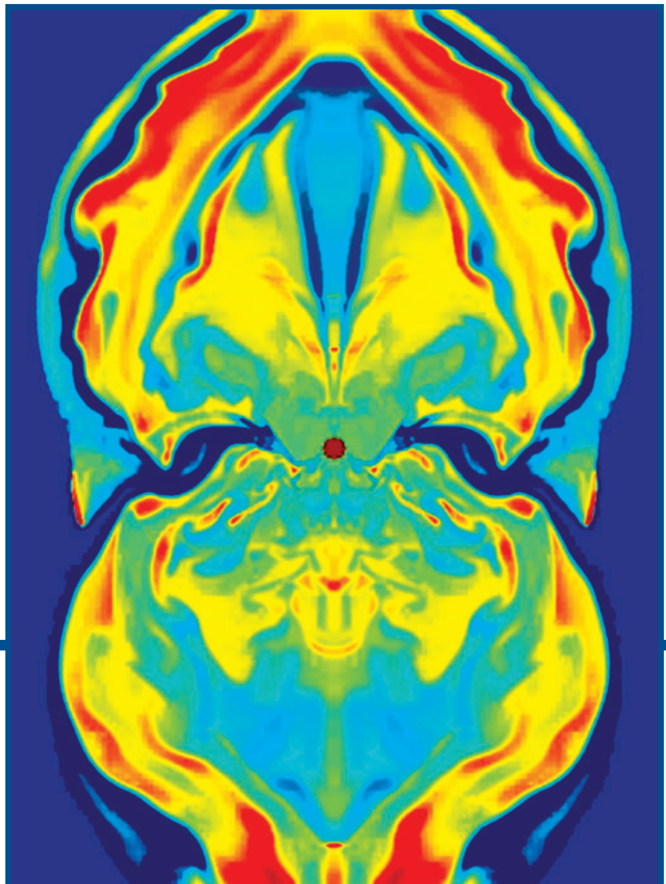

Two-dimensional simulation on a CCS supercomputer uncovers the existence of a supernova shock wave instability.

Robert Harrison received the 2002 Sidney Fernbach Award.

CHEMISTRY: Tomorrow's Molecular and Nanoscale Devices

Large multidisciplinary collaborations that combine chemistry, physics, mathematics, and computer science will tackle future challenges in molecular electronics and nanotechnology.

Chemistry—the study of molecules—is the science of the everyday world. For the world at large, molecules are the fundamental units of matter. An understanding of the structure, interactions, and reactions of molecules is thus of critical importance to a wide range of phenomena, from the fate of contaminants in the environment, through the production of plastics from crude oil, to the occurrence and treatment of genetic diseases.

Chemists are now about to cross a remarkable threshold and expect a dramatic expansion in their ability to make reliable predictions about molecular structure and processes. This sea change is due to the confluence of advances in theoretical, computational, and experimental capabilities, allowing chemists to understand and characterize matter at detailed atomic and molecular levels. By integrating chemists' capabilities of synthesis and characterization with computational modeling and simulation, it will soon be possible to use computation to design molecules to do what we want, and to control how we make them.

The computational chemical sciences group in the Department of Energy's Center for Computational Sciences at ORNL is working to fulfill this vision. Real-world chemical problems are complex and a multi-disciplinary approach embracing chemistry, materials science, solid-state physics, mathematics, and computational science is essential, as is a strong connection with experiment. Excellent examples of the interplay of theory and experiment are advances being made at ORNL in molecular optics and electronics. Using ink-jet printing techniques proposed by computational chemists, an experimental group led by Mike Barnes of ORNL's Chemical Sciences Division, including researchers from the University of Tennessee (UT) and Georgia Tech, coaxed stretched molecules of polyphenylene vinylene to form an array of glowing antennas on a glass substrate. These semiconducting polymer nanostructures have many potential applications including optical wires. Members of the computational chemical sciences group, including Bobby Sumpter, Bill Shelton, and Jack Wells, are performing detailed simulations to better understand this newly discovered nanostructure.

The group is led by Robert J. Harrison, the principal architect of the Northwest Computational Chemistry Software (NWChem) code for massively parallel computers, for which he received the IEEE Computer Society's 2002 Sidney Fernbach Award. DOE's Scientific Discovery through Advanced Computing (SciDAC) program is supporting two of the group's many projects. Harrison is the principal investigator of the project entitled "Advanced Methods for Electronic Structure," which develops and uses new multi-scale methods for fast and accurate numerical solution of electronic structure problems, such as predicting how atoms are arranged in a molecule. These new approaches can deliver greater precision than previous methods, even for very large systems.

The other SciDAC project, entitled "An Integrated Approach to Multi-scale Modeling of Molecular Electronic Devices," is led by Peter Cummings of Vanderbilt University. This multidisciplinary effort involves researchers from four universities (including UT) and ORNL. The goals of this computational project are to develop and apply tools to understand how to construct and control molecular electronic devices by self-assembly of molecules on surfaces. Molecular electronic devices may someday replace silicon-based devices in future computers. High-performance computer programs are being developed by Vincent Meunier to use quantum theory to predict, for instance, the current across a single molecule. Fundamental new theoretical models are also being developed to confront a profound challenge to chemists—the wide range of time and length scales, which cause severe problems as molecular devices are extended to the nanoscale and beyond.

To probe nanoscale devices, design new catalysts, and develop clean energy sources, chemists nationwide see a critical need for multi-scale methods as well as leadership-class computing to provide them with a "computational microscope."

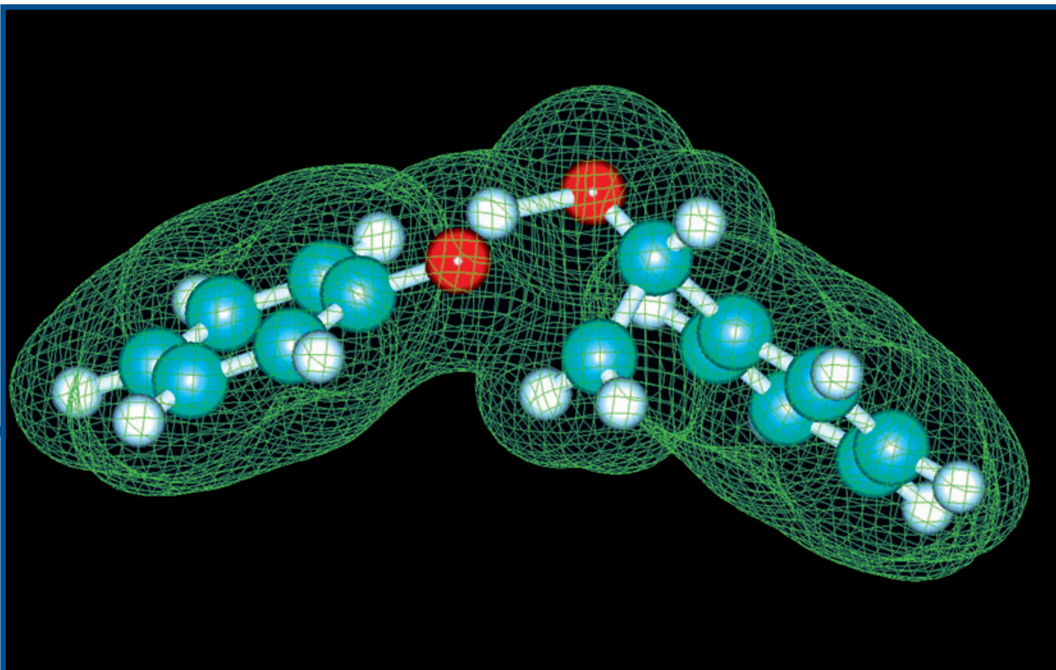

Computer model of a hydrogen bond–mediated transition state of a chemical reaction. The electrostatic potential isosurface (the lines that resemble a green net) is overlain on the molecular structure, all computed using quantum chemistry.

ORNL's $300 million modernization program, funded by a federal, state, and private partnership, includes two new facilities for the Center for Computational Sciences.

ORNL has America's largest unclassified scientific computing facility, with the personnel, infrastructure, and services to deliver leadership-class scientific computing for DOE.

The Department of Energy's Center for Computational Sciences at Oak Ridge National Laboratory is housed in the nation's largest unclassified computing facility. The new facility, which is part of ORNL's $300 million campus modernization project, is a 170,000 sq. ft building that includes 40,000 sq. ft for DOE's unclassified scientific computing resources. The spacious new building houses a state-of-the-art computer center and more than 400 staff, including researchers and computational scientists. The facility provides the networking, visualization, data storage capabilities, and electrical and mechanical infrastructure required to support a leadership-class scientific computing system.

Networking the Nation

CCS is the hub for ORNL's fast network connections. CCS has several major networks in place and under construction, enabling it to connect with DOE laboratories, university partners, and other supercomputing centers. Currently, two networks are operating. CCS links with other DOE facilities through the Energy Sciences Network (ESNet), which has an OC-12 connection—that is, a data transfer speed of 622 megabits per second (mb/s). DOE plans to upgrade ESNet to OC-48, or 2.5 gigabits per second (Gbs).

CCS also connects to the university user community through the National Science Foundation (NSF) Internet-2 network, which has a data transfer speed of OC-192, or 10 Gbs. The high-speed network connection was funded by ORNL as an integral part of its strategy to work more closely with a variety of university partners. CCS staff are building four new networks, including the TeraGrid and the UltraScience Net (see the article starting on p. 22).

Providing the Power

The electrical and mechanical infrastructure of CCS is designed specifically for a leadership-class computing system. The computer center can accommodate the delivery of up to 12 megawatts (MW) of electrical power to the computers and building infrastructure. Initially 4 MW was provided to the computers and an additional 4 MW was made available for the remainder of the building. The building is programmed for expansion with the pad space and conduit in place to easily ramp up to to 8 MW for the computer systems, bringing the CCS total to 12 MW.

The Tennessee Valley Authority's electrical infrastructure provides very reliable power to the CCS. However, CCS also provides a 500-kilowatt (KW) uninterruptible power system that is backed up by a 750-KW diesel generator. The back-up power is insufficient to supply the large computer systems, but is more than enough to power all the networks, disk drives, storage systems, and server infrastructure that support the CCS.

To remove the heat generated by all this power, the facility has three 1200-ton chillers to provide cold water used to cool

the computer systems. The chillers are fully redundant so that if any one fails, the other two can carry the load of the building. The cooling infrastructure is designed to allow easy expansion of the capacity without having to shut down the operating computer center.

Evaluating Machine Performance

In a recent review by DOE and independent panels, the review committee referred to the CCS as the "pathfinder for the nation" in scientific computing. The CCS is the nation's only unclassified center that has deployed large-scale supercomputer systems representing the three major computer architectures. By having these three systems available in one place, the CCS is in a unique position to help computational scientists select the best architecture for each type of scientific application. Scientific applications—computer codes that enable supercomputers to run calculations on huge amounts of data to make scientific discoveries—will run faster or more efficiently on one computer architecture than on other architectures. Because CCS is an evaluation center, researchers test different codes to determine the best architecture for each code. The results guide both hardware vendors and software developers in improving current and next-generation systems.

The Cray X1 system, which has a scalable vector architecture, is the newest system being evaluated by the CCS evaluation team and users. The Cray X1 uses custom-designed vector processors (chips) to get high performance for scientific codes. The Cray-designed processors fabricated by IBM are linked together using a high-performance shared memory interconnect technology. In other words, all the processors share the same memory where they fetch and deposit data. The machine, named Phoenix, has 512 multistreaming processors, each of which can carry out as many as 12.8 billion operations per second, making the performance of the total system as high as 6.4 trillion operations per second. Early evaluation results from the Cray X1 system reveal performance ranges from 5 to 25 times faster than the same number of processors on the IBM Power4 system for applications that can take advantage of the vector processors.

The second supercomputer architecture at CCS is the cluster of symmetric multi-processor systems. These commodity off-the-shelf computers are linked together using a commodity or proprietary interconnect technology across which scientific applications are distributed. In the CCS, this architecture is represented by the IBMPower4 (Cheetah) and its predecessor, the IBMPower3 (Eagle). The IBMPower4 is capable of 4.5 trillion operations per second. The IBMPower3 is capable of 1 trillion operations per second. These supercomputers

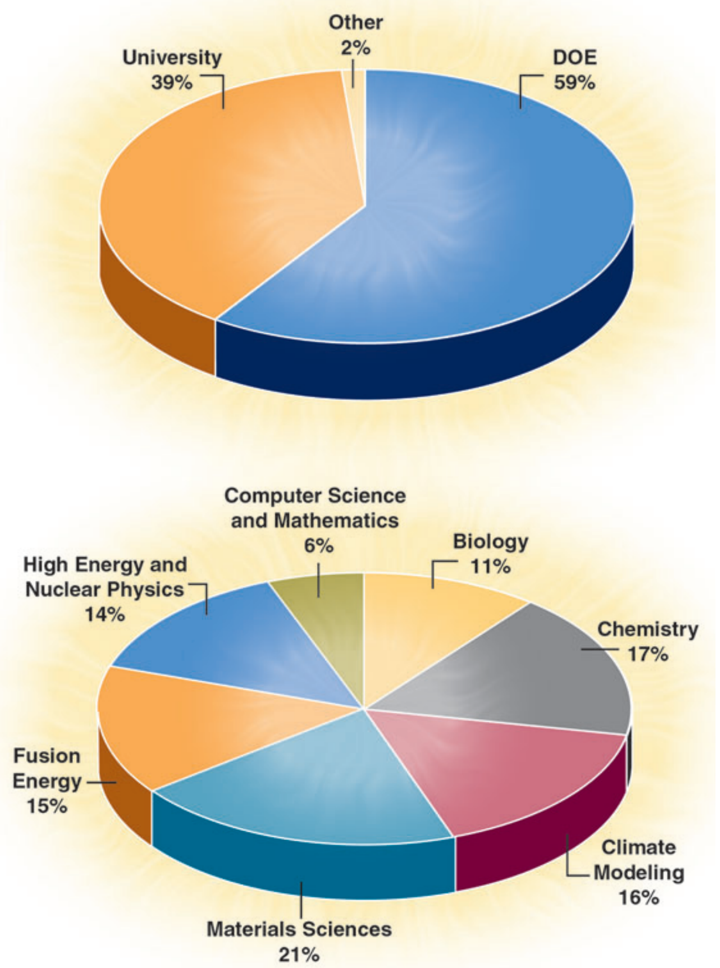

Usage of CCS supercomputers in terms of processor hours by (top) DOE and university researchers, and (bottom) scientific discipline.

CCS offers 40,000 sq. ft for supercomputers and data storage in support of DOE's open scientific computing.

use IBM's custom processor chips. Both of these systems are stable and provide highly reliable supercomputing capability for many projects within DOE.

In 2003 these cluster-based IBM machines were the primary production resource for DOE's Scientific Discovery through Advanced Computing (SciDAC) program. Some 31 million MPP processor hours of IBM machine time were delivered to CCS users in DOE's SciDAC program in the past year. Researchers utilized more than 60 million MPP processor hours of IBM machine time for all DOE science applications.

CCS staff members are evaluating IBM's next-generation interconnect technology, which is the method used to link processors together. IBM Power3 employs the first-generation technology called "SP switch" and the IBM Power4 supercomputer uses the second-generation Colony technology. ORNL researchers are evaluating for IBM its third-generation Federation interconnect technology, which will be used on its Power5 supercomputer now under development. Use of this faster internal network in the IBM Power4 will quadruple the speed of messages inside this machine.

The third supercomputer architecture in the CCS is the shared memory system using commodity microprocessors. The SGI Altix, which has 256 Intel Itanium2 ™ processors, can expand up to 512 processors. The most remarkable feature of this computer, called Ram, is its memory, which holds 2 trillion bytes (2 terabytes) of data. By contrast, the Cray X1 and the IBM systems each have only about 1 terabyte of data in memory. The first supercomputer in Oak Ridge, the Cray XMP installed in 1985, had one-millionth the memory of the SGI Altix.

The shared-memory programming model of the SGI Altix has several advantages. For example, the model makes it conceptually easier for developers of scientific applications to design their codes to run on such a machine. The model also enables certain kinds of applications to run very efficiently. ORNL researchers are seeing excellent results on this system in computational chemistry, global climate modeling, and computational biology.

The SGI Altix supercomputer at CCS has a huge memory. It can hold up to 2 trillion bytes of data.

Cray X1—A New Capability for Science

On August 15, 2002, DOE selected ORNL to test the effectiveness of new Cray architectures in solving important scientific problems in climate, fusion, biology, nanoscale and magnetic materials, and astrophysics. Under the program, ORNL acquired the first few nodes of a Cray X1 supercomputer system.

The Cray X1 system installed at ORNL is the first U.S. computer to offer vector processing and massively parallel processing capabilities in a single architecture. The system is designed to be scaled up to provide scientific applications with performance greater than can be achieved at currently available U.S. computers. Phoenix has 256 multi-streaming vector processors (MSPs). Each MSP has 2 megabytes (MB) of cache memory. Four MSPs form a node that has 16 gigabytes (GB) of shared memory.

In 2002 CCS procured the first 32 processors of the Cray X1 to evaluate the processors, memory subsystem, scalability of the architecture, and software environment to determine the supercomputer's suitability for solving complex scientific problems. CCS and Cray have been evaluating the Phoenix processors to predict its expected sustained performance on key DOE applications codes. The results of the benchmarks show that the Phoenix's architecture is exceptionally fast for most operations and that it markedly improves the performance of several scientific applications.

For example, large-scale simulations of high-temperature superconductors run 25 times faster on the Phoenix system than on an IBM Power4 cluster using the same number of processors. A fusion application requires 16 times more processors on an IBM Power3 cluster to achieve the same performance as that of Phoenix.

The best performance of the parallel ocean-circulation program (POP v1.4.3) for climate simulations is 50% higher on Phoenix than on Japan's Earth Simulator and 5 times higher on Phoenix than on an IBM Power4 cluster. "Even at 256 processors," the evaluation report states, CCS staff evaluate and help users get the most scientific productivity from the CRAY X1, CCS's flagship supercomputer.

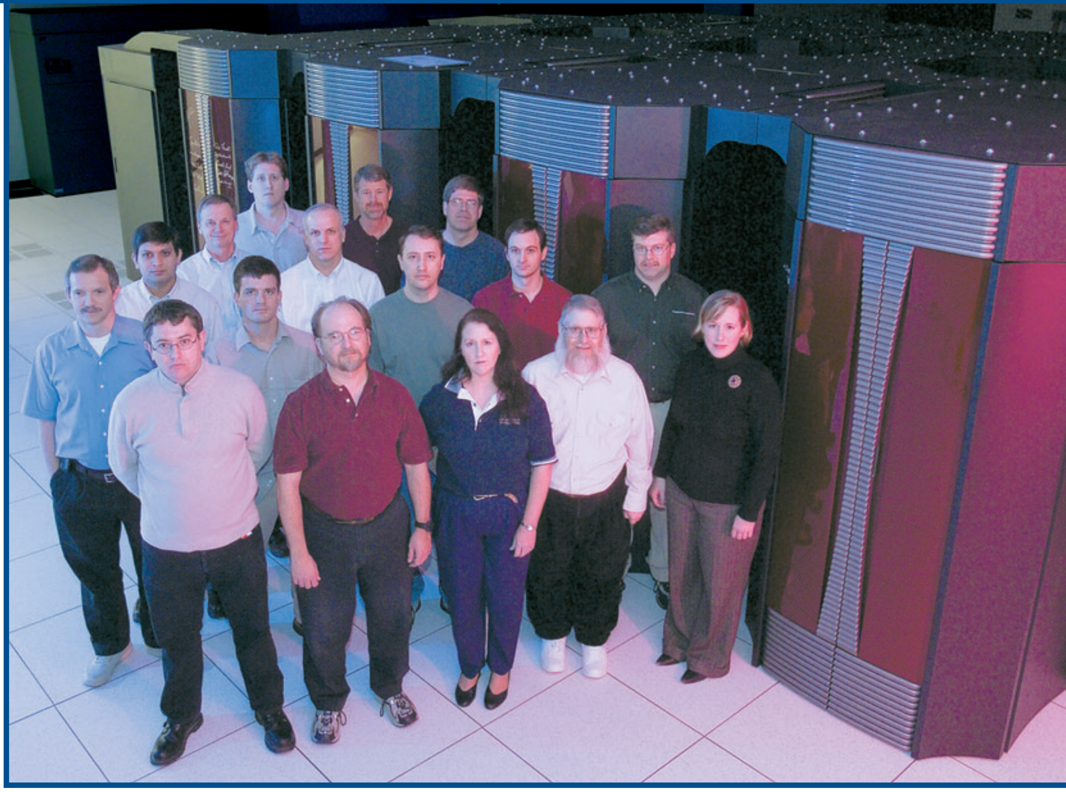

"the Cray X1 system is already outperforming other supercomputers with thousands of processors for certain classes of applications such as climate modeling and some fusion applications."

An astrophysics simulation runs 15 times faster on Phoenix than on an IBM Power4 cluster using the same number of processors. A molecular dynamics simulation related to the phenomenon of photon echo—light from an exploding star scattered off shells of cosmic dust—runs 8 times faster on Phoenix than on other supercomputers.

The best results are shown on large problems, where researchers cannot fit the entire problem into the cache of the processors. These large problems are the types important to DOE and other government agencies and industries interested in world-class simulations. The increased performance enables some simulations to fully resolve questions raised by previous scientific studies.

Based on the extremely positive results, the plan for fiscal-year 2004 is to expand the initial Cray X1 system from 8 cabinets (256 processors) to 10 to 12 cabinets with 640 to 768 processors to develop and deploy the scalable operating system, file system, and other critical features. Along with the deployment of the Cray X1, the CCS staff is partnering with Cray to develop their next-generation supercomputer code named "Black Widow." The system will be capable of 100+ teraflops, specifically designed for DOE science applications. Because financing for this advanced system has been pledged by a combination of government and industry, this advanced supercomputer system will be deployed in fiscal-year 2006.

Visualizing and Storing Supercomputer Results

A key part of the new CCS facility is a state-of-the-art Science Visualization Lab for visualization of multi-terabyte datasets generated by simulation programs run on internal computer systems. The signature capability of this facility is the 35-million pixel, 30-ft.-wide display wall used for viewing high-resolution images. The display is similar to an IMAX movie screen for viewing scientific data. Such visualizations enable researchers to better understand and communicate the significance of the results of supercomputer calculations

In addition to displaying data, ORNL helped to develop a state-of-the-art data storage system that has become the *de facto* standard for many of the largest supercomputing centers in the world. CCS has one of these data repositories, which is aptly named the High Performance Storage System (HPSS). The system received an R&D 100 award in 1995 from *R&D* magazine as one of 100 best innovations of that year.

The HPSS system at ORNL stores files of multiple terabytes for the users of the CCS computers. Other customers include NASA's Distributed Active Archive Center and the Atmospheric Radiation Measurement Program, the largest global change research program supported by DOE.

HPSS is extremely scaleable. The system can support data import rates and data objects of many different sizes. Archive files range in size from a few megabytes up to many terabytes. The CCS is continually upgrading the capability and capacity of HPSS by increasing the bandwidth, upgrading its storage devices, and improving system availability.

With an acre of new space, advanced network capabilities, unlimited infrastructure, and state-of-the-art data storage systems, DOE's Center for Computational Sciences at ORNL is ready to assume the primary role in developing a leadership-class computing program for America's scientific future.

PARTNERSHIPS:
Joining Hands to Address National Problems

This new building, one result of ORNL's $300 million modernization program, houses the Joint Institute for Computational Sciences and the Oak Ridge Center for Advanced Studies.

The CCS has established collaborations and partnerships with researchers throughout the nation. Some collaborations involve university faculty and students.

The mission of the Department of Energy's national laboratories is to address problems of national scope that cannot be easily addressed in the nation's research universities. Such problems require a breadth and depth of expertise not found in any one institution. DOE's Center for Computational Sciences at ORNL addresses such problems through partnerships with academia, other national laboratories, and, as appropriate, industry.

The CCS has established collaborations and partnerships with research scientists throughout the nation. CCS's activities in support of DOE's Scientific Discovery through Advanced Computing program involves collaborations with scientists at 50 universities and 13 national laboratories. Most other major programs also involve partnerships and collaborations.

CCS supports joint appointments at a number of research universities, especially the University of Tennessee, Vanderbilt University, and UT-Battelle's six "core universities"—Duke, Georgia Tech, Virginia Tech, Florida State, North Carolina State, and the University of Virginia. These appointments involve direct participation of university faculty in ORNL research projects that then become part of the educational curriculum at the faculty's home institution. CCS now has about a dozen joint appointments, a number that is growing each year.

Joint Institute for Computational Sciences

To further enhance the Laboratory's research collaborations, ORNL and UT established the Joint Institute for Computational Sciences (JICS) three years ago. In June 2004 JICS began occupying a new 52,000 sq. ft building that was constructed on the ORNL campus using $10 million from the state of Tennessee.

Both the University of Tennessee at Knoxville and ORNL have pledged new joint faculty positions to support JICS's research and educational programs. The facility has space for 30 faculty and 40 to 50 students. Tennessee state officials see JICS as a significant step forward in cementing the close and valuable partnership between UT and ORNL.

JICS has three strategic goals. First, the institute will work closely with CCS to help create major new modeling and simulation capabilities for terascale and beyond computers and to use these new capabilities to solve the science and engineering problems of greatest importance to DOE and the nation. Second, JICS will partner with UT to help train scientists and engineers to model and simulate natural and engineered systems on terascale and beyond supercomputers and to educate a new generation of scientists and en-

Householder Fellow

Jennifer K. Ryan is the latest Alston S. Householder Fellow in a line of outstanding mathematics researchers at ORNL. The Householder Fellowship honors a renowned pioneer in the field of linear algebra who was also the founder of what is now the Computer Science and Mathematics Division at ORNL.

Ryan, a native of Houston, Texas, is an applied mathematician working on numerical algorithms for solving real-world problems. Her primary research focuses on accurate numerical solutions to physical and chemical models represented by partial differential equations. Her current work is on post-processing for the discontinuous Galerkin method, which offers the advantages of high accuracy and the ability to handle complex data. She is collaborating with ORNL researchers on climate change applications and computational chemistry models. She is a co-author of three recent papers, one of which will appear in the *SIAM Journal of Scientific Computing.*

Ryan's interest in math developed during her undergraduate studies at Rutgers University. She received her master's degree from New York University and her doctorate from Brown University, where she received the Stella Dafermos Award for excellence in graduate studies. She previously held a graduate fellowship at NASA-Langley.

Thomas Dunning is director of the ORNL-UT Joint Institute for Computational Sciences.

gineers well versed in the application of computational modeling and simulation in research and education. Finally, JICS will join with CCS, UT, and UT-Battelle's core universities to help create a state-of-the-art cyberinfrastructure for science and engineering in the southeastern United States.

The new director of JICS is Thom Dunning. A pioneer in computational molecular science, Dunning has long been involved in the management of scientific research. Dunning expects to spend the next five years building the institute — the collaborations and partnerships needed to advance computational science in this century.

"Computational science is at a turning point," he says. "Our ability to simulate the real world increases with every advance in computing technology. However, it is becoming more and more difficult to realize the full benefit of these advances. The way to solve this problem is through long-term partnerships—partnerships between physical and biological scientists, computer scientists, and mathematicians, as well as partnerships between the national laboratories and universities. JICS will be the meeting ground for these partnerships. The institute will provide a direct link between scientists at ORNL with those at UT and UT-Battelle's core universities."

One key feature of the JICS facility is that it will provide five incubator suites, each about 1600 sq. ft. The suites can be used as classrooms, small conference rooms, or computer laboratories by ORNL staff, visiting faculty, and graduate students working together on computational research projects.

One of the suites has a raised floor, electrical power, and cooling equipment, making it ideal for a cluster of computers for research. The other four suites also have electrical power and cooling connections, so they can house computers and printers, as well as tables, chairs, and chalkboards.

In addition, a suite could be made available to a commercial firm for proprietary research in collaboration with ORNL researchers under an approved cost recovery basis. Other suites could be used for basic research and nonproprietary technology development collaborations among researchers from ORNL and university partners.

Joint Faculty Appointments

Researchers in JICS will have joint faculty appointments, enabling them to teach at a university and conduct research on site at ORNL. JICS will minimize the problem of geographical non-locality by using the latest electronic collaboratory and distance education tools. Electronic collaboratory tools will allow faculty to continue their research activities when they return to their home institutions. Distance education tools will enable graduate students doing research in JICS to take classes from their home institution. The tools will also allow faculty members working several months on a project at JICS to televise their class from Oak Ridge to their home university, or even to other members of the consortium.

These greatly enhanced communications capabilities will be supported by a new, 10 gigabit-per-second network being constructed to connect ORNL with UT and the core universities.

Research Alliance for Minorities

A critical partnership aimed at identifying and training the next generation of computational scientists is the DOE Research Alliance for Minorities (RAM) Program, which is based on the belief that national laboratories and universities, working hand in hand, offer the best opportunity to improve the quality and diversity of the workforce. The short-term goal is to increase the number of underrepresented minorities pursuing degrees in science, mathematics, engineering, and technology. These minorities include African Americans, Hispanic Americans, Native Americans, and women. The RAM program supports DOE's long-term goal of increasing the number of underrepresented individuals with advanced technical degrees on the staffs of the national laboratories. It is designed to provide collaborative research experiences involving students, faculty, and laboratory researchers in numerous technological disciplines. Students in 10-week summer internships in the computational sciences and applied mathematics learn from the experts and gain hands-on experience in preparing and giving oral and poster session presentations. They are encouraged to seek advanced degrees and consider ORNL the employer of choice upon graduation.

Participants in the 2003 Research Alliance for Minorities summer internship program at ORNL were from Alabama A&M University; City University of New York, York College; Fisk University; Mississippi Valley State University; Tuskegee University; the University of Tennessee at Knoxville; and Wofford College.

NETWORKING: Linking America's Laboratories

ORNL researchers have been tasked by DOE to develop the next generation of scientific networks to address the challenges of large science applications, such as moving large volumes of data at high speeds between science facilities and supercomputers.

Big science requires big computers that are not just scaled-up desktop personal computers. Big computers are fundamentally different from PCs in their ability to model enormous systems, generate immense volumes of data, and, as a payoff, solve uniquely difficult scientific problems. To put this difference in perspective, next-generation science datasets will approach or exceed a petabyte in size. If one of today's desktop PCs had a disk able to hold a petabyte-sized file, the PC would require over three years to read the file.

To make their discoveries, scientists must interact with supercomputers to generate, examine, and archive huge datasets. To turn data into insight, this interaction must occur on human time scales—not over days or weeks, but over minutes.

Large-scale science projects increasingly require close collaboration among domain experts located across the nation or around the globe. However, because even high-performance networks cannot deliver next-generation science datasets to users on human time scales, supercomputers require physical proximity of the scientists who want to interact with the data. Entire teams must travel, sometimes for days, to a supercomputer they will use for as little as an hour. Thus, the development of super-networks scaled to these supercomputers holds enormous promise for improving the efficiency and productivity of big science projects.

The super-networks will require unprecedented capabilities, significantly beyond today's Internet technologies. First, there is the issue of sheer bandwidth. The fastest commercial networks available today have a theoretical bandwidth of 10 gigabits per second (Gbs). If such a network could deliver all its capacity to a supercomputer, it would still take a month to move a petabyte file to the researcher. Bandwidth must be greatly increased.

Second, network technologies do not exist that can stably support interactive visualizations and steer computations of this scale. With current Internet protocols visualizations drop out of sync and computations wander off into unwanted regions, wasting precious computing power. This situation is also unacceptable. Clearly, the next generation of supercomputers will require comparable network capabilities to make them productive.

Several ORNL networking projects, funded by the Department of Energy, National Science Foundation, and Defense Advanced Research Projects Agency, are under way to address these important networking problems.

DOE UltraScience Net

The Center for Computational Sciences at ORNL has been tasked by DOE to develop the next generation of scientific networks to address the challenges of large science applications. The techniques developed at Oak Ridge will eventually filter out into the high end of the business world. Just as yesterday's scientific supercomputers have become today's central business and engineering computers, the same transfer will result in this network, called the DOE UltraScience Net, becoming the core of tomorrow's commercial networks.

Today's commercial networks are optimized for carrying huge numbers of small data streams. What big science needs are networks optimized for small numbers of large data streams and high precision control. High-energy physics, climate modeling, nanotechnology, fusion energy, astrophysics, and genomics are among the research areas that will benefit from the UltraScience Net. Computational steering and instrument control required to run experiments remotely place different types of demands on a network, making this task far more challenging than designing a network system solely for transferring data.

ORNL researchers will take advantage of current optical networking technologies to build the prototype network infrastructure. This infrastructure will enable development and testing of the scheduling and signaling technologies that will be needed to process requests from users and optimize the system. The UltraScience Net, which will operate at 10 Gbs to 40 Gbs, will develop (and demonstrate) the techniques that will allow networks to deliver as much as 200 Gbs. Eventually, the UltraScience Net could become a high-impact, special-purpose network that will connect DOE laboratories, collaborating universities, and institutions around the country. Such a 200 Gbs network will provide scientists with access to the volumes of data produced and demanded by leadership-class systems.

DOE UltraScience Net

While these new networking techniques of tomorrow are being developed, they will need a yardstick against which they can be compared. The yardstick will be provided by the TeraGrid.

Cyber-Science via the TeraGrid

A visiting neutron scattering scientist at ORNL sends data from her experiment to a San Diego supercomputer for analysis. The calculation results are sent to Argonne National Laboratory, where they are turned into "pictures." These visualizations are sent to a collaborating scientist's workstation at North Carolina State University (NCSU), one of the core universities of UT-Battelle, which manages ORNL for DOE. Finally, the significance of an experiment is communicated to interested researchers, thanks to high-tech devices around the nation connected by today's highest-speed network.

The $1.4 billion Spallation Neutron Source is on time for completion in 2006 in Oak Ridge. The High Flux Isotope Reactor is the world's most powerful source of thermal neutrons used to unlock the molecular secrets of materials and to provide radioisotopes for a number of medical, industrial, and academic uses. Data from the two neutron sources will be made available through CCS.

As ORNL becomes the world's foremost neutron science center, researchers from around the nation will be able to analyze its data up close or at a distance. Because of a $3.9 million grant from NSF to the Center for Computational Sciences at ORNL, a network hub and high-performance network connections are being established to support access to ORNL's neutron science instruments across the TeraGrid.

The TeraGrid is part of a high-speed network that will provide scientists with extraordinary amounts of data from ORNL's High Flux Isotope Reactor and the accelerator-based Spallation Neutron Source. When complete, the TeraGrid's network backbone will operate at 40 Gbs, making it the fastest research network in the world today.

The ORNL-led addition to the TeraGrid, called the Southeastern TeraGrid Extension for Neutron Science (SETENS), will allow scientists working at these facilities to use the massive computing and data storage resources on the TeraGrid to rapidly make detailed analyses and visualizations of the data from neutron scattering experiments. The system will provide near real-time feedback.

SETENS represents a major commitment to the region and to economic growth in the Southeast. From an economic development perspective, these resources show a continued commitment to build the intellectual capital of the Southeast—an investment that will reap benefits in terms of new business and research opportunities for decades to come.

Teragrid

eScience

Science projects that are enabled by computational and other resources distributed across networks are often referred to as eScience projects. Both UltraScience Net and TeraGrid projects provide the networking infrastructures in support of such projects—the former provides dedicated channels and the latter provides shared links like the Internet.

ORNL is involved in jointly developing the associated technologies needed to bring the network capabilities to eScientists in a transparent manner. Two joint NSF projects are under way in providing the transport, middleware, and visualization capabilities, all optimized to UltraScience Net and Internet infrastructures. Funded at $3.5 million, the CHEETAH project is a joint venture involving ORNL and the University of Virginia, City University of New York, and NCSU. The NetReact project involves ORNL and Georgia Tech and is funded at $1.2 million.

While both projects specifically target the SciDAC TeraScale Supernova Initiative, led by ORNL, they also provide valuable enabling networking technologies for other DOE large science projects. Using these technologies over UltraScience Net, the goal is to control supernova computations at ORNL by researchers at NCSU and UVA.

Precision Control

Controlling remote devices and instruments over the Internet is extremely challenging. The randomness caused by Internet traffic makes control very complicated, and few technologies are available to help. Under joint $1 million funding from DOE and DARPA, however, technologies are being developed at ORNL for performing such control operations across the networks.

In short, ORNL networking research capabilities, and the projects described here, will advance the state of the networking art, helping position CCS to host leadership-class scientific computing.

SOFTWARE TOOLS: Scalable Systems Software for Terascale Computer Centers

The Scalable Systems Software project coordinated by ORNL is fundamentally changing the way future high-end systems software is developed to make it more cost effective, robust, and scalable to multi-teraflop supercomputers.

System administrators and managers of terascale computer centers are facing a crisis. The nation's premiere scientific computing centers all use incompatible, ad hoc sets of systems tools that were not designed to scale to the multi-teraflop systems being installed in supercomputer centers today. One solution would be for each computer center to take their home-grown software and rewrite it to be scalable. But this approach would incur a tremendous duplication of effort and delay the availability of terascale computers for scientific discovery.

The purpose of the Scalable Systems Software project is to provide a much more timely and cost-effective solution by pulling together representatives from the major computer centers and industry and having them collectively define standardized interfaces between system components. At the same time this group can produce a fully integrated suite of systems software components that can be used by the nation's largest scientific computing centers.

The scalable systems software suite is being designed to support computers that scale to very large physical sizes without requiring that the number of support staff scale along with the machine. This strategy goes beyond just creating a collection of separate scalable components. By defining a software architecture and interfaces between system components, the Scalable Systems Software research is creating an interoperable framework for the components. This makes it much easier and cost effective for supercomputer centers to adapt, update, and maintain the components in order to keep up with new hardware and software. A well-defined interface allows a site to replace or customize individual components as needed. Defining the interfaces between components across the entire system software architecture provides an integrating force between the system components as a whole and improves the long-term usability and manageability of terascale systems at supercomputer centers across the country.

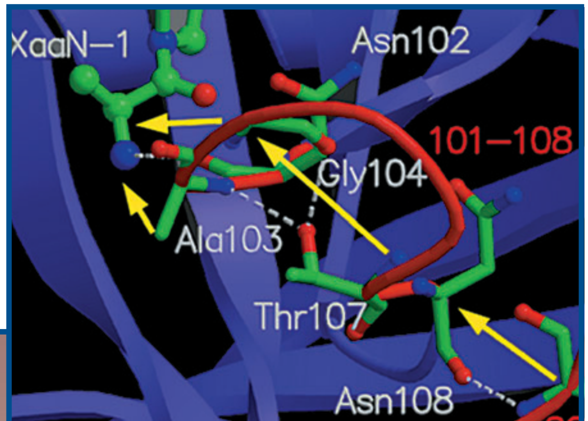

Systems interfaces are being standardized using a process similar to that employed to successfully define the message passing standard (MPI). This process is an open forum of university, lab, and industry representatives who meet regularly to propose and vote on pieces of the standard. The figure at the right represents the significant progress to date on producing scalable components and defining standardized interfaces between them. The bold lines represent working interfaces. The light lines represent interfaces in progress. The colors of the components represent which of the four

Stephen Shevlin (foreground) and Tom Dunigan (standing) discuss with Pratul Agarwal the image on the computer monitor which resulted from a simulation of a protein. The IBM Power4 (Cheetah) supercomputer at CCS was used to perform multi-scale modeling of vibrations in the protein cyclophilin A, which is related to HIV infections.

multi-lab working groups inside the project is responsible for it.

In November 2003 the first release of a complete, integrated set of scalable systems components was made. This distribution utilized the popular OSCAR packaging and install technology. A second release was scheduled in March 2004. This past year the system administrators at Argonne National Laboratory decided to switch their "Chiba City" cluster to use our scalable systems suite exclusively. In January 2004 the suite underwent scale tests on the 5160 processor Titanium cluster at the National Center for Supercomputer Applications. Our research has developed software to provide communication service between components over multiple protocols as well as a flexible authentication scheme to provide security to the overall system. Research continues to harden the working prototypes, improve integration, and increase scalability to the target of 10,000 processor systems.

The coordinator for this project is Al Geist, an ORNL corporate fellow. The participating organizations include seven Department of Energy laboratories, three National Science Foundation supercomputer centers, and five supercomputer vendors. The DOE facilities are ORNL, ANL, Ames, Lawrence Berkeley, Los Alamos, Pacific Northwest, and Sandia national laboratories. The NSF sites are the NCSA, Pittsburgh Super-computer Center, and the San Diego Supercomputer Center. The vendors are IBM, Silicon Graphic, Cray, Hewlett Packard, and Intel.

What is the impact of this project? The Scalable Systems Software project is a catalyst for fundamentally changing the way future high-end systems software is developed and distributed. It will reduce facility management costs by: reducing the need to support home-grown software, making higher quality systems tools available, and providing the ability to get new machines up and running faster and keep them running. The project will also facilitate more effective use of machines by scientific applications by providing scalable job launch, standardized job monitoring and management software, and allocation tools for the cost-effective management and utilization of terascale computer resources.

One of our top award winners associated with software tool development at CCS is Jack Dongarra, who directs UT's Innovative Computing Laboratory. He recently won two R&D 100 Awards, was elected a member of the National Academy of Engineering, and earned a Fernbach Award. He annually compiles a list of the Top 500 supercomputers based on peak performance.

System components presently under development and their interfaces. Dark lines represent working interfaces.

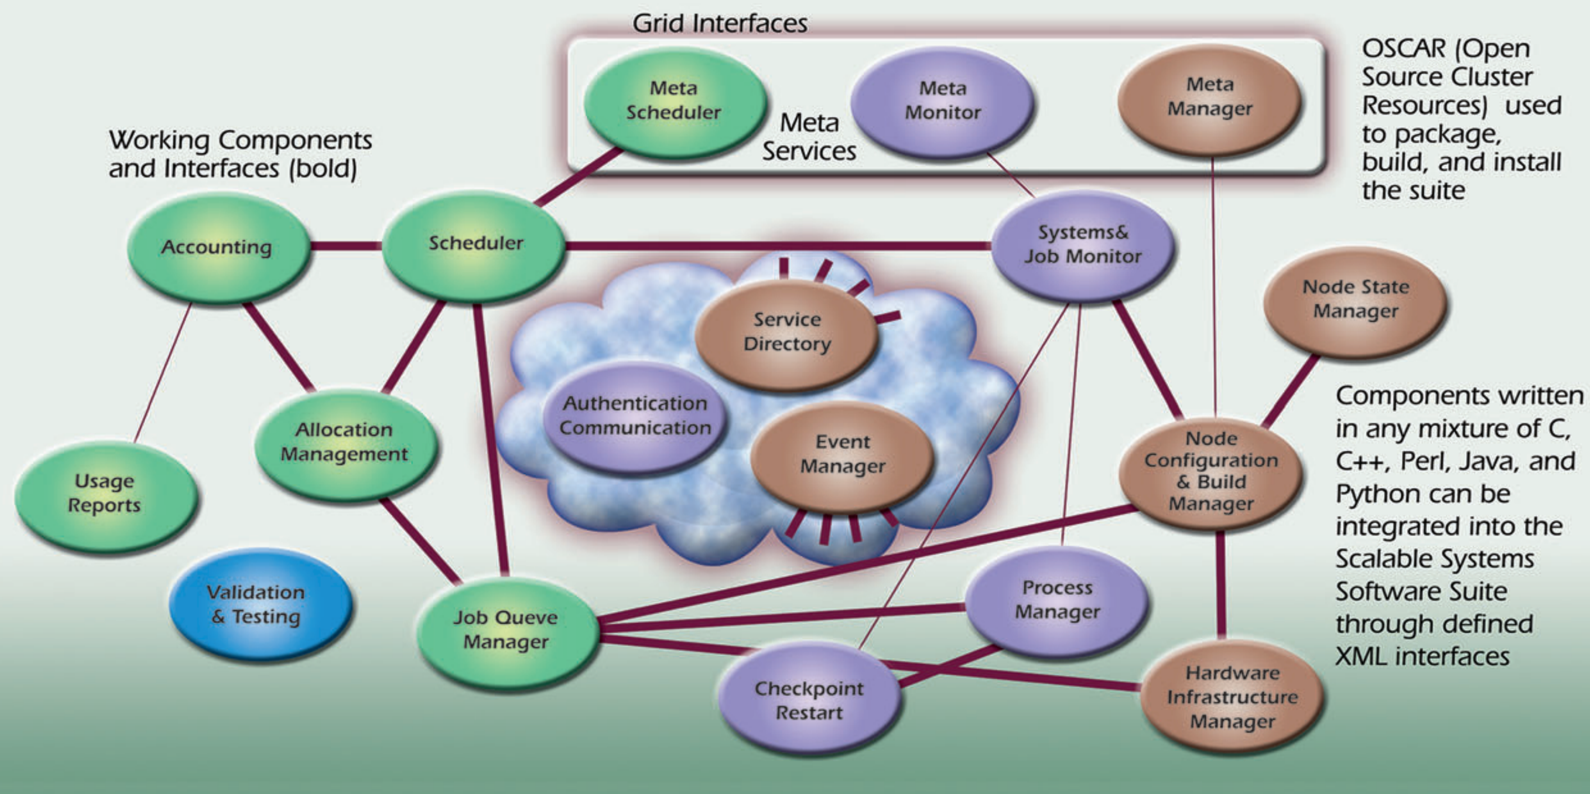

FUTURE COMPUTING TECHNOLOGIES: Pathfinders for Ultrascale Computing

The future technologies team at CCS plays an important role in the development of leadership-class computing.

Scientific applications—computer codes that enable supercomputers to run calculations on huge amounts of data to solve complex problems—will help researchers make important scientific discoveries more quickly if the codes are run on more powerful unclassified supercomputers than exist today. The goal of the Department of Energy is to increase by a factor of 100 the computing capability available to support open scientific research. Such leadership-class computing would reduce the time required to simulate complex systems, such as future climate, from years to days.

The future technologies team at DOE's Center for Computational Sciences at ORNL plays an important role in the development of leadership-class computing. By investigating the continuously evolving core technologies critical to leadership-class systems, the future technologies team led by Jeffrey Vetter is identifying technologies that satisfy the performance requirements of DOE applications. Furthermore, with their intimate knowledge of the applications, software, and hardware, this team works cooperatively with vendors to ensure that the next generation of computing technology meets the requirements of DOE mission applications. Initially, the future technologies team is focusing on four application areas: biology, climate, fusion, and nanoscience.

In collaboration with universities and other government labs, the future technologies team evaluates new computer architectures, tracks and helps design future architectures, gathers contemporary application-driven systems requirements, develops performance prediction capabilities for leadership-class computing, and assesses the state of software for these systems.

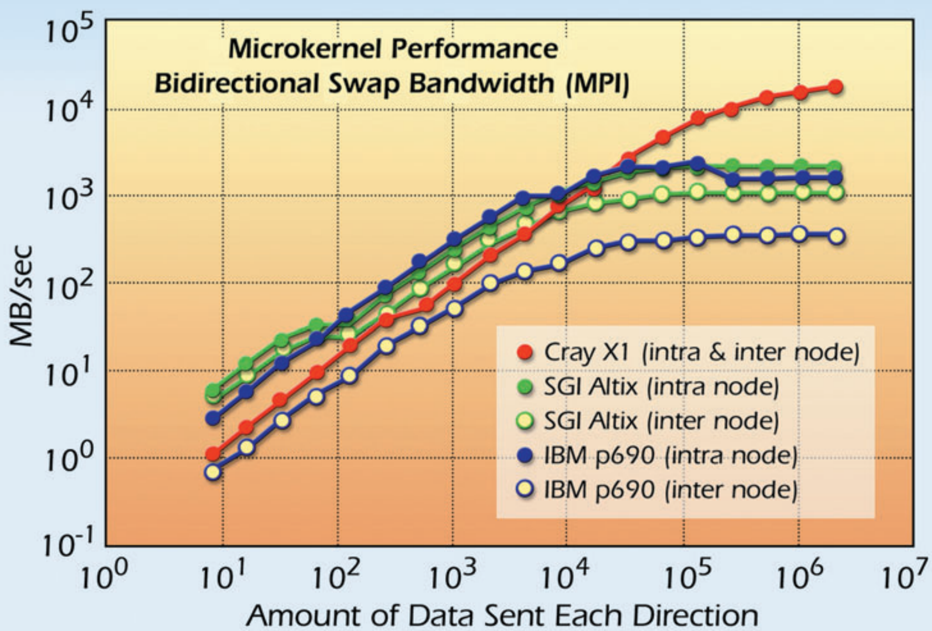

The Cray X1 outperforms the other CCS supercomputers in the amount of data passed both ways between the processors and memory.

Evaluating New Architectures

New computer architectures are evolving at a rapid rate. The future technologies team studies these new architectures and evaluates the suitability of each architecture to important ORNL applications. In particular, the overall goal of this strategy is to identify each system's architectural strengths and weaknesses in the context of these important applications. Then, in cooperation with users and other government agencies, the team uses the gathered evidence to assist vendors in the design of future architectures.

Recently, CCS evaluated several systems, including the Cray X1, the SGI Altix 3000, and the IBM Power4 cluster. In the short term, the future technologies team will evaluate the Cray Red Storm system, the IBM Blue Gene/L system, and the IBM Federation interconnect. Evaluations on the horizon will focus on the Cray X2, the IBM Power5, and the DARPA HPCS systems from Cray, IBM, and Sun Microsystems. In addition, the future technologies team is evaluating individual core technologies for leadership-class computing including processors, interconnects, memory subsystems, storage subsystems, system software, and programming models. Core technologies that the team is considering include the IBM's Federation interconnect, Infiniband interconnect, and processors for reconfigurable computing.

Simply put, CCS constantly surveys the computing landscape for future architectures that may offer orders-of-magnitude improvements in leadership-class computing. Thus, the future technologies team continues to track and evaluate evolving technologies such as system-on-a-chip (SOC), processor-in-memory (PIM), simultaneous multithreading, smart networks, optical interconnects, reconfigurable computers, and streaming

Core technologies for ultrascale computing.

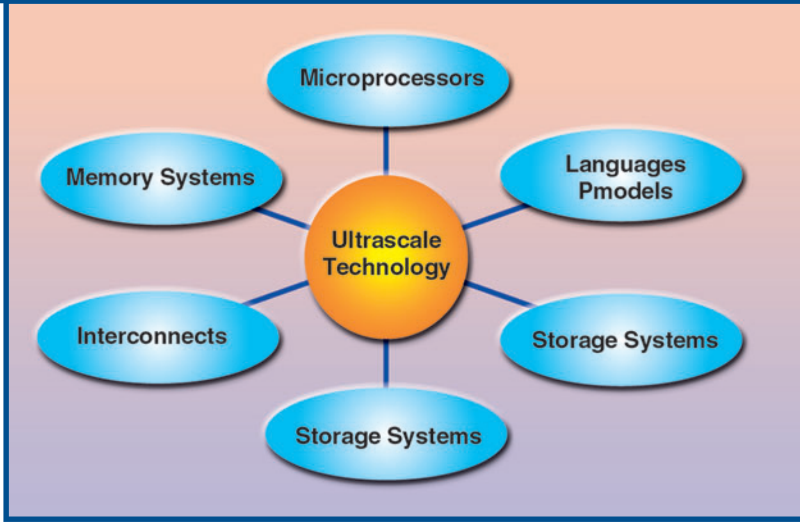

supercomputers. In some cases, these technologies could provide CCS users with tremendous improvements in their scientific productivity because of higher performance, lower costs of ownership, and increased reliability and availability. Using evidence gathered by the future technologies team, CCS makes informed decisions about which speculative architectures to endorse, fund, and procure.

Application-Driven System Requirements

In addition to its evaluation activities, the future technologies team maintains an ongoing, contemporary set of application-driven systems requirements to use for procurements and for feedback to architects. In conjunction with the user communities, the group has identified major classes of scientific applications likely to dominate leadership-class system usage over the next five to ten years. Within each area, the CCS team investigates the machine properties (e.g., floating-point performance, memory, interconnect performance, input/output capability, and mass storage capacity) needed to enable major progress in each application class. This information helps CCS address major hardware, software, and algorithmic challenges to the effective use of leadership-class computing systems.

Predicting Performance

Generating application-driven systems requirements and performing evaluations on existing systems provide some insight into which architectures best match CCS workloads. However, the best possibility for furthering the understanding of performance phenomena and for assisting in intelligent procurement selections may lie in the technique of performance modeling. By having core competencies in the future technologies team for modeling, measurement, and simulation of computer systems themselves, CCS has at its disposal not only performance information about existing systems, but also the capability to speculate about future architectures. With this information, it is possible to predict the performance achieved by a future system even though that system is much larger than systems available today.

Scalable Software

Software infrastructure plays a critical role in reducing the time needed for computational scientists to solve a problem across all phases of application development and use. Simply addressing algorithms and architectures is insufficient to ensure success on a leadership-class computer. First, software programming models and design must efficiently harness the underlying computer hardware to ensure reasonable execution performance. Second, software will play an important role in the reliability and availability of these new architectures. Applications and system software must adapt to and overcome faults in the underlying hardware. Third, scientists' productivity is often at odds with current software development techniques. In particular, the current methods of constructing, optimizing, and using massively parallel scientific applications can be difficult, tedious, and inefficient.

The future technologies team in CCS is in a critical position to evaluate and experiment with alternative technologies on leadership-class computers. Scalable, reliable system software is vital to the operation of a leadership-class platform. This software will provide baseline services such as operating systems, file systems, job and task scheduling, communication (MPI) libraries, resource management, configuration management, security, programming model support, and fault management. Given the possible scale of new systems that will have thousands to millions of processors, it is imperative that these services be efficient, scalable, and reliable. Because the requirements for scientific computing on this scale differ drastically from the requirements for commodity operating systems, the future technologies team is evaluating software environments to judge the scalability and reliability of system software and the associated programming environments and tools. Using this information, the team maintains a list of critical software requirements and undertakes research to solve these problems, in collaboration with universities, vendors, and other DOE national laboratories.

Ultimately, a goal of the future technologies team is to ensure that users can fully capitalize on leadership-class hardware and software to boost scientific productivity in meeting DOE missions. In this role of pathfinder for CCS, the future technologies team explores new technologies and interacts with users, researchers, and vendors to drive promising technologies forward.

VISUALIZATION: New Ways to Understand the Data

The Science Visualization Facility at CCS enables researchers to understand and convey the significance of a supercomputer's calculation results.

Atoms are so close and stars are so far away. In both cases, we cannot really "see" how they behave when various forces act on them. Fortunately, technologies exist to transform supercomputers into microscopes and telescopes that possess ultrahigh resolution. Using experimental data, scientists can make calculations on supercomputers to predict how atoms and stars change under different conditions. Thanks to techniques for visualizing data, suddenly "pictures" worth a thousand gigabytes can be produced. Scientists can "see" in vivid colors images of an alloy's electrons pointing in different directions—like a compass needle—as a metallic alloy's magnetic field shifts because of changes in alloy composition. They can "observe" explosions in a star as it synthesizes and disperses elements that make life possible on our planet. These images give researchers insights into scientific processes and phenomena.

Computational science generates, processes and interprets nearly unfathomable quantities of data from scientific simulations of, for example, galactic supernovae, proteins, nanostructures, and fusion energy devices. Because researchers have vast amounts of data to be analyzed, they often prefer to examine data in a compact form that is accurate, reliable, and readily understandable. Viewing the results as pictures is immensely more efficient. This is where the science of visualization demonstrates its value, and why visualization must keep pace with advances in scientific computing.

Visualization is a discipline that develops computer graphics, interaction devices, three-dimensional (3D) imaging, and realistic rendering techniques. The tools use simulation data as input and create a scientifically accurate representation.

The Science Visualization Facility in the Department of Energy's Center for Computational Sciences (CCS) at ORNL offers Laboratory and visiting researchers the ability to visually analyze and compare results of computer simulations. The ability to see an image of the results may reveal intricacies and flaws otherwise concealed in a haystack of data.

For example, in an image produced computationally from huge amounts of data, researchers can see up close the dimensions and structure of an engineered molecule. They can observe pulses of currents and winds in a simulated weather event or climate map. They can also witness the events leading to a supernova explosion and understand how it may create chemical elements.

At CCS researchers working with high-performance computers and increas-

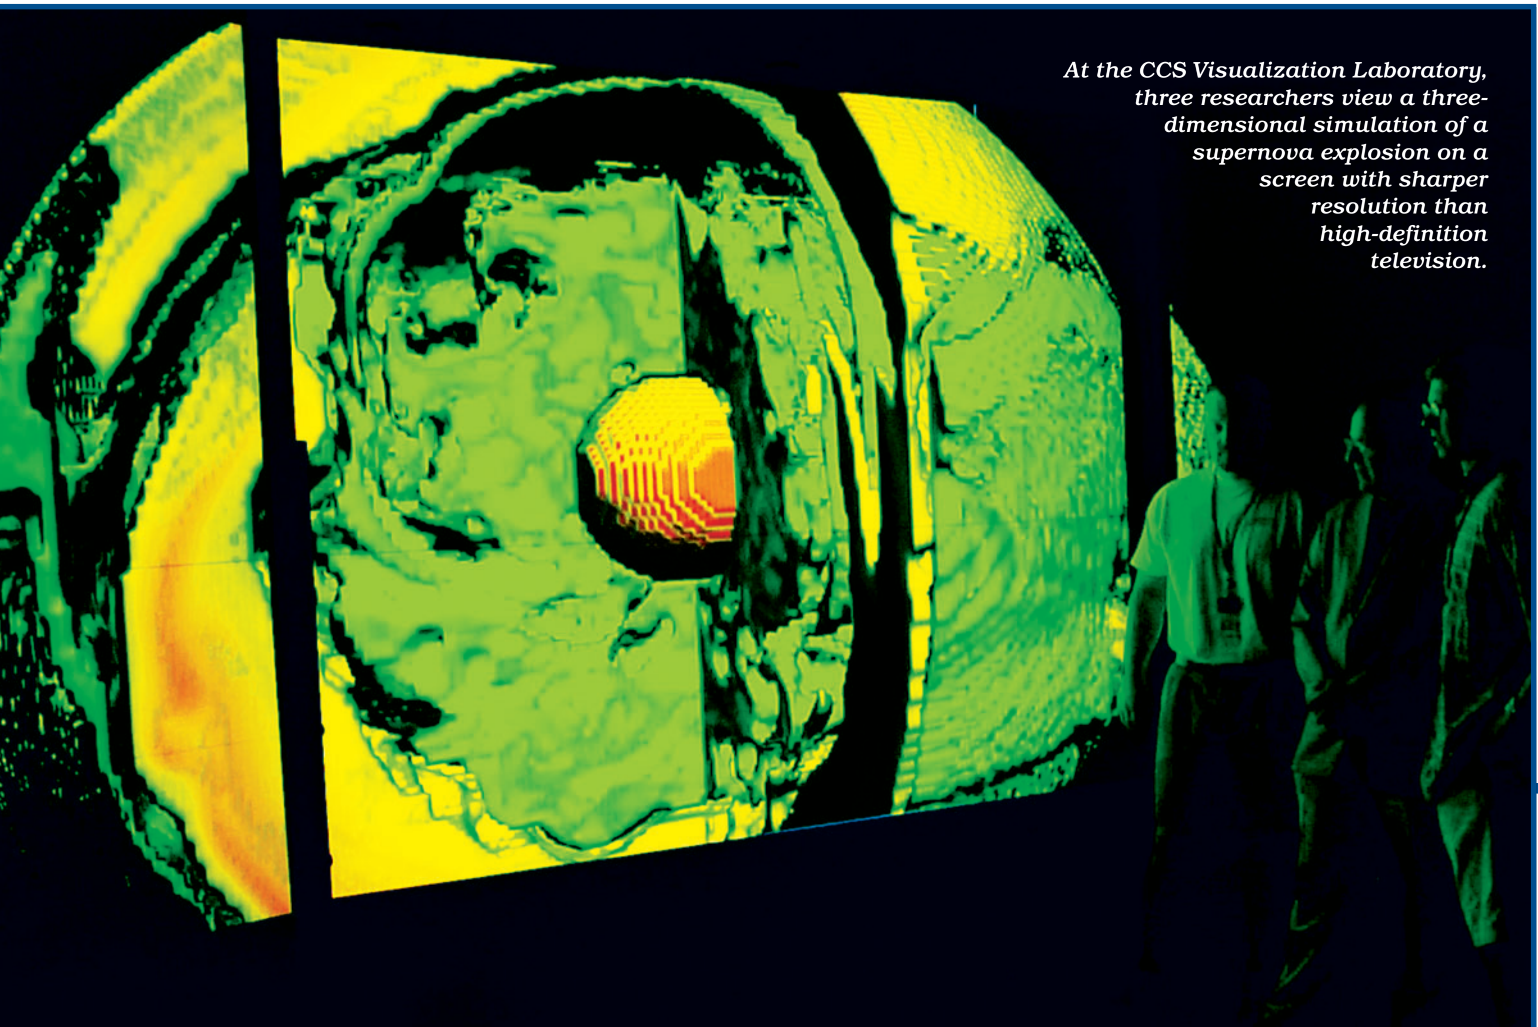

At the CCS Visualization Laboratory, three researchers view a three-dimensional simulation of a supernova explosion on a screen with sharper resolution than high-definition television.

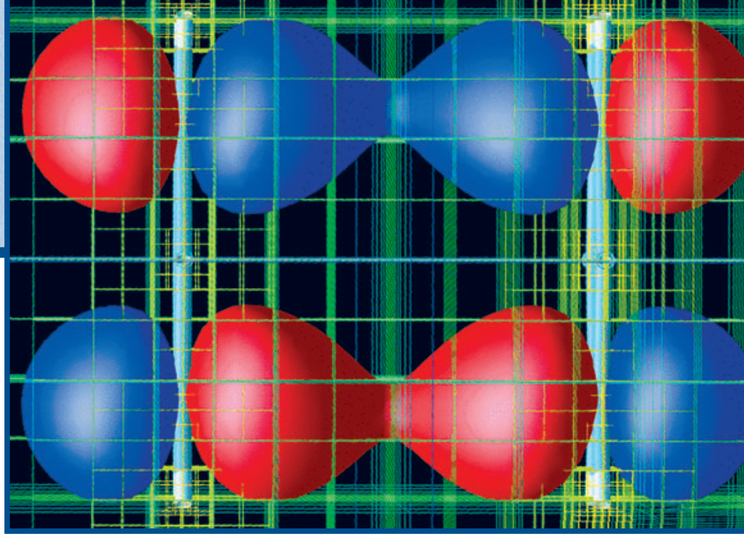

Graphical representation of a molecular orbital of the benzene dimer with the underlying adaptive computational mesh.

ing volumes of data will be able to view quickly the results of their work in a lab, office or conference room setting. With virtual-lab capabilities, these collaborations can involve researchers in the same room or in offices throughout the world. Researchers can compare their latest work with previous data through applications that are intuitively designed and customized to the user.

The Science Visualization Facility of CCS is a large-scale, immersive venue for data exploration and analysis. The new research tool is a 30-ft wide by 8-ft high display comparable in size to 150 standard computer displays with a resolution of more than 11,000 by 3,000 pixels, creating a total space of 35 million pixels. The immersive qualities of this environment create a powerful discovery tool for research groups and collaborations. The Visualization Facility provides a premier data analysis and visualization capability and facility in DOE's Office of Science.

The Visualization Facility houses display venues of different sizes and capabilities. The reCAVE offers a seamless integration of multiple projections into a large, reconfigurable, and inexpensive virtual environment. Because screens are mounted on movable frames with their respective projectors and mirrors, scientific images can be projected and viewed in different geometries, from a long wall, to "L"-shapes, boxes, or other shapes that best present the data.

An array of liquid-crystal display, flatscreen monitors, currently 6 tiles in a 3 × 2 grid, can fit in an office—either hung from the ceiling or on a wall. This portable array can display 12 megapixels of information. Another display, the IBM T221 monitor, can display images in photographic resolution of 9 megapixels, much higher than high-definition television, which has a 2-megapixel capability.

The visualization facilities are integrated with CCS through a high-bandwidth network between large-scale, high-performance computers and large-scale data visualization computers and displays. Images displayed on each venue are rendered using a group of commodity PCs with high-end graphics capabilities—a low-cost, high-performance Beowulf-type visualization cluster. Workload is distributed among the PCs in a parallel fashion for high-performance and real-time interactions.

Viewing Changes in Climate, Stars, and Materials

The Science Visualization Lab is used to investigate the intricacies of global climate change simulations. Through projected climate simulations portrayed as maps of the globe, color-coding helps deliver an understanding of parameters such as surface temperature and carbon dioxide concentrations. Such simulations may involve hundreds of parameters. Because of highly interconnected and complex interrelationships among variables, subtle changes in one variable can have significant long-term effects. Large, tiled displays provide an ideal environment for observing such sensitivities between multiple simulations and variables and can be experienced by as many as 20 participants.

Researchers can examine virtual stellar explosions—complex hydrodynamics from near the stellar core to the turbulent shock front. As the models are refined, scientists can observe parameters such as temperature, density, and velocity in the supernovae simulations and compare results from both theoretical and observational data. Eventually, it may be possible to understand how the elements that make life possible are synthesized and dispersed by such explosions.

Materials science research is conducted at the atomic level as researchers build and study the dynamics of custom-engineered materials. ORNL's visualization facilities offer scientists the ability to zoom in to the atomistic scale and observe spin dynamics and vibrational modes or zoom back out to determine bulk material properties, say, in a magnetic alloy whose composition is being changed. Visualization venues such as those in the Science Visualization Facility enable researchers to learn what works, as well as what fails—filling in the missing pieces of scientific understanding.

Visual Mural for Visitors

In addition to providing researchers with an indispensable visual tool for comprehending and communicating their research, ORNL's visualization facilities have repeatedly proven their effectiveness at communicating and demonstrating the value of ORNL research to the public. Collectively, the facilities comprise a visual mural for Laboratory visitors ranging from the nation's highest elected officials to local school groups. The Science Visualization Facility was a highlight of the recent Laboratory tour by science writers attending the New Horizons in Science Briefing sponsored by the Council for the Advancement of Science Writing.

Scientific visualization is a rapidly moving discipline as research and technology evolves. As computers advance in power and speed and scientific models become more complex, the art and science of visualization and its underlying analytical process must advance in step. As scientific computing closes in on its 100-teraflop goal, visualization researchers will be pressed to keep pace with the progress. But the resulting discoveries from images on the screen and the understanding they provide will be worth the investment.

INDUSTRY: A Critical Connection

Industry can gain access to DOE's supercomputers through the Computational Center for Industrial Innovation at ORNL.

Industrial firms in the Southeast have an opportunity to use the computational resources of a Department of Energy national lab to solve their difficult scientific problems. Through the Computational Center for Industrial Innovation (CCII), part of DOE's Center for Computational Sciences at ORNL, industry can gain access to the Laboratory's high-performance computing resources by doing collaborative research with ORNL researchers who use CCS supercomputers.

Established in 1994, CCII initially offered a place for ORNL and industrial researchers to work together, often face to face, to solve complex industrial problems. For example, in 1997 collaborative projects at CCII involved simulations of advanced aircraft, aluminum production processes, internal combustion processes, and the buildup of ice on airplane wings. CCII's industrial users included Reynolds Metals, Lockheed Martin Skunk Works, and Eastman Chemical.

Thomas Zacharia led CCII researchers in developing computer codes to facilitate the design of a stronger beverage can using less aluminum and to help car companies and the Department of Transportation model car collisions using supercomputers. Because a fuel-efficient, pollution-free car made of lightweight materials has been a national goal, car companies are interested in using computer modeling to determine whether a car made of materials lighter than steel could be designed to be as resistant to damage in a collision as today's heavier vehicles.

In 1998, when Zacharia moved from his position as head of CCII to director of ORNL's Computer Science and Mathematics Division (CSMD), he focused on developing a close relationship between CCS and supercomputer vendors. His goals were to bring the fastest supercomputers to ORNL for evaluation and production runs and to guide vendors in developing next-generation supercomputer architectures to solve "grand challenge" science problems. Zacharia has cultivated close relationships with IBM, SGI, Cray, and other vendors.

Today Alex Fischer, director of ORNL's Office of Technology Transfer and Economic Development, and Jeff Nichols, current director of CSMD, believe that CCS could give Tennessee a competitive advantage, similar to what the state enjoyed during the mid-20th century when manufacturing prospered because of low labor costs and the ability to ship goods cheaply on railroads. When the interstate highways were built, Tennessee had a similar goods distribution advantage because of its central location in the nation; the state is within a day's drive of two-thirds of the U.S. population.

The American economy in the 21st century may be driven by information and knowledge. Industry and businesses will seek greater access to tools for making better use of information. Just as railroads and interstate highways provided Tennessee and the Southeast with ways to ship large volumes of products quickly to customers, the telecommunications infrastructure will be essential for moving large amounts of data at high speeds.

CCS will be a hub linked by high-speed networks with Chattanooga and Atlanta to the south, Nashville and Memphis to the west, the Research Triangle to the east, and Chicago to the north. A contract has already been signed with Qwest to build the network south from CCS to a university simulation center. CCS will have high-speed connectivity around the nation.

The biomedical industry in Memphis and the health care industry in Nashville have expressed an interest in using the powerful supercomputers and electron microscopes at ORNL. These industries will likely have access to a high-speed network to ORNL because the Laboratory is partnering with the Tennessee Valley Authority (TVA) to use the excess "dark" optical fiber that the nation's largest public utility installed with its power transmission system. ORNL would provide the lasers and repeaters to turn TVA's unused fiber into a high-speed network connecting Oak Ridge with Nashville and Memphis.

According to Fischer, "If we can connect the dots of industry in the Southeast to ORNL's powerful supercomputers, Tennessee would have a strategic advantage for economic development for the 21st century."

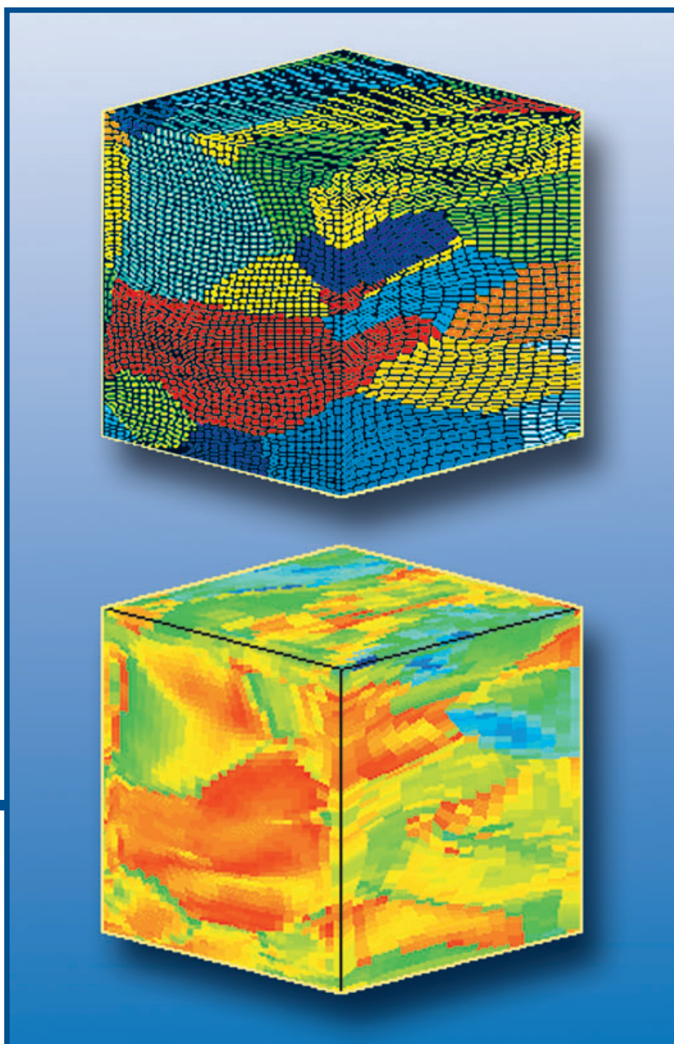

Model of deformed aluminum sheet showing structural changes at the atomic level.

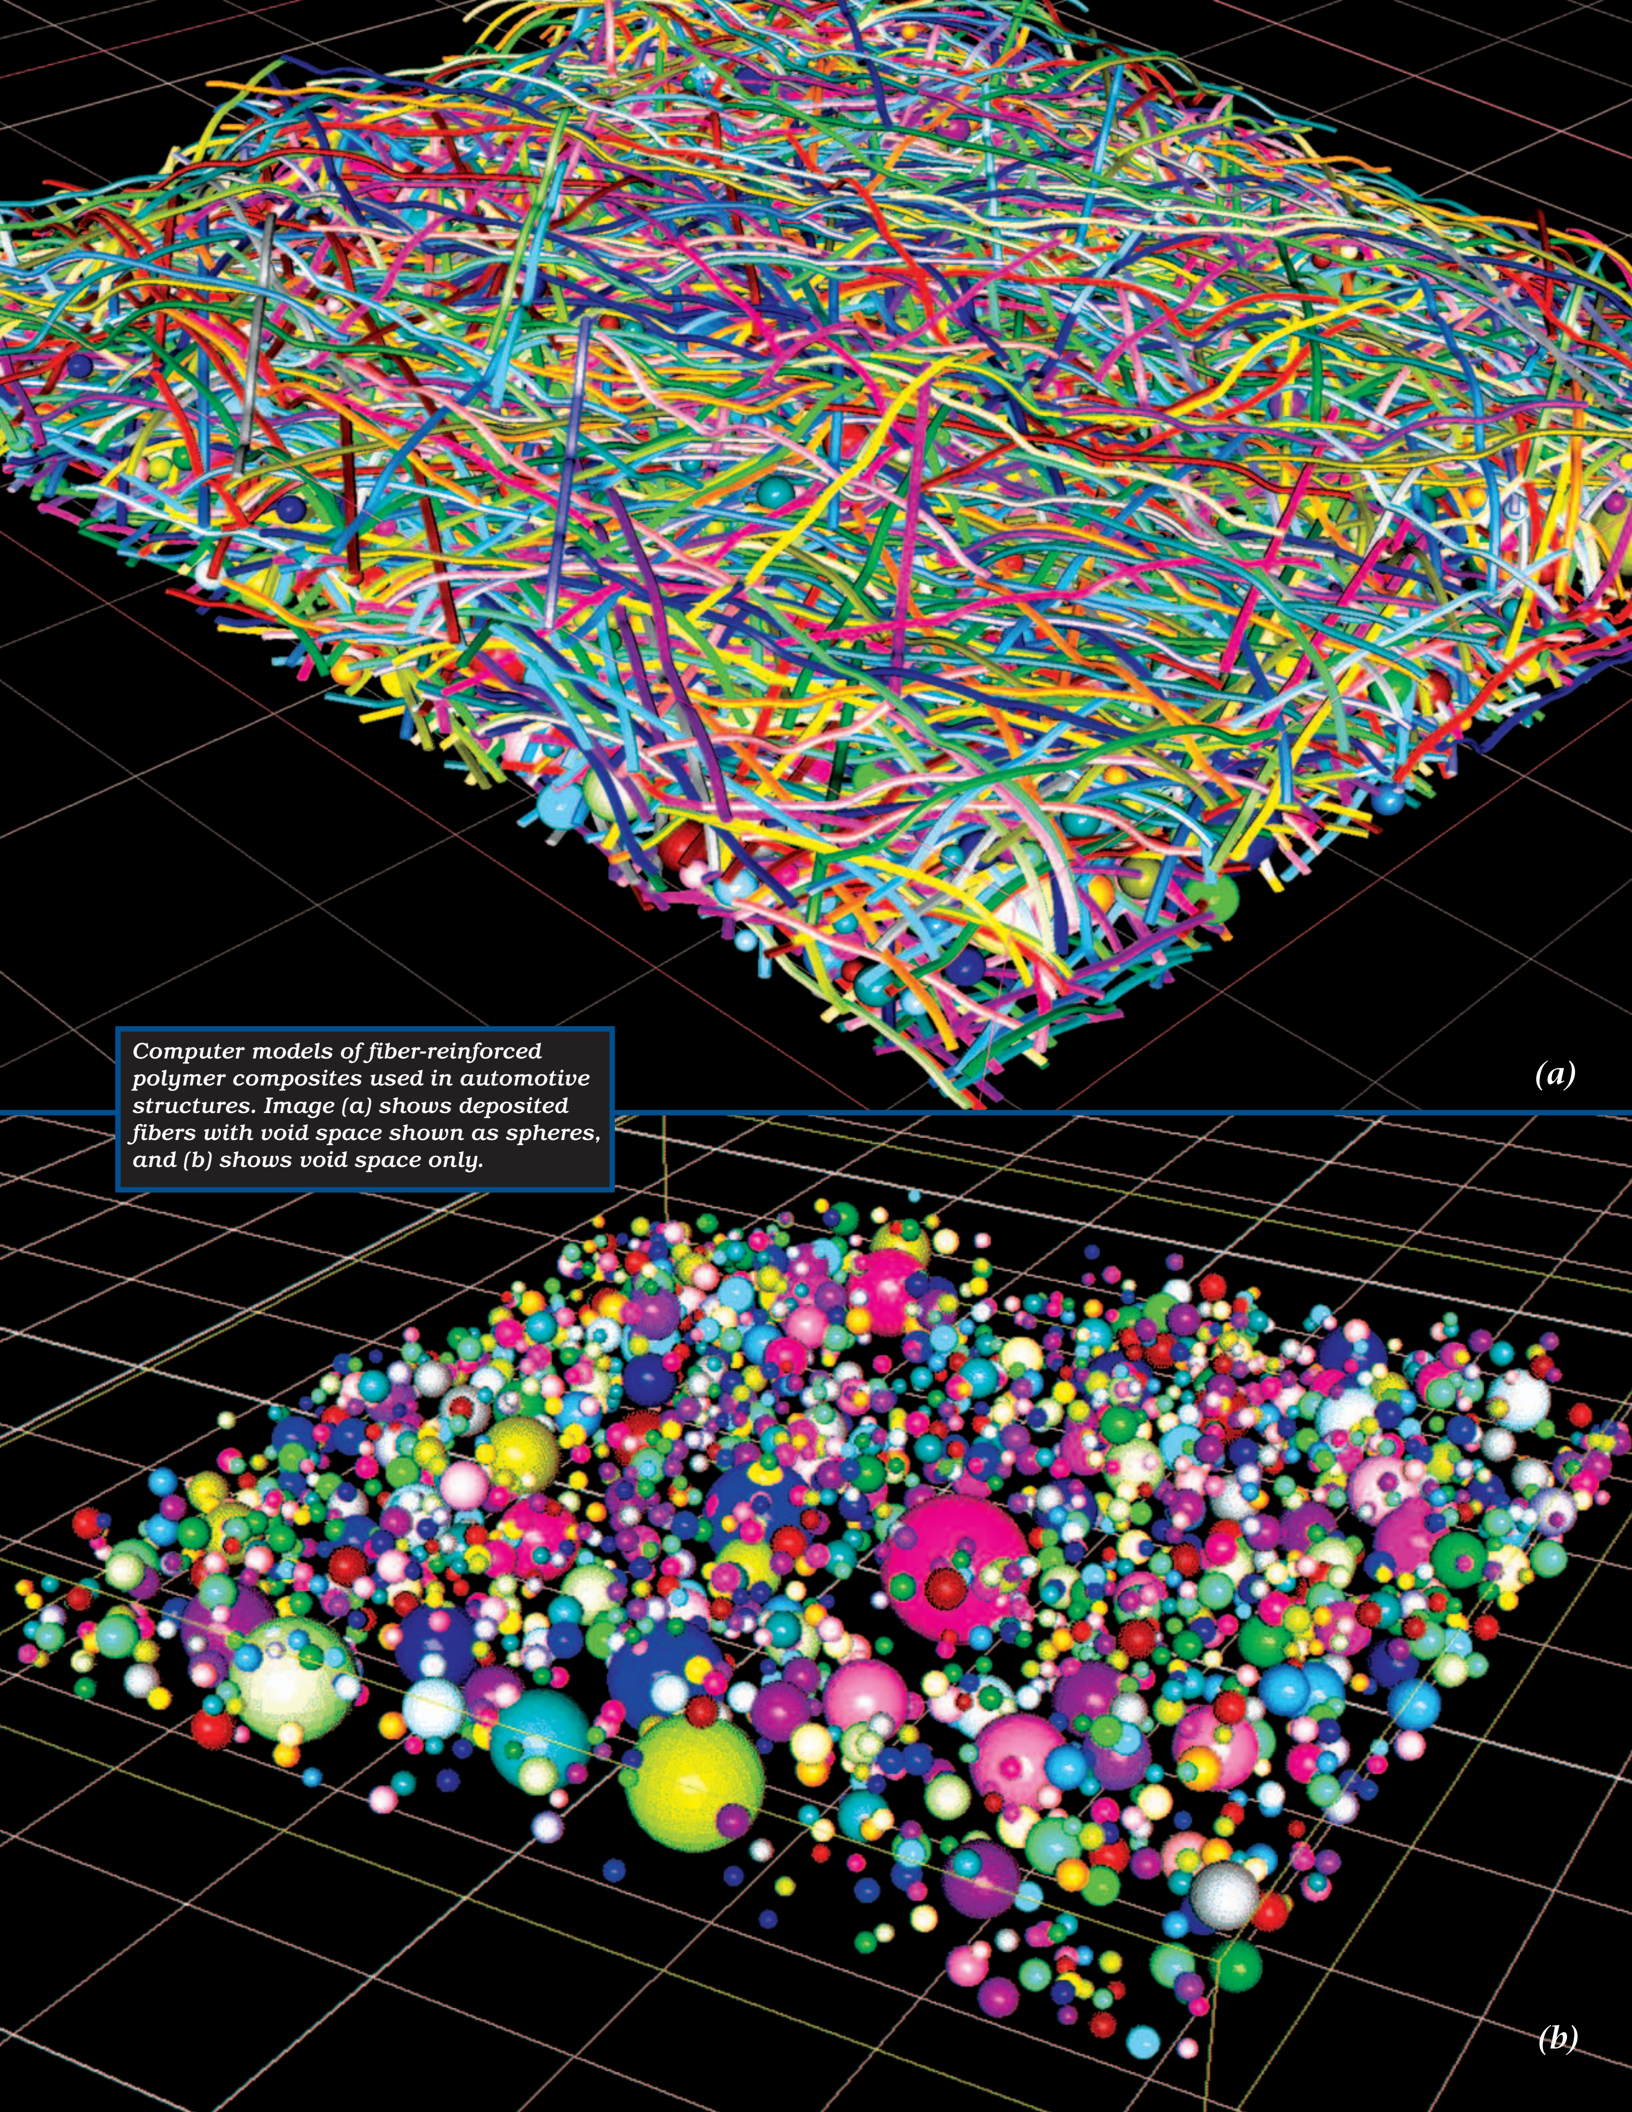

Computer models of fiber-reinforced polymer composites used in automotive structures. Image (a) shows deposited fibers with void space shown as spheres, and (b) shows void space only.

...and the WINNERS are...

Accomplishments of Distinction at Oak Ridge National Laboratory

C. T. Liu has been elected to the National Academy of Engineering. He was cited "for advancing ordered metallic compounds from the laboratory to practice." Thirteen other researchers previously affiliated with ORNL have been elected to the academy.

Jeff Blackmon is one of five Department of Energy recipients of the latest Presidential Early Career Awards for Scientists and Engineers (PECASE). He was cited for his pioneering work at ORNL's Holifield Radioactive Ion Beam Facility that advanced understanding of the nuclear processes behind stellar explosions. Eleven ORNL researchers have received PECASE awards.

Marilyn Brown has earned the 2004 James R. Anderson Medal of Applied Geography from the Association of American Geographers. She was cited for her national leadership in the analysis and interpretation of energy futures in the United States, particularly noting her publication *Scenarios of a Clean Energy Future*, which illustrates how technologies and policies could address national energy needs over the next two decades.

ORNL corporate fellows **Tom Wilbanks** and **David Greene** have been designated lifetime national associates of the National Academies, which consists of the National Academy of Sciences, the National Academy of Engineering, the National Research Council, and the Institute of Medicine. The designation of national associate is in recognition of extraordinary service to the National Academies and advising the nation in matters of science, engineering, and health.

Sang-ho Kim has won the 2004 Outstanding Young Researcher Award of the Association of Korean Physicists in America for his outstanding contribution to superconducting cavity research.

Ted Barnes and **Tom Ferrell** have been named fellows of the American Physical Society. Barnes was cited for his research on hybrid and exotic hadrons, and Ferrell was cited for his work in developing and understanding the photon scanning tunneling microscope

Emory Collins received the Glenn Seaborg Actinide Separations Award for 2004.

Pat Parr shared the Hinote Award of the Southern Appalachian Man and the Biosphere for her energy and leadership in promoting stewardship and ecosystem management at the Oak Ridge Reservation and in addressing the Oak Ridge Reservation's invasive species issue.

Jeff Blackmon is a recipient of a Presidential Early Career Award for Scientists and Engineers.

The ORNL Center for Computational Sciences is sponsored by
the Mathematical, Information, and Computational Sciences Division,
Office of Advanced Scientific Computing Research, U.S. Department of Energy

Contact Thomas Zacharia: Telephone: 865-574-4897 • Fax: 865-574-4839 • E-mail: zachariat@ornl.gov

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

PRSRT STD
U.S. Postage
PAID
Permit No. 3
Oak Ridge, TN

Editorial office telephone: (865) 574-7183;
Editorial office FAX: (865) 241-9787;
Electronic mail: krausech@ornl.gov
Web addresses: www.ornl.gov
www.ornl.gov/ORNLReview/

The *Review* is printed in the United States of America and is also available from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Road, Springfield, VA 22161

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

Oak Ridge National Laboratory is the Department of Energy's largest science and energy laboratory, with research competencies in neutron science, energy, biological systems, advanced materials, national security, and high-performance computing.

UT-BATTELLE
OAK RIDGE NATIONAL LABORATORY