

REINVENTING THE SUPERCOMPUTER CENTER AT NERSC

Horst D. Simon

The power of high-performance computers grows exponentially with each new generation, but making productive use of that power for scientific research also becomes increasingly difficult. It took the scientific community about seven years to make the transition from single-CPU supercomputers of the Cray-1 variety to vector parallel computers such as the Cray Y-MP and C90, even though the shared-memory architecture of those machines was relatively straightforward. Distributed-memory, massively parallel machines such as the IBM SP, Intel Paragon, and Cray T3D and T3E present burdensome data distribution problems for users to overcome. The next generation of symmetric multiprocessors will add another level of memory hierarchy and new barriers to effective use.

Helping scientists overcome such barriers to using high-performance computers was a key objective in last year's relocation of the National Energy Research Scientific Computing Center to Ernest Orlando Lawrence Berkeley National Laboratory. The relocation makes it easier for NERSC to collaborate closely with scientists at Berkeley Lab and throughout the energy research community, with scientists at the adjacent University of California, Berkeley, campus, and with major computer and data communications companies throughout the San Francisco Bay area.

Founded in 1974 and funded by the US Department of Energy's Office of Energy Research, NERSC was the first unclassified supercomputer center and was the model for those that followed. Its mission is to accelerate the pace of scientific discovery in the energy research community by providing high-performance computing, information, and communications services. Our primary goal is to enable computational science of scale, in which large, interdisciplinary teams of scientists attack fundamental problems in science and engineering that require massive calculations and have broad scientific and economic impacts.

Computing resources and services

In August 1997, NERSC took delivery of a 512-processor Cray T3E-900, with a Linpack rating of 264.8 Gflops. When compared with the June 1997 TOP500 list, this is the largest unclassified supercomputer in the US and the fourth largest in the world. Our six Cray Research computers—the T3E-

900, four J90SEs, and a C90—have a combined computational capacity of 495 Gflops. We are currently developing a single user interface to access all of our Cray systems. Our objective for the end of the decade is to offer computational capacity in the teraflops range and storage capacity in the petabyte range.

The T3E-900 is one of the largest Cray systems available at an openly accessible facility. The system has 512 user-accessible processing elements and 32 system PEs, each with 256 Mbytes of memory, as well as approximately 1.5 Tbytes of disk storage. The J90 cluster has a total of 96 PEs distributed among the four machines, with 32 Gbytes of cluster memory and 752 Gbytes of disk storage. The C90 has 16 PEs and 2 Gbytes of memory. These computing platforms are supported by an array of multilevel mass storage facilities that provide approximately 1 Tbyte of intermediate disk storage and 120 Tbytes of tertiary storage. The computers and storage systems are interconnected by a HIPPI interface with an 800-Mbit data transfer rate.

NERSC consultants work with users to enhance the speed of their applications, respond to questions, and conduct online, video, and classroom training. The center's scientific computing staff collaborates with users to develop new computational approaches and applications. The Energy Research Supercomputer Users' Group provides guidance to NERSC and OER regarding NERSC services and the direction of future development.

Computer science research

To stay at the forefront of HPC and the underlying technologies, NERSC participates in computer science research as both a consumer of state-of-the-art technologies and an activist in communicating the needs of scientists to computer science R&D organizations. In the consumer role, we evaluate and test next-generation software and hardware, such as new mass-storage systems, compilers, languages, and visualization tools, for potential introduction into NERSC systems. Our users' needs for computational robustness, reliability, and efficiency provide a solid testbed for new ideas and designs. In our activist role, we strive to influence the direction of technology development and shape the national technology agenda by focusing on technologies that are two or three generations in the future.

To jump-start the computer science program, NERSC offered joint appointments to several UC Berkeley Computer Science faculty members who are conducting scientific and high-performance computing research. Current joint research efforts include Comps, NOW, and Millennium.

The Cluster of Multiprocessor Systems project, or Comps, is a joint research effort of NERSC, UC Berkeley, Sun Microsystems, and the Berkeley Lab Information and Computing Sciences Division and Materials Sciences Division. The project's goal is to install a networked cluster of shared-memory multiprocessors 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 problems: numerical computation, laboratory experiment control, and high-volume data analysis. If the system meets expectations, it will provide an incrementally scalable common platform that is architecturally similar to the next generation of massively parallel supercomputers.

The Berkeley NOW (Network of Workstations) project is harnessing the power of clustered machines connected via high-speed switched networks. The NOW cluster is a supercomputer in its own right, with 100 Sun Ultrasparcs and 40 Sun Sparcstations totaling about 1 Tbyte, with a Myrinet-based interconnect providing about IBM SP2-level latency and bandwidth. On April 30, 1997, the NOW team achieved over 10 Gflops on the Linpack benchmark.

Millennium, a joint project of UC Berkeley and Intel, will expand the NOW concept into a campus-wide hierarchical cluster architecture with five layers:

- ◆ single-processor workstations,
 - ◆ SMP servers,
 - ◆ clusters of SMPs connected by high-speed switches,
 - ◆ a campus-wide NOW for the most demanding computational modeling, and
 - ◆ a global access layer connecting the other four levels.
- Millennium will make parallel processing available to more Berkeley researchers than ever before. NERSC will receive one of the SMP clusters and will collaborate in the area of high-performance parallel software.

Computational science research

The old model for applications groups at many supercomputer centers has been to assemble a group of experts in the relevant application disciplines who also happened to have relevant parallel programming experience. This model of the early 1990s is no longer efficient, since many users of highly parallel machines have acquired baseline parallel programming knowledge themselves. A group structured this way would no longer add value to a center.

NERSC has evolved the role of the scientific computing staff member away from the application-specific parallel programmer towards the computational technology expert. Our scientific computing staff collaborates with strategic users to port and develop scientific applications, as well as to evalu-

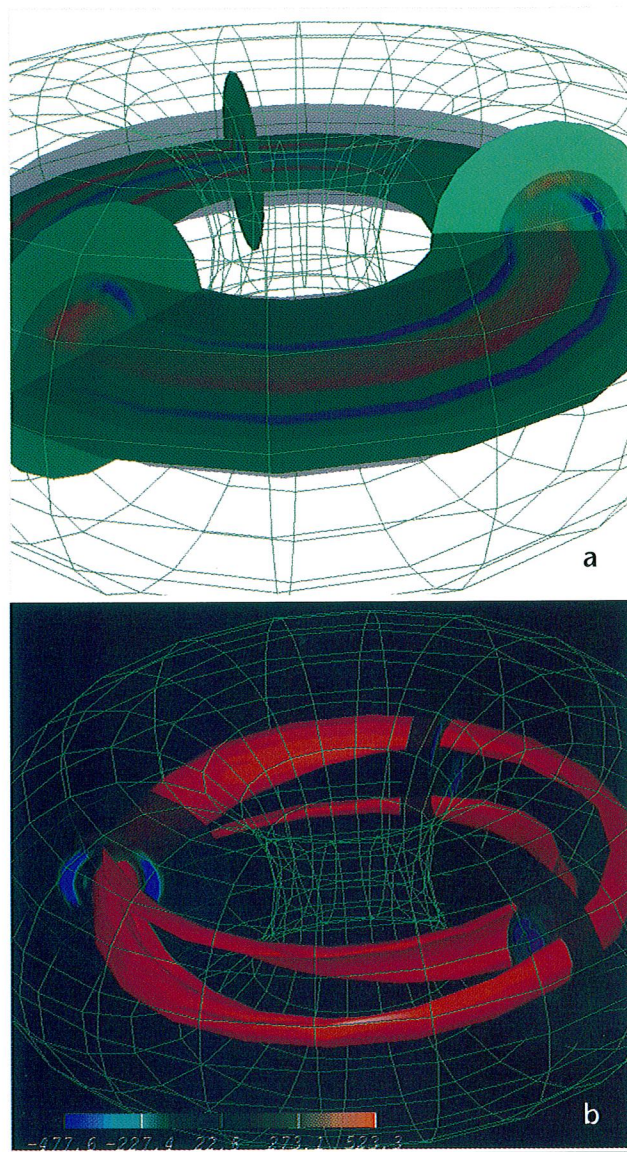


Figure 1. This visualization helped researchers demonstrate inaccuracies in standard computer models of tokamaks. (a) Current density on the midplane of a tokamak and in three cross-sectional slices. (b) An isosurface of the same data set. (Research: Linda Sugiyama, MIT; Wonchull Park, Guoyung Fu, and Elena Belova, Princeton Plasma Physics Laboratory; and Hank Strauss, New York University, Courant Institute of Mathematical Sciences. All visualizations are by the NERSC/Berkeley Lab Visualization Group.)

ate, integrate, and create new hardware and software technologies and new numerical and non-numerical algorithms.

The NERSC/Berkeley Lab Visualization Group creates innovative simulations and images that help scientists understand the results of their research. They have developed virtual reality interfaces that allow remote viewing and control of computer simulations, and they are exploring new methods for visualizing extremely large data sets.




Figure 2. This lattice QCD simulation created on NERSC's Cray T3E shows a single 3D slice of a 4D space. Colored volumes and surfaces depict the probable location densities of quarks. (Research: Gregory Kilcup, The Ohio State University)

Grand Challenge projects

NERSC's approach to computational science research is being applied in seven Grand Challenge projects supported by DoE, in partnership with investigators from national labs and universities across the US.

- ◆ We hope to advance the computational study of materials by providing a better understanding of the relationship between the atomic-scale structure of materials and their macroscopic structural and electrical properties. Future researchers would use this knowledge to design materials on the computer with specific properties needed for engineering applications.

- ◆ Another project expects to improve researchers' ability to understand and predict the experimentally observed turbulent transport of heat and particles out of the core of the tokamak, a toroid-shaped magnetic fusion reactor (see Figure 1). This understanding would be critical to the success of the International Thermonuclear Experimental Reactor.

- ◆ By testing the theory of quantum chromodynamics through simulation, we expect to improve physicists' understanding of the fundamental structure of subatomic particles. Quantum chromodynamics is a theory about how quarks interact with gluons to form hadrons—the strongly interacting particles that we observe in nature. How quarks are confined in space determines the size of protons and ultimately the properties of matter (see Figure 2).

- ◆ The implementation of relativistic quantum-chemical methods on massively parallel computers will, for the first time, enable modeling of heavy-element compounds similar to models currently available for light-element compounds. Radioactive waste involves actinides, whose large atomic number implies that relativistic effects have important chemical consequences.

- ◆ Computational tools developed in another project should enable end-to-end modeling of beam dynamics and engineering design for a new generation of particle accelerators with improved efficiency, performance, and reliability. New accelerator-based technologies are an important component in strategies for treating radioactive waste, disposing of plutonium from nuclear warheads, and producing energy.

- ◆ Another project centers around developing tools for analyzing and managing the massive amounts of data that will be generated by the next generation of physics experiments (see Figure 3). The results of this work can also be applied to other governmental and commercial enterprises faced with massive amounts of data.

- ◆ We are also developing a computational genome analysis tool required in the next phase of the Human Genome Project. The tool will also benefit other human and microorganism genome research as well as related research in structural biology, genetic medicine, biotechnology, pharmaceuticals, and bioremediation.

The Berkeley Lab Computational Sciences Program

With the arrival of NERSC, the lab initiated a broad-based Computational Sciences Program. NERSC is currently collaborating with Berkeley Lab scientists on 21 projects in disciplines ranging from fluid dynamics to biology to geochemistry. For example, one project is developing the first full-scale simulation of the geologic evolution of an actual ore deposit through reactive porous media flow, subject to unique geologic constraints (see Figure 4). The results will assist in future exploration for porphyry copper deposits; management of acid mine waters, leach dumps, and mine tailings; and remediation of environmental contamination.

The goal of another project is to discover enough Type Ia supernovas to measure the rate at which the expansion of the universe is slowing down—the rate that will determine the ultimate fate of the universe. Computational requirements include both real-time image processing and simulation of radiation transport from the distant exploding stars. More accurate simulations will result in better estimates of the rate of deceleration.

A third project is attempting to construct new predictive models of the (macroscopic) properties of real materials starting only with the atomic numbers and masses of constituent atoms. This research is intended to exploit advances in computer hardware and computational techniques, in the fundamental theory of bonding in solids, and in the theory of plastic deformation.

While grounded in the reality of offering production-level scientific computing services, NERSC is committed to constantly rethinking the four elements of scientific computing—mathematical modeling, algorithmic design, software implementation, and system architecture. We intend to achieve significant scientific discoveries through computation. We also intend to make a fundamental change in the way scientists do their work.

By promoting the development of a powerful computational community at Berkeley Lab and throughout the DoE, in which there is rapid sharing of technological developments, we hope to make high-performance computational tools easier to use, bringing computational science of scale from the cutting edge into the mainstream. Our vision is that 10 years from now, giving computational instructions to a high-performance computer will be as easy as writing equations and algorithms on a blackboard. We're doing what we can to make that vision a reality.

For more information see the NERSC home page at <http://www.nersc.gov>. ♦

Acknowledgments

This report is based on an earlier white paper, "The New NERSC Program in FY 1997: An Overview of Current and Proposed Efforts" by H.D. Simon, C.W. McCurdy, W.T. Kramer, and A.X. Merola (<http://www.nersc.gov/research/whitepaper/whitepaper.html>). I thank the NERSC staff and group leaders who contributed many of the ideas expressed here, and John Hules of Berkeley Lab's Technical and Electronic Information Dept. for his help in preparing this report. This work was supported by the US Dept. of Energy Office of Computational and Technology Research under contract no. DE-AC 03-76SF00098.




Figure 3. In the STAR experiment at Brookhaven National Laboratory, a large particle detector will aid in the search for the quark-gluon plasma, a state of matter that has not existed since a few microseconds after the Big Bang. The detector will record the locations of ionization events (blue or red squares in the simulation) as the particles given off by ion collisions pass through the gas filling the chamber (Figure 3a, detail in Figure 3b). The events will then be reconstructed into tracks (red lines) by pattern recognition software. (Research: Herb Ward, University of Texas, Austin; Iwona Sakrejda and Doug Olson, Berkeley Lab)




Figure 4. The vertical plane shows a section of lithologic units (rock types) in a Chilean copper ore deposit. The green grid shows the topographic surface. The white-to-red lines show borings within 50 meters on either side of the section; color indicates ore grade. The model will attempt to replicate 16 million years of geologic evolution resulting in these conditions. (Research: Tim Mote and George Brimhall, University of California, Berkeley; Karsten Pruess, Berkeley Lab)

Horst D. Simon was named director of NERSC in early 1996. Previously he was with the Advanced System Division of Silicon Graphics and with Computer Sciences Corp. at the NAS Division of NASA Ames Research Center. E-mail, simon@nersc.gov.