Oak Ridge National Laboratory

REVIEW

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

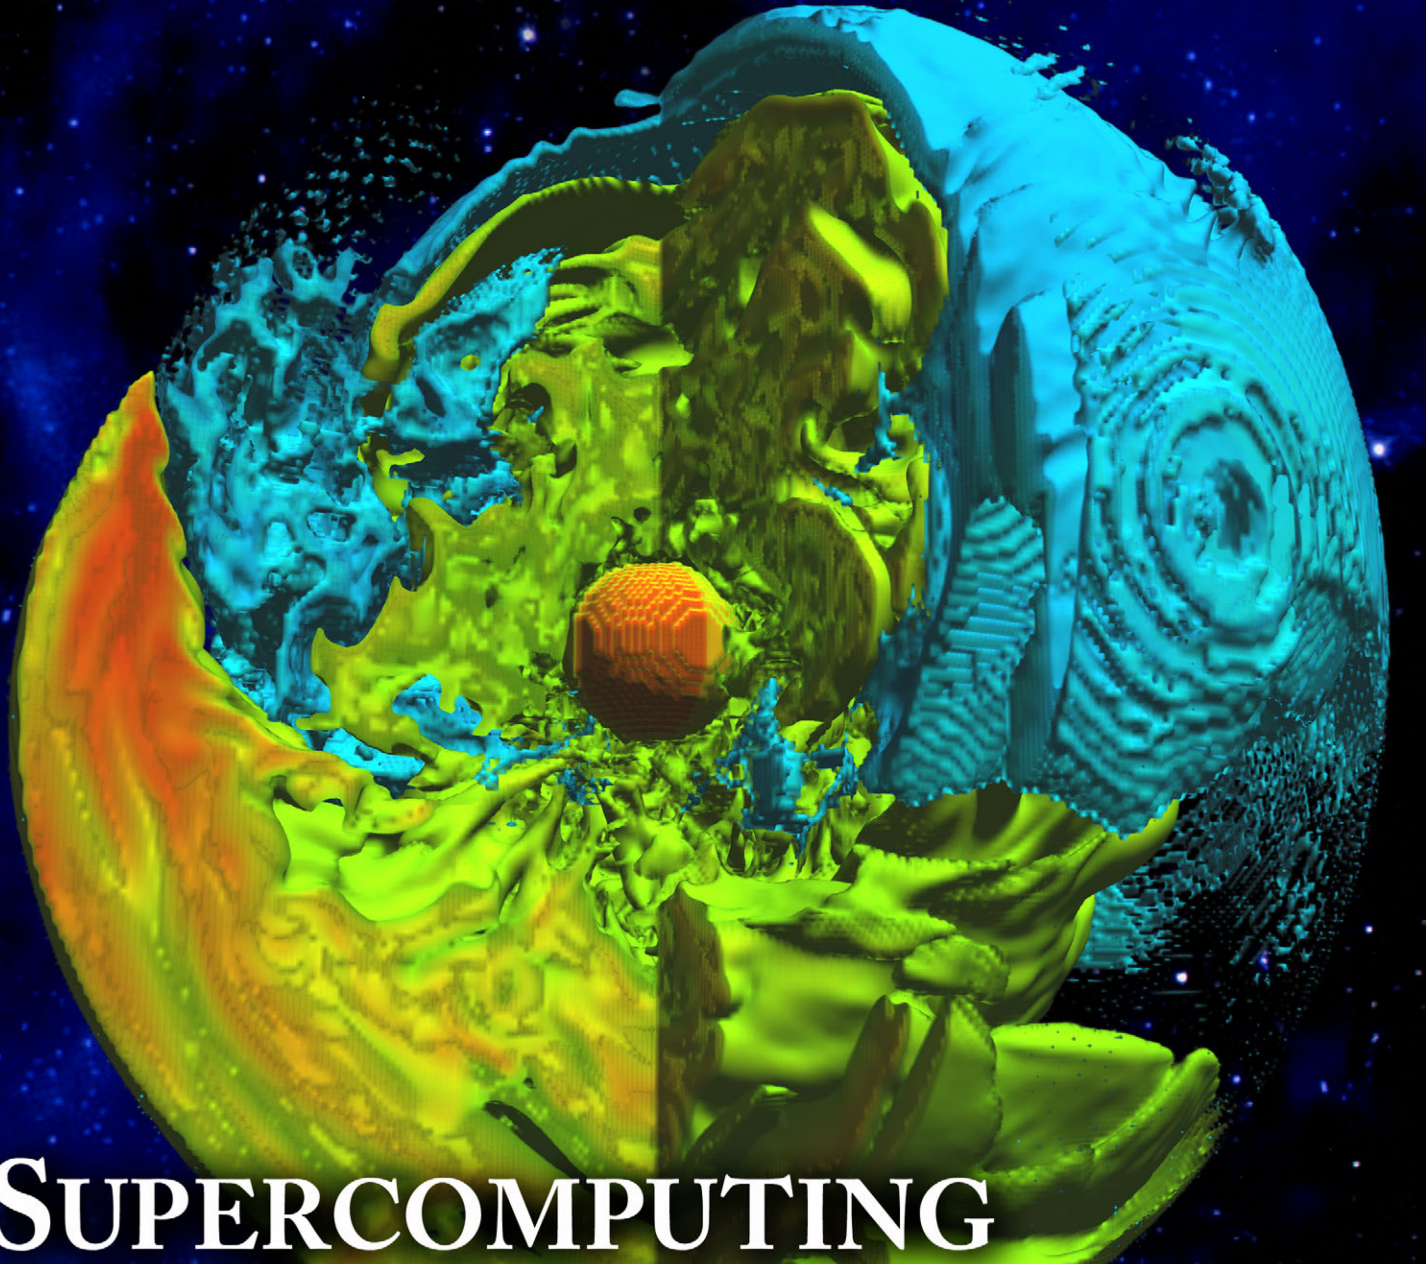

Supercomputing for Science

Also: Computer Modeling and Homeland Security

Vol. 35 • No. 1 • 2002
www.ornl.gov/ORNLReview

OAK RIDGE NATIONAL LABORATORY
REVIEW Vol. 35, No. 1, 2002

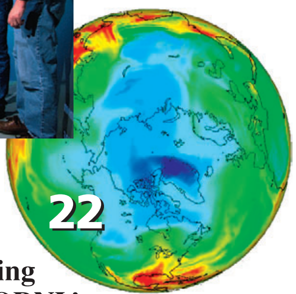

Features

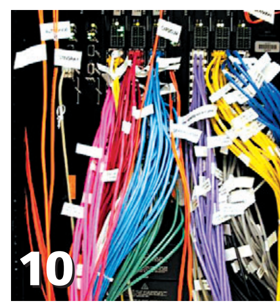

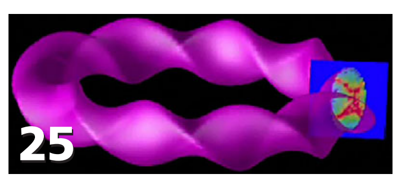

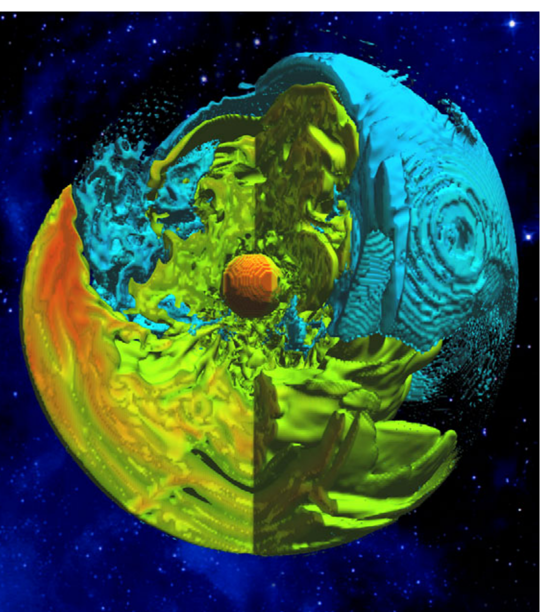

COVER: *This is a volume rendering of a three-dimensional stellar explosion, produced by ORNL's Ross Toedte. The three-dimensional simulation was carried out on the IBM SP supercomputer at DOE's Center for Computational Sciences by John Blondin of North Carolina State University as part of the SciDAC TeraScaleSupernova Initiative, which is led by ORNL's Tony Mezzacappa. Shown are two outflows in green, directed in nearly opposite directions. The outflows occur below a shock wave, the surface of the volume rendered. The organized flow below the shock helps drive the shock outward, and in turn the shock causes the explosion. Astrophysics simulation is one of many areas of computational science that involve ORNL researchers, as discussed in this issue.*

OAK RIDGE NATIONAL LABORATORY
REVIEW

Editor and writer—Carolyn Krause
Assistant editor—Deborah Barnes
Editorial Board—Lee Riedinger (chair), Fred Bertrand, Eli Greenbaum, Russ Knapp, Stan Milora, Thomas Zacharia
Designer and illustrator—Jane Parrott
Photographer—Curtis Boles

The Oak Ridge National Laboratory *Review* is published two or three times a year and distributed to employees and others associated with or interested in ORNL.

Editorial office address:
Building 4500-North, M.S. 6266,
Oak Ridge, TN 37831-2008
Telephone: (865) 574-7183;
FAX: (865) 241-6776;
Electronic mail: krausech@ornl.gov
ORNL *Review* Web address:
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 Department of Energy under contract DE-AC05-00OR22725

ISSN 0048-1262

Oak Ridge National Laboratory is a multiprogram, multipurpose laboratory that conducts research in energy production and end-use technologies; biological and environmental science and technology; advanced materials synthesis, processing, and characterization; computing and computational sciences; and the physical sciences, including neutron-based science and technology.

EDITORIAL: Supercomputing for Science—ORNL's Commitment to Scientific Discovery


Curtis Boles

Thomas Zacharia

From every corner of science a revolution is under way because of the growing amount of data being generated and the rapid increase in scientific understanding resulting from applying advanced computational science tools to these data. John Marburger, science adviser to President Bush, echoed this message when he spoke recently at ORNL. He observed that advanced computing combined with advanced instrumentation is enabling scientific discovery at an unprecedented pace—a pace that will only increase in the future.

Advanced computing and computational science add a new dimension to the more traditional experimental and theoretical approaches to scientific research. The use of computational tools has become vital to most fields of science and engineering and to many parts of the educational enterprise. ORNL is well positioned to play a key role in making breakthrough advances in important national priorities such as anti-terrorism research and development (R&D), networking and information technology R&D, nanotechnology R&D, biotechnology R&D, and climate change R&D, by integrating advanced computing with physical, chemical, and biological sciences.

By creating the new Computing and Computational Sciences Directorate, ORNL reaffirms its commitment to the critical area of scientific computing. In the words of Bill Madia, director of Oak Ridge National Laboratory, "ORNL has established expertise and ongoing research in developing new materials, studying global climate change and the effects of pollution, mapping human chromosomes and safety testing automobiles of the future in virtual reality. As we move those important national objectives forward, high-performance computing is essential. The creation of this directorate will improve the Laboratory's ability to advance in this area which is critical to the Laboratory's future."

DOE's Office of Advanced Scientific Computing Research has launched an exciting program on Scientific Discovery through Advanced Computing (SciDAC) to develop the scientific computing software and hardware infrastructure that will dramatically extend our exploration of the fundamental processes of nature, as well as advance our ability to predict the behavior of a broad range of complex natural and engineered systems. ORNL is a full partner in the SciDAC Program, collaborating with 13 laboratories and 50 universities in helping to create a new generation of scientific simulation codes that take full advantage of extraordinary computing capabilities such as those at DOE's Center for Computational Sciences at ORNL.


Jim Richmond, enhanced by Vickie Conner

ORNL recently received the first IBM Power4 Regatta H system for early evaluation.

Recently, ORNL took delivery of a 4-teraflops IBM Power4 supercomputer that can perform 4 trillion arithmetic operations per second. This addition to our existing 1.5 teraflops of computing represents an exciting tool of discovery in support of the new SciDAC program. The computer is already enabling breakthrough science in nanotechnology, global climate change prediction, and systems biology.

Within the next five years, computers 1000 times faster than those available to the scientific community today will be operating. These dramatic boosts in supercomputing power must be matched by corresponding increases in the capabilities of scientific modeling and simulation codes. Researchers across the Laboratory have teamed together to carry out the rigorous interdisciplinary effort of designing, building, tuning, integrating, and using codes to accelerate our solutions of complex scientific problems.

In this issue of the *Review*, we cover many areas in which ORNL researchers and our many collaborators are helping the Department of Energy meet its research needs now and well into the 21st century. We have several advantages on our side as we pursue this bold scientific agenda. We have an exciting research program in support of the SciDAC Program, we have a talented staff, and our work has applications in nanoscale science, biology, and climate change prediction. In addition, we are benefiting from the University of Tennessee's strong commitment to ORNL's computational sciences agenda and the state of Tennessee's financial support for the Joint Institute for Computational Sciences. As we build state-of-the-art computing facilities, bring new computers on line, and support exciting new science, we are creating a leading scientific enterprise that will help advance the revolution in science.

Thomas Zacharia
Associate Laboratory Director for Computing and Computational Sciences

A National Resource at ORNL: Supercomputers Support Superb Science

Supercomputers at ORNL are enabling scientists in a number of fields to make discoveries that could not be made through either theoretical or experimental research.

Jim Richmond

ORNL's Richard Alexander shows the inside of a dual-processor Pentium Xeon computer (which has the same layout as a supercomputer) to Boy Scouts (from left) Morgan Alexander, Liam Holland, Miles Pacheco, and Evan Meredith. Educating youth about computers may attract more talent into the growing area of computational sciences.

Oak Ridge National Laboratory is home to some of the most powerful open, or unclassified, computers in the nation. These recently acquired supercomputers will soon give the Department of Energy's Center for Computational Sciences (CCS) at ORNL a total computing speed approaching 6 teraflops, or 6 trillion arithmetic calculations per second. They have also advanced ORNL's leadership role in computational science and enabled scientific discovery that is not possible without high-performance computing.

The IBM Power4 supercomputer (dubbed Cheetah by ORNL researchers), delivered to ORNL in stages in late 2001, has a computing speed of 4 teraflops. The IBM RS/6000 SP (Eagle) supercomputer and Compaq AlphaServer SC (Falcon) system together offer 1.5 teraflops of computing power. As a result, researchers can solve complex scientific problems in a virtual computational environment. ORNL researchers also are involved in "collaboratories," in which scientists from different sites share data over computer networks as if they were working side by side.

The development of new algorithms by ORNL and University of Tennessee computer scientists will allow researchers to solve even more complex scientific problems more efficiently through simulations of experiments. Terascale computing is a powerful tool for analyzing, understanding, and predicting scientific phenomena because it serves as a bridge between theoretical understanding and experiment.

In this issue of the *ORNL Review*, the first section (Computing Infrastructure) describes the supercomputers at the heart of CCS and the computational infrastructure that is in place to support breakthrough computational science. We also describe our tools for monitoring and evaluating performance—for example, identifying the type of supercomputer on which a numerical code performs best and guiding decisions on which type of supercomputer to purchase next. We discuss our development of computer tools to enable scientists to run their codes more efficiently on supercomputers and the creation of a Scalable Systems Software Center for the preparation of software that will effectively manage terascale computational resources. We also discuss our contribution to the development of the computer industry's leading data-storage system—in terms of capacity and transfer speed—and Probe, our new storage research facility to improve data storage and transfer for terascale systems.

John Drake is participating in a climate simulation in a CAVE at ORNL.

Curtis Boles

Another article discusses our work in devising ways to improve our ability to send large data files over the Internet so that supercomputers are not idle because of delays in data delivery. ORNL is building a high-speed fiber-optic link that will connect Laboratory supercomputers with those in Atlanta, Georgia, and Research Triangle Park near Durham, North Carolina. Through this network, huge volumes of data and calculation results will be transferred among these supercomputers and, later, the new 100-teraflops Blue Gene supercomputer that ORNL is helping IBM to develop. The Blue Gene machine will be used to relate protein shapes to diseases. ORNL researchers will contribute their expertise in developing fault-tolerant algorithms and predicting protein structures.

Visualization tools for supercomputers provide insight into physical phenomena, help scientists verify calculated results, highlight the unexpected, and enable scientists to communicate their results more effectively. In this issue, we show that some visualization tools are being used in a CAVE™ virtual reality theater at ORNL, to enable scientists to interact with predicted phenomena that relate, for example, to stellar explosions and climate changes.

Our plans for a new building to accommodate all our research groups and our next-generation supercomputer are described, along with our partnerships with universities and industry and how they will benefit from the ORNL-University of Tennessee Joint Institute for Computational Sciences, which will also be housed in a new building.

The second section of this issue (Discovery by Computing) shows why computers are needed to enable scientific discovery and why scientific research that depends on high-performance computing is important to the nation. Supercomputer-supported research at ORNL falls into six areas of importance to DOE: astrophysics, biology, chemistry, climate prediction, fusion energy, and materials. Other computer-driven research at ORNL addresses questions concerning energy efficiency and health, as well as electronic devices that may result from advances in nanotechnology. In this issue you will learn that

- Studies of magnetic materials using supercomputers at ORNL are paving the way for the next generation of information technologies, including better digital cameras.
- A proposed molecular memory cell that would enable laptop computer batteries to last 100 times longer than today's batteries is being modeled computationally on an IBM supercomputer at ORNL. This machine is also being used to simulate carbon nanotubes in contact with other components that may be used in tomorrow's nanoscale electronic devices.
- Multidimensional simulations of core-collapse supernovae by ORNL and its partners could help determine how and why stars explode and how elements are formed and disseminated into space. This work in astrophysics could also advance the understanding of combustion, future climate, fusion energy, and radiation therapy.

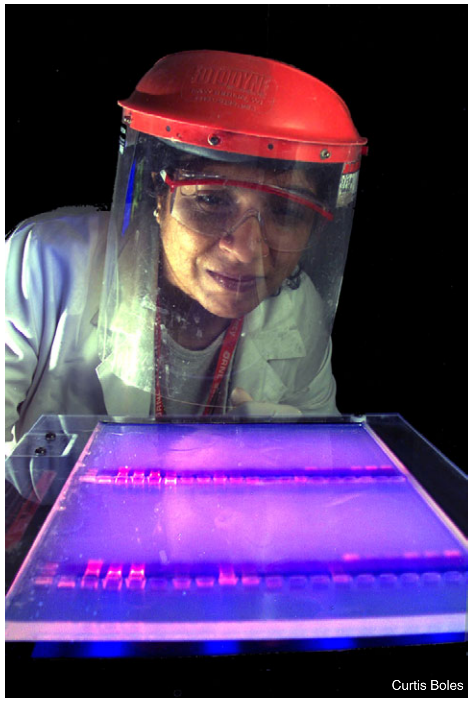
Curtis Boles

Madhu Dhar examines a gel used in DNA analysis at ORNL. Information on the order of chemical bases in DNA strands obtained through gel electrophoresis and other techniques is available in data bases accessible by computer. ORNL offers the latest information on DNA and protein analyses through its Genome Channel (http://compbio.ornl.gov/channel/index.html).

- A computational analysis of human and bacterial genomes by ORNL researchers provides insights into what our genes do. ORNL researchers will soon be computationally predicting 100 protein structures a day and will evaluate which compounds could make highly effective therapeutic drugs.
- Thanks to computer modeling, a scientific discovery was made that might lead to a way to save victims of cardiac arrest.
- Some of the world's largest global climate models are being run on ORNL's supercomputers, providing insights for national and international assessments of the effects of global warming caused by human activities.
- Supercomputers can be used to simulate chemical reactions, saving time and money and improving safety.
- ORNL fusion researchers are using supercomputers to understand plasma turbulence, design a device that could eliminate plasma disruptions, and find ways to get radio waves not only to heat but also to control the plasma to allow sustained energy-producing fusion reactions.
- In support of energy efficiency research, ORNL researchers are building computer models of vehicles made of aluminum, high-strength steel, regular steel, and carbon-fiber composites. Other researchers are developing software tools for supercomputers to simulate engine exhaust from various lean-burn diesel and gasoline engine designs as it flows through envisioned catalytic converters designed to chemically transform pollutants into harmless emissions. This research could lead to safer and cleaner, energy-efficient cars.
- Researchers in ORNL's Computational Sciences and Engineering Division are developing software in support of efforts to ensure our homeland security.

This issue provides a window into supercomputing at ORNL and shows how supercomputing helps scientists better "see" various phenomena and material structures that could lead to more detailed scientific understanding and improved devices and drugs. Clearly, ORNL is advancing science through advanced computing.

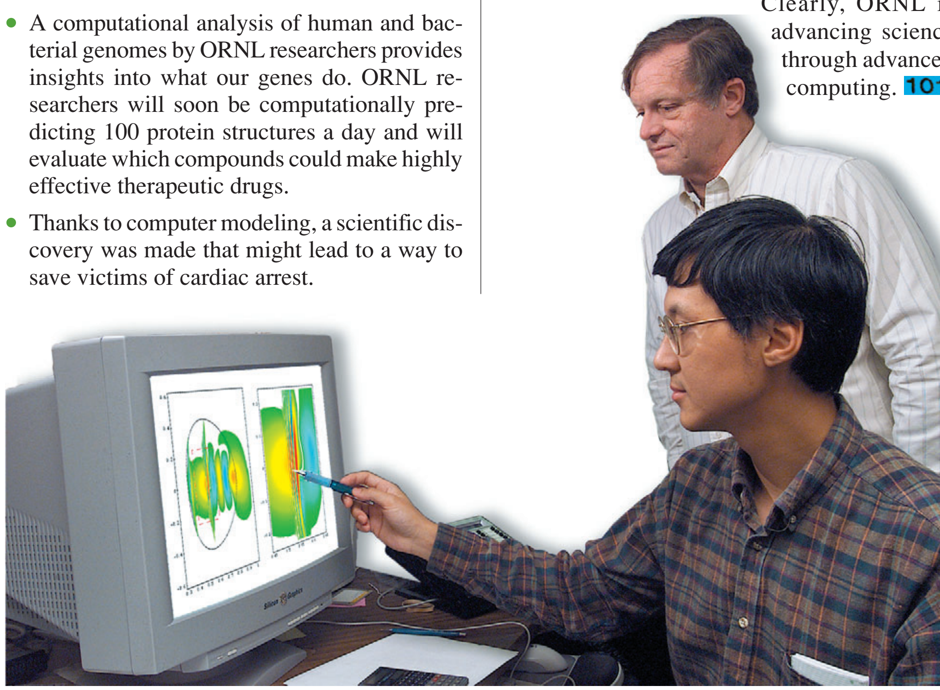
Jim Richmond; enhanced by Gail Sweeden

Fusion theorist Don Batchelor (standing) and mathematical consultant Ed D'Azevedo have collaborated on a project in which a supercomputer is used to help understand how electromagnetic waves can be used to heat and control high-temperature plasmas in a fusion research device.

ORNL's Powerful Tools

ORNL offers 5.5 teraflops of computing to advance scientific discovery. Some ORNL computational research is funded by DOE's Scientific Discovery through Advanced Computing Program.

Speed in supercomputers is an elusive commodity. Only seven years ago, in January 1995, ORNL installed what was then the world's fastest computer. The Intel XP/S 150 was capable of an astounding 150 billion arithmetic operations per second (150 gigaflops) and at $16 million, seemed like a bargain. Today ORNL has three large supercomputers with a total capacity of 5.5 trillion calculations per second (5.5 teraflops). The three massively parallel supercomputers at DOE's Center for Computational Sciences (CCS) at ORNL are dubbed Eagle, Falcon, and Cheetah (the fastest mammal). A single node of ORNL's new IBM Power4 computer, Cheetah, is faster than the Intel machine of 1995.

Eagle is an IBM RS/6000 SP supercomputer and Falcon is a Compaq AlphaServer SC computer. Both were installed and up and running in April 2000. Eagle performs 1 trillion arithmetic operations per second, making it a one-teraflop machine. Falcon has a peak performance of 0.5 teraflop. Both supercomputers have four central processor units (CPUs) per node, which includes memory, disk drives, and input-output capabilities. Cheetah is an IBM Power4 supercomputer; it arrived at ORNL in late 2001. It provides an amazing 4 teraflops of computing power. Cheetah is the first system to use IBM's new Power4 "Regatta-H" nodes. Each of Cheetah's 24 nodes has 32 processors. The total system has more than a terabyte of memory and 40 terabytes of disk space. A comparison of the capabilities of ORNL's three supercomputers is shown below:

	CHEETAH (IBM)	EAGLE (IBM)	FALCON (Compaq)
Teraflops	4	1.08	0.5
Nodes	24	184	64
Disk space (TB)	40	9.2	4.5
Memory (GB)	1020	368	128

Eagle, which was the tenth-fastest supercomputer in the world in April 2000, is used for running codes that address complex problems in astrophysics, biology, climate prediction, computational chemistry, fusion energy, and materials science. Falcon is used to evaluate code and machine performance (see p. 9), as well as for large scientific runs in the research areas of astrophysics, computational chemistry, and materials science.

As Ernie Moniz, then Department of Energy undersecretary, said at the June 20, 2000, dedication of ORNL's new supercomputers, these machines are "extraordinary tools for extraordinary science." He noted further that "simulation using teraflop computers will be a tool of scientific discovery. Simulation will play an important role in the bridging from the molecular level to engineering systems to get the needed efficiencies..." to solve energy, environmental, materials, and medical problems. ORNL is providing supercomputing resources to researchers supported by DOE's Scientific Discovery through Advanced Computing Program.

Already calculations performed using ORNL computers and codes have led to the location of disease-causing genes (including the gene that causes the disease suffered by the title character in the movie *Lorenzo's Oil*). Computer modeling at ORNL helped IBM better understand the use of giant magnetoresistance (GMR) for reading data while the company was developing quarter-size disk drives for digital cameras. GMR simulation at ORNL has also influenced Seagate's new designs of disk drives for desktop computers. Our codes and nodes may help predict which lightweight materials considered candidates for future, highly-efficient cars will likely hold up as well in a crash as today's heavier steel cars; this approach saves

New Home Planned for 10-Teraflops Supercomputer

Rendering by John K. Jordan

Artist's conception of ORNL's planned new Computational Sciences Building (right), which will be connected to the Research Office Building (left). These privately funded buildings will be completed in 2003.

The next big thing in the supercomputer world will, indeed, take up more space. You can't go from a trillion calculations to ten trillion calculations per second without building a bigger machine that has many more parallel processors. Because ORNL's goal is to acquire a 10-teraflops supercomputer by 2003, a new, larger building is required.

Buddy Bland of ORNL's Computer Science and Mathematics Division, says, "To build a many-teraflops parallel machine, tens of thousands of processors must be grouped together. Such a machine will need a vast amount of space. Additionally, the computer will require larger amounts of power and more air-conditioning equipment."

For example, the Intel Paragon at ORNL, which was the world's fastest supercomputer in early 1995, took up 500 square feet of space. Eagle takes up 1000 square feet. The space in Building 4500-North available for Eagle, Falcon, Cheetah, and the necessary data storage units amounts to 7500 square feet. The fastest supercomputer in 2001 takes up 12,000 square feet.

"Our new building at ORNL to house DOE's Center for Computational Sciences will have 40,000 square feet for the 10-teraflops supercomputer and other equipment that we are planning to have in 2003," Bland says. "We will need 2.5 megawatts of power and 750 tons of air conditioning for the computer center.

"To meet DOE's future requirements for doing science using high-performance computing, we must have state-of-the-art facilities. Fortunately, UT-Battelle is committed to building these facilities as part of its modernization plan for ORNL."

for Scientific Discovery

Falcon: The Compaq AlphaServer SC supercomputer at ORNL. Buddy Bland stands in back.

Both photos by Curtis Boles; enhanced by Mark Robbins and Gail Sweeden

money and reduces waste since crunching numbers is much cheaper than crunching cars in real crash tests.

Because ORNL's supercomputers will generate a mind-boggling number of results from a trillion calculations per second, systems must be in place to obtain, store, catalog, retrieve, and transmit over long distances these huge quantities of data. ORNL and its collaborators have developed a state-of-the-art data storage-and-retrieval system called the High Performance Storage System (see p. 10). In addition, ORNL researchers are devising ways to move large chunks of data (up to 2 terabytes) more quickly and efficiently over high-speed network links (see p. 11) between ORNL and the National Energy Research Scientific Computing Center (NERSC) at DOE's Lawrence Berkeley National Laboratory in California. A high-speed fiber-optic link 10,000 times faster than today's fastest networks is also being set up to connect ORNL, Atlanta, and the Research Triangle in North Carolina. This network may connect with the first IBM Blue Gene supercomputer now being developed with help from ORNL researchers. When operational in 2005, this 100-teraflop supercomputer will help biologists understand the complex rules by which proteins assume their shapes (which are related to disease), as well as predict future climate as the earth's atmosphere is loaded with increasing amounts of carbon dioxide, and aid in the design of nanoscale electronic devices.

Trey White (front) and Tim Jones work at the supercomputer.

Eagle: The IBM RS/6000 SP supercomputer at ORNL.

Both photos by Curtis Boles; enhanced by Reneé Balogh and Gail Sweeden

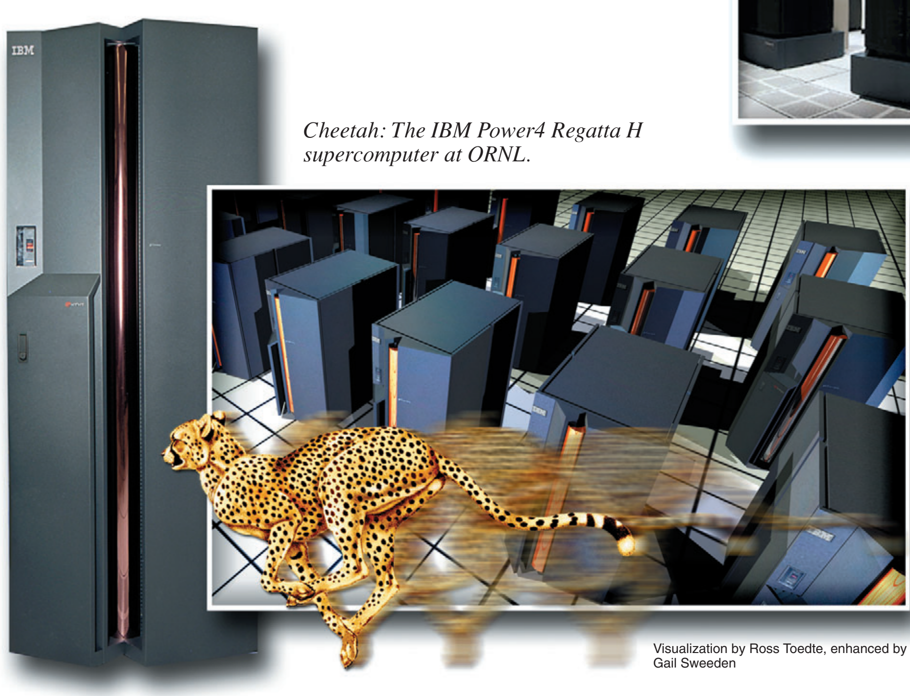

Cheetah: The IBM Power4 Regatta H supercomputer at ORNL.

Jim Richmond; enhanced by Vickie Conner

Visualization by Ross Toedte, enhanced by Gail Sweeden

Climate modelers use ORNL and NERSC supercomputers to simulate the earth's climate for past, present, and future greenhouse-gas scenarios. Climate models solve complex mathematical equations that describe atmospheric and oceanic circulation, temperature, pressure, and many other variables over the entire earth. These calculations must simulate time in 20-minute increments, and many simulations are for hundreds of years. This large number of calculation results requires supercomputers, as well as high-capacity data storage and networking capabilities.

Because of their expertise and equipment, ORNL and other researchers are complementing theory and experiment as they use CCS's powerful tools for scientific discovery. 101

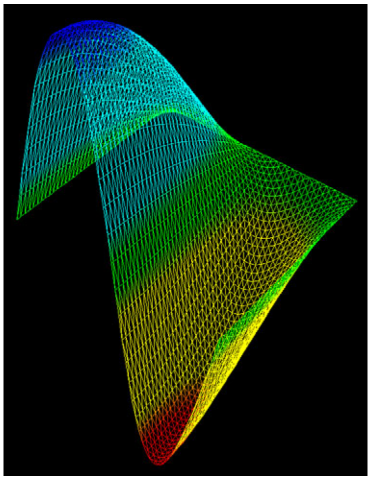

Developing Computer Tools for Scientists

ORNL researchers and their university and national lab colleagues are developing tools to enable scientists to run simulation codes more efficiently on massively parallel supercomputers and clusters of personal computers.

"Having a supercomputer that doesn't have any software that lets you use it is like having a fast car that you have locked your keys inside," says Al Geist, a group leader in ORNL's Computer Science and Mathematics Division (CSMD). "Supercomputing tools are the keys that help scientists unlock the speed inside the nation's fastest computers."

To help unlock this speed, the Department of Energy recently started the Scientific Discovery through Advanced Computing (SciDAC) Program to help create a new generation of scientific simulation codes. The codes will take full advantage of extraordinary terascale computer resources that can perform trillions of calculations per second and handle trillions of bytes of data to address complex scientific problems. These codes for massively parallel supercomputers will be used to address increasingly complex problems in climate modeling, fusion energy sciences, chemical sciences, nuclear astrophysics, high-energy physics, and high-performance computing. ORNL is involved in several SciDAC projects aimed at developing supercomputer tools for scientists.

The performance evaluation project focuses on finding the best ways to execute a specific application on a given platform (see article on p. 9). The tools from this effort will answer three fundamental questions: What are the limits of performance for a given supercomputer? How can we accelerate applications toward these limits? How can this information drive the design of future applications and high-performance computing systems? ORNL has a long history of evaluating early prototype systems from supercomputer vendors. The most recent ORNL acquisition is an IBM Power4 system that arrived so new it didn't even have an IBM product name. ORNL has already determined how this system will perform on a variety of scientific applications.

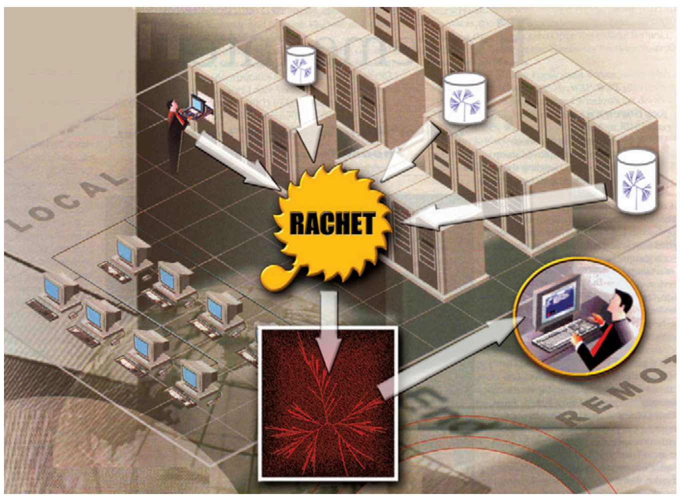

Nagiza Samatova has developed Rachet, a petascale distributed-data-analysis suite. It is designed for scientific data that are massive, distributed, dynamic, and high dimensional. This highly scalable approach allows users to make computations from local analyses, merge information with minimum data transfer, and visualize global results. Rachet can be applied to analyses and predictions in the scientific areas of climate, genomics, astrophysics, and high-energy physics.

A growing trend among scientists is to buy a bunch of personal computers (PCs) and "cluster" them together to run their applications. But just as the right key is needed to run the fast car, cluster computing software is required to make the PCs work as one computer. The Scalable Systems Software Center (see sidebar below) leverages a lot of the work that ORNL has done in cluster

ORNL Leads Effort to Improve Supercomputer Centers

Supercomputers provide researchers with powerful tools, but managing them can be quite difficult, says an ORNL researcher who heads a team working to fix the problem.

Through an $11-million, ORNL-led project, a team composed of people from DOE labs, several universities, and industrial firms has created the Scalable Systems Software Center. The center, funded by DOE's Scientific Discovery through Advanced Computing Program, will address the lack of software for effective management of terascale computational resources, such as the ones being installed at ORNL and other sites around the country.

"DOE operates many of the world's largest computers and computer facilities," says Al Geist of ORNL's Computer Science and Mathematics Division. "Today, each computer facility uses ad hoc and homegrown systems software solutions to, for example, schedule jobs and monitor the health of the supercomputers.

"The Scalable Systems Software Center provides the opportunity to create and support a common set of systems software for all the large computer facilities across the country. It's a problem that the computer industry isn't going to solve because business trends push the industry toward smaller systems aimed at Web serving, database farms, and departmental-sized systems."

The vision and goal of the Center are to bring together a team of experts who, with industry involvement, can agree on and specify standardized interfaces between system components. The Center will also produce a fully integrated set of systems software and tools to effectively use terascale computational resources.

Researchers also plan to study and develop more advanced versions of the system tools to meet the needs of future – and even larger—supercomputers.

computing. For instance, ORNL initiated and leads the Open Source Cluster Application Resources (OSCAR) project. "The interest in this software has been phenomenal," says CSMD's Stephen Scott, who leads the project. "In the first two months after the OSCAR toolset was released, more than 12,000 people downloaded it!"

OSCAR is a snapshot of the best-known methods from across the nation for building, programming, and using clusters. It consists of a fully integrated, easy-to-install software bundle designed for high-performance cluster computing. Everything needed to install, build, maintain, and use a modest-sized Linux cluster is included in the suite, making it unnecessary to download or even install any individual software packages on a cluster. OSCAR team members are now busy working on the Scalable Systems Software project, for which they plan to build the same kind of easy-to-use tools for supercomputers.

"Sure, computers can run fast and make lots of calculations, but if you don't have the tools to analyze the terabytes of data they produce, you are still going nowhere," says CSMD's Nagiza Samatova, who is one of the investigators on the SciDAC Scientific Data Management project. This project's goal is to optimize and simplify access to very large, distributed, heterogeneous datasets and to use data mining to extract meaningful data from these datasets. Samatova has developed a new algorithm for analyzing biological data to determine metabolic pathways that cuts the run time from 3 days to 2 minutes.

"This innovative algorithm is a perfect example of how computer science and mathematics expertise can make breakthrough tools available to the scientists," says Thomas Zacharia, ORNL's associate laboratory director for Computing and Computational Sciences. "It is one of the things that makes ORNL's Center for Computational Sciences so successful." 101

From the Stone Age to the Lego Block Age of Computing

The Center for Component Technology for Terascale Simulation Software (CCTTSS), another of ORNL's projects to receive funding from DOE's Scientific Discovery through Advanced Computing (SciDAC) Program, may well revolutionize the way terascale software simulations are developed. Traditionally, large software codes have been chiseled from a single, monolithic "rock." The CCTTSS staff strive to pull scientific software development out of the "stone age" and build high-performance simulations, using flexible and powerful "component" building blocks that can be assembled in myriad ways—like Lego blocks—while maintaining high levels of performance.

As an offshoot of the ongoing Common Component Architecture (CCA) Forum effort, the CCTTSS [a $16-million joint project including Sandia, Argonne, Lawrence Livermore, Los Alamos (LANL), and Pacific Northwest national laboratories, as well as Indiana University and the University of Utah] is creating high-performance, component-based software specifications and frameworks for use in next-generation scientific simulations. "While there is a variety of component systems available in industry, from Corba to DCOM to JavaBeans, they just aren't built for *speed* to suit the needs of scientific computing," says James Kohl, a co-principal investigator (Co-PI) and ORNL leader for the center.

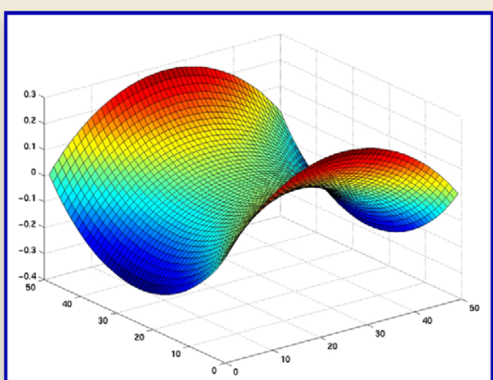

In this scientific simulation using high-performance software components at last November's Supercomputing 2001 conference in Denver, Colorado, independently developed software modules work together in this "Common Component Architecture" demonstration of an unconstrained minimization problem (minimal surface area given boundary constraints) that combines Argonne National Laboratory's TAO solver component with ORNL's CUMULVS visualization and "MxN" components.

Components encapsulate the functionality of one or more traditional software "objects" to make it easier to plug in components as large software systems are assembled. The use of components allows each code module in a complex software simulation to be "swapped out," even at run time, to change the capabilities or behavior of the simulation. For example, a particular linear solver component might be replaced with another one having the same functional interface but with an algorithm better suited to the given matrix properties.

CCTTSS and CCA staff members are developing tools that support the high-performance component interactions required for many new terascale scientific simulation projects for DOE's Office of Science. Special components and services are being constructed to assist in the description and manipulation of parallel datasets (where individual data elements are decomposed and distributed across a number of parallel processors).

Because many CCA-compliant components are themselves parallel entities, advanced technology is required to share and exchange their data in parallel. ORNL leads the CCTTSS effort in this so-called "MxN" (pronounced "M by N") parallel data redistribution technology, which allows one parallel model running on "M" processors to transfer data to another parallel model on "N" processors. ORNL is working closely with researchers at LANL to generalize existing technology for this parallel model coupling; two initial MxN component implementations are based on the CUMULVS system (ORNL) and the PAWS system (LANL).

The key to success is getting real scientific applications to exploit the advantages of a high-performance component environment. "The CCTTSS is reaching out to scientists and teaching them how to use this new CCA technology, and we're also actively supporting the integration of CCA into their simulation codes," says David Bernholdt, ORNL Co-PI and CCTTSS applications liaison.

The CCTTSS has already formed a number of collaborations, with national efforts in key application domains, including global climate modeling, computational chemistry, astrophysics, and fusion. Many other SciDAC centers are also working closely with the CCTTSS to incorporate component technology into the development of their simulation code and to use CCA-compliant components to deploy their tools and systems.

The CCTTSS had its first big public showing at the Supercomputing 2001 conference in November 2001 in Denver, Colorado. ORNL's booth hosted the main multi-lab CCA/CCTTSS presentation, which demonstrated three different scientific applications built from a variety of high-performance components. Several prototype components developed by ORNL, including "MxN," "Distributed Data Descriptor,"and "Visualization Proxy," were used in each of three simulations, illustrating the potential "re-use" and interoperability of DOE technology via CCA components. A public release of this initial software will be available as a tutorial in early 2002. (See the www.cca-forum.org web site for more details.)

ORNL is working with IBM to develop the Blue Gene supercomputer for relating protein shapes to disease.

ORNL, IBM, and the Blue Gene Project

Advanced cellular architecture in the next-generation supercomputer will help scientists better understand the makeup and purpose of different genes and proteins in living cells. Massive computing power and the intricacies of biological matter at the molecular level will be colliding through a cooperative research and development agreement (CRADA) announced August 22, 2001, by ORNL and IBM and funded by IBM and the Department of Energy.

At the heart of the agreement is IBM's Blue Gene research project, which combines advanced protein science with IBM's next-generation cellular architecture supercomputer design. Unlike today's computers, cellular servers will run on chips containing "cells," which are processors that contain memory and communications circuits. Cellular architecture will help scale computer performance from a teraflop (1 trillion calculations per second) to a petaflop (1000 trillion calculations per second).

The new supercomputer will be a petaflop machine. The fastest existing computer, ASCI White, unveiled by IBM in early August 2001, can perform about 12 trillion calculations per second, or 12 teraflops. That computer is being used for nuclear weapons stockpile stewardship research at DOE's Lawrence Livermore National Laboratory. IBM, also known as Big Blue, began its five-year, $100 million Blue Gene project at the end of 1999; its goal is to create a supercomputer that can handle large-scale computing projects.

Supercomputing power of this magnitude (1 petaflop) will improve scientists' ability to predict future climate, advance the field of nanotechnology, and gain a better understanding of how gene sequences and the folding of proteins relate to diseases.

"Proteins control all processes occurring in the cells of the body," says Joe Jasinski, manager of the Computational Biology Center for IBM Research. "These proteins are made up of a vast array of different combinations of amino acids that fold and bend into very complex, three-dimensional shapes that determine the exact function of each protein. If the shape of a protein changes because of some environmental, physical, or biological factor, the protein may turn from being beneficial to one that causes a specific disease."

The understanding of the protein-folding phenomenon is a recognized "grand challenge problem" of great interest to the life sciences. The scientific knowledge derived from research on protein folding can potentially be applied to a variety of problems of great scientific and commercial interest, including protein-drug interactions, enzyme catalysis, and refinement of protein structures created through other methods.

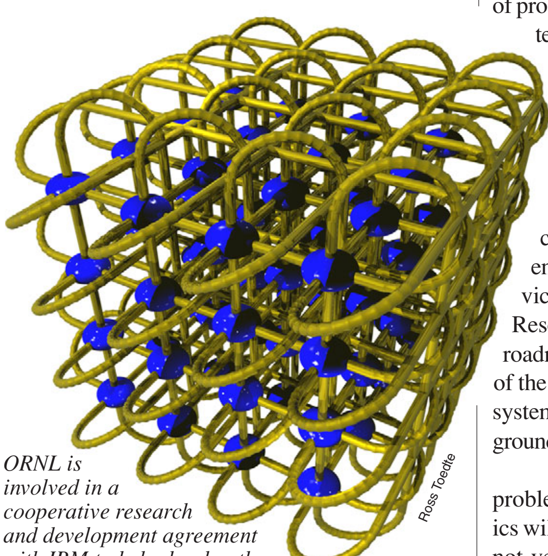

ORNL is involved in a cooperative research and development agreement with IBM to help develop the Blue Gene supercomputer that will improve our understanding of how living cells work. This supercomputer will use advanced cellular architectures to allow 1000 trillion calculations per second (petaflops computing). ORNL researchers will write programs to help the machine run effectively.

"Our collaboration with Oak Ridge National Laboratory is vital to IBM's work to extend the boundaries for applications of large-scale computing, focusing on the combination of IBM and ORNL's deep scientific capabilities," says David McQueeney, vice president of Emerging Business for IBM Research. "Together we have built a common roadmap for an ambitious, multi-year evolution of the simulation and modeling of many complex systems. We are confident that we will break new ground in several domains, including life sciences."

"The complexity of the protein-folding problem, nanoscale science, and climate dynamics will require computational resources at a scale not yet achieved by any scientific application," says Thomas Zacharia, ORNL's associate laboratory director for Computing and Computational Sciences. "This is an exciting next step in ORNL's history of evaluating new computational architectures and pushing the computational science envelope." Before it will be possible to solve problems in biology, climate, and nanotechnology, scientists must devise methods to run applications that use tens of thousands of processors in the Blue Gene supercomputer. Each processor forms a cell with memory, communication, and input/output built in. This approach departs from past designs and offers a glimpse of what's to come in high-performance computing.

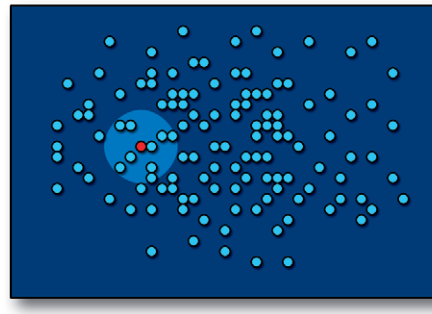

A Blue Gene IBM supercomputer with 100,000 processors could fail at a rate of one processor every few seconds. ORNL researchers are establishing a theoretical foundation for a whole new class of fault-tolerant algorithms that are scalable beyond 100,000 processors—that is, they allow the supercomputer to proceed with calculations and work around the processors that have failed. Each program is a cell in a larger job where each cell interacts with a fixed number of other cells.

"The world of supercomputing is rapidly changing," says Ed Oliver, associate director in the Department of Energy's Office of Advanced Scientific Computing Research. "We need to develop approaches to solving computational problems that are able to scale to thousands of processors and at the same time be tolerant of failures of some of these processors."

Working with IBM, ORNL researchers led by Al Geist of the Computer Sciences and Mathematics Division will develop fault-tolerant algorithms to allow the Blue Gene supercomputer to work around processors that fail, as well as other capabilities, to ensure that the machine operates effectively. ORNL scientists led by Ying Xu of the Life Sciences Division will collaborate with IBM on how the supercomputer should be programmed to analyze proteins and predict their structures.

IBM and ORNL hope to use this enormous computing power to explore numerous other areas, as well. This effort merely represents the beginning of what is expected to be a long relationship. 101

Evaluating Supercomputer Performance

In selecting the right machine on which to run a code or the next machine to buy, evaluators of supercomputers focus on many parameters ranging from speeds to feeds, teraflops to terabytes.

Like buildings, supercomputers have different architectures. Picture four computer processing units (CPUs) and four data-storage units (computer memory). Give each processor its own memory unit and then connect the processors. If one processor wants to read data in memory attached to another processor, it must ask the other processor for the data. This arrangement is called distributed memory, and the collection of processors is called a cluster. If, instead, each processor is connected to each of the four memory units and can access data directly, this arrangement is called shared memory. Now, put these four processors and their memory units in one box in a shared memory arrangement and call it a node.

Eagle, the IBM RS/6000 SP supercomputer at ORNL, is a cluster of 176 four-processor nodes, combining both distributed and shared memory in a single system. In 1998 when DOE's Center for Computational Sciences (CCS) at ORNL was planning to purchase a next-generation supercomputer for ORNL, it signed a contract with IBM that called for 16-processor nodes. At the time, a performance evaluation team led by Pat Worley of ORNL's Computer Science and Mathematics Division (CSMD) compared the 4-processor and 16-processor IBM nodes, to determine which architecture would work best for the codes that were to be run on the machine.

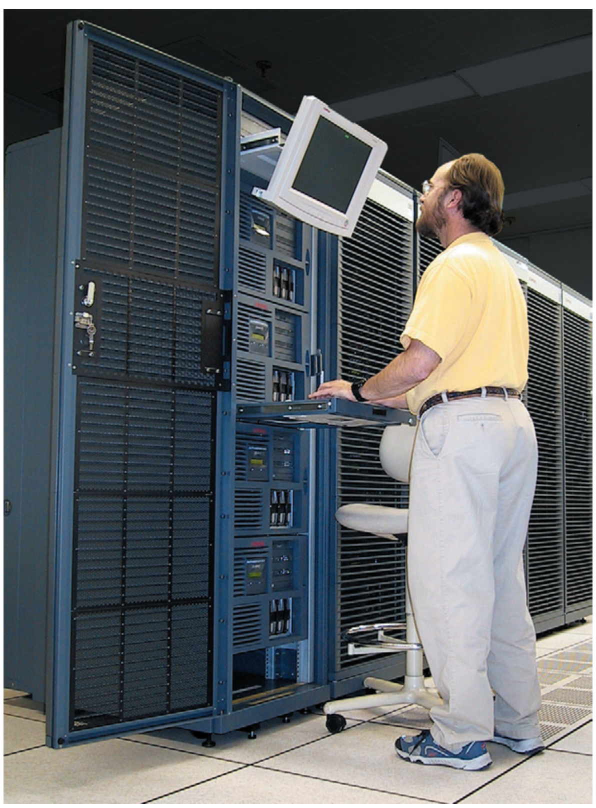

Buddy Bland

Pat Worley works at the console of the Compaq AlphaServer SC supercomputer.

"We found that smaller nodes work better for our science applications," says Buddy Bland, head of CSMD's Systems and Operations Group. "So we changed the contract with IBM from 16-processor nodes to 4-processor nodes. As a result, we obtained Eagle eight months earlier at a cost 20% less than the total in the original contract. Now, we must decide which architecture will work best for the 10-teraflop supercomputer we want built for 2003. Already climate modelers are writing codes that will port to this future supercomputer."

Worley and his team have focused their recent performance evaluation efforts on the Compaq AlphaServer SC machine at ORNL known as Falcon and on a new IBM machine at ORNL known as Cheetah. Falcon uses 4-processor nodes, like Eagle. Cheetah uses the new IBM p690 nodes, which each have 32 processors. Worley's team has found that IBM Power4 processors used in a p690 node are two-and-a-half times faster than Eagle's processors and twice as fast as Falcon's processors for a variety of application codes.

Unlike the earlier 16-processor IBM nodes, the 32-processor, p690 node has up to 4 times better bandwidth than Eagle for communication within a node. Hence, a larger volume of messages and other data can be passed more quickly among Cheetah's processors than among Eagle's. As Bland puts it, "If you have a really fast water pump, you want a fire hose, not a straw, to increase the speed and volume of flow. Cheetah has the bandwidth equivalent of a fire hose."

As part of their performance evaluation, Worley and his team do "benchmarking." They test existing parallel-computing codes to determine whether each code runs faster on, for example, the IBM or Compaq machines. Then they "diagnose" the performance of the code.

"We try to determine why a code runs faster on one machine than another," Worley says. "We investigate whether a code may run more slowly on one machine because of the coding style—the way a computer program is written. If so, we can advise code developers on how to alter their style so the code will run faster on a particular machine."

ORNL team members also do performance engineering. They can tune a code to improve its performance on a specific machine. In addition, Worley's team tells vendors which problems they need to solve in designing their next-generation machines so that certain codes will run faster.

"Our customers are code writers and users, vendors, and system administrators," Worley says. "We provide advice on how to configure and run their systems and on what machines they should buy next. We guide the development of both codes and supercomputers.

"In our most recent efforts we have focused on evaluating the performance of Falcon and Cheetah in running climate, car crash, computational chemistry, human genome analysis, and materials codes. We measure how fast each code runs and predict how much time and how many processors are needed to get the computing job done."

The ORNL team was the first to show that a supercomputer made in the United States (Falcon) could exceed a performance goal (5 seconds per model day) for modeling the global climate. Later the team also showed that Eagle can exceed that goal.

Without the input of the CCS performance evaluation team, ORNL's supercomputers would not have nearly as good an output. 101

Retaining and Retrieving Data More Effectively

ORNL is a co-developer of and customer for the computer industry's leading data-storage system in terms of capacity and transfer speed. The ORNL data-storage program also includes the Probe Storage Research Facility.

A scientist needs data about how different types of clouds reflect, absorb, and transmit the energy of sunlight. The data, based on measurements taken by instruments on the ground and aboard airplanes and satellites, will help the scientist improve the accuracy of a computer model in predicting the influence of industrial emissions of greenhouse gases on global warming.

The scientist accesses a Web-based interface and requests 100 files of data from the Department of Energy's Atmospheric Radiation Measurement (ARM) data archive, located at ORNL. In this archive are tape drives (for slower-speed but higher-capacity storage) and disks (for high-speed access.) They contain more than 4 million files representing more than 25 terabytes of data. Three robots retrieve the tapes on which the requested files are stored and load them for copying on the disk drive of the ARM Web site server. Within an hour, the scientist can access the requested files.

Randy Burris examines the archive of tape drives and disks of a High-Performance Storage System (HPSS) at ORNL. The HPSS, which was developed by ORNL and several partners, is storage-system software that leads the computer industry in data capacity and transfer speeds. At ORNL the HPSS is used for DOE's Atmospheric Radiation Measurement, or ARM, data archive. This archive contains more than 4 million files representing more than 25 terabytes of data.

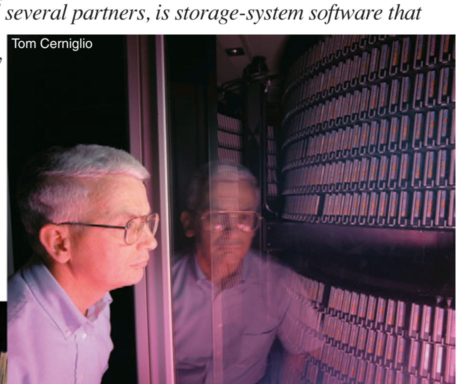

For the past four years, the ARM data archive has been using the High-Performance Storage System (HPSS), storage-system software that leads the computer industry in data capacity and transfer speeds and is standard for storage systems in the high-performance computing community. The ARM project is one of two large customers for HPSS at DOE's Center for Computational Sciences (CCS) at ORNL, where Laboratory researchers provide and support the data archive. The HPSS system manages the hierarchy of devices, storing more than 3.5 billion measurements. It can place 12,000 new files a day into storage. It will eventually be able routinely to find and retrieve up to 5000 files an hour to meet the growing requests for information related to global change.

The other large customer is the group of climate prediction modelers using ORNL's supercomputers. They can produce a run of calculations generating 1 terabyte of data that needs to be stored. These results may also be sent from ORNL to the data archives at DOE's National Energy Research Scientific Computing Center (NERSC) in California in chunks of 250 megabytes.

HPSS was developed by a consortium of DOE national laboratories and IBM. The DOE participants are ORNL, Sandia, Lawrence Berkeley (LBNL), Los Alamos, and Lawrence Livermore national laboratories. HPSS, which received an R&D 100 Award in 1997, is marketed by IBM. ORNL researchers Deryl Steinert, Vicky White, and Mark Arnold have been developing the graphical user interface between the operator and the HPSS for running, monitoring, and otherwise managing the system. More than 70 terabytes are now stored in ORNL's production HPSS installation, managed by Stan White, Nina Hathaway, and Tim Jones.

The ORNL mass-storage program also includes the Probe Storage Research Facility, operated by Dan Million. In one probe project, researchers Nagiza Samatova and George Ostrouchov investigate the use of data mining to extract meaningful information from massive scientific datasets.

Probe resources are also used for developing new software to send larger chunks of data more rapidly over the network to such facilities as the CAVE virtual reality theater at ORNL (see p. 12).

"Our Probe staff recently accomplished one of our goals," says Randy Burris, manager of data storage systems for CCS. "Thanks in part to work by ORNL researchers Tom Dunigan and Florence Fowler on network protocols, we are now using the bandwidth between CCS and NERSC more effectively. We are now transmitting more than 12 megabytes per second over ESnet, DOE's semiprivate portion of the Internet."

Probe researchers also have a role in several projects funded by DOE's Scientific Discovery through Advanced Computing (SciDAC) program. For the Scientific Data Management Integrated Software Infrastructure Center, a SciDAC project led by Arie Shoshoni of LBNL, Probe resources will be used to develop ways to improve data access and transfer and to test and implement other concepts. Probe resources are also being used in the DOE Science Grid and the Earth Systems Grid II projects. The SciDAC project on climate prediction, led by John Drake, will be using the Probe facility to determine how to transfer bulk amounts of data over the wide-area network. In work for the SciDAC project on astrophysics modeling, led by Tony Mezzacappa, Ross Toedte will be using Probe resources as he develops an effective visualization of the details of a stellar explosion. Finally, Net100 project researchers will use Probe resources as they seek to improve computer operating systems so excellent network throughput will be easily achievable without extensive application-specific tuning.

The production and research elements of ORNL's mass-storage program are providing and promising valuable services to computational scientists throughout the Laboratory. 101

Networking: Making Faster Connections among Supercomputers

ORNL researchers are devising ways to move large data files faster over computer networks and to reduce delays in data delivery so supercomputers are not idle.

Illustration by Gail Sweeden

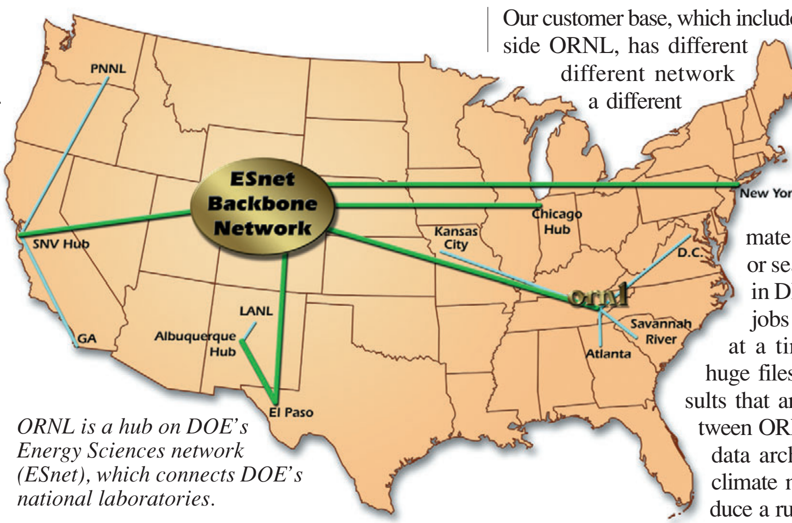

ORNL is a hub on DOE's Energy Sciences network (ESnet), which connects DOE's national laboratories.

Some computational scientists' high-performance computing jobs are getting done even while they are not working, thanks in part to networking. Networks, particularly high-speed networks, allow supercomputer nodes to "talk" to each other, send messages to other nodes asking for data, and transmit large data files across the country. In addition, networks allow computational scientists to keep tabs on the progress of long-running jobs that can often run for days at a time.

DOE's Center for Computational Sciences (CCS) at ORNL has supercomputers, as do the National Science Foundation center in Pittsburgh, the National Center for Supercomputing Applications facility at San Diego, and DOE's National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL) in California. But, according to Bill Wing of ORNL's Computer Science and Mathematics Division, the similarities stop there.

"Other supercomputer centers slice up their resources on a fine scale and run hundreds of jobs for thousands of users," he says. "We are different. We focus our computational resources on a few high-end users who need massive computing capacity for climate prediction, human genome analysis, and materials science simulations. Our customer base, which includes many users outside ORNL, has different needs, including different network needs, so we use a different model."

At CCS the computational scientists model future climate or exploding stars or searching for genes in DNA sequences run jobs for days or weeks at a time and generate huge files of calculated results that are transmitted between ORNL and NERSC's data archives. Sometimes climate modeling can produce a run of data amounting to 1 trillion byes (1 terabyte). These data are sent between ORNL and NERSC in chunks of 250 million bytes (megabytes).

"We focus on moving large files of data," Wing says. "ORNL and LBNL are writing computer programs to ensure that these data packages slide through the network rather than clog it. In addition, we are developing the ability to allow users to monitor the progress of these long-running jobs—and steer them if necessary—from a variety of portable access points, including laptops and personal digital assistants like Palm Pilots or iPAQs."

Data are sent over the network mostly using the transmission control protocol (TCP), a predefined protocol that computers use to communicate over a network. LBNL and ORNL researchers are devising ways to improve the ability to send large files so that supercomputers are not idle because of delays in data delivery.

To reduce delays in data delivery, Nageswara Rao of ORNL's Computer Science and Mathematics Division has developed a computer program called NetLet that is being tested on 12 free telnet and university sites serving as monitors and routers. "NetLet allows computers to efficiently talk with each other, 'predict' the delay in getting the message to the receiver, and suitably route the message," Rao says. "This algorithm enables the computers to measure connection speeds and the delays of pathways and then identify the best combination of pathways to get the information delivered efficiently in the time or at the rate guaranteed."

Demonstrations of NetLet have shown that the algorithm has improved the speed of data delivery by about 40% without any additional support from the Internet routers. "Some of our data files used to take 10 seconds to get from our computer to a destination computer," Rao says. "Those same data files can now get there in 6 seconds. That means that a huge data file that took 10 hours to arrive at a destination computer can now get there in 6 hours."

The data files transmitted from ORNL's Eagle (IBM RS/6000 SP supercomputer) to the NERSC data archive fly over DOE's Energy Sciences Network (ESnet), a semiprivate part of the Internet. Currently, DOE facilities such as ORNL and LBNL are using the new ESnet (OC12), a high-speed link operated by Qwest that supports data transmission at 622 megabits per second—4 times faster than the old ESnet. 101

Networking for More Powerful Supercomputers

ORNL is building a high-speed fiber-optic link that will connect Laboratory supercomputers with those in Atlanta, Georgia, and Research Triangle Park near Durham, North Carolina. The network, which will be about 10,000 times faster than a traditional T-1 connection, is expected to be operating in early 2002.

According to ORNL's Bill Wing, this new network will connect to the Atlanta GigaPOP (a regional aggregation point at which high-speed networks can exchange data with each other), giving ORNL the ability to link with several other high-performance networks, including the Internet2 and perhaps the Supernet managed by the U.S. Defense Advanced Research Projects Agency. In addition, the new network will serve as the core of a regional research network.

Visualization Tools: Interacting with Data in Many Dimensions

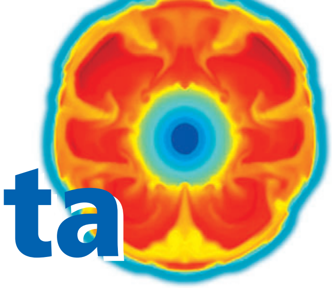

Visualization tools used in the CAVE™ virtual reality theater at ORNL let scientists interact with predicted phenomena such as stellar explosions and climate changes.

Imagine walking into a CAVE virtual reality theater and feeling as if you're flying in a kaleidoscope toward the red-hot core of an exploding star. As you meet up with the shock wave in this supernova, you see blobs of hot, reddish, rising material pounding on the shock wave and cooler, yellowish material moving inward. As you continue your flight to the core, some of this material falls back toward the center with you to be reheated by core-generated neutrinos, which drive the explosion. The yellow blobs accompanying you eventually turn red and move outward again.

Such an experience is available at ORNL's CAVE, where four projectors throw full-color, computer-generated video streams onto three walls and the floor. Originally developed by the Electronic Visualization Laboratory at the University of Illinois at Chicago in the early 1990s, the CAVE (CAVE Automatic Virtual Environment) was designed to provide an environment in which multiple users could experience and interact with data in "natural" ways.

The CAVE uses a special blend of software and hardware to enable users not only to see but also potentially to hear and touch data. The CAVE floods the senses with information so users feel "immersed" or coexistent with data. Tracking hardware, such as headgear and a programmable joystick or "wand," provides information about the user's head position and orientation. This information is relayed from the hardware to a computer and used to calculate multiple realistic perspectives of a scene from the point of view of the user. Views are calculated for each wall and each eye to simulate a surrounding three-dimensional visual environment. The user wears special glasses that rapidly alternate between the left and right eye, thereby mimicking human depth perception. By moving and pressing on the wand's programmed buttons, the user can choose an interactive experience such as flying or can select an object to learn more about it.

Ross Toedte and Ray Flanery, both of ORNL's Computing and Computational Sciences Directorate, are programming the CAVE to provide interactive experiences for users in many applications areas, such as astrophysics and climate prediction. Toedte says his goal for ORNL's CAVE is to find the ideal blend of data features and performance "so that meaningful images can be seen interactively. I am exploring various visualization tools to get the ideal mix of resolution and speed for each CAVE application."

The CAVE is being used to help understand the intricacies of global climate change simulations. In the photograph below, ORNL researchers are looking at climate data calculated by colleagues at the National Center for Atmospheric Research. Color-coding helps deliver an understanding of parameters such as monthly surface temperature for a climate simulation between 1870 and 2170. Such simulations involve potentially hundreds of parameters. Because of the highly complex interrelationships among climate variables, subtle changes in just one can have magnified long-term effects. The CAVE provides an ideal environment for observing such sensitivities between multiple simulations and variables.

The CAVE at ORNL is being programmed to let viewers experience virtual stellar explosions. Scientists will be able to watch stellar explosions evolve from near the core or outside the shock wave. They will also be able to observe parameters such as temperature, density, and velocity in the supernovae simulations and compare findings based on both theoretical and experimental data. Eventually, it may be possible to understand how the elements that make life possible are synthesized and dispersed by such explosions.

Toedte and Flanery have used a number of visualization tools to provide insight into physical phenomena, help scientists verify calculated results, highlight the unexpected, and enable scientists to communicate their results more effectively. The immersive environment of the CAVE represents the high end of the visualization environment available at ORNL. This plethora of software and hardware tools offers large challenges and even larger potential for fostering meaningful scientific understanding through visualizing computational results.

Computation is now seen as an equal partner with theory and experimentation in the advancement of science; visualization helps bring this one-time vision to reality.

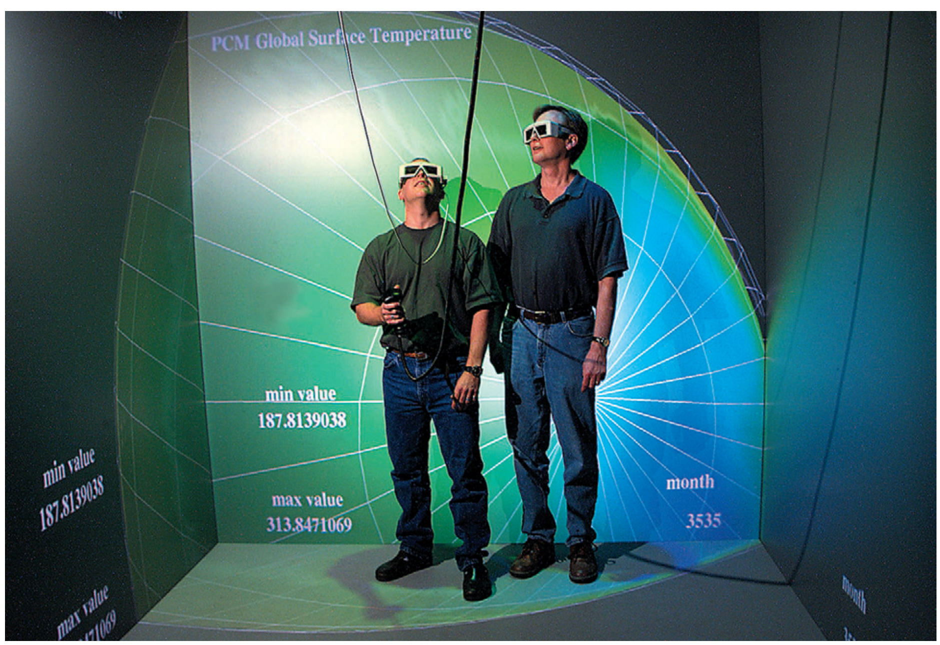

Curtis Boles

In ORNL's CAVE, Ross Toedte (left) and John Drake examine the earth's predicted surface temperature for a future month in a global climate simulation.

Collaborations and Partnerships

Through the ORNL-UT Joint Institute for Computational Sciences and DOE's Center for Computational Sciences (both of which will be housed in new buildings at ORNL), ORNL's new Computing and Computational Sciences Directorate has many partnerships with universities, industrial firms, and other research laboratories.

Artist's conception of the new building to house the Joint Institute for Computational Sciences, as well as the new Oak Ridge Center for Advanced Studies.

A new $10-million, 45,000-square-foot building on the ORNL campus will be constructed for the Joint Institute for Computational Sciences (JICS), as well as the new Oak Ridge Center for Advanced Studies (ORCAS). JICS was established in 1991 through the Science Alliance, a University of Tennessee (UT) Center of Excellence that seeks to promote collaborative relationships between research groups in the UT system and ORNL. JICS was created to encourage and facilitate the effective use of high-performance computing resources in the state of Tennessee. In pursuit of its main goal, JICS directs activities in the areas of research, education and training, and technology transfer.

In research, JICS personnel consult with researchers on individual application codes and recommend target architecture(s), tools for writing parallel codes, visualization toolkits, graphical user interfaces, and homogeneous or heterogeneous parallel implementations.

In education and training, JICS staff conduct classes, organize seminars and workshops, and participate in one-on-one and small group educational activities.

"An important feature of JICS is the strong collaboration between researchers at ORNL and faculty and students at UT," says ORNL's Buddy Bland. "These collaborations are under the new umbrella of the ORNL-UT Computational Sciences Initiative. For example, Jack Dongarra's group of 40 or so researchers and graduate students at UT is developing performance tools to enable researchers to run their scientific simulation codes at the most efficient supercomputing level of power, using resources at ORNL. One application in which there is a strong collaboration is climate prediction." UT is making funds available to support tool development research for which ORNL provides access to its supercomputers.

"The Computational Sciences Initiative makes collaborators more competitive for research grants," says Science Alliance Director Jesse Poore, "because researchers can draw on the combined strength of ORNL and UT. In addition, they gain ready access to terascale computing facilities." The initiative is open to all facets of science and engineering research and education. Research selected for the program is of strategic importance to both UT and ORNL. JICS will oversee the initiative. Seminars to acquaint interested UT and ORNL researchers with the opportunities available have already begun.

Future goals of JICS include partnerships with government, private industry, and other universities, emphasizing expansion of collaborative research, training in high-performance computing, expanded use of clusters of heterogeneous computing environments, and technology transfer and consulting to benefit the private sector. Through these types of collaborations, JICS is working to earn for Tennessee an international reputation in the use of massive parallel processing.

DOE's Center for Computational Sciences (CCS) at ORNL has partnerships with universities, governmental institutions, and major industrial firms, including IBM and Compaq corporations. These partnerships are essential to solving complex scientific problems using supercomputers and to advancing computer science and scientific research.

CCS has partnerships with the core universities of UT-Battelle, which manages ORNL for DOE. These universities are Duke University, Florida State University, North Carolina State University, Virginia Tech, the University of Virginia, and UT.

CCS is part of ORNL's first scientific collaboration with the IBM Research Division in Zurich, Switzerland. The goal of this collaboration is to develop a quantum-dot array to do innovative computations. (See *ORNL Review*, Vol. 34, No. 2, 2001, pp. 18–19 for more details).

Thanks to David McQueeney, vice president of IBM Communication Technology, IBM fellowships will allow five postdoctoral scientists to split their time between ORNL and IBM. One of these "shared postdocs" will work on the quantum-dot array project.

In August 2001, ORNL signed a cooperative research and development agreement with IBM to help develop its next-generation supercomputer—the 100-teraflops Blue Gene, which will be used to relate gene sequences and protein structures to human diseases. In this collaboration, ORNL and IBM researchers will explore the limits of scalability in high-performance computing. Because this machine will have hundreds of thousands of processors, ORNL researchers will help develop fault-tolerant algorithms to allow the supercomputer to work around processors that fail and will provide their expertise in using computing to predict protein structures.

Clearly, ORNL's supercomputing personnel have good connections.

Modeling Magnetic Materials

Materials research using supercomputers is paving the way for the next generation of information technology.

For many years, information technology has been marked by an exponential growth in the ability to store and retrieve data—so much so that the industry is reaching its limits on further improvement in data storage and retrieval with the materials at hand. Developing new materials requires research at the molecular and atomic level to understand their properties. Such studies require calculations involving even the complexities of quantum physics. Somewhat ironically, the ability to advance this type of research has, until recently, required computational power that has been constrained by the very limits the research is trying to overcome.

Yet the demand to handle more data faster keeps growing—for instance, stuffing more data into smaller devices, such as digital cameras, and retrieving that information faster. Meeting this demand requires creating new materials when the demands on current materials are more than they can effectively meet. These are tough materials sciences problems whose complex solutions will depend on calculations made using the latest and most powerful supercomputers.

ORNL's Malcolm Stocks, Thomas Schulthess, Xiaoguang Zhang, Don Nicholson, Bill Shelton, and the University of Tennessee's Balazs Ujfalussy, along with other collaborators, are performing complex materials science calculations using ORNL's growing power in supercomputing. "We have been developing computational methods to model magnetic states of various materials and to try to unravel the intricacies of magnetism," says Stocks.

Stocks, Shelton, Schulthess, and Ujfalussy have been exploring the role of antiferromagnetic materials in contact with ferromagnetic materials in magnetic multilayer storage and read-head devices, or, simply put, the tiny yet complex components that make a computer's hard drive tick. Data are stored, in binary fashion, through multitudes of yes/no or on/off commands captured in the magnetic orientation of small regions of magnetism, which are, in turn, made up of the magnetism of thousands of atoms. Individual atomic "magnetic moments" result from an imbalance between the number of spin "up" and spin "down" electrons associated with the atomic site. If the atomic magnetic moments all point in the same direction, that material is ferromagnetic. ("Ferro" connotes iron, the most common magnetic element.) It will stick to a refrigerator. It will make an electric motor run. Or it can be used to store data.

In *anti*ferromagnetic materials, the individual atomic magnetic moments are oriented in the materials' structure in such a way that the total magnetic moment is zero; for example, the magnetic moments associated with individual atoms in a crystal can point "up" and "down" in a regular manner, with as many "up" moments as "down" moments.

The read head in a computer disk drive is a multilayered component, known as a spin valve (because it functions similarly to the way a mechanical valve might turn water on or off). A spin-valve device consists of a ferromagnetic material, a nonmagnetic layer, and another ferromagnetic material that is held constant by an adjoining layer of antiferromagnetic material. Through an effect called "exchange bias" the magnetic orientation of the ferromagnetic layer closest to the antiferromagnetic layer is held constant, or "pinned," enabling the other, outermost ferromagnetic layer to "switch" from one orientation to the other, depending on the bit of information that is stored on the disk. When this process cycles several million times, the word-processing document, spreadsheet, or digital photo stored a few days ago appears on a computer screen.

To develop lighter, smaller devices for faster access to and storage of more data, new and better materials are needed. Stocks, Schulthess, and colleagues at ORNL and across the country have been exploring the incredibly complex interactions at the atomic levels that can produce insights into the development of new, advanced materials for the next generation of information technology.

Using computer simulations performed at ORNL, Stocks, Schulthess, Shelton, and Ujfalussy have recently explored magnetism in iron manganese (FeMn), an alloy used as the anitferromagnet in spin-valve devices. FeMn is an antiferromagnetic material used to fix magnetic moments in those pinned layers of ferromagnetic material, to allow detection of differences in electrical resistance (between the pinned and moving layers) that represent 1's and 0's sensed on magnetic media. In addition, they are studying the magnetic structure of interfaces between FeMn—an antiferromagnet—and cobalt (Co)—a ferromagnet—to shed light on the mechanisms responsible for exchange bias.

In a collaboration with the Department of Energy's National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory, the Pittsburgh Supercomputer Center (PSC), and DOE's Center for Computational Sciences (CCS) at ORNL, the ORNL team has run simulations of FeMn and FeMn-Co interfaces at unprecedented computational speeds. Calculations involving 2016 atoms ran at a computational speed of 2.26 teraflops (2.26 trillion additions, subtractions, multiplications, and divisions per second). The calculations involved 126 sixteen-processor nodes on the NERSC IBM SP3 supercomputer. A subsequent job involving 2176 atoms ran at 2.46 teraflops on 136 nodes; this run was done at about 75% efficiency, a figure much higher than has been seen in typical large-scale production codes.

What have these computational experiments shown? First, the FeMn alloy has a different magnetic structure than previously thought. Second, at interfaces between ferromagnetic Co and antiferromagnetic FeMn, the magnetic structure of the antiferromagnetic material changes radically from that associated with bulk FeMn.

Stocks explains that, previously, the magnetic structure of FeMn was thought to be one of three states of antiferromagnetic ordering, called 1Q, 2Q, and 3Q. This system was studied experimentally by various techniques, including neutron scattering, but experiments were unable to unambiguously distinguish among these three possible magnetic structures. The 3Q structure

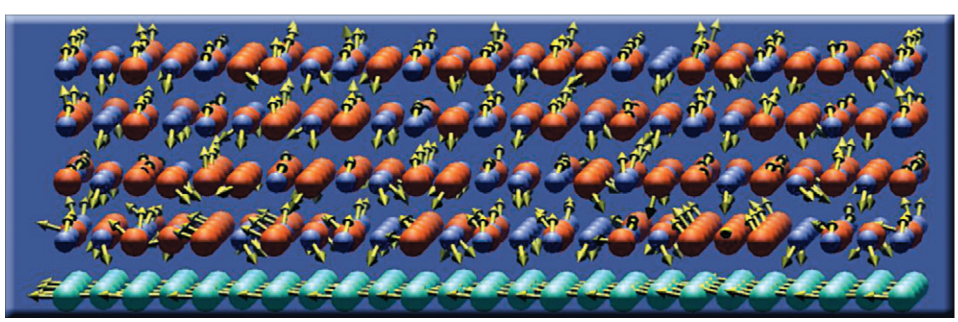

This section shows the final configuration of magnetic moments in five layers at the interface between iron-manganese (red and purple balls with yellow arrows showing magnetic moments) and cobalt (blue balls with yellow arrows.) The simulation of 2016 atoms had 15 iron-manganese and 6 cobalt layers.

was thought to have the lowest known energy and, therefore, to be the most stable, when compared with the 1Q and 2Q structures. However, Bill Shelton found, through hours of high-performance computer modeling using the 176-node IBM RS/6000 SP machine at ORNL, that a magnetic configuration exists that is even more stable than the 3Q structure.

"We did calculations using constrained density functional theory and spin dynamics," Stocks says. "Our goal was to understand the alloy's noncollinear, antiferromagnetic 3Q magnetic structure. We predicted there is a relaxed, even lower-energy, and therefore more stable, magnetic state that we call $3Q_R$, in which R stands for relaxed. According to our simulation, in this state, the magnetic moment orientations are pointed in slightly different directions from those of the 3Q structure. We thought that this insight could be an important ingredient in the explanation of exchange bias."

Schulthess explains that with read heads, for instance, issues of practicality stand in the way of improving their performance and further miniaturizing them. Improvements in the spin-valve function could be obtained by creating a larger magnetic field to fix the magnetic moments of the ferromagnetic layer. But exposing components to such a magnetic field would likely also wreck the disk and the data stored on it. The other solution is to increase the amount of antiferromagnetic material to enhance "the magic hand that holds"—or pins—the magnetic moments of the ferromagnetic material. However, doing so with currently available materials would require disproportionately more antiferromagnetic material, thwarting any further attempts at miniaturization.

The ORNL team believes the key to making an improved, miniaturized spin-valve device could lie in changes in the magnetic moments of the antiferromagnetic and ferromagnetic interfaces, which surfaced in the computer models. "Details of the magnetic structure appear to be very rich and could not have been predicted without the full machinery of first-principles spin dynamics and the use of massively parallel computing resources," Stocks says. "Although the origins of exchange bias are currently not fully understood, it is known that the magnetic structure of the ferromagnetic-antiferromagnetic interface is a key ingredient. The ORNL group's prediction of disorder in the magnetic moment orientation and the large changes that occur when the antiferromagnet is proximate to a ferromagnet have a number of interesting implications for the mechanism controlling exchange bias."

"It's all in how the magnetic moments arrange themselves," says Schulthess.

Stocks, Shelton, Schulthess, and Ujfalussy performed their calculations using a program called the Locally Self-Consistent Multiple Scattering (LSMS) code, which was developed at ORNL on the Laboratory's first supercomputer, the Intel Paragon. The 45,000-line LSMS code continues to be developed and extended to treat larger and more complex systems. Currently, the code is modeling magnetic materials on ORNL's latest supercomputers, including the Lab's new IBM Power4 machine at CCS. The researchers have run their models on the most powerful computers they can line up, including supercomputers at NERSC and PSC. Supercomputing high points came when they won the prestigious Gordon Bell Prize in 1998 and were the first to run a real application code at a speed of greater than 1 teraflop. Most recently, the code ran at a staggering 4.46 teraflops on a Compaq supercomputer at PSC.

Advances in supercomputing and advances in materials technology could develop into a symbiotic and self-perpetuating relationship. Strides in one area may very well lead to strides in the other. Supercomputers are helping materials researchers understand what's happening with individual atoms, as well as the incredibly complex relationships that exist between those atoms and at the interfaces of the materials they compose. As both the researcher and the casual Web surfer can already attest, the sky may be the limit for what research into new materials for computers can produce and the speed and amount of information the resulting computers can process.—*Bill Cabage, editor of the* ORNL Reporter, *an employee newsletter*.

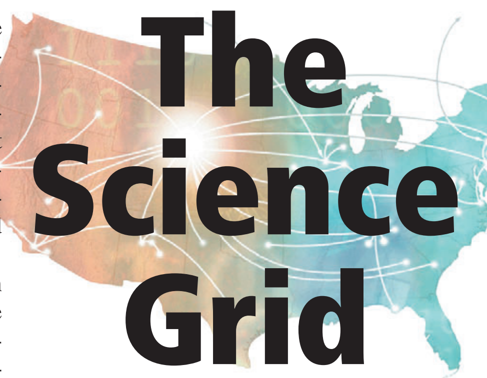

The Science Grid

Science grids are being established to connect scientists, instruments, computing, and data.

Scattered across the globe is a wide range of unique scientific resources: large experimental facilities (which will include the Spallation Neutron Source under construction at ORNL), supercomputer centers, petabyte data archives, the high-speed Internet, and, most importantly, the expertise of scientists at laboratories and universities around the world.

"The easy science problems have been solved," says Thomas Zacharia, ORNL's associate laboratory director for Computing and Computational Sciences. "Tomorrow's scientific breakthroughs in biology, nanotechnology, and physics now require large, multidisciplinary teams and the ability to effectively use the scientific resources distributed around the country." In fact, all scientific fields have the need to bring resources to the fingertips of area experts.

To connect scientists, instruments, computing, and data, the Department of Energy is establishing a DOE Science Grid, one of several grids being constructed by science agencies, including the National Science Foundation and the National Aeronautics and Space Administration. According to the DOE Science Grid Web site, "The vision for 'grids' is to revolutionize the use of computing in science by making the construction and use of large-scale systems of diverse resources as easy as using today's desktop environments." Collectively, all these grids are part of a global science grid.

This science grid creates a "virtual laboratory" environment that allows scientists to tackle tomorrow's problems more effectively. It enables innovative approaches to scientific computing through secure remote access to online facilities, high-speed Internet access, collaborations among scientists separated by distance (e.g., at different DOE labs in "collaboratories"), shared petabyte datasets, and large-scale distributed computation.

Scientific communities must have the capability for multiple users to remotely access high-performance computing resources and large data archives to both perform simulations and analyze the results of experiments. Researchers also need to collaborate with scientists at different sites who are involved in the simulations or experiments. They must coordinate access to and use of the resources.

The complex and evolving nature of scientific discovery requires general services that can be combined in many different ways to support different problem-solving approaches and the ability to evolve along with the scientific understanding of the problem. Resource management for such dynamic and distributed environments requires global naming and authorization services, scalability, fault tolerance, data management, security, authentication, and protection of proprietary data.

"The goal of the science grid," says Al Geist, who is overseeing ORNL's efforts in developing the DOE grid, "is to provide a common and supported set of services across all the scientific resources so that scientists can easily access, use, and share these resources and results more efficiently with the larger scientific community."

The design and deployment of large, multi-site grids are still evolving. You can learn about the state of the art and stay abreast of new developments by accessing this ORNL Web site (http://www.csm.ornl.gov/sciencegrid) and the links listed there.

Designing Electronic Devices Using Supercomputers

A proposed molecular memory cell that would allow laptop computer batteries to last 100 times longer than today's batteries is being modeled computationally on an IBM supercomputer at ORNL. This machine is also being used to simulate electron transport in carbon nanotubes in contact with other components, for future nanoscale electronic devices.

A large supercomputer at ORNL is being used to learn more about the best ways to design electronic device components on a very small scale.

DESIGNING NANOCIRCUITS

The successor to the silicon chip may be a nanoscale device—a self-assembled monolayer of organic molecules of benzene (a ring of six carbon atoms bonded with four hydrogen atoms) attached to sulfur atoms at each end. Because sulfur (thiol) has an affinity for gold, a single layer of these benzenedithiol molecules can be sandwiched between thin gold contacts. Scientists at Rice and Yale universities have induced self-assembly of such a device by dipping a gold surface into a beaker of benzenedithiol molecules. In large numbers these molecules attached themselves to the gold surface.

The scientists added nitrogen-containing (nitro) groups to the molecule's center benzene ring. The resulting perturbed electron cloud made the asymmetric molecule twist when an electric field was set up by applying a voltage between the gold contacts. When the molecule twisted, current flow through the "molecular wire" was blocked. When the voltage was removed, the molecule adopted its original shape, allowing current to flow again.

Such a device, if fabricated on a large scale, could be used as an ultrafast on-off switch, a key to creating ultrasmall, highly dense computer circuits required to make computers fast and powerful enough to mimic the human brain. Or the device could be used to make superior computer memory elements. A charge can be stored on the nitro group to prevent electrical conduction (binary 0), or the group can have no charge, allowing conduction (binary 1). Such a molecular memory cell retains a stored bit for nearly 10 minutes. By comparison, today's silicon-dynamic, random-access memories must be refreshed by an electrical current every 20 milliseconds. The new type of memory would save energy, allowing laptop computer batteries to last 100 times longer.

Density functional theory calculation of the molecular structure of three benzenedithiols sandwiched between two gold surfaces.

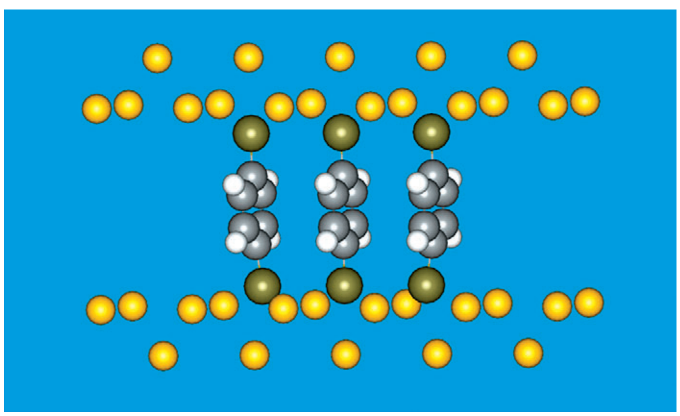

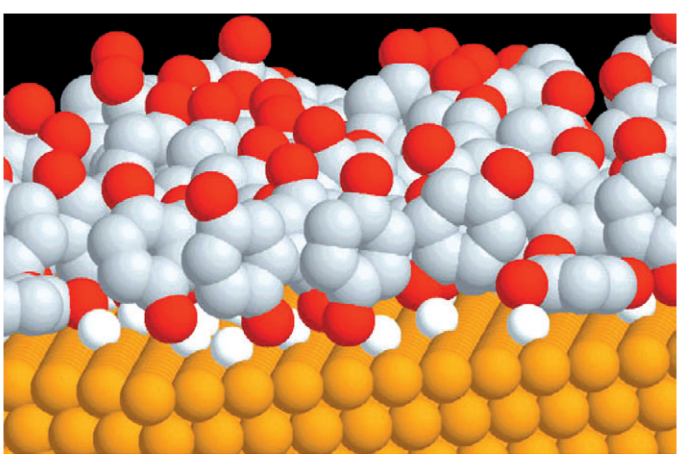

Snapshot from molecular dynamic simulation of 144 benzenedithiol molecules adsorbed on gold [Au(111)]. Color code: sulfur (red); carbon (gray); gold (yellow), anchor point on gold surface (white).

Such a concept is being modeled computationally on ORNL's IBM supercomputer, dubbed Eagle, in a Laboratory Directed Research and Development project. ORNL's David Dean, Bill Butler (now at the University of Alabama), Peter Cummings (an ORNL-UT Distinguished Scientist), Predrag Krstic, David Schultz, Mike Strayer, Jack Wells, and Xiaoguang Zhang are running the calculations using a modified version of NWChem, a computational chemistry code.

"Using ab initio methods, we modeled the self-assembly and electrical conductivity of five benzenethiol (BT) and benzenedithiol (BDT) molecules on a gold surface," says Dean. In the example shown in the illustration, two gold lattices are shown on the top and bottom. Three BDT molecules are seen in the middle area. This particular configuration has 70 atoms and 590 active electrons. A single calculation of this type requires 46.67 hours on 80 nodes of Eagle, or 14,930 processor hours. The single-particle wave functions resulting from this calculation will be used in a conductance calculation to determine the current-voltage characteristics of this molecular device.

The researchers have also performed preliminary molecular simulations of self-assembled monolayers composed of BT molecules on the [111] surface of gold. They included state-of-the-art force fields generated through electronic structure calculations. Both molecular dynamics (MD) and Monte Carlo (MC) simulations are being used. Gibbs ensemble MC simulation is being used to establish the equilibrium between adsorbed monolayers of BT and BDT with a low-density solution. MD is then being used to equilibrate the structures thus found. “The structure of the adsorbed monolayers appears to be consistent with available experimental results,” Dean says.

To produce a useful device, self-assembly must be combined with fabrication methods such as photolithography. The ORNL scientists will model how best to assemble these molecules and align them with the gold contacts to optimize electrical conductivity.

NANOSCALE DEVICE DESIGN AND PERFORMANCE

Marco Buongiorno Nardelli, who holds a joint position with ORNL and North Carolina State University (NCSU, one of UT-Battelle’s core universities), has been exploring the feasibility of using carbon nanotubes in nanoscale electronic devices. He is currently using Eagle to run his own suite of codes simulating electron transport in carbon nanotubes in contact with other materials.

In one project in which he provided computer modeling, experiments at the University of North Carolina (UNC) at Chapel Hill have shown that it is possible to build a nano-rheostat, similar to a dimmer light switch. In such a device, a carbon nanotube—a cylinder resembling rolled-up chicken wire because its carbon atoms are arranged in a hexagonal configuration—is placed on a sheet of graphite whose carbon atoms also have a hexagonal arrangement.

“If you place the carbon cylinder on the graphite sheet so that the carbon atoms of both are aligned, a current will flow at the interface,” Buongiorno Nardelli says. “As you rotate the carbon cylinder on the graphite sheet, changing the angle between the atoms in the system, you get increased electrical resistance and reduced current flow. As the atoms become aligned, you get low resistance and high current flow.”

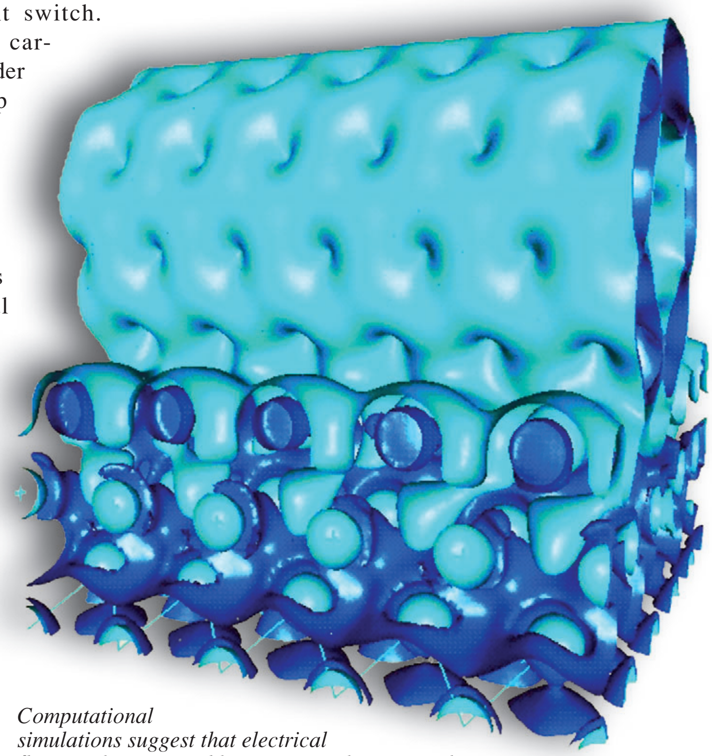

Illustration enhanced by LeJean Hardin

Computational simulations suggest that electrical flow can be improved between a carbon nanotube and an aluminum contact by mechanically deforming the tube.

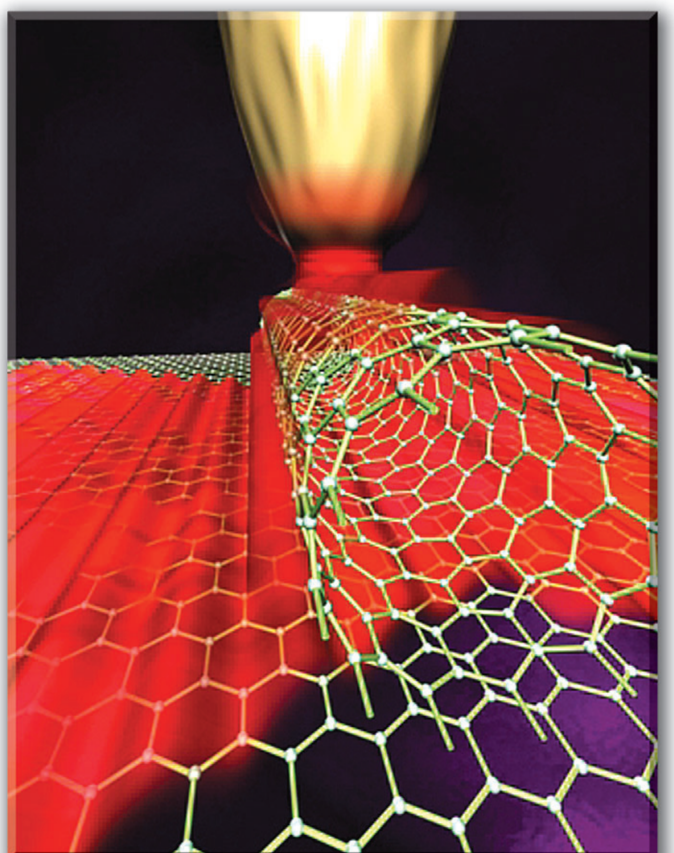

Illustration enhanced by LeJean Hardin

A carbon nanotube whose atoms are aligned with the carbon atoms of a graphite sheet exhibits good current flow. By rotating the tube on the sheet, tunable resistance results, making possible the creation of a nano-rheostat that acts like a dimmer light switch.

Computational simulations by Buongiorno Nardelli verified that the interface between a carbon nanotube and graphite gives tunable resistance (as in a dimmer switch). His theoretical predictions on the feasibility of a nano-rheostat agreed with the UNC experimental results. The work was published in *Science* magazine in 2000.

If carbon nanotubes are to be used as nanowires or other components in nanoscale devices, electrons must flow between these nanotubes and metal contacts in the device. In some experimental devices, high resistance at the tube-contact interface can make the mechanism of electron transfer quite inefficient. Buongiorno Nardelli and his NCSU colleagues have used computer modeling to address the question of why some nanodevices have better performance than others.

“In some devices,” Buongiorno Nardelli says, “electrons in the carbon nanotube stay in the tube and electrons in the aluminum stay in the metal. Our simulations suggest that contacts can be improved by mechanical deformations. For example, if a carbon nanotube sandwiched between two aluminum contacts is squeezed and deformed, new bonds form between the carbon and aluminum atoms, increasing electron flow at the tube-contact interface.”

The strength of carbon nanotubes is also of interest to Buongiorno Nardelli. Of all materials, carbon nanotubes have the highest tensile strength. They are 100 times stronger than steel but have one-sixth its weight. Scientists propose using carbon nanotubes as fibers in a polymer composite to form stronger structural materials for aircraft, spacecraft, and suspension bridges.

Computational simulations by Buongiorno Nardelli and his colleagues have shown that the geometry—the arrangement of the carbon hexagons along the nanotube—influences tube strength. “Our simulations,” he says, “predicted that whether a nanotube is brittle or ductile depends on the temperature at which it was deformed, the orientation of the hexagons with respect to the tube’s axis, and the amount it is stretched—that is, strain.”

Carbon nanotubes are very small, but simulations of their behavior in nanoscale electronic devices require a large amount of computer capacity. 101

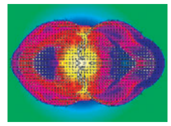

Simulating Supernovae on Supercomputers

Researchers using an ORNL supercomputer have found that the organized flow beneath the shock wave in a previous two-dimensional model of a stellar explosion persists in three dimensions, as shown here. A quasi-"bipolar" outflow can easily be seen in this snapshot taken from a three-dimensional simulation. Although this simulation is not a full-blown supernova simulation because all the needed physics is not yet included, it does illuminate the important interaction between the three-dimensional fluid flow that develops beneath the shock and the shock wave itself. In this model, this interaction leads to an explosion.

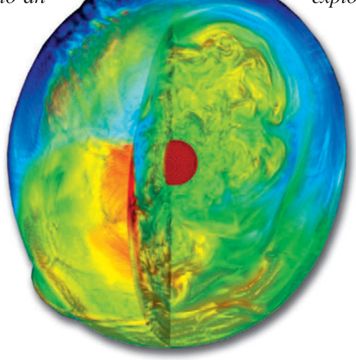

Multidimensional simulations of core-collapse supernovae will answer important questions about the creation and dissemination of elements that make life possible. They may also be important in the development of "enabling technologies" for other applications, such as combustion, climate, fusion, stockpile stewardship, and nuclear medicine.

Billions of years ago, the universe became the stage for spectacular stellar explosions. These momentous celestial events marked the end of a star's life but planted the seeds that would ultimately result in the explosion of life on the earth. After millions of years of evolution, massive stars were disrupted in a creative process. These core-collapse supernovae spewed forth lightweight, life-giving elements such as carbon and oxygen that eventually reached the regions out of which our solar system formed. They also synthesized heavy elements and disseminated them to interstellar space. In addition to the elements they created, supernovae also left their mark in the form of neutron stars and black holes.

"Life as we know it would not exist if not for these incredible explosions of stars," says Tony Mezzacappa, task leader for astrophysics theory in ORNL's Physics Division. "When stars die in these explosions that generate energy at the rate of billions upon billions of watts, elements necessary for life are strewn throughout our galaxy and become part of the 'soup' from which our solar system formed."

A major challenge for computational astrophysicists is to solve the "supernova problem," using massively parallel supercomputers to model stars greater than 10 times the mass of the sun. One goal is to predict whether these stars will explode like Supernova 1987A (a much observed supernova in a nearby galaxy). The other is to predict all of the observed phenomena associated with such stellar explosions. Three-dimensional (3D) simulations that can be run only on terascale supercomputers (such as the IBM supercomputers at ORNL) are being developed to do just that.

A core-collapse supernova explosion occurs in only a few hours in a star that has evolved over millions of years. This event is thought to be 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. Astrophysicists believe the shock wave stalls while trying to propagate through the stellar core and is re-energized by neutrino heating. Neutrinos are particles with no charge and small mass that interact very weakly with matter. At the core of a supernova is a neutrino "bulb" that radiates heat and energy at the staggering rate of 10^{45} watts.

ORNL leads the field in simulations of neutrino transport and must now apply this ex-

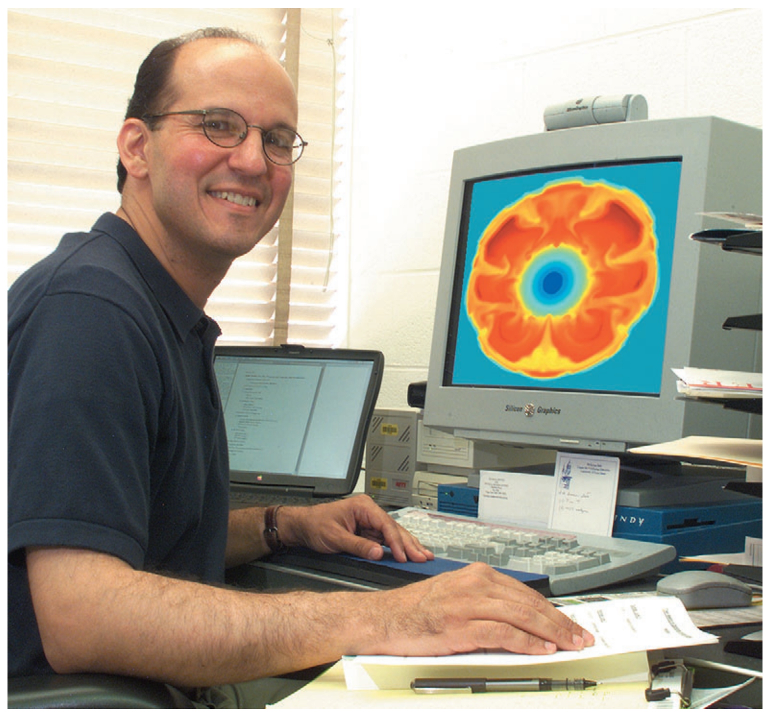

Curtis Boles

Tony Mezzacappa sits before a computer visualization of an exploding star. The visualization illustrates convection in a core-collapse supernova explosion. The red plumes, composed of material that has been heated by neutrinos, are rising, and the yellow fingers, composed of cooler material, are moving inward.

These five computer visualizations from a 2D model show convection in a core-collapse supernova explosion and its role in aiding the explosion.

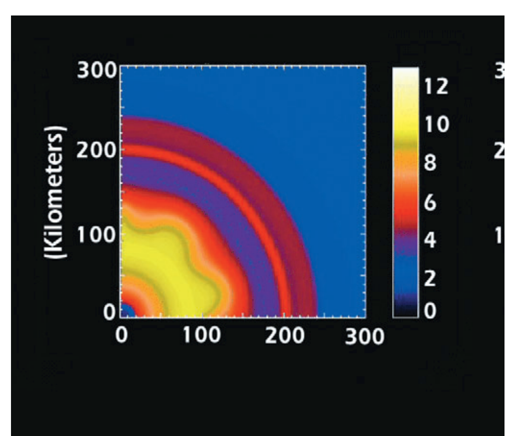

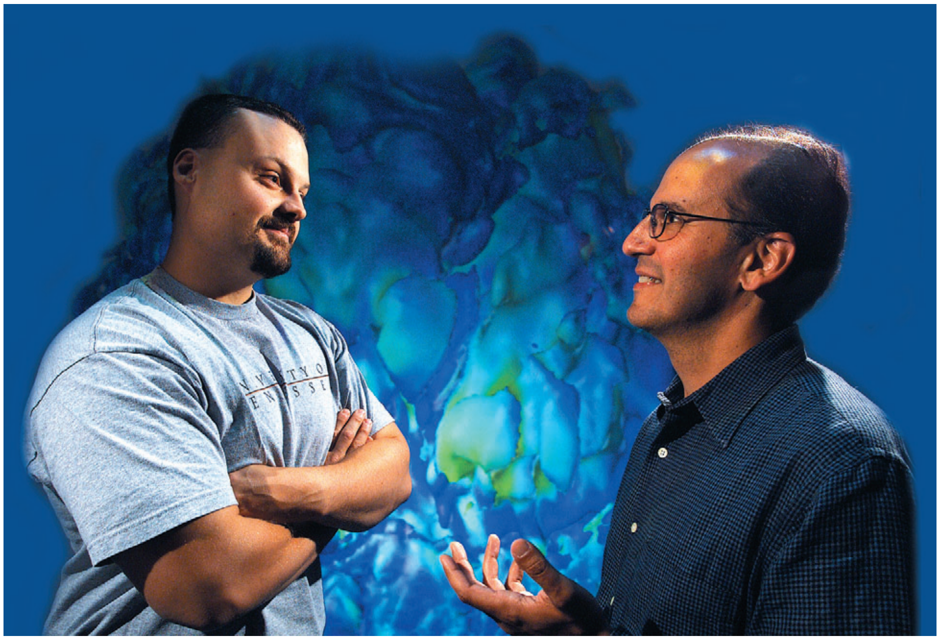

Curtis Boles, enhanced by Tina Curry

Tony Mezzacappa (right) discusses with Bronson Messer the results of a three-dimensional supernova simulation shown in the ORNL CAVE. The visualization shows turbulent fluid flow below the supernova shock wave.

pertise to developing 3D simulations that will explore the role played by convection—transfer of heat by the circulation of the core's proton-neutron fluid—in aiding this shock revival process, as well as the role played by the star's rotation and magnetic field. ORNL researchers are also interested in using these simulations to predict the emitted gamma rays and gravitational waves (ripples in space-time) from supernovae.

Supported by $9 million in funding over five years from the Department of Energy's Office of Science initiative called Scientific Discovery through Advanced Computing (SciDAC), ORNL, the National Center for Supercomputer Applications (NCSA), and eight universities plan to obtain a detailed understanding of how a star explodes. Their approach is to perform 3D simulations of the radiation of the enormous amounts of neutrino energy and the resulting turbulent fluid flow (hydrodynamics) that together may propel material into outer space. Computational predictions will be made consistent with data obtained from astronomical observations.

"Thanks to the growing wealth of observational data from ground- and space-based facilities and the growing computing power afforded by massively parallel supercomputers at ORNL and elsewhere, we are presented with a unique opportunity to finally solve one of nature's most important mysteries," says Mezzacappa.

The simulations will uncover how supernovae synthesize elements and disseminate them into interstellar space for processing by other astrophysical systems. Additionally, they hope to determine how the "neutrino-driven wind," which arises from the proto-neutron star left behind after the explosion, synthesizes elements heavier than iron in a process of rapid neutron capture.

The simulations will also be important in ultimately understanding how the cooled-down remnants of supernovae give rise to neutron stars and black holes, creating the basic building blocks of rotating neutron star (pulsar) and X-ray binary systems.

"We will try to predict whether a black hole or a neutron star will be left behind in the next supernova explosion in our galaxy," says Mezzacappa. "These events occur in our galaxy two or three times each century."

One of the collaborating institutions is the University of Tennessee (UT), where Jack Dongarra and research faculty member Victor Eijkhout will be working on mathematical solutions (algorithms) to help solve the equations that govern the motion of neutrinos through the stellar material.

In addition to UT, ORNL's collaborators for the project are the State University of New York at Stony Brook, the University of Illinois at Urbana-Champaign, the University of California at San Diego, the University of Washington, Florida Atlantic University, North Carolina State University, and Clemson University. Mezzacappa's co-investigators at ORNL are David Dean and Michael Strayer, both of the Physics Division, and Ross Toedte of the Computer Science and Mathematics Division (CSMD).

The ORNL-centered SciDAC team led by Mezzacappa is also working with five other SciDAC teams ("ISICs") on issues that include scalable solution of large sparse linear systems of equations, code architecture and design, management and analysis of terascale datasets, code performance, and adaptive meshes, as well as with a supporting "base project" to address networking issues for this distributed collaboration. These collaborations involve a number of ORNL staff in CSMD.

This "mother of all applications," as Mezzacappa calls the core-collapse supernova problem, should provide new insights into radiation transport and fluid flow relevant to many phenomena. "Our work addresses very broad themes important to DOE's national mission," Mezzacappa says. "Our ability to model the movement of radiation through matter may help advance DOE's energy and basic research missions."

Examples of interest to DOE are combustion processes in internal combustion engines, effects of increased atmospheric greenhouse gas concentrations on future climate, simulated nuclear weapon explosions (stockpile stewardship), production of fusion energy reactions in hot plasmas, and the effects of radiation therapy on tumors and normal tissue. Supernova simulations on supercomputers will likely be a shining star in the astrophysics and other scientific communities. 101

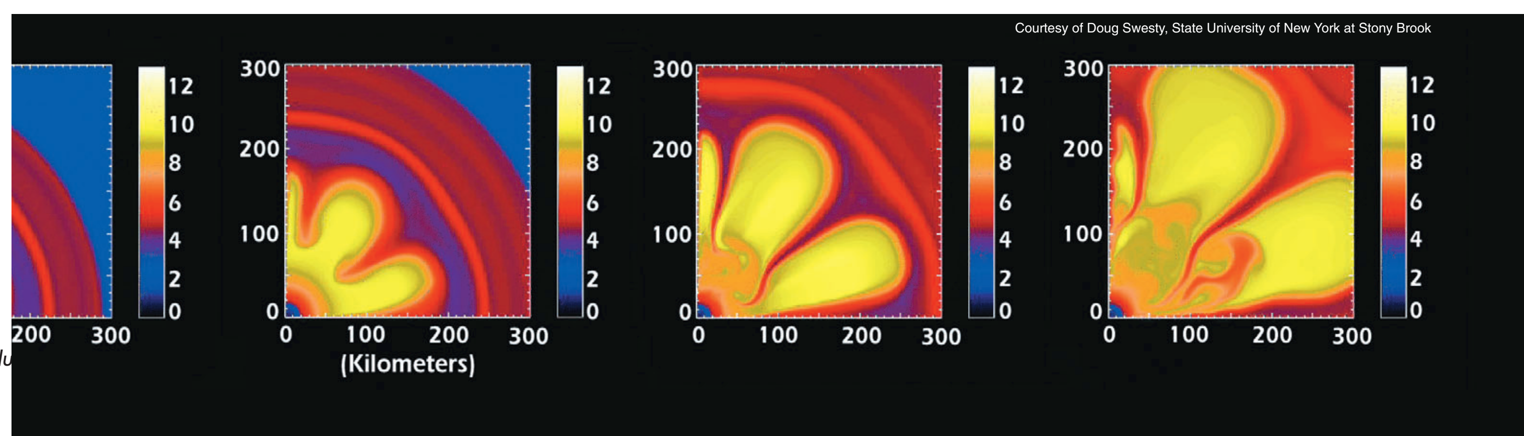

Probing Cells by Computer

A computational analysis of human and bacterial genomes by ORNL researchers provides insights into what our genes do. ORNL researchers will soon be predicting 100 protein structures a day and evaluating which compounds could make highly effective therapeutic drugs.

Supercomputers are being used at ORNL to increase our knowledge about the structure and function of genes and proteins in living cells.

ANALYZING GENOMES COMPUTATIONALLY

In 2001 scientists using supercomputers suggested we should say goodbye to some common beliefs in biology. No longer was it considered true that the human genome has 100,000 genes, that each gene makes only one protein, and that humans and bacteria have entirely different genes in their cells.

These tenets were tossed out in response to findings of the International Human Genome Sequencing Consortium, including the Department of Energy's Joint Genome Institute (JGI), to which ORNL contributes computational analysis. On February 15, 2001, the consortium published the paper "Initial Sequencing and Analysis of the Human Genome" in the journal *Nature*. The paper states that the human genome has "about 30,000 to 40,000 protein-coding genes, only about twice as many as in a worm or fly"; each gene codes for an average of three proteins; and hundreds of genes may have been transferred from bacteria to human genes.

Ed Uberbacher, head of the Computational Biology Section in ORNL's Life Sciences Division (LSD), was one of the hundreds of authors who contributed to this landmark paper. Using

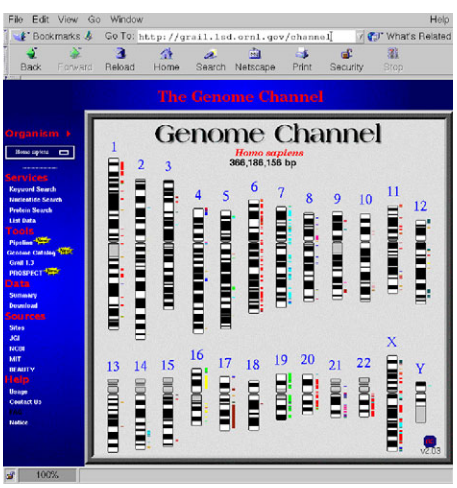

A view of the Genome Channel home page on the World Wide Web.

the IBM RS/6000 SP supercomputer (Eagle) at ORNL, he and his ORNL colleagues performed computational analysis and annotation of the human genome to uncover evidence of the existence of genes about which little or nothing was known—until this study. To perform their analysis, Uberbacher et al. used the latest version of the Gene Recognition and Analysis Internet Link (GRAIL), which was developed by Uberbacher and others in 1990 at ORNL and was rewritten as GrailEXP for parallel supercomputers. Use of GrailEXP helped provide evidence for alternative splicing—different ways of combining a gene's protein-coding regions (exons) to produce variants of the complete protein. The evidence suggests that some genes when expressed produce up to 10 different protein products.

Researchers in LSD's Computational Biology Section have identified many genes in bacterial, mouse, and human genomes. For the JGI they have created and used assembly programs and analysis tools to produce draft sequences of the 300 million DNA base pairs in chromosomes 19, 16, and 5. They have analyzed 25 complete microbial genomes and many JGI draft microbial genomes.

The section's researchers have written algorithms and developed other tools that make it easier for biologists to use computers to find genes and make sense out of the rising flood of biological data. Through ORNL's popular, user-friendly Genome Channel Web site (150,000 sessions per month) and its Genomic Integrated Supercomputing Toolkit (developed by ORNL's Phil LoCascio and commonly called GIST), the international biology community, including pharmaceutical industry researchers, have readily obtained meaningful interpretations of their DNA sequences. With help from its supercomputers, ORNL is on the genome analysis map.

COMPUTATIONALLY PREDICTING PROTEIN STRUCTURES

In the summer of 2000, an LSD group led by Ying Xu participated in an international competition to predict the three-dimensional (3D) structures of 43 proteins, using computational tools. Of the 123 groups competing in the fourth Critical Assessment of Techniques for Protein Structure Prediction competition, this group placed sixth, putting ORNL in the top 4% and placing it ahead of all other DOE national laboratories in the contest.

The actual structures of the 43 target proteins had been determined experimentally by using nuclear magnetic resonance spectroscopy and X-ray crystallography. The computational groups were provided with the identity and order of amino acids making up each protein and the length of the one-dimensional amino-acid sequence. Their predicted structures (obtained in a few weeks) were compared with the experimentally determined structures (obtained in about a year).

Protein structure is the key to protein behavior. Because the function of a protein is related to its shape, it is essential to find or predict correctly the 3D structures of proteins that make us ill or keep us well. Using the details of a protein's shape, a chemical compound can be custom designed to fit precisely in the protein, like a hand in a glove, blocking or enhancing the protein's activity. In this way, a highly effective drug with no side effects could be created for an individual.

To speed up drug development, the goal is to predict computationally the structures of 100,000 proteins by aligning different amino-acid sequences along 1000 unique structural folds that are being determined experimentally.

ORNL researchers will soon be predicting 100 protein structures a day and evaluating which

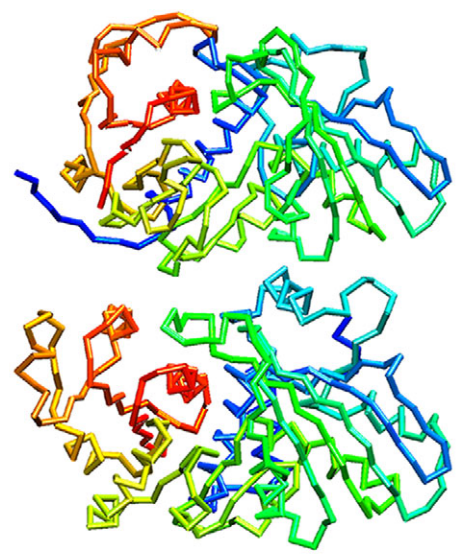

A protein structure, predicted at ORNL (top) and the actual structure, determined experimentally (bottom).

potential drug molecules dock well with specific proteins by running various automated tools on the Eagle supercomputer. One of those tools is PROSPECT, the Laboratory's copyrighted protein-threading computer program that brought the group a high world ranking and an R&D 100 Award in 2001. It is giving ORNL good prospects in a field that could shape future health care.

Modeling Blood Flow during CPR

Thanks to computer modeling, a scientific discovery was made that might lead to a way to save victims of cardiac arrest.

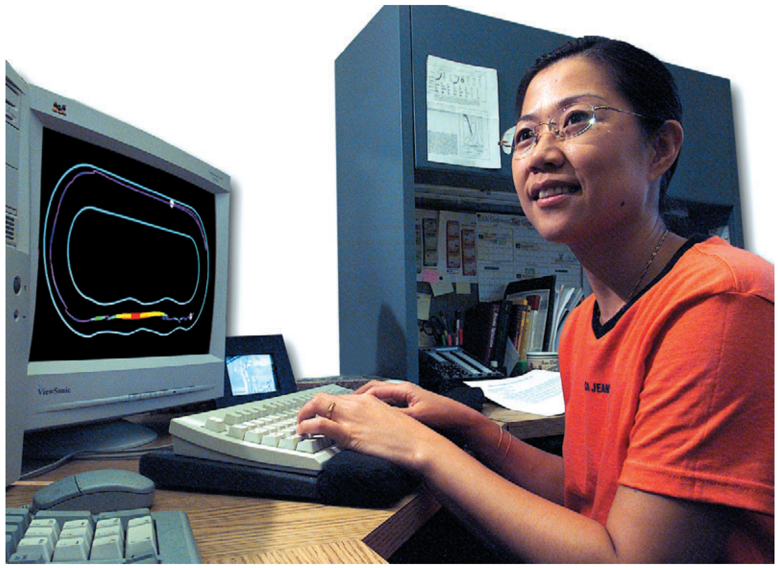

Jim Richmond, enhanced by Jane Parrott

Eunok Jung conducts a computer experiment in which the flow of a liquid changes direction when the frequency of pulses is changed. The black and white fluid markers represent the position at the initial and final time respectively.

Frank, 42, fell to the floor at home, a victim of cardiac arrest. His brother Jim immediately put his ear to Frank's chest. Frank was not breathing; his heart had stopped beating. Jim called 911. Because he had been trained in cardiopulmonary resuscitation (CPR), he began chest compressions and mouth-to-mouth airflow in the hope of restoring his brother's heartbeat and breathing. Tragically, Jim's heroic efforts failed; Frank died.

CPR has been successful in restarting the hearts of people who have been electrically shocked, badly injured, or frozen, but in few instances has CPR revived victims of cardiac arrest. More than 250,000 people die from cardiac arrest each year in the United States. Yet CPR, despite its high failure rate, is used by physicians and rescue workers to preserve blood flow during cardiac arrest. If the mechanisms of blood flow in the body during CPR were better understood, it might be possible to improve CPR techniques and save more lives of victims of cardiac arrest.

At least that's the hope of Eunok Jung, a staff member in ORNL's Computational Mathematics Group in the Computer Science and Mathematics Division. She recently made a scientific discovery using computational simulation that is relevant to CPR. Jung conducted a computational experiment using a two-dimensional model of a rigid, doughnut-shaped tube in which one section is replaced with a flexible membrane. Earlier laboratory experiments with this fluid-filled device, which has no valves, showed that periodic squeezing of the membrane caused a flow in one direction.

Jung discovered that changing the frequency of squeezing affects not only the amount of flow but also its direction. She verified this computational finding with a physical apparatus. "I found that you can reverse the flow of fluid simply by varying the frequency of squeezing," she says. If during CPR the heart valves remain open, then Jung's results suggest that the rate of chest compression may partly determine whether CPR saves a life. Jung's experiment took advantage of her Ph.D. dissertation adviser's "immersed boundary method" for modeling the fluid dynamics of the heart. Her adviser is Dr. Charles Peskin, one of the world's leading experts on heart modeling, who works at New York University's Courant Institute.

Whether Jung's finding is important to CPR may depend on which theory about the heart is correct. The cardiac compression theory says that during CPR the heart works as an active pump. The thoracic compression theory argues that during CPR the heart is a passive conduit that allows blood to flow, as a result of periodic squeezing and pressure differences between the external and internal thoracic compartments, through the cardiac valves that remain open (valveless pumping). Some imaging data suggest that the valves of the heart remain open during CPR in some instances. "These results," Jung says, "imply that the heart is acting at least partly as a passive conduit."

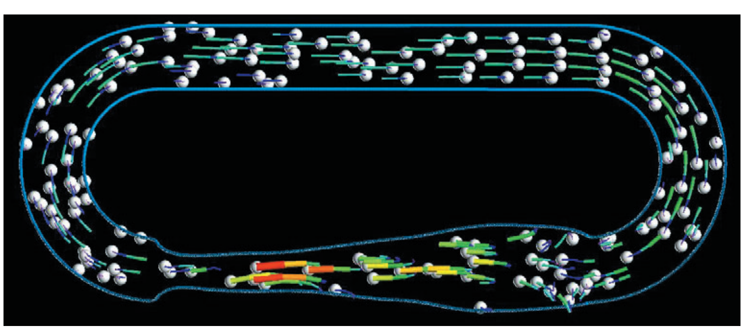

Despite the absence of valves, a clockwise net flow of the liquid around the loop is observed during a computer experiment.

To better understand blood flow in the heart during CPR and valveless pumping in general, Jung proposes to computationally simulate the heart as a pump, as a passive conduit, and as a combination of both. She will write a number of differential equations to create three-dimensional heart models, coupled with a lumped parameter circuit model (ordinary differential equations) of the circulation that can be solved using ORNL's supercomputers. Her results could get at the heart of how to modify CPR techniques to save victims of cardiac arrest.

Other Uses Envisioned for Valveless Pumping

Valveless pumping may also occur in the human embryo until the third week of gestation. At this stage of development, the valves of the heart have not yet formed; however, there is a net flow in the circulatory system that is somehow generated by the beating of the heart.

Modeling of valveless pumping may also aid in the design of valveless pumps for microelectromechanical system (MEMS) devices. In these devices, fluid motion must be produced without moving anything inside the fluid. A MEMS device, for example, might be used to control when and how much insulin is released into the body of a diabetic.

Some of the world's largest global climate models are being run on ORNL's supercomputers, providing insights for national and international assessments of the effects of global warming caused by human activities.

World-Class Climate Modeling

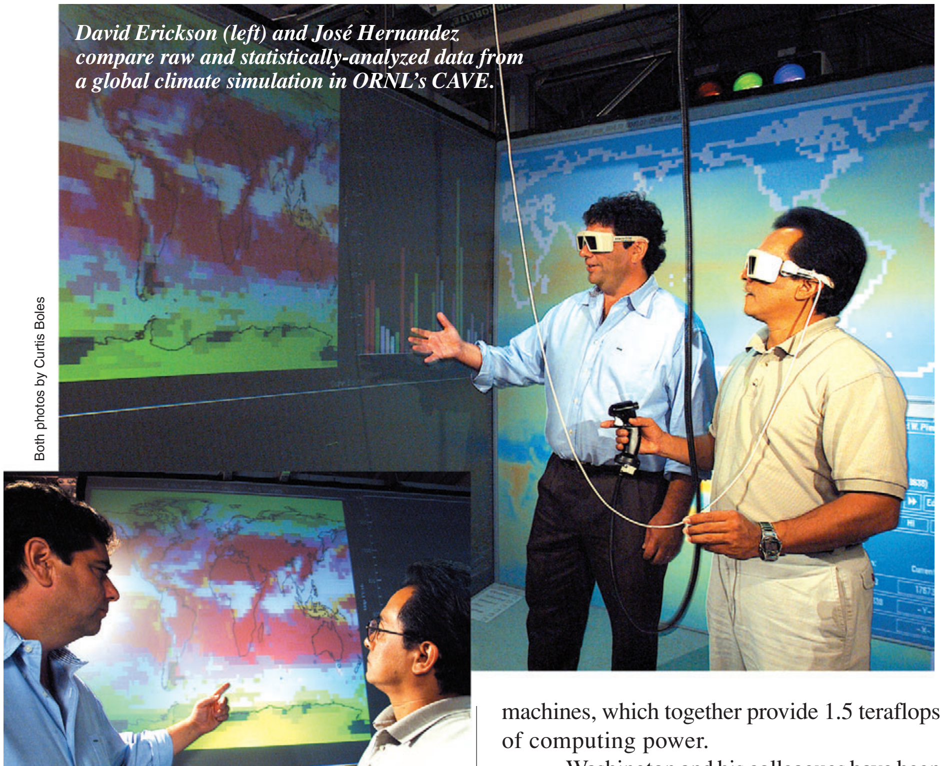

David Erickson (left) and José Hernandez compare raw and statistically-analyzed data from a global climate simulation in ORNL's CAVE.

Both photos by Curtis Boles

When Warren M. Washington, a member of the National Science Board, which advises the Executive Branch and Congress on science-related matters, conducts his research, he ponders what will happen 10, 50, or 100 years from now. As head of the Climate Change Research Section at the National Center for Atmospheric Research (NCAR) in Boulder, Colorado, Washington, along with his colleagues, tries to predict how the climate will change in the next century under various conditions. And for Washington, the most interesting hardware right now is the new supercomputer at ORNL named after the cheetah, the fastest land mammal.

Cheetah is the new IBM Power4 supercomputer, a machine rated at 4 teraflops (4 trillion calculations per second) that has 24 "regatta" nodes, each having 32 processors. This supercomputer has 1 terabyte of memory and 24 terabytes of disk space. It is located at DOE's Center for Computational Sciences (CCS) at ORNL, which is also home for the IBM RS/6000 SP and Compaq AlphaServer SC machines, which together provide 1.5 teraflops of computing power.

Washington and his colleagues have been using the ORNL supercomputers for century-long climate model runs to simulate changes in the world's climate from 1870 to 2170 under several different greenhouse-gas scenarios. In the business-as-usual case, atmospheric carbon dioxide (CO_2) levels steadily increase, trapping heat and causing global warming. Under the "stabilization" scenario, atmospheric CO_2 concentrations rise and then level off in response to various nations' carbon management strategies. The stabilization scenarios assume that CO_2 emissions from fossil fuel power plants are reduced and that enhanced absorption and sequestration of CO_2 in land plant life and ocean waters result in a significant slowing of the accumulation of CO_2 in the atmosphere.

The DOE-sponsored Parallel Climate Model being run on ORNL supercomputers results from a joint effort involving NCAR, ORNL, DOE's Los Alamos National Laboratory (LANL), the Naval Postgraduate School, and the U.S. Army Corps of Engineers' Cold Regions Research and Engineering Laboratory. John Drake in ORNL's Computer Science and Mathematics Division (CSMD) led the effort that enabled the climate model to be modified so that it could be run on massively parallel supercomputers at ORNL. The Parallel Climate Model brings together the NCAR Community Climate Model version 3, the LANL Parallel

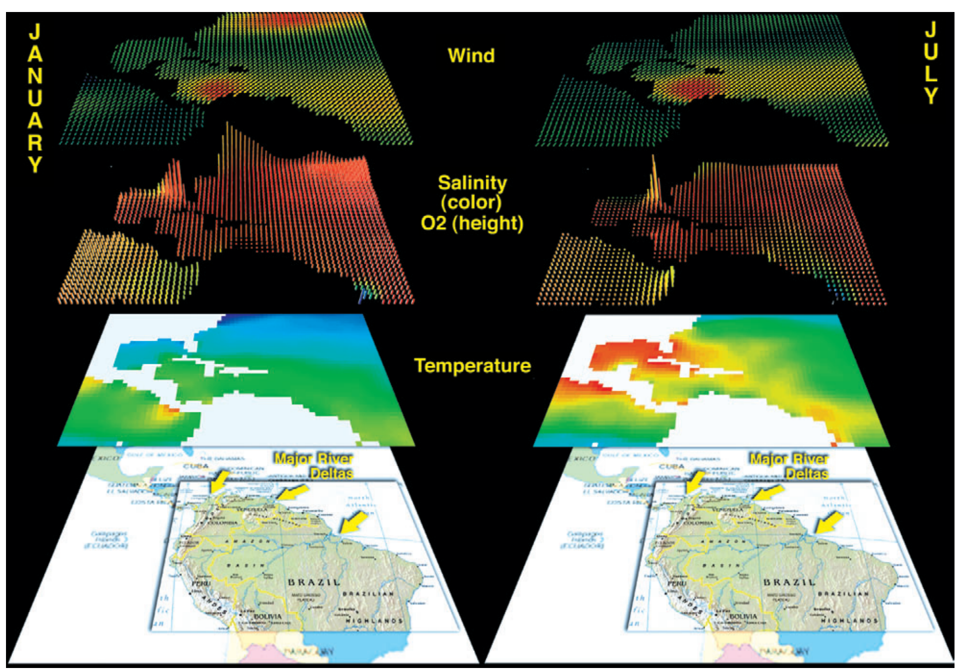

These images, produced by computer simulations, show how wind, salinity, and temperature vary between January and July in the northern portion of South America.

Ocean Program, and a sea-ice model from the Naval Postgraduate School in a massively parallel computer environment.

Drake and David Erickson (formerly with NCAR, now with CSMD) are using the ORNL supercomputers for climate predictions in which the interactions of the atmosphere with land and with the oceans are simulated. They plan to couple carbon and climate models together with help from researchers Mac Post and Tony King in ORNL's Environmental Sciences Division. The project will be multi-institutional, involving colleagues from other DOE labs, as well as a number of universities. Because CO_2 is emitted by energy production and because atmospheric increases in the greenhouse gas contribute to global warming, combining carbon and climate modeling fits well into the DOE mission of finding ways to produce energy in an environmentally acceptable fashion. The ORNL researchers plan to perform simulations using a coupled climate-carbon model to determine how the carbon budget would change in a greenhouse-warmed world.

In addition to the prescribed increases in CO_2, the model will be allowed to "run free" with the climate and carbon cycle evolving in unison to reach different future greenhouse-gas, or climate, states. The combined models will address these questions: How will increased atmospheric carbon dioxide affect desert size? How will it affect the frequency of droughts, precipitation, and severe storms—such as tornadoes and hurricanes—in different regions?

"Eventually, we hope to do climate modeling that is useful to public health service planners," Erickson says. "Our modeling could indicate when and where temperature and moisture conditions are likely to be conducive to various insect-borne diseases."

When climate modeling experts look into the details of the climate system, what they see is often cloudy. Emissions of particles from fossil fuel plants can have confounding effects on the warming of the earth's surface, making predictions less certain. While CO_2 and other greenhouse gases in the atmosphere absorb infrared radiation from the earth's surface and prevent the escape of heat, sulfate aerosols from coal-fired power plant emissions can have a cooling effect, moderating the temperature signal and changing weather patterns. Other challenges faced by the climate modelers are the uncertain effects on the radiative signal of clouds, the motion of sea ice, airborne dust from African deserts, and variability in Pacific Ocean surface temperatures. Research areas include the response of the oceans to rising temperatures, the effect of aerosols on cloud formation, and the feedback between the climate and carbon and water cycles.

The availability of the ORNL supercomputers allows U.S. climate researchers—for the first time—to make an ensemble of predictions for each future climate scenario. This capability enables a more detailed assessment of the variability and error estimates in the simulations, thus reducing uncertainties of model predictions. These climate predictions by Oak Ridge researchers and by NCAR's Washington and others using ORNL supercomputers will provide timely information for national climate change assessments and reports compiled by the Intergovernmental Panel on Climate Change (IPCC).

Based partly on results from global climate models, the IPCC concluded in 2001 that "there is new and stronger evidence that most of the warming observed over the last 50 years is attributable to human activities" and projected that, by the end of 2100, the global average temperature of the earth could rise by 1.5 to 5.8°C. Policymakers are now paying attention to what the IPCC sees as the future of the earth's climate.

SciDAC Funding for Local Climate Modelers

ORNL co-leads a team that will receive $20 million in funding over the next five years from the Department of Energy to speed the development of computer models to predict climate change and to improve the representation of physical, chemical, and biological processes. The funding comes from DOE's Scientific Discovery through Advanced Computing (SciDAC) Program. The project, awarded jointly to ORNL and Los Alamos National Laboratory, will allow the labs to develop a climate model that will provide a scientific basis for evaluating policy alternatives.

"It's difficult to assess the effects of policies if we do not have accurate models," says John Drake of ORNL's Computer Science and Mathematics Division. "This funding will allow us to develop tools to provide accurate information to policymakers looking at the potential effects of increased greenhouse gases."

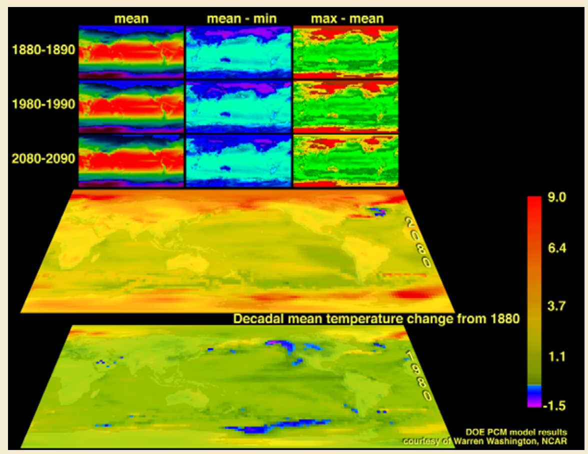

ORNL and other groups involved in climate research are using supercomputers to address this key science question: How will the earth's climate respond to physical, chemical, and biological changes produced by global alterations of the atmosphere, the oceans, and land? The computer image above shows climate variations over several different decades. These variations indicate long-term climatological factors. Comparisons of decadal averages clearly show polar warming trends, especially in the north.

A climate model solves the complex physical equations that balance atmospheric and oceanic flows with incoming solar radiation and energy absorption. The project focuses on model development and the software design and engineering required by the climate research community to develop and maintain a high-quality climate prediction capability that effectively uses high-performance computers.

High-performance computers that have thousands of processors acting simultaneously require scalable software to take advantage of unprecedented computer power. The focus of the DOE effort is on joining the atmospheric, oceanic, and sea-ice submodels in a way that efficiently uses the supercomputers at ORNL and other national laboratories.

"We have two primary goals," Drake says. "First, we want to restructure and redesign the climate-model components with a view to enhancing performance on a range of computing platforms. Second, we hope to extend the model for more realistic climate simulations by including hydrological and chemical processes with interactive carbon cycles."

Other institutions involved in the project are DOE's Argonne National Laboratory, Lawrence Livermore National Laboratory, and Pacific Northwest National Laboratory, as well as the National Center for Atmospheric Research and the National Aeronautics and Space Administration (NASA).

In support of the rapidly growing field of climate-simulation science, DOE's Center for Computational Sciences (CCS) at ORNL has also initiated a new focus area called Climate and Carbon Research (CCR). Its role is to conduct world-class computational climate research and use the hardware and software tools in the CCS to assess future climate simulations. The CCR director is David Erickson. CCR staff will work closely with researchers at other DOE labs, NASA, the National Science Foundation supercomputer centers, and universities on important issues regarding climate simulation.

Chemical Experiments and Predictions by Computer

Supercomputers can be used to simulate chemical reactions, saving time and money and increasing safety.

The term "chemistry" conjures up images of people in white lab coats, pouring liquids from test tubes into a beaker. But an increasing number of chemists do most of their chemistry on computers, partly to save money and increase safety. Simulating chemistry experiments on the computer instead of doing them in the lab can produce chemical reactions and properties with a faster turnaround and better accuracy. Computational chemistry also can aid understanding of actual lab results.

David Bernholdt, who helped develop the NWChem computational chemistry software package while a postdoctoral scientist at DOE's Pacific Northwest National Laboratory, leads the chemistry initiative at DOE's Center for Computational Sciences at ORNL. Bernholdt knows the value of computational chemistry, so he is working to develop scaling techniques that will allow chemists to work with even larger molecules on ORNL's new supercomputers.

"Computational chemistry allows you to predict which chemical compounds are more likely to give the desired property or result," he says. "For instance, computational chemistry was used by Kodak's color film developers to predict which chemicals would produce the right colors yet still hold up during chemical processing to develop the film. It saves you from having to synthesize lots of different chemicals and then screen them for the needed properties, such as what might be effective in a therapeutic drug."

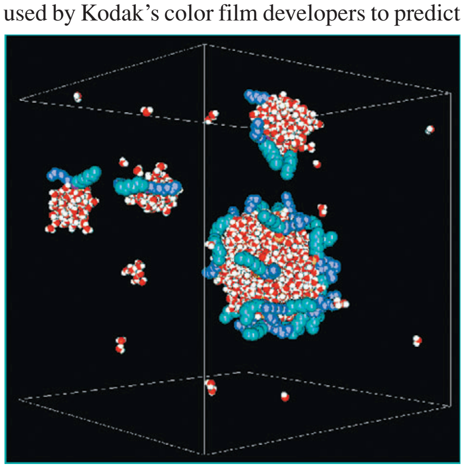

Snapshot of a droplet of water (red and white), too small to be seen, which has formed spontaneously in carbon dioxide (invisible for clarity) and is stabilized by a novel surfactant (blue and turquoise). The size of the droplet formed in the simulation agreed with experimental measurements.

At ORNL several researchers have been using supercomputers to work on computational chemistry projects. Don Noid and Bobby Sumpter, both in ORNL's Computer Science and Mathematics Division, were the first to computationally model the dynamics of fluid flow inside carbon nanotubes—tiny cylinders resembling rolled-up chicken wire. Their molecular dynamics simulations showed that argon slowed down more quickly than helium and that both fluids slowed down faster in a flexible nanotube than in a rigid one.

More recently, they found in their simulation studies with Clemson University researchers that hydrocarbons break up more slowly when heated in carbon nanotubes than in a furnace under vacuum. They believe the rate of breakup of these polymers (similar to crude oil cracking under pyrolysis) is altered by the chemistry of confined spaces.

Working with Mike Barnes of ORNL's Chemical Sciences Division (CSD), Noid and Sumpter showed both computationally and experimentally that a nanosized polymer droplet can be formed from materials that don't normally mix and that the droplet can be forced through a micron-sized orifice. The resulting particles exhibit unique properties that could make them useful for optical displays and industrial coatings. The researchers also helped develop a software tool to calculate how far and which way thousands of atoms move relative to their neighbors (vibrational modes), providing insight into the structure and behavior of various materials.

Hank Cochran, Peter Cummings, and Shengting Cui, all with both CSD and the University of Tennessee at Knoxville, are doing molecular simulations of microdispersions stabilized by surfactants in supercritical carbon dioxide (CO_2). Surfactants act like detergent by reducing surface tension; they are being used in several new dry cleaning establishments where CO_2 replaces traditionally used carcinogenic solvents. DuPont is building a big plant in North Carolina to produce Teflon and other fluoryl polymers using other new surfactants with supercritical CO_2 instead of ozone-destroying chlorinated fluorocarbons.

The CSD group is also simulating the effects of velocity gradients (shear flow) on the arrangement and behavior of long-chain molecules in such situations as during the extrusion of filaments and the melting of polymers. In addition, they are looking at the behavior of lubricants in the narrowest nanoscale separations between motor vehicle components, which can be substantially different from their behavior in bulk.

Cochran and Cui are simulating the behavior of water that contains salts and DNA or proteins in channels thousands of times smaller than a hair in a "nanofluidic lab on a chip." When developed, under the leadership of lab-on-a-chip inventor Mike Ramsey of CSD, such a device might be used in a doctor's office for ultrafast DNA sequencing of blood drops from individual patients for rapid disease diagnosis. Liquids containing DNA or protein could be moved by the influence of electric fields through ultrasmall channels, which might enable the increase of sequencing and diagnostic speeds by a million to a billion times. In their simulations, the researchers take into account the electric-field and surface forces that extend through the liquid inside the nanoscale channel. These simulations help to guide and interpret Ramsey's experiments.

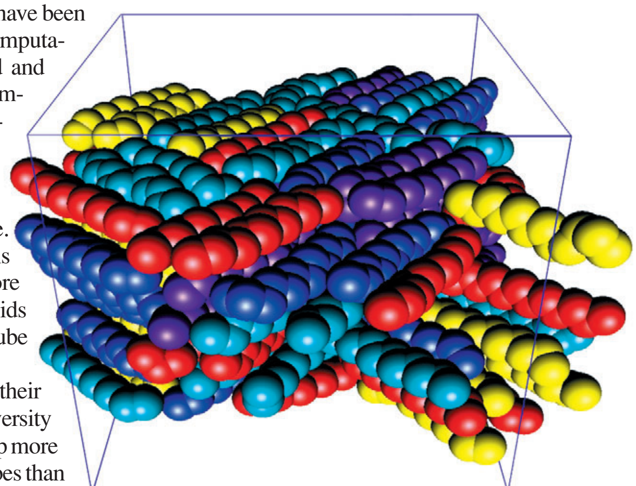

Solid-like structure (with short-range order in three dimensions) of n-dodecane narrowly confined between solid surfaces is induced by interfacial forces. This effect may explain the orders-of-magnitude higher viscosity observed in confined-fluid experiments compared with bulk fluid values.

As devices get smaller, computational chemistry is likely to play a bigger role in figuring out how to make them work.

Computer Modeling Aids Understanding of Plasma Physics

ORNL fusion researchers are using supercomputers to understand plasma turbulence, design a device that could eliminate plasma disruptions, and find ways to get radio waves to not only heat but also control the plasma to allow sustained energy-producing fusion reactions.

Fusion energy research has a long history of employing supercomputers to solve highly complex mathematical equations. Fusion researchers have long used the National Energy Research Scientific Computing Center (NERSC) at the Department of Energy's Lawrence Berkeley National Laboratory in California, which started life in the late 1970s as the Magnetic Fusion Energy Computer Center at DOE's Lawrence Livermore National Laboratory. "The ORNL Fusion Theory program uses the large computers at both NERSC and ORNL," says Don Batchelor, head of the Plasma Theory Group in ORNL's Fusion Energy Division (FED).

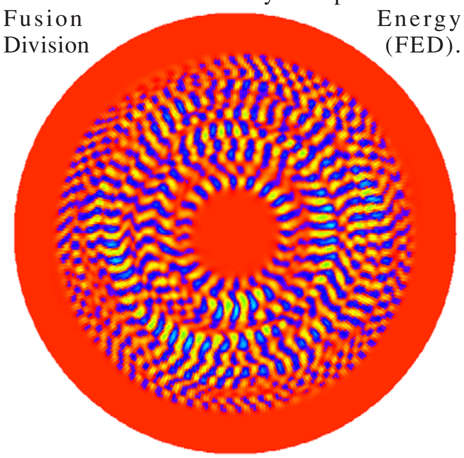

A large-scale three-dimensional turbulence calculation resulted in this color contour of constant T_i fluctuations in the "steady state" of the turbulence. The evolution to fully developed turbulence took weeks to simulate on NERSC's Cray T3E using 64 processors.

"The ORNL capability has dramatically increased our progress in developing large-scale computing applications for fusion research."

Fusion, the energy that powers the sun and stars, is a long-term energy source that could provide an environmentally acceptable alternative to fossil fuels. Achieving fusion energy requires that the fuel, heavy isotopes of hydrogen, be heated to hundreds of millions of degrees, much hotter than the sun. The atoms in matter at such temperatures are torn apart into electrons and ions to form a "fourth" state of matter called plasma (which makes up over 99% of the visible universe). The fuel particles and their energy must then be confined by magnetic fields for a long enough time to produce more energy by fusion reactions than was needed to establish the plasma.

"A great challenge is to understand the physics of how plasma and plasma energy leak out of the carefully constructed magnetic fields used for confinement," says Batchelor. "The dominant transport process in most cases is turbulent motion due to plasma instabilities."

Ben Carreras of FED and Vickie Lynch of ORNL's Computational Science and Engineering Division are carrying out massively parallel computations on the Cray T3E and IBM supercomputers at NERSC and the IBM RS/6000 SP supercomputer at ORNL. These calculations simulate the evolution of certain instabilities that occur in plasma devices, resulting in turbulent fluctuations and greatly increased transport of energy away from the plasma center.

"We are finding that turbulent transport is not following traditional laws of diffusion or heat conduction," says Carreras. "We find that evidence of a non-linear process called self-organized criticality exists and that transport may be described by more complicated integro-differential or 'fractional' differential equations. We are also applying these techniques to fields outside fusion. Methods of self-organized criticality can provide insight into the very timely topic of vulnerability of complex systems such as power grids or communication networks."

The design of very complex, nonsymmetrical magnetic systems to minimize plasma losses in fusion devices is another area in which FED scientists are using supercomputers. ORNL supercomputers are being used in the analysis and design of a new type of magnetic fusion device called the Quasi-Poloidal Stellarator (QPS); see the *Review*, Vol. 34, No. 2, 2001, p. 16, for more details. This device may result in a much smaller and more economically attractive fusion reactor than existing stellarators and would eliminate the potentially damaging plasma disruptions that plague conventional research tokamaks. It is hoped that QPS will be built at ORNL starting in 2003.

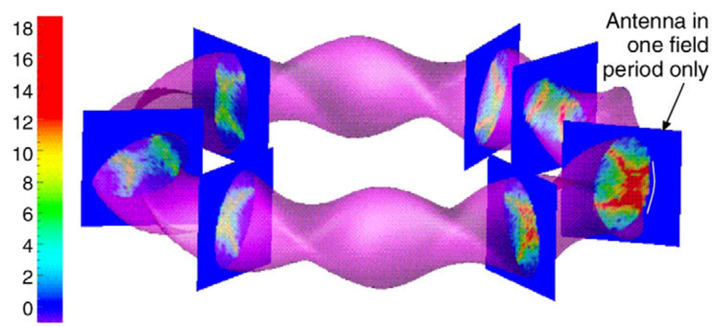

All-orders spectral calculation in three dimensions of minority ion cyclotron heating for a 10-field period stellarator. Individual cross sections show the intensity of ion heating at various toroidal angles.

With the help of ORNL supercomputers and new funding from DOE's Scientific Discovery through Advanced Computation (SciDAC) Program, Batchelor, Fred Jaeger, and Lee Berry, all of FED, and Ed D'Azevedo of ORNL's Computer Science and Mathematics Division are investigating the heating of plasmas to the astronomical temperatures needed for fusion by electromagnetic waves. "Besides heating the plasma in the way that a microwave oven heats food, experiments show that radio waves can drive electric currents through the plasma and force the plasma fluid to flow," says Batchelor. "These waves have even been seen to improve the ability of the applied magnetic field to hold the energetic particles and plasma energy inside the device."

"In 2000 we had a very significant breakthrough in developing a computational technique we call the all-orders spectral algorithm in two dimensions," says Jaeger. "This algorithm eliminates a number of restrictive mathematical approximations to the theory that were previously necessary. Simultaneously, it enables us to study essentially arbitrarily small-scale wave phenomena, limited only by the size and speed of the computer, not the approximations in the theory."

Using 576 processors on the IBM SP computer at ORNL, FED researchers obtained the first converged wave solutions in 2D for an important wave process in fusion called fast wave to ion Bernstein wave mode conversion. According to Berry, "As soon as we heard about the SciDAC award, we pressed ahead as rapidly as possible to implement a three-dimensional version of the all-orders spectral algorithm."

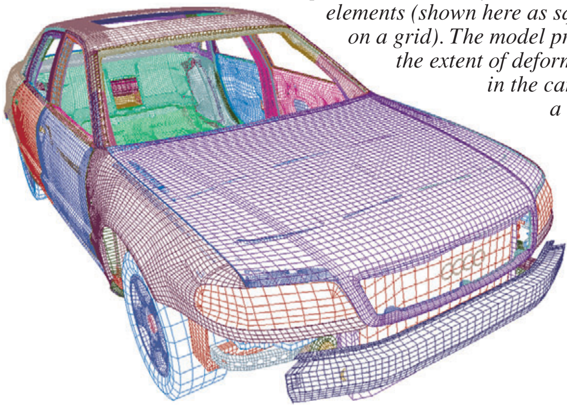

This Audi A8 car-crash model contains numerous materials and structural components modeled by 290,000 finite elements (shown here as squares on a grid). The model predicts the extent of deformation in the car after a crash.

Car Crash Simulations May Improve Vehicle Efficiency

ORNL researchers are building computer models of vehicles made of aluminum, regular steel, high-strength steel, and carbon-fiber composites. This research could lead to safer, energy-efficient cars.

The lighter the structural materials that make up a car, the less fuel it requires per mile of travel. But, if a car is lighter than the steel cars most of us drive, will it hold up as well in a crash?

To help answer this question, the Computational Material Sciences Group in ORNL's Computer Science and Mathematics Division has completed a computational simulation of an all-aluminum Audi A8 car crashing against a rigid barrier at 35 miles per hour (mph). The group, led by Srdan Simunovic, built the computer model after disassembling an Audi A8 car and scanning its structural components into a computer. The model contains equations and numbers representing approximately 290,000 finite elements and 200 different material components. The model has been run on the IBM RS/6000 SP supercomputer at DOE's Center for Computational Sciences at ORNL.

While the ORNL supercomputer was crunching the numbers, an actual crash test of an Audi A8 car was performed by the National Highway Traffic Safety Administration (NHTSA). The test results were used to tune the computer model and determine whether its predictions about the extent of deformation throughout the car were correct.

"The deformation predicted by our model was verified against the crash test data," Simunovic says. "That suggests that this model can be used as a low-cost method for testing new design concepts and materials-processing technologies without the need for building and crashing expensive prototypes. The model can be used to test crashworthiness and analyze stiffness of structural components."

The ORNL group also has developed computer models of vehicles whose bodies are made of regular steel and high-strength steel. With funding from NHTSA, the group recently developed detailed computer models of the Ford Explorer. One material model predicts how the Explorer's body material will behave as the vehicle collides from different angles with a rigid barrier at 35 mph.

"We are also working on the computational analysis of a concept car made of high-strength steel," Simunovic says. The high-strength-steel, UltraLight Steel Auto Body (ULSAB) design and the computational crash models were developed by Porsche Engineering Services, Inc., for the ULSAB Consortium. "We use the models to predict the effects on new advanced materials of various collisions, such as two cars colliding with each other," he says. "Because these new steel alloys have such high strength, less steel is needed for the body of the car, making it lighter. We found that the ability of the high-strength steel vehicles to hold up in a crash can be even better than that of today's heavier steel vehicles."

To explain what happens when cars made of metal crash, Simunovic squeezes an empty beverage can until it folds and collapses like an accordion. "This is what you want to happen to a car during a collision with a rigid barrier or another car," he says. "Metals tend to bend and deform as they absorb the energy of the impact. It is this simple plasticity of metals in response to sudden impacts that we can simulate using our materials modeling codes."

For Simunovic an even bigger challenge is modeling fiber-reinforced polymer composites, a project he has been working on since 1993. These composites, which are lighter than steel and aluminum, consist of glass or carbon fibers embedded in a polymer matrix.

"We are developing constitutive models to predict how the material will behave during an impact at 35 mph," Simunovic says. "Composites don't act like metals and dissipate energy by bending and deforming plastically in response to a blow. Although composites have higher specific strength, they tend to be brittle, making them less likely to give as easily. They are more likely to shatter like glass.

"The impact could cause fibers to break away, or de-bond, from the polymer matrix. The goal is to develop a composite that exhibits controlled progressive fracture during impact. Such a material could dissipate a large amount of impact energy and gradually decelerate the vehicle. We must learn how to model these effects and accurately predict how they improve the ability of the material to resist breaking catastrophically in a crash."

For computer simulations of crashes involving cars made of carbon-fiber composites, the ORNL group will use data from the intermediate strain rate crush test station, which will be installed in 2002 at the National Transportation Research Center, where ORNL and University of Tennessee researchers work. The station will compress samples at speeds up to 15 mph, providing information on changes in the number of small and long cracks produced as the impact velocity varies.

"Our goal," Simunovic says, "is to provide the material models and computational tools that designers need to develop highly efficient, low-emission, lightweight vehicles that have improved safety features." 101

Images of an all-aluminum Audi A8 car during an actual crash test. The car hit a rigid barrier at 35 miles per hour. The crash data are used to help ORNL researchers improve their car-crash computer models.

Evaluating Vehicle Emissions Controls

ORNL researchers are developing software tools for supercomputers that will simulate engine exhaust from various lean-burn diesel and gasoline engines as it flows through envisioned catalytic converters designed to chemically transform pollutants into harmless emissions.

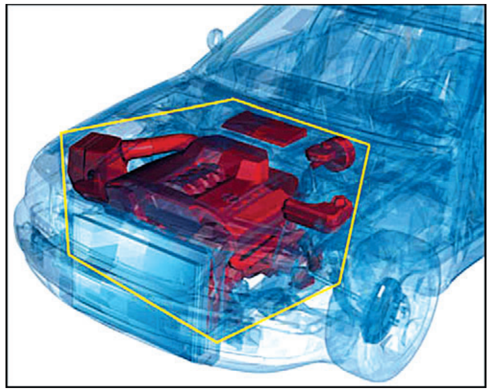

ORNL researchers are developing specialized supercomputer software tools to simulate the transformation of harmful compounds in lean-burn engine exhaust into harmless emissions.

A major stumbling block to putting 80-mile-per-gallon cars on the road within this decade is the lack of effective emissions controls for lean-burn engines. Lean-burn engines, whether diesel- or gasoline-fueled, are designed to carry out combustion with an excess of air. Such combustion achieves increased energy efficiency and reduced emissions of greenhouse gases, such as carbon dioxide. If lean-burn engines could be used for passenger cars, the Department of Energy estimates that fuel economy increases of over 30% could be readily achieved. Such improved efficiencies would be quite a step forward in the move to reduce U.S. dependency on foreign oil. However, in spite of higher efficiency, lean-burn emissions continue to be a problem because cleanup technologies are not available for lean exhaust. Conventional catalytic converters are unable to simultaneously reduce the nitrogen oxides, carbon monoxide, hydrocarbons, and particulates from lean-burn engines to required levels. Thus, the development of new exhaust cleanup technologies is critical.

ORNL's Bill Shelton and Stuart Daw are developing supercomputer software for simulating the physics and chemistry of lean exhaust cleanup devices. Most of the promising cleanup technologies involve complex chemical reactions between the gaseous exhaust species and special solid-phase catalytic materials coated on ceramic substrates. Up to now, the level of complexity involved has restricted the development of new cleanup systems to the construction and testing of experimental prototypes. This empirical approach proved adequate in previous decades for developing the automotive catalysts used today, but it is simply too slow and costly for current needs.

In focusing on detailed simulations of the underlying physical processes of cleanup devices, Shelton and Daw are joining a new generation of researchers who are trying to apply the power of high-performance computing to go beyond empiricism. Specific goals of the ORNL researchers include the simultaneous description of atomic-scale interactions on the surface with models for heat and mass transport between the surface and gas. By accurately modeling the dynamics of cleanup devices at multiple scales, it is expected that development of new technology can be greatly accelerated.

"Simulations based on fundamental physics and chemistry can reveal previously unanticipated approaches for formulating the catalytic materials, and better ways to link them with the engine exhaust can be identified and exploited," Shelton says. "Realistically, some degree of empiricism will always be necessary, but even then accurate simulations can be used to more effectively plan and interpret experiments."

ORNL is emerging as a leader in this field because of its experience in the experimental evaluation of emissions control devices, its considerable expertise in fundamental surface physics and chemistry, and its world-class facilities for high-performance computing. "Through workshops and direct collaborations, we have been bringing together a broad range of experts from national labs, universities, auto makers, emissions control manufacturers, and engine companies who are extremely interested in addressing the problems of lean-exhaust simulation," says Daw. "We hope to help DOE's Office of Transportation Technologies set priorities for research and construct a coordinated approach to overcoming this hurdle in developing a clean, efficient car."

"By combining models for the dominant physical processes at multiple scales, we are obtaining previously unavailable insights into the coupling of local surface transformations and chemical reactions with global heat and mass flow through the devices," Shelton adds. "This approach could lead to more innovative solutions for improving performance of emissions controls. For example, the multi-scale approach is crucial to understanding the durability of the catalytic material—that is, how long the catalyst will function before it must be replaced. Our goal is to understand how global heat flow in the gas and ceramic substrate produces coarsening of catalyst nanoparticles. This information would be used to determine options for reducing coarsening to delay degradation in catalyst performance."

Shelton and Daw emphasize that their long-term objective is to produce simulations of lean-exhaust cleanup that are directly relevant to realistic driving conditions. The availability of simulations at this level will allow industry and DOE to calculate cost-benefit ratios for different lean-exhaust-cleanup technologies so they can make more informed decisions. 101

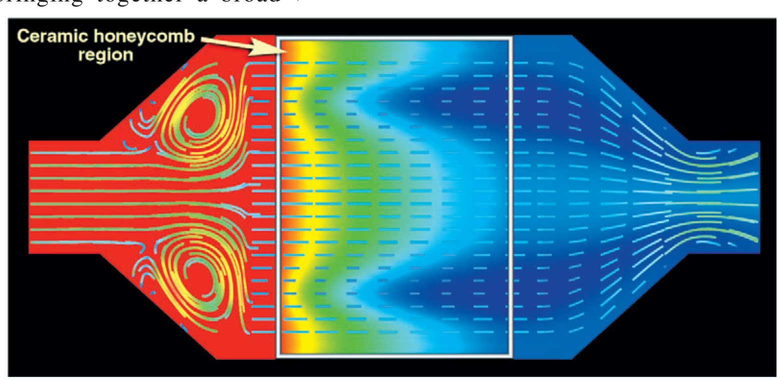

In this computational visualization, the temperature contours and flow streamlines represent the combined effects of exhaust gas flow, heat transport, and chemical reactions in a typical automotive catalytic converter during a cold start. Cold start performance is critical to reducing harmful emissions from automobiles.

Computer Modeling and Homeland Security

ORNL researchers have developed computer-based products that could provide information to help Americans better protect themselves from natural, accidental, or deliberate threats.

Suppose that a group of terrorists releases anthrax disease spores from a crop duster. Could these airborne spores be detected? How many people might be exposed?

What if a group of terrorists detonates a "dirty bomb" in which radioactive materials are dispersed by a conventional explosive? How many people could be exposed to hazardous radioactivity from such an explosion?

Suppose that a Boeing 767 airplane hijacked by terrorists crashes into a nuclear power plant, destroying the building in which spent nuclear fuel is being stored under water and causing the release of large amounts of radioactivity. How many people could be exposed to dangerous amounts of radiation?

These questions have been raised in the news media in the wake of the September 11, 2001, terrorist attacks on America. Research at ORNL using computational simulation and modeling has resulted in products that can help America and other countries address these questions concerning the use of weapons of mass destruction. Computer software and modeling and simulation techniques developed at ORNL can provide American citizens with the information they need to protect themselves from terrorist threats, improving homeland security.

The chemical biological mass spectrometer (right) developed by ORNL researchers will help U.S. Army soldiers detect biological and chemical warfare agents. A CBMS can easily be carried in a military patrol vehicle (shown below).

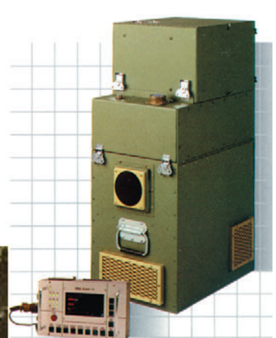

COMBINING TECHNOLOGIES

The only device in the world capable of detecting both chemical and biological threats in minutes has been built at ORNL with help from its partners. This system won an R&D 100 Award in 2000 as one of the most significant technological advances of the year. This Block II chemical-biological mass spectrometer (CBMS) is the first integrated system capable of detecting and identifying chemical warfare agents, such as mustard gas and sarin, and biological warfare agents, such as anthrax spores.

ORNL researchers in the newly formed Computational Sciences and Engineering Division (CSED), Chemical Sciences Division (CSD), and Engineering Science and Technology Division (ESTD) worked with several partners to develop the CBMS for the U.S. Army Soldier Biological Chemical Command. The goal of this five-year program, now in its fifth year, has been to develop a technology to rapidly detect and identify chemical and biological agents in the mobile battlefield. The CBMS achieves this goal by combining a high-sensitivity mass spectrometer, integrated electronics, and an expert system (using software developed by CSED).

The $45-million CBMS program, headed by CSD's Wayne Griest, is responsible for designing, building, and demonstrating prototypes and pre-production units of the Block II CBMS. Tests performed thus far at the Dugway Proving Ground in Utah and by an independent contractor at ORNL have shown that ORNL is on track for meeting program goals.

"The CBMS can be easily operated by soldiers in the field," says Griest. "It also can be used for homeland security by detecting and identifying chemical and biological threats to civilian populations and determining when it is safe for people to return to evacuated areas."

Atmospheric dispersion of hazardous materials released accidentally or deliberately can be predicted by Hazard Prediction and Assessment Capability software.

CSED researchers Bob Morris and Ron Lee are working on improving the Version 4.0 software architecture for the Hazard Prediction and Assessment Capability (HPAC) Program. This program is funded by the Defense Threat Reduction Agency and Strategic Command of the U.S. Department of Defense (DoD).

HPAC is an integrated system of codes and data that provide useful information on the dispersal of hazardous nuclear, biological, and chemical materials released to the atmosphere. HPAC includes detailed, three-dimensional (3D) information on transport and dispersion of hazardous materials through the atmosphere, and their deposition to the ground. All this information is linked to worldwide population information at a one-kilometer (1-km) resolution (see next section), to predict the consequences of various releases (e.g., how many people are being exposed to potentially dangerous concentrations of these materials). The HPAC system provides estimates of acute and long-term doses of radioactive, chemical, or biological materials to the population.

HPAC is currently being used at most military command centers throughout the world. It was used during the Bosnia conflict in the 1990s; the 1996 Summer Olympic Games in Atlanta, Georgia; the inauguration of President George Bush in 2001; and studies of illnesses suffered by soldiers who participated in the Persian Gulf War of 1991. It was used at the 2002 Winter Olympic Games in Salt Lake City, Utah.

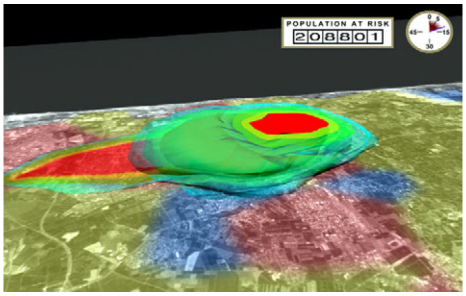

HPAC software developed at ORNL models atmospheric transport and ground deposition of hazardous releases (shown as colorful plume).

In March 2001, Jim Kulesz of CSED chaired an internal quality-assurance surveillance of the HPAC project. Familiar with the CBMS development, he believed that CBMS could be modified to measure the plume concentrations of biological and chemical warfare agents it detects. This information could then be entered into modified HPAC air-dispersion models to predict future migration rates and concentrations of hazardous materials in the plume (prognostics). The models could also use this information to locate the source of the hazardous plume (forensics).

After the incidents of September 11 and subsequent anthrax releases in the U.S. mail system, it became apparent that the combined CBMS-HPAC technologies could be used to enhance homeland security. Sponsors of ORNL work on the CBMS and HPAC projects are now discussing the possibility of funding a collaborative effort at ORNL to develop improved products for both battlefield and homeland security. ORNL's National Security Directorate is actively discussing the concept with military and homeland security agencies.

ASSESSING POPULATIONS IN CRISIS

When a hazardous material is released into the atmosphere, emergency responders to the crisis need to know who is being exposed, where aid is needed the most, how many people are leaving the affected area, and where the "refugees" are going. To help answer these and other questions, CSED's Geographic Information Science and Technology (GIST) Group has developed a world population database that has the world's finest resolution, as well as computer models and 3D visualizations that use this unique database. For this project, the group has applied and improved geographic information systems (GIS) techniques.

The GIST LandScan Population Distribution Project, funded by DoD, developed the LandScan 1998 and LandScan 2000 databases for the entire world. The LandScan database allows the user to assess how many people are present in any given 1-km^2 area. CSED researchers Eddie Bright, Phil Coleman, and Budhendra Bhaduri are currently developing a very high-resolution population distribution database (LandScan USA) for the United States. LandScan USA indicates the number of people in any 90-m cell. The unique feature of LandScan USA is that in addition to residential (nighttime) population, it predicts a daytime population at the same spatial resolution. Because natural and human-induced disasters often strike during the day, the LandScan USA daytime population distribution data is critical in emergency management and response applications.

LandScan uses population numbers provided by the U.S. Census Bureau's International Program and geographic data derived from satellite imagery. The LandScan Population Distribution Model takes into account various factors—such as land cover (e.g., buildings, trees, grassy areas), nighttime lights, the slope of the ground, and the proximity to roads—to determine the likelihood that a cell has a population.

Recent terrorist attacks on the United States have made emergency responders, legislators, and the general public more aware of their need to know where people are and where they are going. "Most natural and manmade disasters strike very unexpectedly, placing vast populations at risk," says Bhaduri, the leader of the GIST group. "The lack of any efficient advanced warning system compels emergency responders to quickly assess how far and in what direction will a contaminant release disperse, how many people are at risk, who are they, and where are they? Geographic information can significantly aid in quickly answering these critical questions for emergency planning, response, and recovery activities."

To answer these questions more accurately, CSED researchers are combining the capabilities of the LandScan population database with those of the HPAC system. In an HPAC simulation, a plume containing a disease-causing agent travels from its source through the air over a rural area toward a large city. Using its imagery and data, the LandScan database generates 3D animations of potentially affected populations. This visual tool used in conjunction with HPAC allows investigators and planners to "see" where a chemical or biological agent will disperse and which populations will be exposed.

The recent terrorist attacks also aroused concerns about resources being at risk. CSED researchers could provide useful information on the risks to the population of contaminated food and water. Currently, Bhaduri is leading a national pesticide usage impact modeling project. He and his group are assessing which of 10,000 drinking water intakes are most likely to receive pesticides used to protect agricultural crops from attack by insects and pathogens. The project makes its analyses using National Hydrography Data, National Land Cover Data, county-based pesticide usage data, and the national agricultural census.

CSED researchers also can aid authorities in reducing possible terrorist threats to U.S. transportation systems. For the Department of Energy's National Transportation Program, Paul Johnson and Richard Michelhaugh, both of ORNL's Nuclear Science and Technology Division, have developed the Transportation Routing Analysis GIS (TRAGIS) model. This routing tool can be used to determine the highway, railroad, or waterway transportation routes that pass through areas having the least population in the unlikely event of an accident that could release hazardous materials. TRAGIS also has the capability to determine alternative routes (e.g., to avoid areas considered at risk of an imminent terrorist attack).

INFORMATION PROCESSING AGENT

CSED's Collaborative Technologies Group, led by Thomas Potok, has created and demonstrated a new way to integrate, manage, and discover information. Called the Virtual Information Processing Agent Research (VIPAR) system, this tool gathers and integrates information from a number of open sources on the Internet. It then rapidly searches, clusters, and analyzes the combined data and presents the results visually to information analysts. This project, sponsored by DoD's Office of Naval Research, was conducted jointly with the Advanced Technologies Program at DOE's Oak Ridge Y-12 National Security Complex.

The VIPAR system is installed at the U.S. Pacific Command Virtual Information Center, which calls it a "grand-slam home run," as well as at ORNL and the Pacific Disaster Center. VIPAR users no longer have to read through large quantities of information to find the answers they seek. VIPAR graphically displays relevant new information relationships that may suggest trends of interest to military, intelligence, and law-enforcement personnel. This technology could also be applied to other text-intensive processes, such as message routing and organizing electronic mail.

Development of a virtual information processing agent was possible because of several research breakthroughs in managing newspaper information, dynamically adding and clustering new information entering the system, and graphically representing the organized information so that it is easily understandable. Thanks to the VIPAR system, meaningful information can be rapidly extracted from the vast repository of data on the Internet.

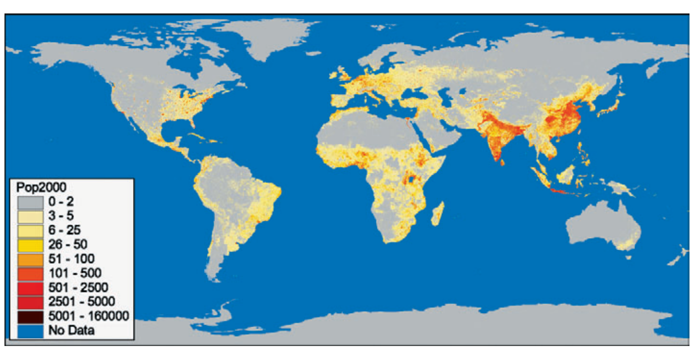

ORNL's Landscan Population Distribution Model produces the finest available global population data.

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

Artist's conception of the new East Campus to be constructed at ORNL by 2003. The blue overlay (not the color of the actual roofs) indicates the two buildings where computer and computational scientists from ORNL and the University of Tennessee will work. The building on the left is the privately funded Computational Sciences Building, which will house the Department of Energy's Center for Computational Sciences. The building on the right, which will be constructed using funds from the state of Tennessee, will house the ORNL-University of Tennessee Joint Institute for Computational Sciences. The buildings behind these new buildings (center and top right) are part of the existing ORNL facility. *Rendering by John K. Jordan.*