

NERSC

NATIONAL ENERGY RESEARCH
SCIENTIFIC COMPUTING CENTER

PUB-779/10-96

RECEIVED

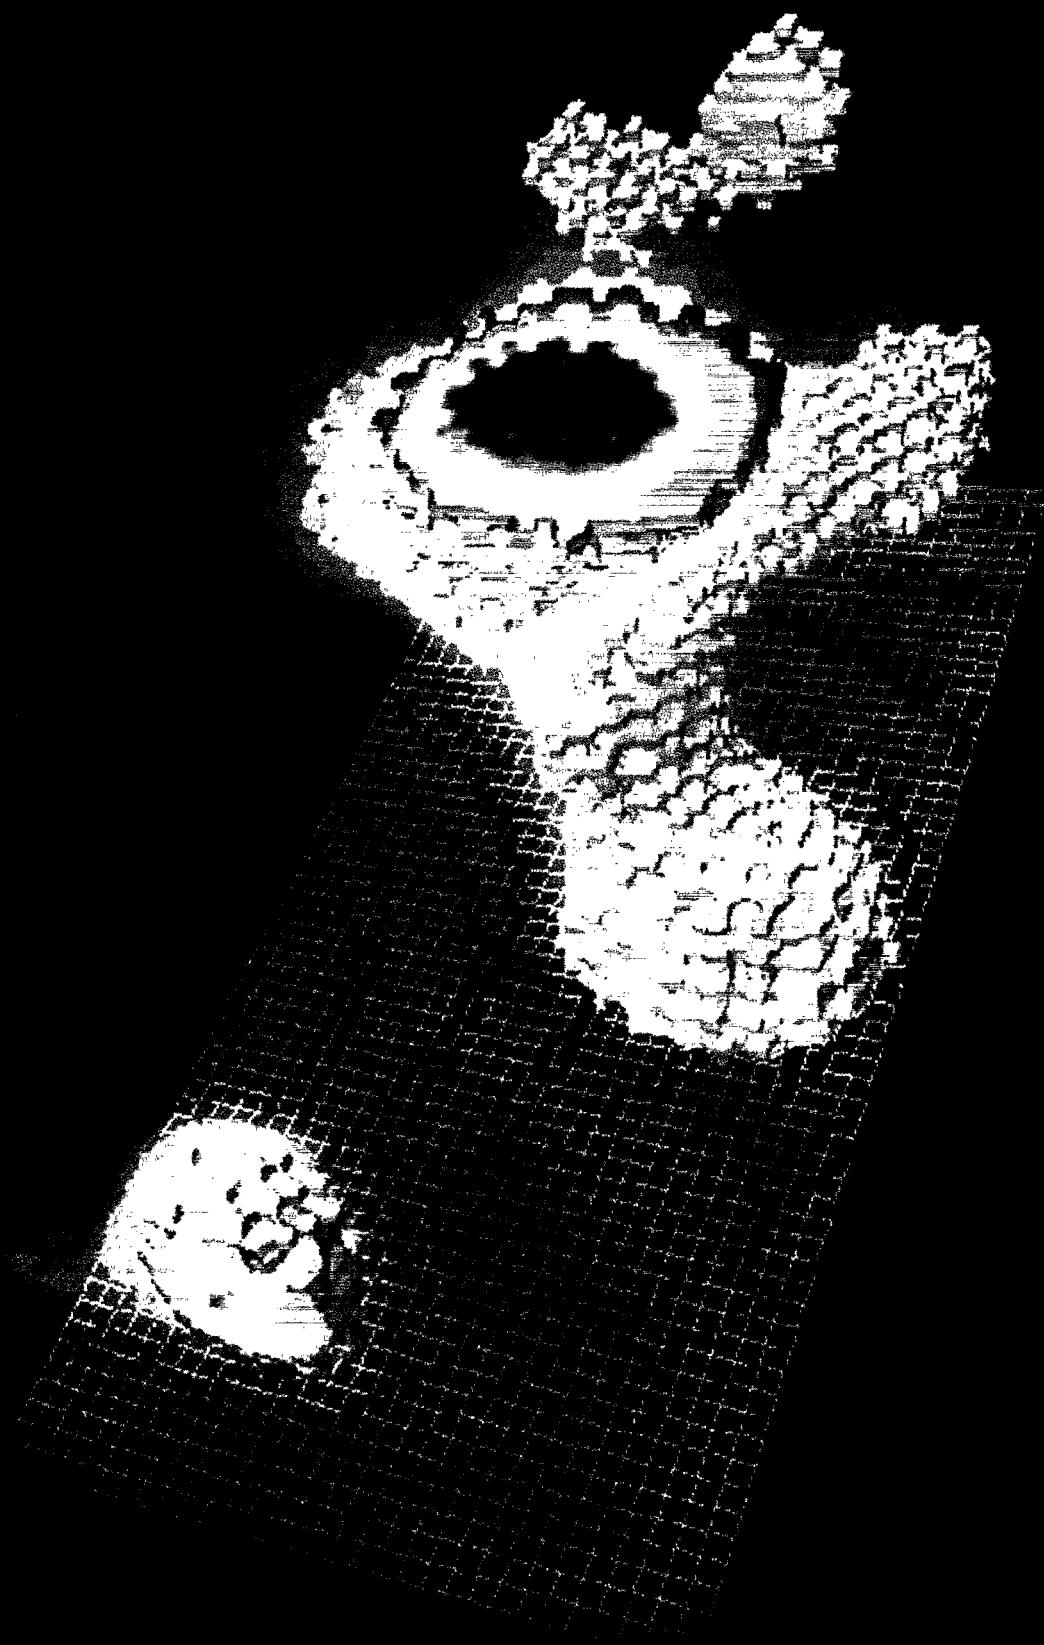
NOV 14 1996

OSTI



MASTER

ADVANCING THE FRONTIERS OF COMPUTATIONAL SCIENCE & TECHNOLOGY



DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

A UNIQUE RESOURCE

Computation has joined experiment and theory as one of the three principal elements of modern scientific research. Questions from the frontiers of science not only require high-performance computers, they also drive the evolution of computer hardware and software.

The National Energy Research Scientific Computing Center (NERSC) provides researchers with high-performance computing tools to tackle science's biggest and most challenging problems. In addition to providing high-end computing services, our goal is to play a major role in advancing large-scale computational science and computing technology.

NERSC works closely with the scientific community we serve to improve the tools of computational science, to demonstrate and advance the value of computation as a complement to traditional theory and experiment, and to be the engine through which significant scientific progress is achieved.

***NERSC works closely with the scientific community
to improve the tools of computational science and to be the engine
through which significant scientific progress is achieved.***



Founded in 1974 by the Department of Energy, Office of Energy Research (DOE/OER), the Controlled Thermonuclear Research Computer Center was the first unclassified supercomputer center and was the model for those that followed. Over the years the center's name was changed to the National Magnetic Fusion Energy Computer Center and later the National Energy Research Supercomputer Center (NERSC). The current name was adopted with NERSC's relocation to the Ernest Orlando Lawrence Berkeley National Laboratory (Berkeley Lab), adjacent to the University of California campus in Berkeley, and reflects NERSC's merger with Berkeley Lab's Computing Sciences program.

NERSC is one of the largest unclassified scientific computing resources in the world and is the principal provider of general-purpose computing services to OER programs—Magnetic Fusion Energy, High Energy and Nuclear Physics, Basic Energy Sciences, Health and Environmental Research, and the Office of Computational and Technology Research. NERSC funding and program guidance is provided by the Mathematical, Information, and Computational Sciences Division (MICS) of the Office of Computational and Technology Research.

A technical staff of 65 and a \$30 million budget enable NERSC to serve several thousand researchers as they collaborate to solve a wide variety of problems related to OER's missions in science, technology, and the environment. NERSC users are a diverse community located throughout the U.S. and in several foreign countries. Thirty percent of our usage comes from researchers in universities, 65% from national laboratories, and 5% from private industry. Nationwide and international high-speed access to NERSC is provided by the Energy Sciences Network (ESnet), whose management is co-located with NERSC in the Computing Sciences Directorate at Berkeley Lab.

OPPOSITE PAGE:
This visualization from calculations run on the NERSC T3E shows one slice of a four-dimensional (32³ x 64 lattice) simulation of typical quantum fluctuations of the quark field. The number and size of the fluctuations ultimately determine the properties of matter.
(Gregory Kilcup, The Ohio State University; NERSC Visualization Group.)



THE NERSC ADVANTAGE

Advances in the physical sciences and in computational science are happening at the interface between these two fields. Traditionally, scientific theory and experiment have relied on advances in mathematics and computing to enable them to address ever larger and more complex problems. In recent decades, new strategies for scientific research have appeared, stimulated by the development of powerful new technologies by the computer science community. Many modern "computational experiments" would have been unthinkable when computing was based on submitting decks of computer cards to an operator. Computational exploration and analysis involving simulation, sophisticated numerical algorithms, and innovative imaging and visualization techniques have become known collectively as computational science.



ESnet and NERSC operators collaborate on high-speed data transfers.

The growing importance of computational science has led Berkeley Lab to create a major new program in Computing Sciences, parallel to its Energy Sciences, General Sciences, and Biosciences programs. The objectives of the Computing Sciences program are:

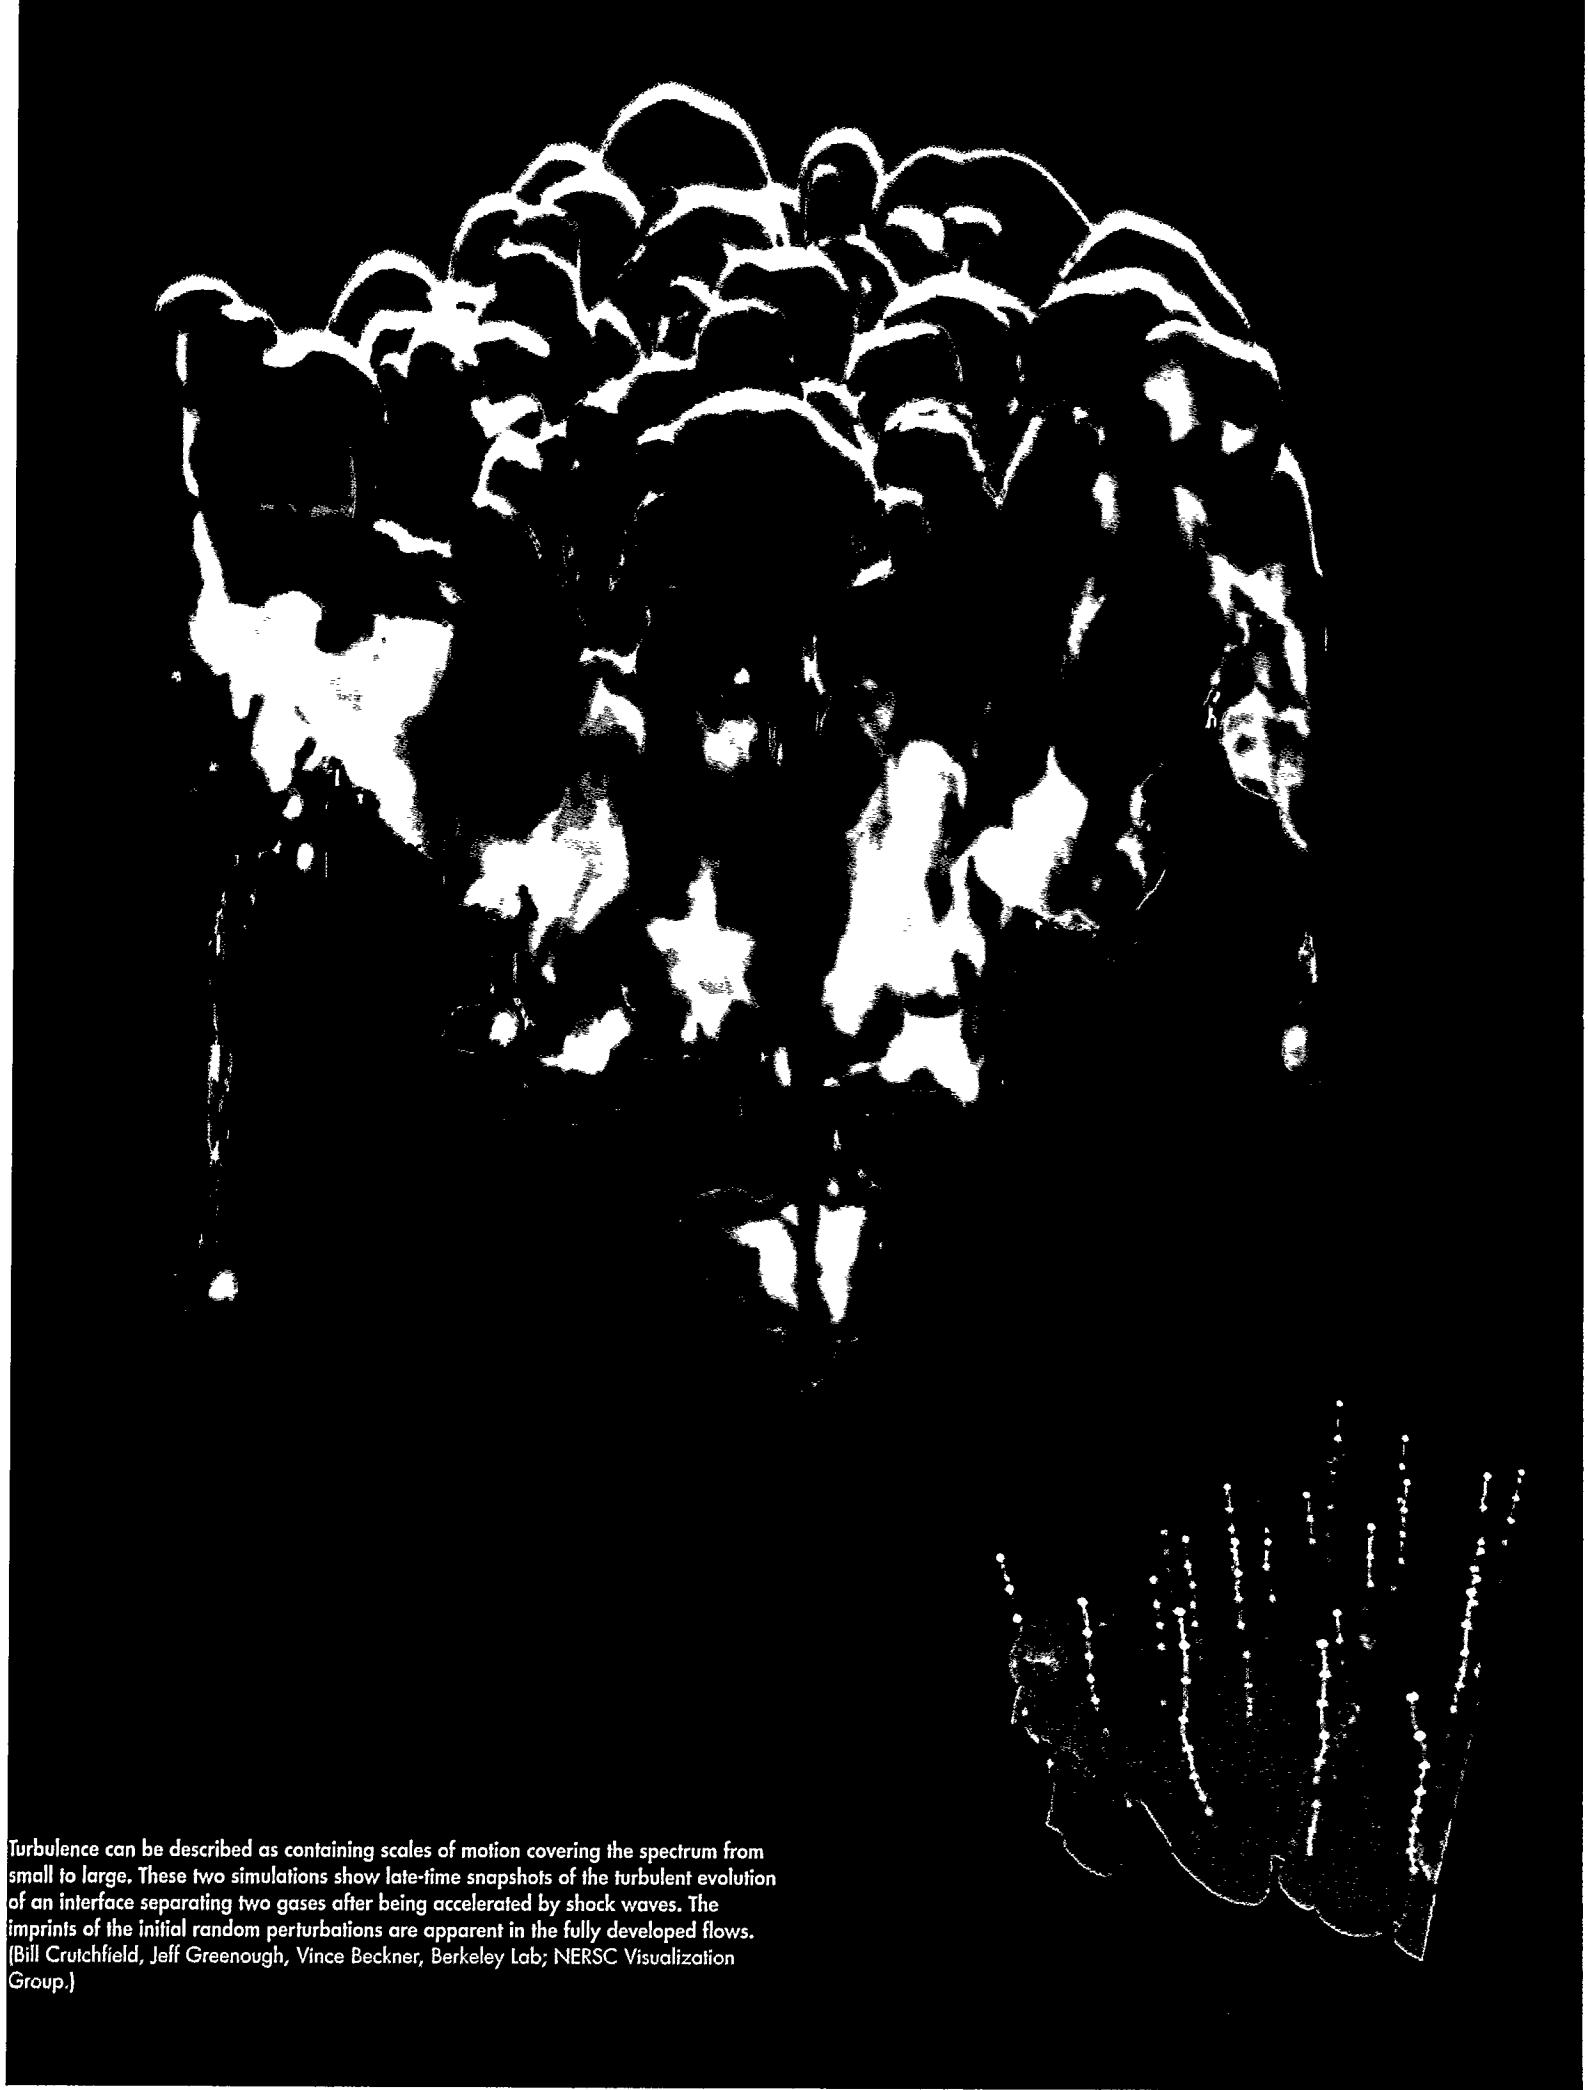
- Promote advances in science through the close integration of computation and telecommunications with all aspects of the scientific endeavor.
- Provide stable, advanced, cost-effective computational and distributed collaborative capabilities to the Energy Research science community.
- Promote mathematics, computer science, and advanced industry products that are relevant to advancing the state of the art in scientific computing.

***NERSC has become an active partner
in the scientific programs that both shape and benefit
from high-performance computing.***

By collaborating with the Berkeley Lab Computing Sciences program, NERSC has become an active partner in the scientific programs that both shape and benefit from high-performance computing. Our new environment provides additional opportunities to quicken the pace of scientific progress by promoting and facilitating the use of the newest computational technologies.

Given our proximity to the UC Berkeley campus, the Mathematical Sciences Research Institute, major computer and data communications companies in the San Francisco Bay Area, and the diverse scientific research programs at Berkeley Lab, NERSC expects to achieve a level of scientific and technological collaboration unequaled by any other high-performance computer center. This collaboration will result in new approaches and solutions to scientific problems and will help shape the next generation of high-performance computer technology. The following are typical examples of research conducted with NERSC resources:

- Using simulations of chemical systems, theoretical chemists are investigating molecular-level processes that may prove useful in the development of new approaches to cleaning up radioactive and toxic pollutants.
- Structural biologists are using computer models to learn how proteins take the shapes that make them work in our bodies.
- Other OER scientists are modeling the interplay of fluid flow and chemistry in combustion to help build cleaner, more efficient engines.
- The Atmospheric Model Intercomparison Project evaluates competing approaches to global atmospheric studies.
- The international Numerical Tokamak Project attempts to solve basic design problems for future commercial generation of fusion energy.
- High-energy and nuclear physicists are successfully studying the smallest known particles.
- The Particle Detector Simulation Facility (PDSF), hosted by NERSC, is being used to develop prototype technologies that look for rare and important events buried in mountains of data from accelerator experiments.



Turbulence can be described as containing scales of motion covering the spectrum from small to large. These two simulations show late-time snapshots of the turbulent evolution of an interface separating two gases after being accelerated by shock waves. The imprints of the initial random perturbations are apparent in the fully developed flows.
(Bill Crutchfield, Jeff Greenough, Vince Beckner, Berkeley Lab; NERSC Visualization Group.)

COMPUTATIONAL RESOURCES & SERVICES

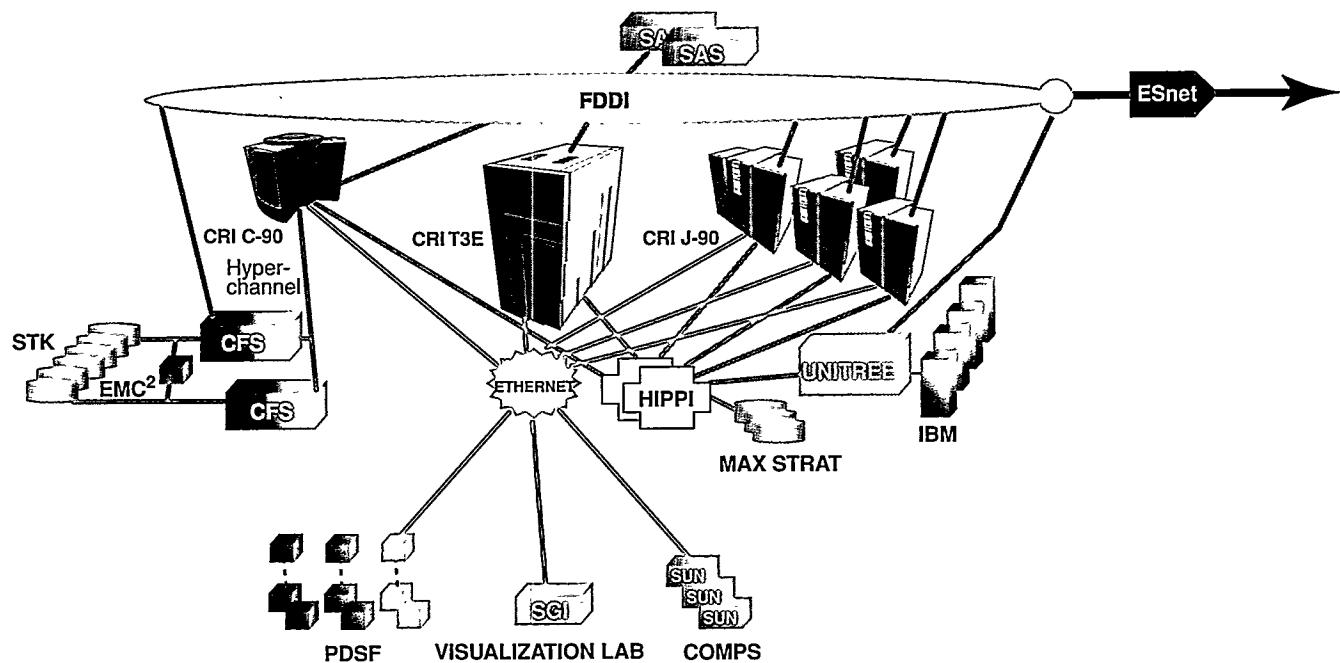
NERSC offers unsurpassed high-end capability to support aggressive science-of-scale projects. In 1997, six high-performance Cray Research computer systems—four J-90s, a C-90, and a T3E—will make up the heart of NERSC's computer hardware capability. Collectively, the four J-90s will have 96 central processing units (CPUs) with a peak aggregate speed of 19 gigaflops (Gflops), 24 gigabytes (GB) of memory, and 750 GB of primary disk storage capacity. The C-90 has 16 CPUs, peak speed of 16 Gflops, and 2 GB of memory. The fully configured T3E will have 512 CPUs with a peak speed of 300 Gflops, 131 GB of memory, and 1.5 terabytes (TB) of disk storage. The PDSF (hosted by NERSC), a multiplatform farm of 124 Hewlett-Packard, Sun, and Silicon Graphics UNIX workstations, provides an additional 12,000 MIPS (million instructions per second) of computing power.

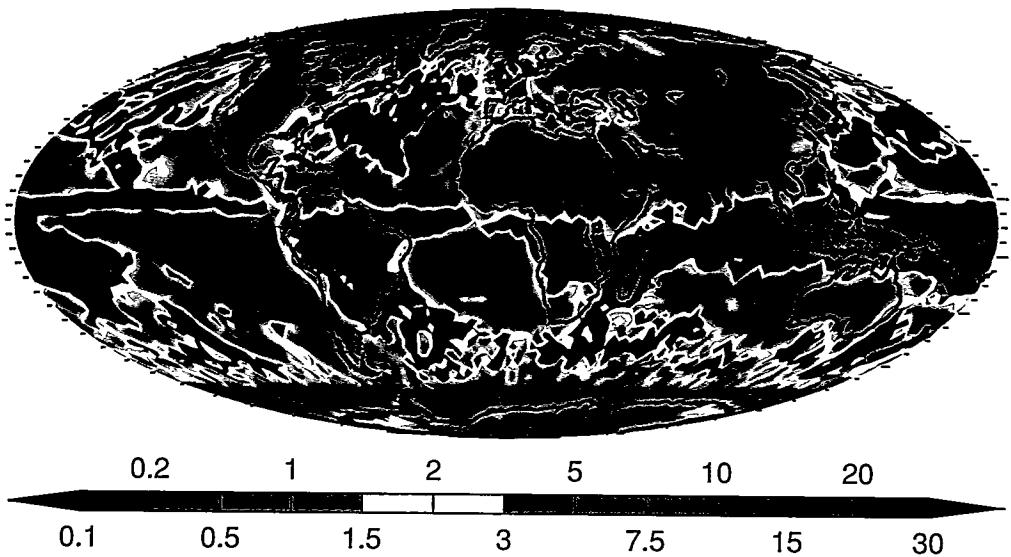
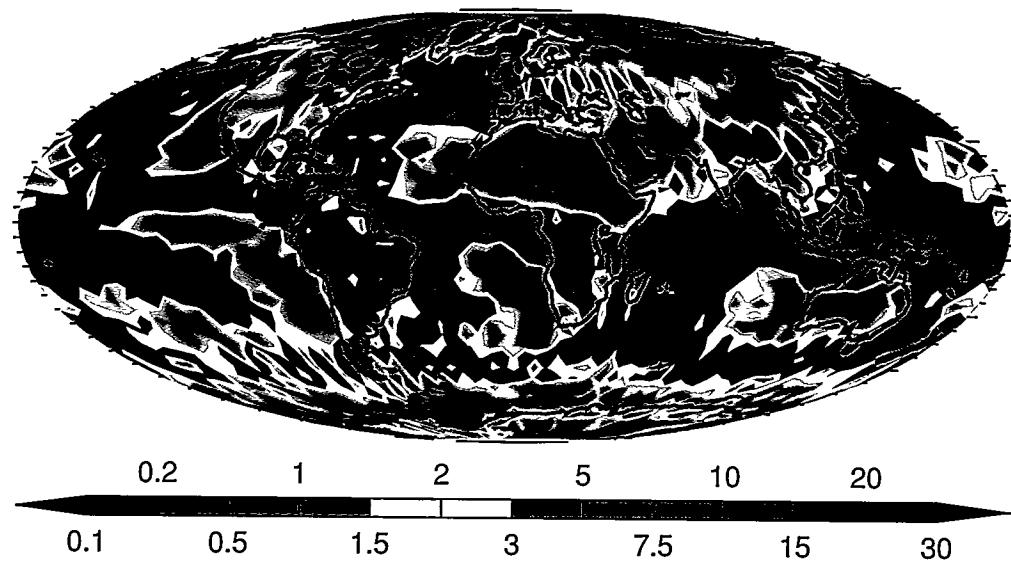
NERSC is committed to responding to clients' problems within four working hours and to resolving at least 90% of problems within two working days.

These computing platforms are supported by an array of mass storage facilities which enable users to store large volumes of data at the highest possible transfer rates with high reliability and availability. These platforms use multi-level storage systems that provide approximately 1 TB of intermediate disk storage and 120 TB of tertiary storage. The computing platforms and storage systems are interconnected by a High Performance Parallel Interface (HIPPI), which provides an 800 megabit data transfer rate. NERSC plans to provide a global filing system directly linked to the storage systems.

To manage and optimize the efficiency of scientific computation on NERSC systems, we maintain dynamic scheduling software that balances computer loads to maximize system efficiency and availability, ensuring that jobs get the capabilities they need. We work with users to enhance the speed of their applications and involve users in decision making regarding future services and system upgrades. Whenever possible, we incorporate new technologies to enhance system capacity and capabilities.

Computational Resources, 1997





NERSC's success is measured by the productivity of our clients' work. We are committed to responding to clients' problems within four working hours and to resolving at least 90% of problems within two working days. Our computers, operating systems, and communications networks are monitored 24 hours a day, and information on system status and scheduled downtime is available to users on line. Our staff consultants respond to user questions, provide one-on-one consulting, and conduct on-line, video, and classroom training. NERSC technical documentation is available on the World Wide Web and in hard-copy publications.

Access to NERSC from anywhere in the U.S. or the world is available through ESnet, which provides T3 bandwidth on major backbone links and T1 links over much of the rest of its coverage area. In addition to providing access to the NERSC computing engines through a Fiber Distributed Data Interface (FDDI), ESnet also offers communications support services (e.g., x.500 directory services and video conferencing).

Users from national laboratories, universities, and industries are represented by the Energy Research Supercomputer Users' Group (ERSUG), which provides guidance to NERSC and OER regarding NERSC services and the direction of future development.

ABOVE: The Atmospheric Model Intercomparison Project evaluates competing approaches to global atmospheric studies. The top map shows the precipitation rate from one climate model, while the bottom map shows the observed precipitation. (Gerald Potter, Lawrence Livermore National Laboratory [LLNL])

FUTURE TECHNOLOGIES

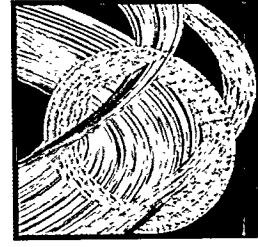
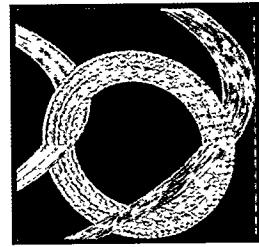
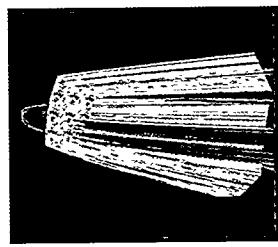
NERSC's goal for the end of the decade is to be a "tera-peta" facility, i.e., to offer computational capacity in the teraflop range (10^{12} floating-point operations per second) combined with storage capacity in the petabyte range (10^{15} bytes of accessible storage). To achieve this goal, we will work closely with computer science research groups at Berkeley Lab, the University of California, and industrial collaborators in Silicon Valley and elsewhere. We will communicate the needs and requirements of NERSC and its users to the computer science community, investigate and evaluate innovative technologies in all areas of computer science (architecture, storage, operating systems, compilers, tools, algorithms, etc.), and recommend and introduce new technologies at NERSC.

With NERSC located at Berkeley Lab, we have the unique combined perspective of a major production facility, computer science programs, computational science programs, and a diverse scientific community. This perspective, combined with the expertise of academic and commercial computer science researchers, will be used to produce prototype facilities that can be tested in a scientific production environment.

***NERSC is producing a prototype facility that will test
the next generation of computer architecture by addressing three different types
of scientific problems in a production environment.***

An example of such a prototype is COMPS (Cluster of Multiprocessor Systems), a joint research project between NERSC, UC Berkeley, Sun Microsystems, and the Berkeley Lab Information and Computing Sciences Division (ICSD) and Materials Sciences Division (MSD). The goal of COMPS is to install a networked cluster of symmetric multiprocessors (SMPs) at Berkeley Lab and evaluate its utility to the scientific community by examining issues encountered during its use as a prototype production facility. The COMPS system will be tested by using it to address three different types of scientific problems: numerical computation, laboratory experiment control, and high-volume data analysis. If the system meets expectations, it will provide an incrementally scalable common platform for all three types of applications that is architecturally similar to the next generation of massively parallel supercomputers.

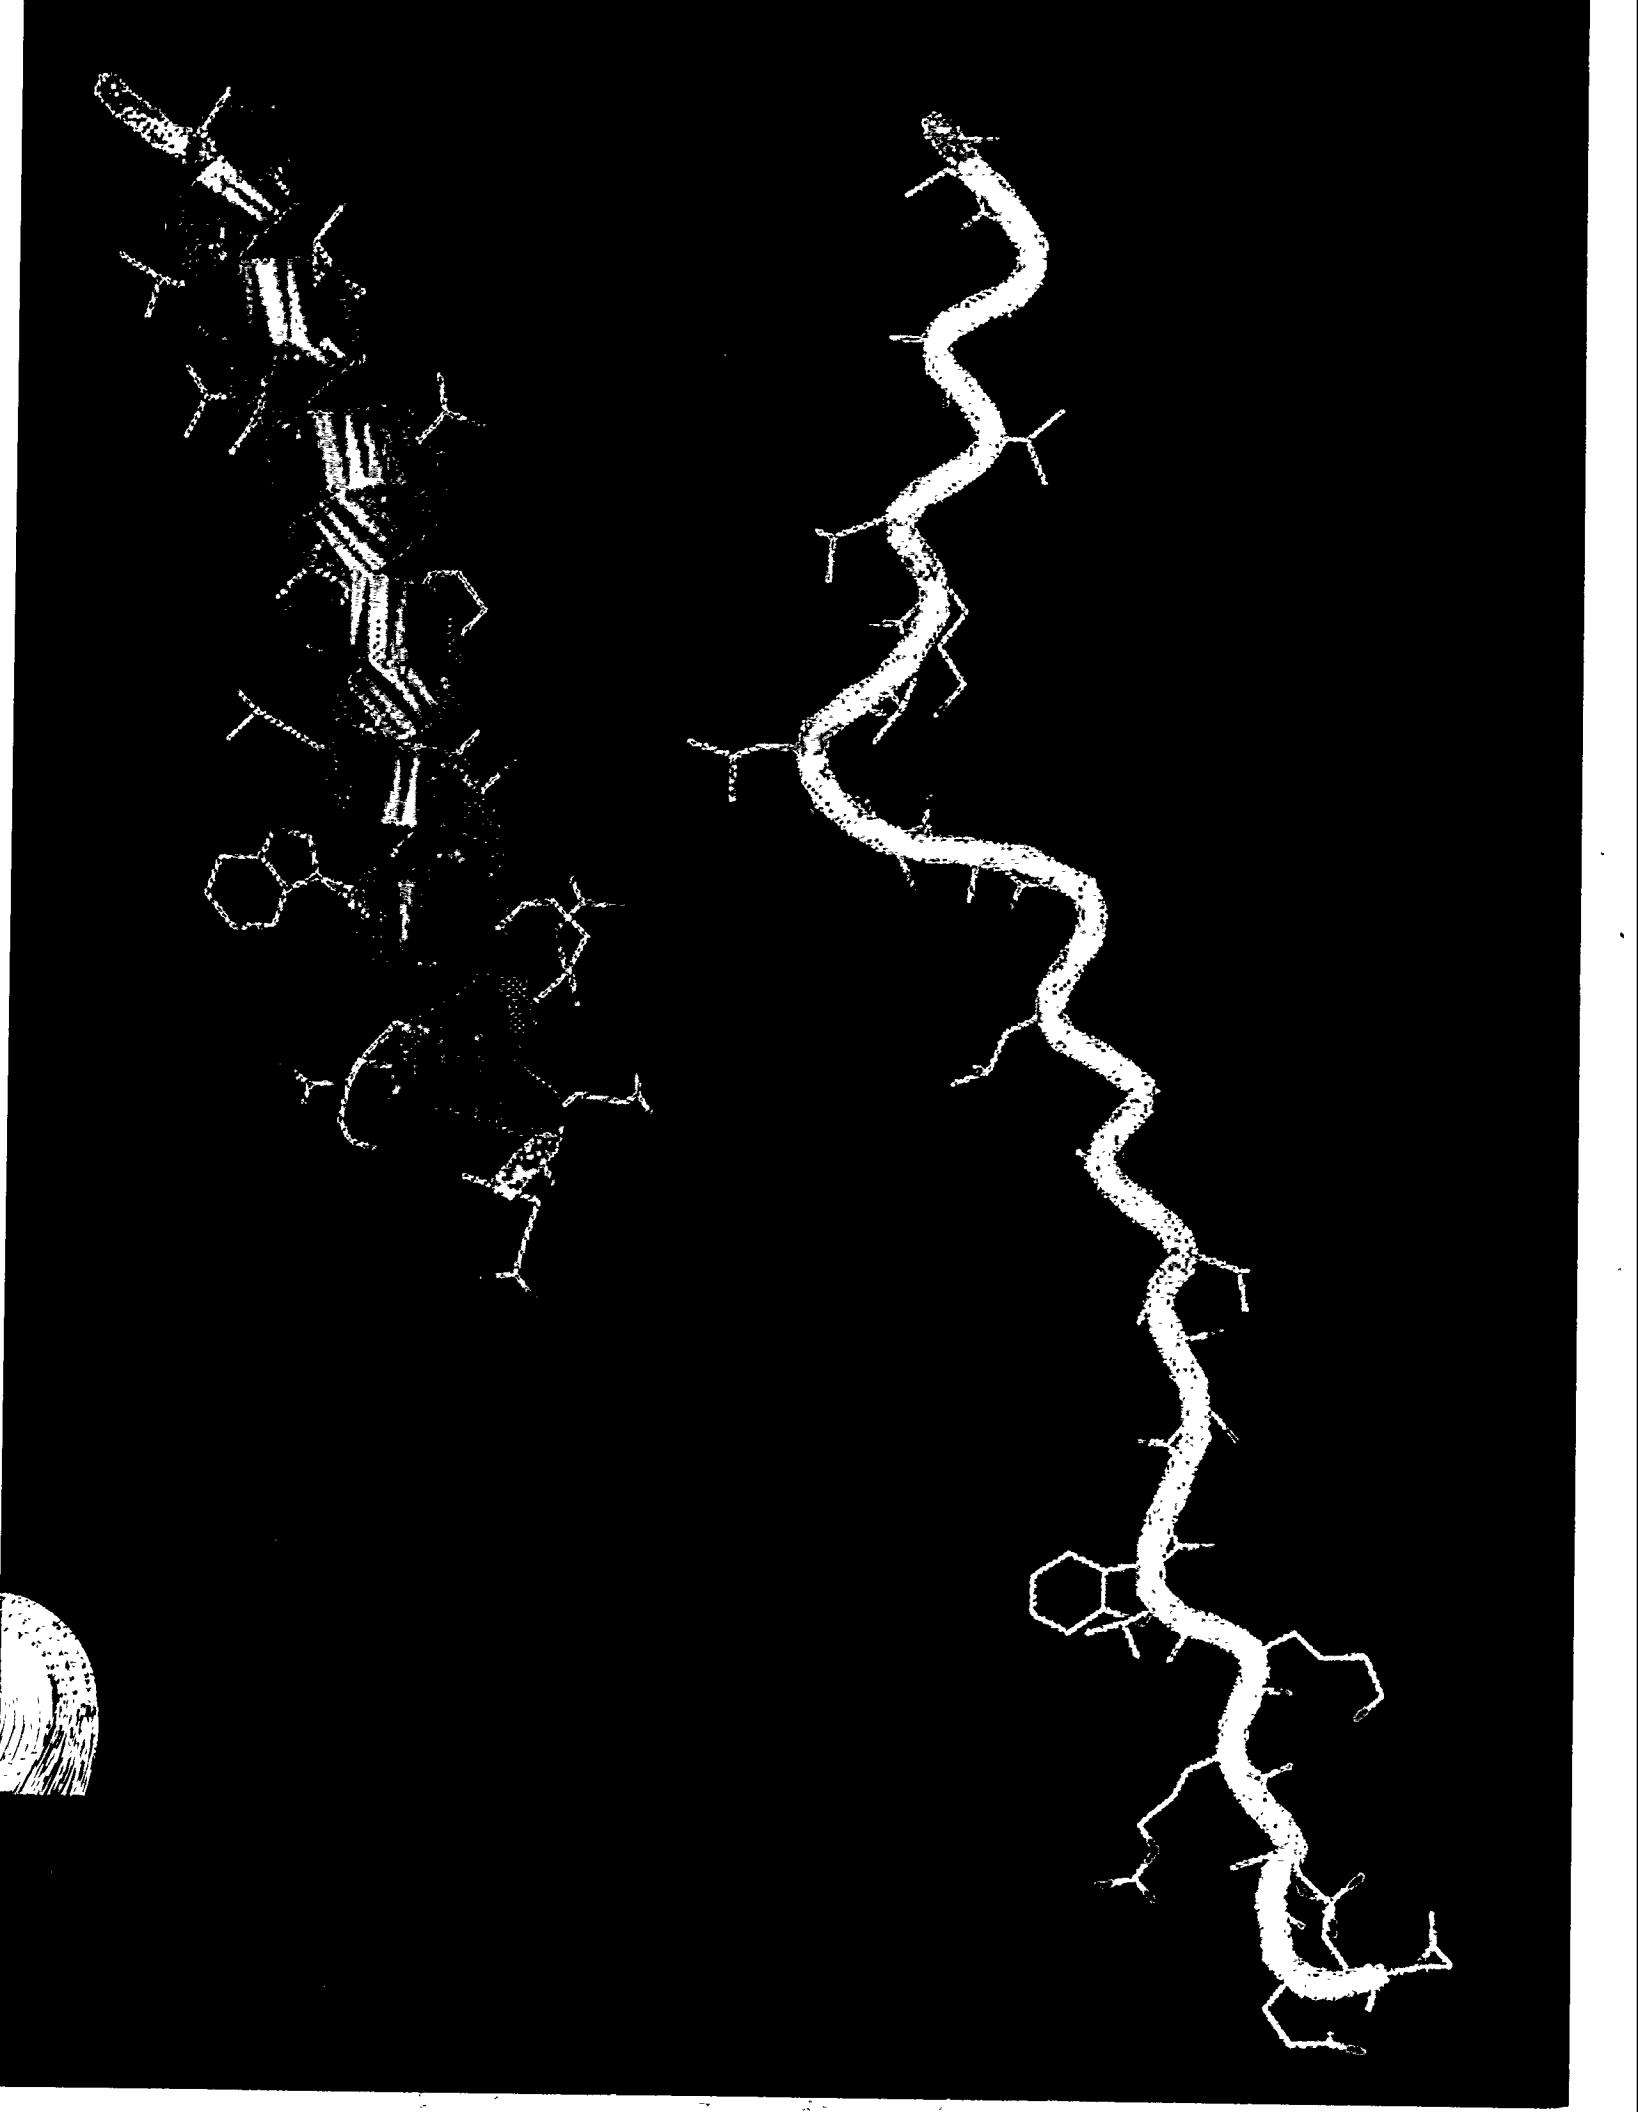
NERSC and ICSD are providing the facilities, network interconnections, numerical algorithms, visualization, and system administration for COMPS. The UC Berkeley Network of Workstations (NOW) Project is supplying systems software and the results of their systems architecture research. Sun Microsystems is providing the workstation hardware and engineering support. And MSD is providing scientific applications and access to experiments at Berkeley Lab's Advanced Light Source. As a result of this collaboration, COMPS will enable NERSC to investigate and contribute to the next generation of computer architecture.



ABOVE: One of the fundamental issues in magnetic fusion research is understanding turbulent heat transport in tokamaks (doughnut-shaped fusion reactors). Visualizations like these, from the Numerical Tokamak Turbulence Grand Challenge Project, enable physicists to examine "slices" of three-dimensional data from turbulence simulations.*

RIGHT: Using sophisticated optimization methods to predict the structure of proteins with 50 or more amino acids is one of the grand challenges in computational biology. This figure shows the successful prediction of the crystal structure form of the 26-residue polypeptide melittin (bee venom). On the right is the simple polypeptide chain; on the left is its folded structure.†

(*Andris Dims, LLNL. †Teresa Head-Gordon, Berkeley Lab.)



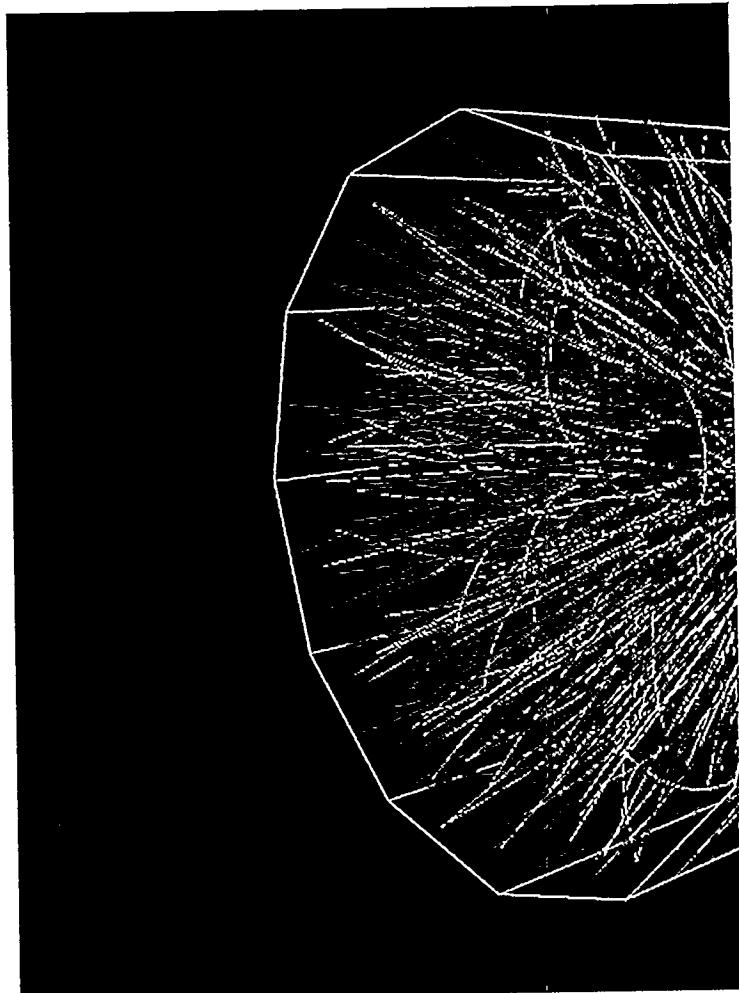
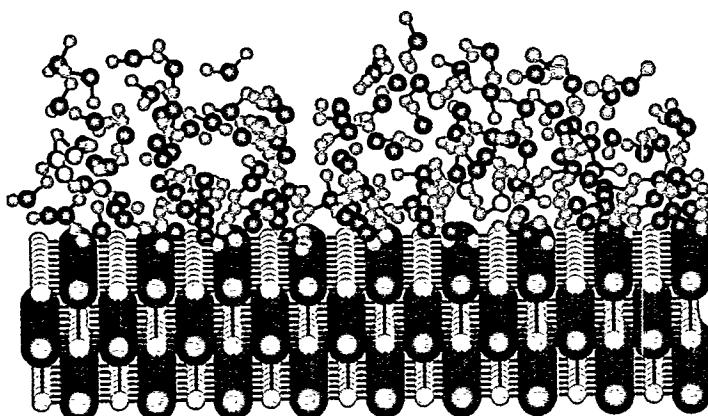
SCIENTIFIC RESOURCES

NERSC's new computational science infrastructure is rich in resources—not just high-performance hardware, but also scientific expertise and research projects that push the limits of computational technologies and methodologies. Using these resources, NERSC will make significant contributions toward integrating high-performance computing into the scientific disciplines by developing and deploying the necessary methods, algorithms, and technologies.

Our scientific computing specialists are taking the lead in achieving this integration by bridging the gap between new computing technology and scientific research. Their mission is to facilitate development of scientific applications that run on NERSC capability platforms and to promote the optimal use of NERSC computing resources. Their responsibilities include collaborating with strategic users to port and develop scientific applications, as well as evaluating, integrating, and creating new hardware and software technologies and new numerical and non-numerical algorithms—in short, developing new computational approaches to scientific problems.

Because the various scientific disciplines require different computational methodologies and algorithms, NERSC's scientific computing specialists have strong backgrounds in both computer science and another discipline, such as physics, chemistry, materials science, or biology. Their dual backgrounds enable them to understand scientific problems and to devise the optimum computational approaches to assist users in developing solutions.

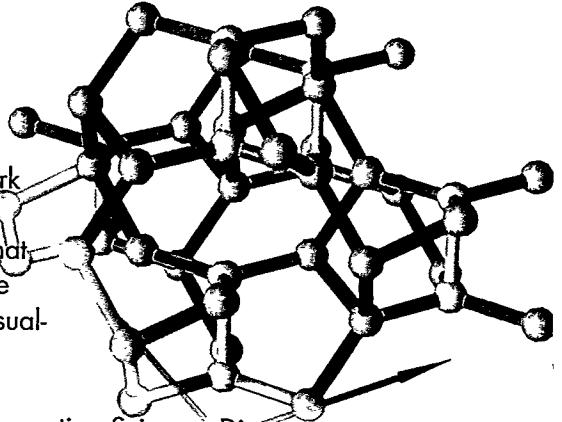
NERSC offers specialized visualization and imaging support to help scientists create simulations and understand the results of their research. We provide state-of-the-art still images, movies, and virtual reality interfaces for scientific data. We also create videos or CDs for presentations or distribution. We keep abreast of current graphics technology and maintain a collection of broadly useful graphics and visualization software for multiple platforms.



NERSC is collaborating with a community of scientists committed to constantly rethinking the four elements of scientific computing—mathematical modeling, algorithmic design, software implementation, and system architecture.

LEFT: Simulations like this one, an interface between water and magnesium oxide, may help chemists understand how to model a variety of industrial and environmental processes (Maureen McCarthy, Pacific Northwest National Laboratory)

We are developing techniques that will allow users on different network bandwidths and hardware platforms to visualize their data and simulations remotely. For example, we developed tools for the Numerical Tokamak Project that allow interactive visual exploration of gyrofluid tokamak turbulence and remote viewing and control of simulations. We are also exploring new methods for visualizing extremely large data sets.



As part of the Computing Sciences Directorate at Berkeley Lab, NERSC draws on the intellectual resources of several other departments in the Directorate which are involved in high-performance applications. These departments include:

Computer Science Research. Several groups within this department have begun working with NERSC. Their research includes:

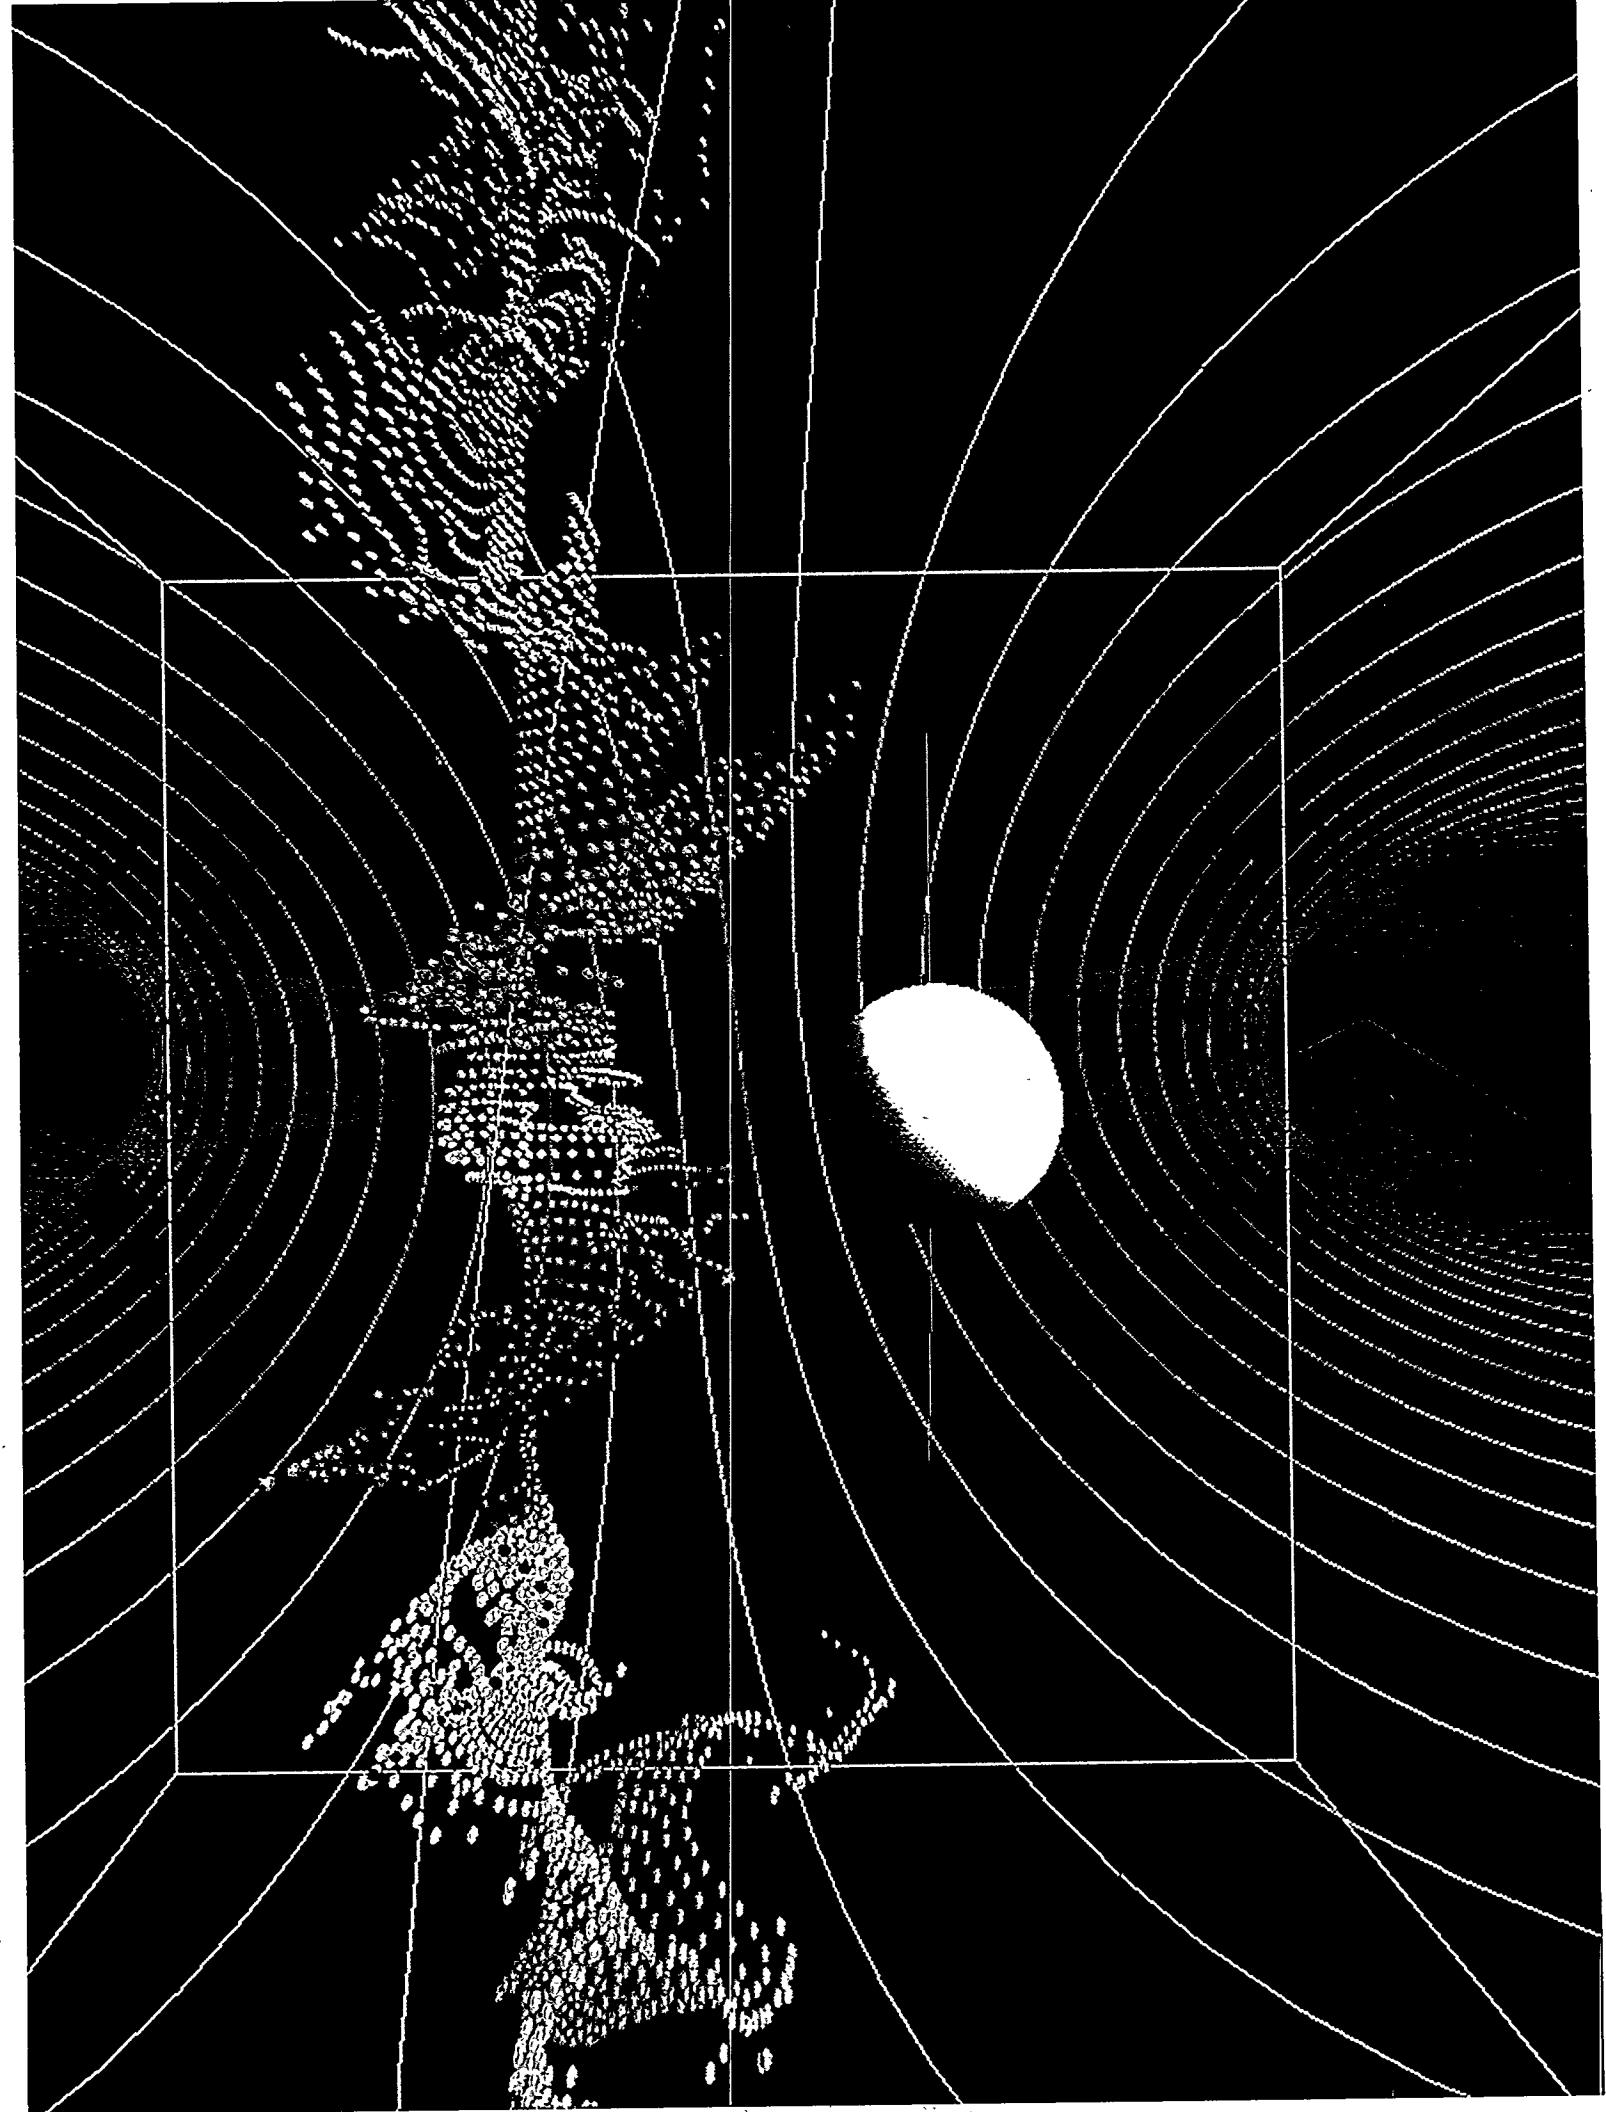
- Developing data management techniques suited to scientific and statistical applications, which have different requirements from commercial database management systems. Researchers are developing new techniques for the efficient storage organization of scientific data, new algorithms to access and manipulate the data efficiently, and new data modeling methods to capture the semantics of scientific data.
- Developing distributed systems to support remote, on-line scientific experiments that include real-time direct control and data analysis. Research includes high-speed distributed computing and networking, computer vision and robotics applications, and image processing and manipulation.
- Network research addressing issues related to protocols and traffic congestion in wide-area networks. This group has developed tools to facilitate multiparty video and audio conferencing over the Multicast Backbone (MBone), a virtual network that sends traffic to multiple remote locations as one data stream.

Mathematics. Located at the intersection of applied mathematics and physical modeling, the Mathematics Department has played a central role in developing and implementing state-of-the-art numerical algorithms used throughout the DOE community.

The Center for Computational Science and Engineering. CCSE develops and applies advanced computational methodologies to solve large-scale scientific and engineering problems. CCSE's application-driven mathematical and numerical research enhances DOE's ability to use high-performance computing as a scientific and engineering tool.

As local users of NERSC services, these departments provide valuable advice and feedback on client support services and the ongoing development of computing resources. Collaboration with a community of scientists committed to constantly rethinking the four elements of scientific computing—mathematical modeling, algorithmic design, software implementation, and system architecture—ensures that NERSC stays committed to new ideas. Our users' needs for computational robustness, reliability, and efficiency provide a solid test bed for new ideas and designs.

ABOVE RIGHT: This figure shows the arrangement of atoms in a recently discovered high-pressure phase of silicon that forms upon decompression in a diamond anvil cell.* CENTER: A simulated collision of two gold nuclei releases thousands of elementary particles. Such reactions help physicists understand the development of the early universe.‡
(*Bernd Pfrommer, Michel Côté, Steven Louie, University of California, Berkeley. ‡Iwona Sakrejda, Berkeley Lab; Herb Ward, University of Texas, Austin.)



COMPUTATIONAL SCIENCE OF SCALE

Computational science of scale is generally characterized by large, interdisciplinary teams of scientists attacking problems that require them to work closely together for a decade or more. These are fundamental problems in science and engineering with broad scientific and economic impacts whose solutions can be advanced by applying the most powerful computing techniques and resources. Science-of-scale applications perform massive calculations, frequently creating high-resolution, three-dimensional simulations of complex physical phenomena.

By continually raising the standards of precision, reliability, and time to solution, science-of-scale projects drive the development of improved software and algorithms.

Researchers throughout the DOE community and academia are collaborating on a variety of science-of-scale projects that utilize NERSC resources. For example:

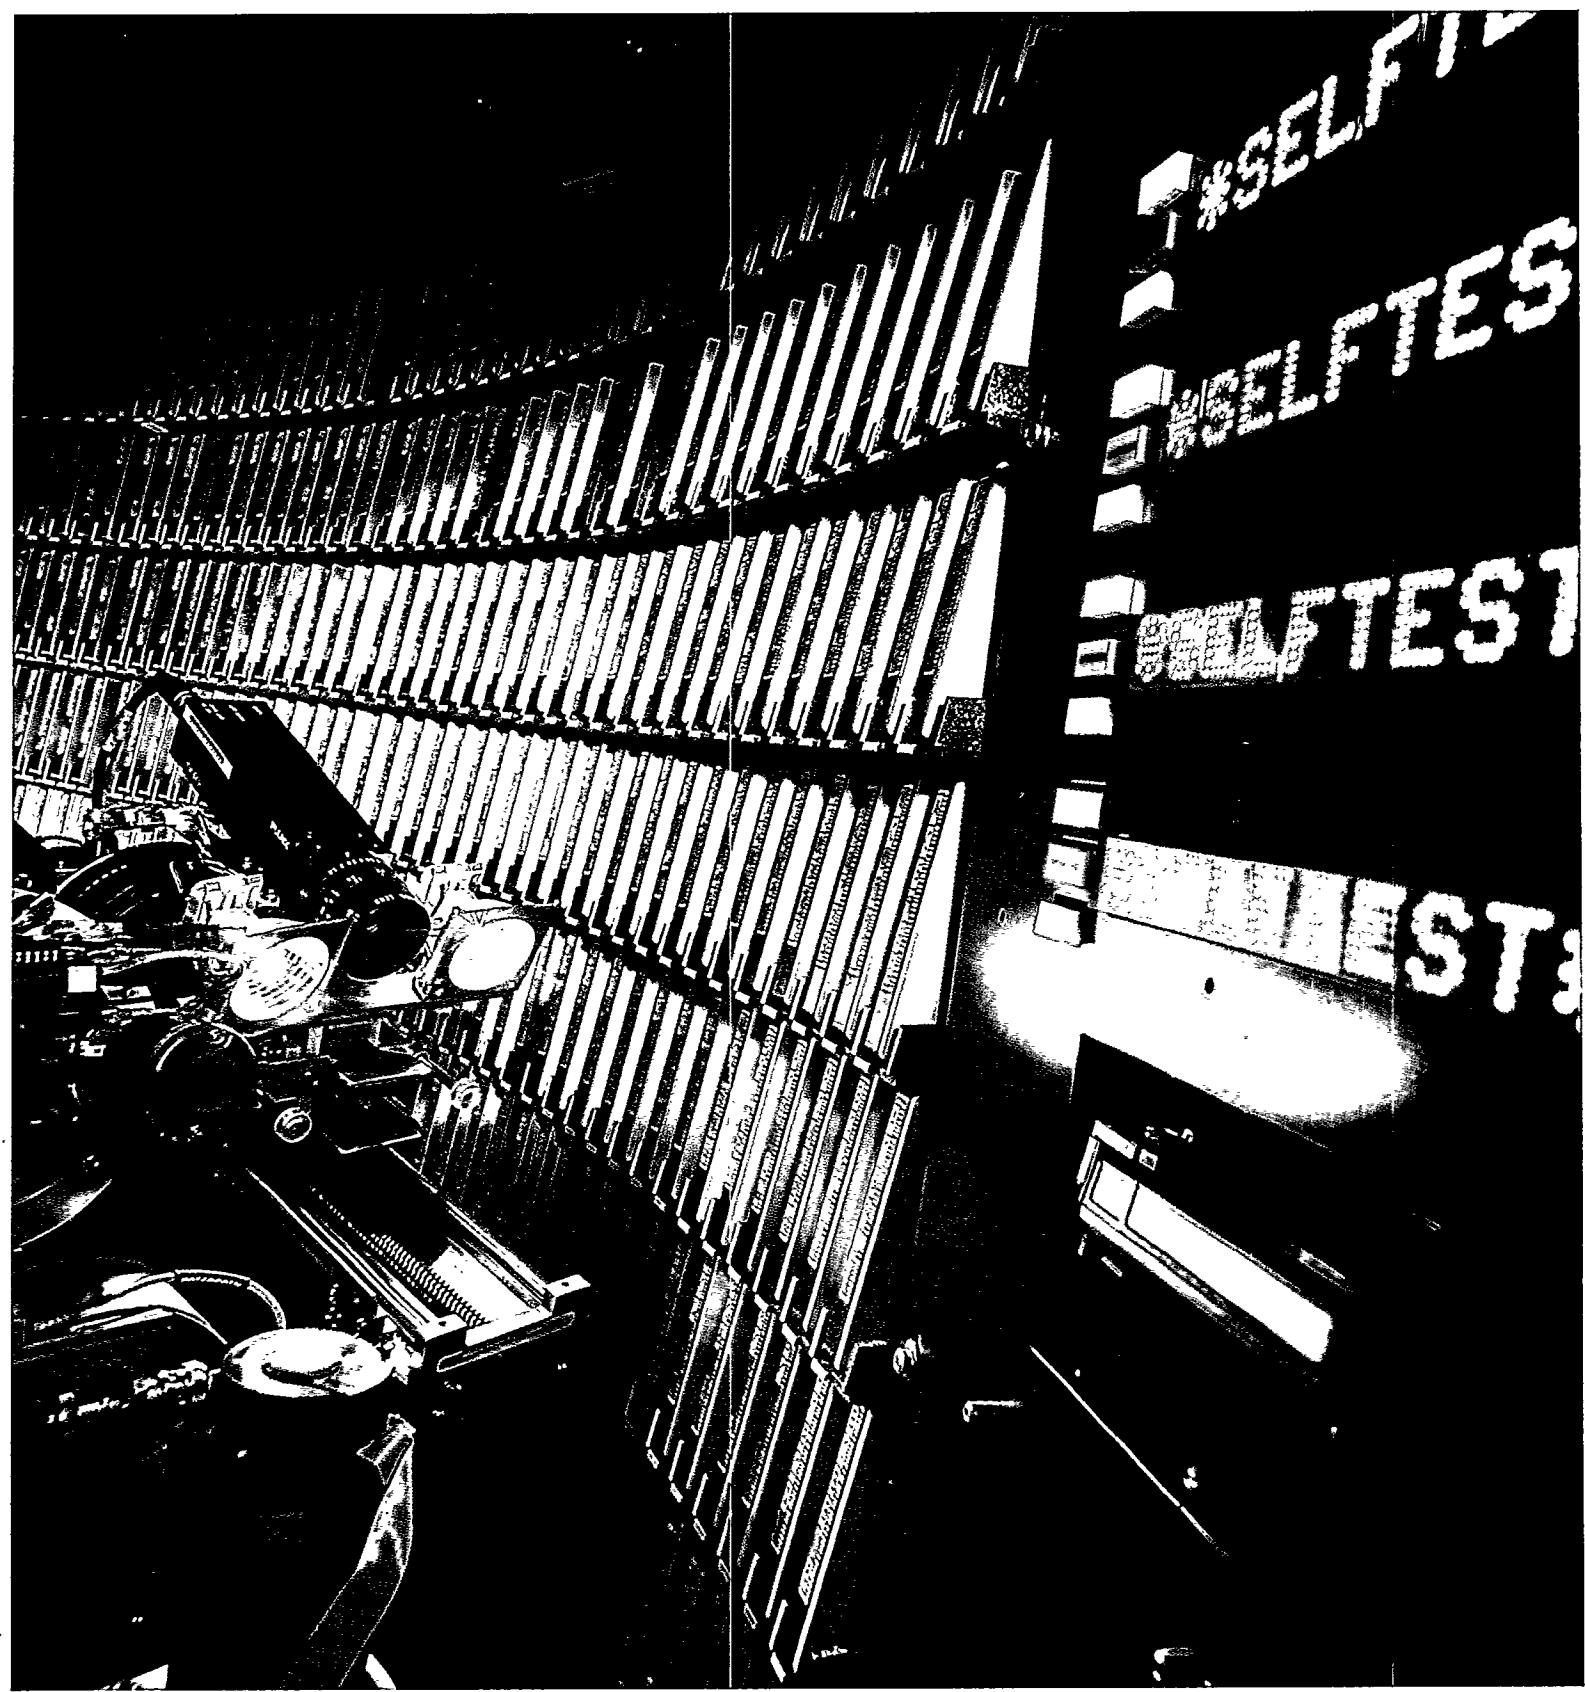
Combustion dynamics for better engines. The drive to build a less polluting and more efficient engine is a prime example of the NERSC user community's commitment to solving complex and practical problems. Industry has already used computer simulations to design better combustion devices. However, existing computer simulation codes cannot faithfully represent the devices' detailed physical processes because of limitations in computational capability. Taking advantage of NERSC's high-performance systems, researchers are developing advanced numerical methodologies for modeling realistic engineering geometries in combustion and other areas of fluid dynamics. The goal is to develop a new generation of computational tools that accurately simulate combustion in applications such as natural gas burners and reciprocating engines. The results of this research are expected to be useful in designing a better engine for a cleaner environment.

Computational chemistry for waste management. The most challenging environmental issues confronting DOE and the energy industry are the safe and cost-effective management of highly radioactive wastes in underground tanks and the remediation of soil and groundwater contamination. A key ingredient to the success and cost-effectiveness of any remediation effort is a fundamental understanding of the chemical properties, such as reactivity and thermodynamics, of the waste components. Such properties are difficult to determine by traditional experimentation due to safety concerns and prohibitive costs. To address particular aspects of these problems, researchers plan to develop new, scalable computational methodologies to simulate these complex chemical systems on high-performance computers. These simulations should yield new information that can be used in assessing the feasibility of various approaches to waste management and remediation.

High-resolution modeling of global climate change. Existing global climate change simulations have demonstrated that higher-resolution atmosphere, ocean, and sea-ice models give more realistic results. But the long computer processing time required for modeling scenarios at high resolution has limited the range of experiments to date. Now, a group of atmospheric scientists and oceanographers plan to use the massively parallel processing capabilities of NERSC's new Cray T3E to couple state-of-the-art atmosphere, ocean, and sea-ice models, producing a new generation of high-resolution global climate simulations. The T3E is expected to compute one calendar year of global climate in only ten hours, allowing the researchers to complete century-long scenarios of climate change. The improved precision of these simulations will enable scientists to produce more realistic representations of various global climate change scenarios. These investigations will serve an important DOE research mission—to assess how various energy use strategies may affect the climate and the environment.

By continually raising the standards of precision, reliability, and time to solution, science-of-scale projects drive the development of improved software and algorithms. In turn, advances in computational power and visualization result in the discovery of new phenomena that enrich those areas of science.

As one of the few facilities that make computational science of scale possible, NERSC is dedicated to helping scientists take advantage of our high-performance computing capabilities, as well as continuing to improve the tools of computational science. In active partnership with the researchers we serve, we hope to participate in major contributions to the future of science.



Interior of NERSC's STK robotic storage silo.

For more information:

Visit NERSC on the World Wide Web:

<http://www.nersc.gov>

Contact NERSC via email:

support@nersc.gov for account, allocation, and password assistance

consult@nersc.gov for consulting

operator@nersc.gov for computer operations

Contact NERSC by phone:

800-66-NERSC (666-3772): U.S. and Canada

510-486-8600: local and international

510-486-4300: fax

Support staff and consultants are available from 8 a.m. to 5 p.m. Pacific time, Monday through Friday. Computer operators are available 24 hours a day.

For network questions, contact ESnet:

trouble@es.net to report problems

info@es.net to request information

800-33-ESnet (333-7638): U.S. and Canada

510-486-7600: local and international

510-486-4300: fax

Published by the Berkeley Lab
Technical and Electronic Information Department

Manager: Ruby Tebelak

Writer/Editor: John Hules

Art Director/Illustrator: Stefan J. Hodes

Illustrator: Bill Bero

Photographers: Roy Kaltschmidt and Don Fike

Print Coordinator: Faye Jobes

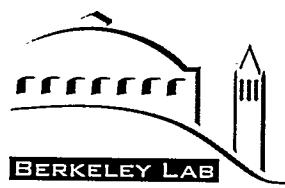
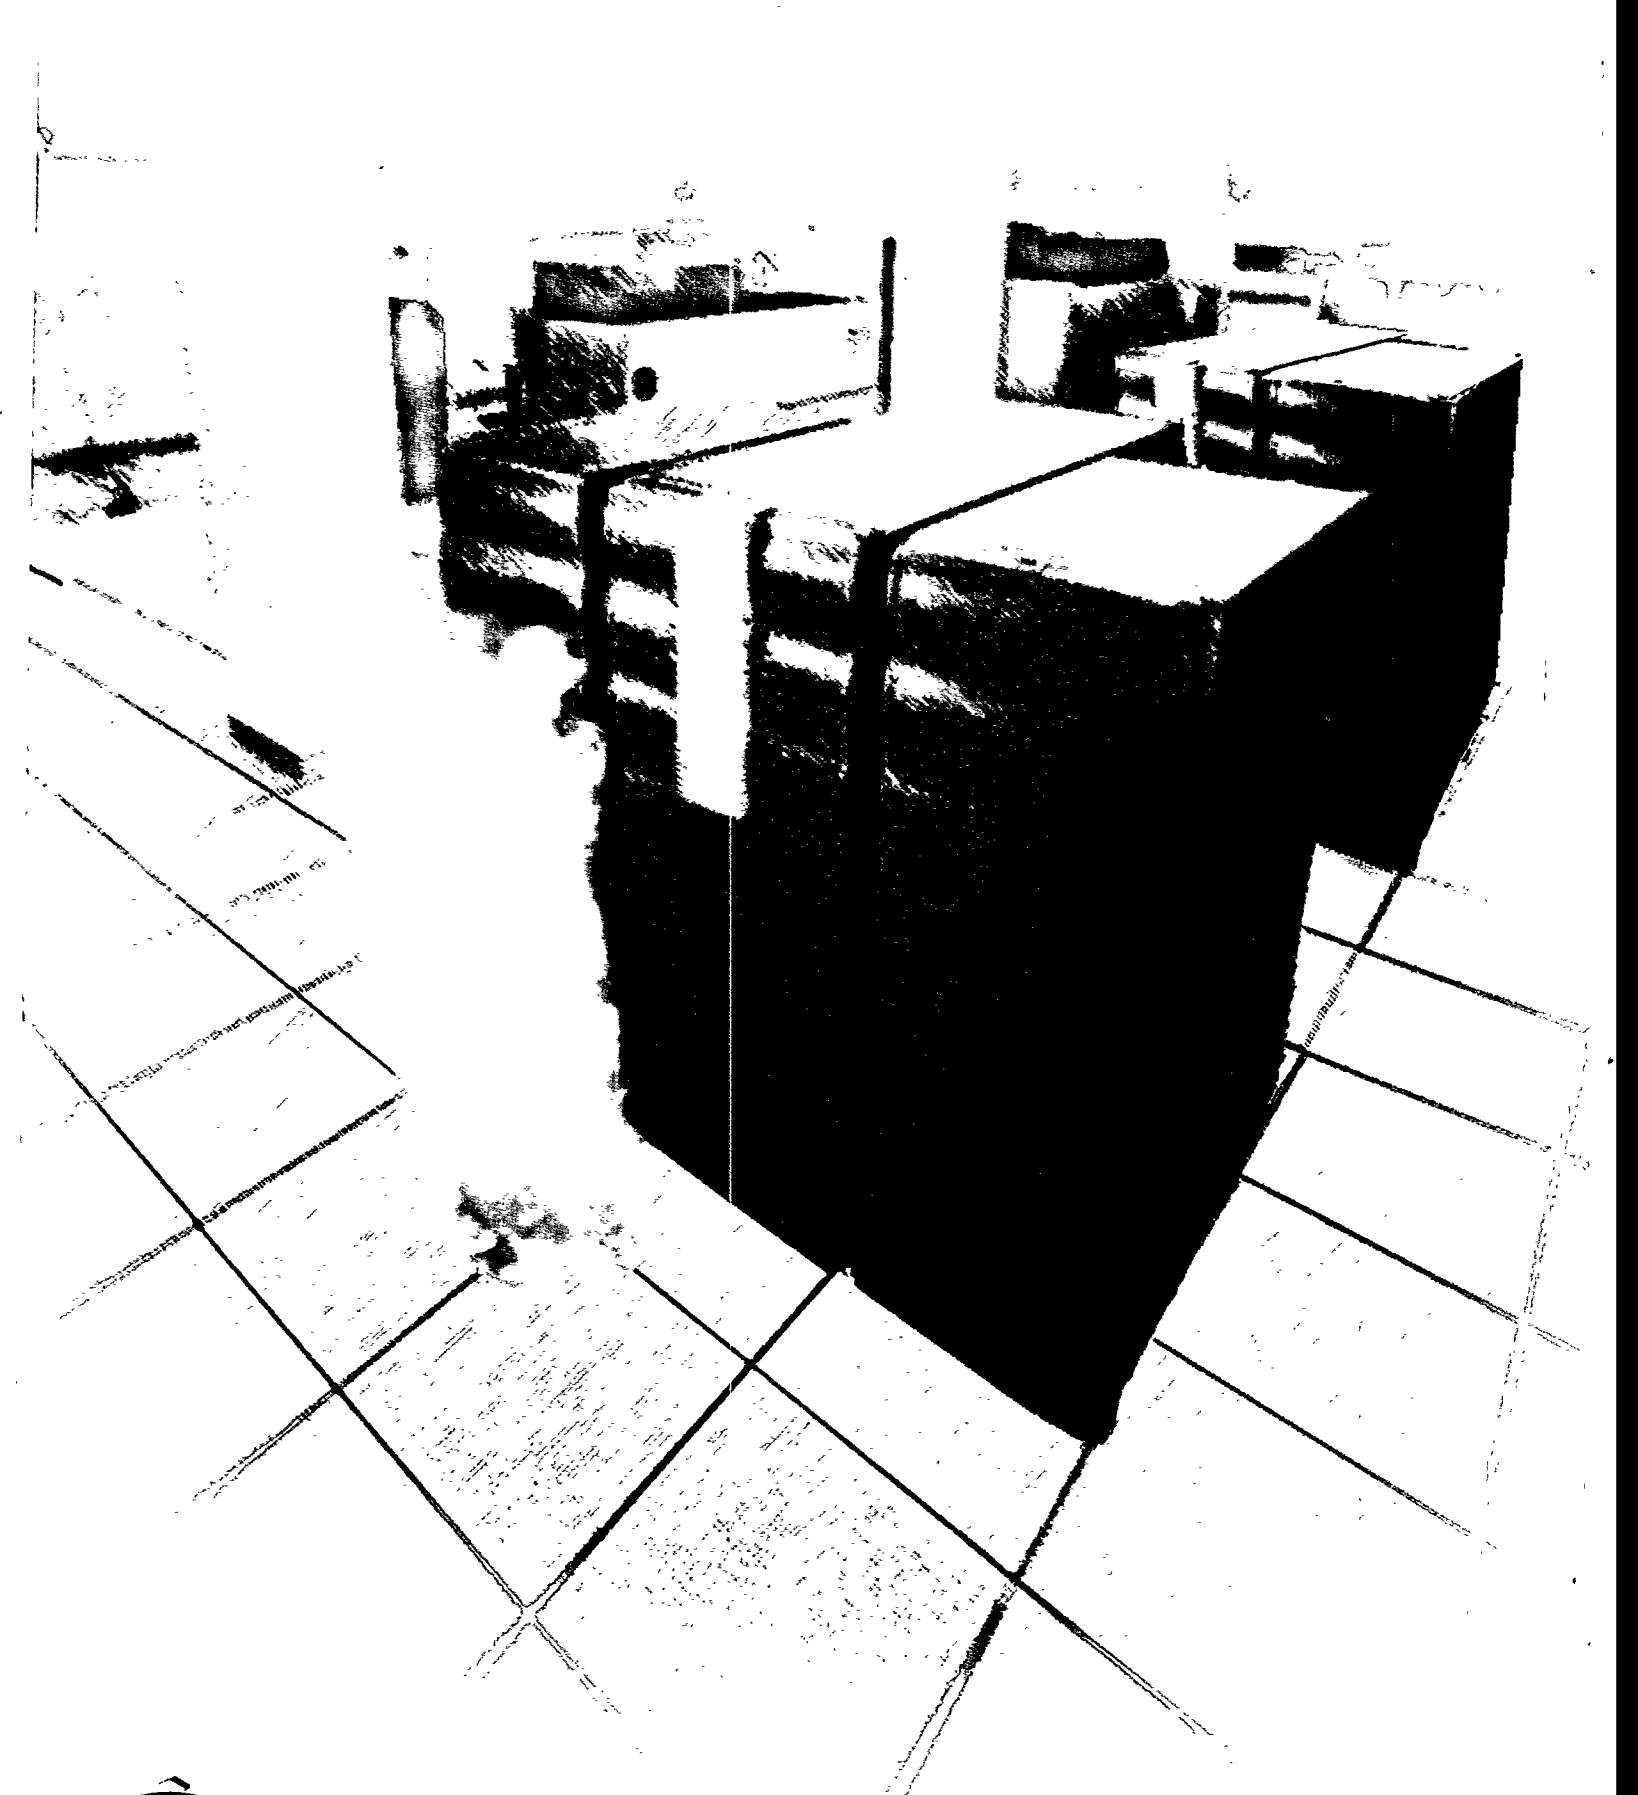
Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

Prepared for the U.S. Department of Energy under Contract DE-AC 03-76SF00098.

PUB-779/10-96



ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY