Advanced Scientific Computing Research

Program Mission

The primary mission of the Advanced Scientific Computing Research (ASCR) program, which is carried out by the Mathematical, Information, and Computational Sciences (MICS) subprogram, is to discover, develop, and deploy the computational and networking tools that enable researchers in the scientific disciplines to analyze, model, simulate, and predict complex physical, chemical, and biological phenomena important to the Department of Energy (DOE). These tools are crucial if DOE researchers in the scientific disciplines are to maintain their world leadership. To accomplish this mission, the program fosters and supports fundamental research in advanced scientific computing – applied mathematics, computer science, and networking – and operates supercomputers, a high performance network, and related facilities. The applied mathematics research efforts provide the fundamental mathematical methods and algorithms needed to model complex physical, chemical, and biological systems. The computer science research efforts enable scientists to efficiently implement these models on the highest performance computers available and to store, manage, analyze, and visualize the massive amounts of data that result. Networking research provides the techniques to link the data producers, e.g., supercomputers and large experimental facilities, with the data consumers, i.e., scientists who need the data.

In fulfilling this primary mission, the ASCR program supports the Office of Science Strategic Plan's goal of providing extraordinary tools for extraordinary science as well as building the foundation for research in support of the other goals of the strategic plan. In the course of accomplishing this mission, the research programs of ASCR have played a critical role in the evolution of high performance computing and networks. The accomplishments of researchers funded by ASCR, which are listed later in this budget, amply demonstrate the world leadership of this program.

In addition to this primary mission, the ASCR program is also responsible for the Laboratory Technology Research subprogram in the Office of Science. The mission of this subprogram is to foster and support high-risk research in the natural sciences and engineering in partnership with the private sector leading to innovative applications relevant to the Nation's energy sector.

The high quality of the research in the entire ASCR program, supporting both of its missions, is continuously evaluated through the use of merit-based peer review and scientific advisory committees.

The research and facilities supported by ASCR are critical to the success of all the missions of the Office of Science (SC) because computational modeling and simulation have become an important contributor to progress in all SC scientific research programs. Modeling and simulation is particularly important for the solution of research problems that are insoluble by traditional theoretical and experimental approaches, hazardous to study in the laboratory, or time-consuming or expensive to solve by traditional means. All of the research programs in the U.S. Department of Energy's Office of Science—in Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, and High Energy and Nuclear Physics—have identified major scientific challenges that can best be addressed through advances in scientific computing.

Advances in computing technologies during the past decade have set the stage for a major step forward in modeling and simulation. By 2005, computers 1,000 times faster than those generally available to the scientific community in 2000, *i.e., terascale* computers (computers that can perform trillions of operations per second, or over one teraflop), will be at hand. However, to deliver on this promise, these increases in "peak" computing power, *i.e.*, the maximum theoretical speed that a computer can attain, must be translated into corresponding increases in the modeling and simulation capabilities of scientific codes. This is a daunting problem that will only be solved by increased investments in *computer software*—the scientific codes for simulating physical phenomena, the mathematical algorithms that underlie these codes, and the computing systems software that enables the use of high-end computer systems. These investments in software research must be made by DOE and other government agencies whose missions depend on high-end computing because technology trends and business forces in the U.S. computer industry have resulted in radically reduced development and production of high-end systems necessary for meeting the most demanding requirements of scientific research. The U.S. computer industry has become focused on the computer hardware and software needs of business applications, which dominate the market, and cannot justify sufficient investments in software research, focused on the computational needs of the scientific community.

An example that shows the dramatic difference between the requirements of scientific computing and business computing is the calculation that won the 1998 Gordon Bell Prize awarded by the Supercomputing 1998 Conference sponsored by the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronic Engineers (IEEE). This calculation resulted from a research project supported by the Mathematical, Information, and Computational Sciences (MICS) subprogram. This calculation, the first to achieve a teraflop on a real scientific application, was undertaken to better understand and predict the magnetic behavior of atoms in a magnetic material. Such calculations are important for understanding the properties of magnetic storage devices, such as computer hard drives, and permanent magnets. The actual calculation that was performed involved a cube of iron approximately two millionths of a centimeter in each dimension containing 1,458 atoms. In order to predict the magnetic behavior of this cube of iron, using the most efficient mathematical tools available today, approximately 30 trillion operations, i.e. additions, subtractions, multiplications and divisions **per atom**, were required. In addition, each atom required 48 megabytes of memory to store critical data. Therefore, the entire calculation required approximately 70 gigabytes of memory and 45 quadrillion (45×10^{15}) operations to complete.

To put this in context, if you could buy a 500Mhz PC with 70 gigabytes of memory (approximately 1000 times the memory on a typical PC), this problem would take approximately 3 years to solve. In addition, while this calculation represents a major scientific advance, to really understand the behavior of materials in real systems, scientists estimate that they need to solve problems with at least 27,000 atoms, almost twenty (20) times larger than this calculation.

In addition to the development of software to address scientific challenges on high performance computers, the Office of Science faces significant challenges in enabling its research community to have effective access to the computational and experimental data resources at its facilities and in enabling geographically-distributed research teams to work effectively together. Many of the DOE facilities produce thousands to millions of gigabytes of data per year that must be managed and analyzed by scientists at universities, government laboratories, and industrial laboratories across the nation, and in some cases across the world. In addition, remote users increasingly need to be able to control experiments from their home institutions. This enables university professors to integrate their research activities at DOE's facilities into their home research and educational activities. In addition, it decreases both the time and cost of performing the research at remote facilities. These geographically-distributed teams of scientists and engineers along with the supporting experimental and computational facilities, tied together with high performance networks are called National Collaboratories. The current technologies that underlie National Collaboratories, many of which were developed by ASCR, only provide support for basic services and must be significantly enhanced to enable collaborative technologies to achieve their full potential as tools for scientists.

In National Collaboratories it is again the scale of the scientific problems that drives SC's need for advances in information technology. For example, the current state of the art in simulations of the earth's climate uses a grid that measures 0.5 degrees of latitude by 0.5 degrees of longitude. A single snapshot of the data from these simulations may contain 300 megabytes of data (over 3,000 times larger than the Netscape home page), and the output of a full simulation may total 75,000 gigabytes (or 75 terabytes). Sharing this data with a distributed scientific team represents a major technical challenge. It should also be noted that there is considerable interest in using grids much finer than 0.5 degrees which corresponds to distances of approximately 50km on a side, (e.g., Washington DC, Baltimore MD, and Frederick MD could be in the same grid element). Again, because of the scale of DOE's requirements, DOE must support significant network and collaboratory research to satisfy its missions. This research complements commercial research investments.

Program Goals

- Maintain world leadership in areas of advanced scientific computing research relevant to the missions of the Department of Energy.
- Integrate the results of advanced scientific computing research into the natural sciences and engineering.
- Provide world-class supercomputer and networking facilities for scientists working on problems that are important to the missions of the Department.
- Integrate and disseminate the results of high-risk research in natural sciences and engineering to the private sector through the Laboratory Technology Research subprogram.

Program Objectives

- *Advance the frontiers of knowledge in advanced scientific computing research.* — Foster research to create new fundamental knowledge in areas of advanced computing research important to the Department, e.g., high performance computing, high speed networks, and software to enable scientists to make effective use of the highest performance computers available and to support National Collaboratories.
- *Apply advanced computing knowledge to complex problems of importance to DOE.* — Promote the transfer of results of advanced scientific computing research to DOE missions in areas such as the improved use of fossil fuels, including understanding the combustion process; the atmospheric and environmental impacts of energy production and use, including global climate modeling and subsurface transport; and future energy sources, including fusion energy, as well as the fundamental understanding of matter and energy.
- *Plan, construct, and operate premier supercomputer and networking facilities.* — Serve researchers at national laboratories, universities, and industry, thus enabling new understanding through analysis, modeling, and simulation of complex natural and engineered systems and effective integration of geographically distributed teams through national collaboratories.
- *Transfer results of fundamental research to the private sector.* — Provide tangible results of research and development activities through cost-shared partnerships with industry.

Evaluation of Objectives

The Advanced Scientific Computing Research (ASCR) program evaluates the progress being made toward achieving its scientific and management objectives in a variety of ways. Regular peer review and merit evaluation is conducted for all activities, except those Congressionally mandated, based on procedures set down in 10 CFR 605 for the extramural grant program and under a similar modified process for the laboratory programs and scientific user facilities. Facilities, including the National Energy Research Scientific Computing Center (NERSC) and ESnet, will be operated within budget and successfully meet user needs and satisfy overall SC program requirements. All new projects are selected by peer review and merit evaluation for their scientific excellence. Beginning in FY 2001, the Advanced Scientific Computing Advisory Committee (ASCAC) will provide advice on subprogram portfolios. As part of these evaluations, the international leadership of ASCR's advanced scientific computing research programs will be assessed. Specific performance measures are included within the detailed program justification narratives as appropriate.

The overall quality of the research in the ASCR program will be judged excellent and relevant by external review by peers, and through various forms of external recognition.

Leadership in key ASCR disciplines that are critical to DOE's mission and the Nation will be measured through external review and other mechanisms.

At least 80% of all new research projects supported by ASCR will be peer reviewed and competitively selected, and will undergo regular peer review merit evaluation.

ASCR will keep within 10%, on average, of cost and schedule milestones for upgrades and construction of the scientific user facilities it manages.

The ASCR scientific user facilities will be operated and maintained so that unscheduled operational downtime will be less than 10%, on average, of scheduled operating time. This includes the National Energy Research Scientific Computing Center and ESnet.

Ensure the safety and health of the workforce and members of the public and the protection of the environment in all its program activities.

Significant Accomplishments and Program Shifts

The ASCR program builds on decades of leadership in high performance computing and has many pioneering accomplishments. Building on this long history, principal investigators of the ASCR program have received recognition through numerous prizes awards and honors. A sample of pioneering results, recent accomplishments and current awards is given below.

SCIENCE ACCOMPLISHMENTS

Mathematical, Information and Computational Sciences

- *A Fundamental Problem of Quantum Physics Solved.* For over half a century, theorists tried and failed to provide a complete solution to scattering in a quantum system of three charged particles, one of the most fundamental phenomena in atomic physics. Such interactions are everywhere;

ionization by electron impact, for example, is responsible for the glow of fluorescent lights and for the ion beams that engrave silicon chips. Collaborators at Lawrence Berkeley National Laboratory (LBNL), Lawrence Livermore National Laboratory (LLNL), and the University of California at Davis used NERSC's Cray T3E supercomputer to obtain a complete solution of the ionization of a hydrogen atom by collision with an electron, the simplest nontrivial example of the problem's last unsolved component. Their breakthrough employs a mathematical transformation of the Schroedinger wave equation that makes it possible to treat the outgoing particles as if their wave functions simply vanish at large distances from the nucleus instead of extending to infinity if treated conventionally.

- *Biggest Dataset Ever from Cosmic Microwave Background*. Since ancient times, the geometry of the universe has been a topic of speculation and inquiry. An international team of scientists reported in Nature that the universe is, as Euclid thought, flat. They obtained their results by combining cosmic microwave radiation data collected during an Antarctic balloon flight and extensive analysis, amounting to 50,000 hours of computer time, using NERSC's Cray T3E supercomputer and software written at NERSC. Called BOOMERANG for "Balloon Observations Of Millimetric Extragalactic Radiation And Geophysics," the collaboration includes over two dozen researchers from seven countries. Supercomputers at NERSC, along with software developed there, were crucial to extracting fundamental cosmological parameters from the data, the largest and most precise set of cosmic microwave background (CMB) data yet collected. From the dataset, the BOOMERANG team was able to make the most detailed map of the CMB's temperature fluctuations ever seen. From a CMB map, cosmologists derive a "power spectrum," a curve that registers the strength of these fluctuations on different angular scales, and which contains information on such characteristics of the universe as its geometry and how much matter and energy it contains. The power spectrum derived from the BOOMERANG Antarctic flight data is detailed enough to allow the determination of fundamental cosmic parameters to within a few percent, indicating that the geometry of the universe is flat and that the universe will expand forever. The calculation required would have taken almost six years to complete if run on a desktop personal computer. On the NERSC Cray T3E, processing time over the life of the project totaled less than 3 weeks.

- *Access Grid: A New Collaborative Environment*. Argonne computer scientists have developed a new collaborative environment, the Access Grid, that can be used not just for science and engineering applications but for lectures, seminars, tutorials, and training. An Access Grid comprises a large-format multimedia display system about 18 feet by 6 feet; several video streams (for example, a wide audience shot, a close-up shot of the presenter, a wide area shot of the display screen, and a roving camera) projected onto the display; numerous microphones and speakers; and several computers for audio capture, video capture, control, and display. Each Access Grid "node," or site, involves 3-20 people, providing a compelling new environment far beyond desktop-to-desktop collaborative systems.

- *Parallel Computational Oil Reservoir Simulator*. To meet the Nation's energy needs, the United States oil and gas industry must continue to advance the technology used to extract oil and gas from both new and old fields. Until recently, most drilling and recovery activities were based on past practices that often lacked a sound scientific basis. Computer scientists at Argonne National Laboratory, in collaboration with petroleum engineers at the University of Texas at Austin, have recently developed a software package capable of simulating the flow of oil and gas in reservoirs. These codes, which are based on software tools designed at Argonne, are able to run on a variety of computer platforms, including massively parallel systems with hundreds and even thousands of processors. The software codes will enable the oil and gas industry to lower exploration and drilling costs and enhance the yield of oil from new and old fields alike.

- *Cracking the Chemistry of Nuclear Waste Characterization and Processing*. Many DOE research areas require accurate and efficient calculations of molecular electronic structure (e.g., catalysis, combustion, and chemistry in the environment or atmosphere). However, the single most challenging environmental issue confronting the DOE is the safe and cost-effective management of highly radioactive mixed wastes generated by four decades of nuclear weapons production. Modeling the chemistry of these heavy elements (e.g., the actinides which include uranium and plutonium) present an enormous challenge, since relativistic effects must be included, and these greatly increase the cost and complexity of the calculations. The MICS project, Computational Chemistry for Nuclear Waste Characterization and Processing: Relativistic Quantum Chemistry of Actinides, has developed major new capabilities for understanding actinide processes. The seven institution team created new chemistry codes for massively parallel computers (Pacific Northwest, Argonne, and Lawrence Berkeley National Laboratories; Stevens Institute; Eloret; Ohio State University; and Syracuse University) and formed liaisons with experimental research at four DOE labs (Pacific Northwest, Argonne, Lawrence Berkeley, and Los Alamos). This project broke new ground in a number of areas, implementing several models that include relativistic effects. In addition, new models were developed, that provide a very significant improvement in accuracy over previous approaches. This is the first time that any of these methods are available on massively parallel computers.

- *High Performance Algorithms for Scientific Simulation*. Many problems in science and engineering involve the complex interplay of forces and effects on different time and length scales. Two significant examples of such nonlinear multi-time, multi-scale phenomena are the interaction of the atmosphere and the oceans in the creation of the global climate and the burning of fossil fuels in engines and other devices. The size and complexity of such problems require the development of fast and efficient algorithms and software that can take advantage of the resolution power of today's massively parallel computing platforms. Applied mathematicians at the Lawrence Berkeley National Laboratory, working in collaboration with applied mathematicians at the Lawrence Livermore National Laboratory and New York University, have developed adaptive mesh refinement algorithms capable of automatically redistributing grid points in computational regions where significant physics is occurring over small time scales. At finer and finer length scales the continuous flow solver is replaced by a particle method such as Monte Carlo, thus allowing the researchers to accurately resolve phenomena over a broad range of length and time scales. Applications of this work to the Accelerated Strategic Computing Initiative (ASCI) and Office of Science problems are many, although the primary focus of the research is the accurate simulation of diesel combustion in realistic, three-dimensional geometries. Laboratory and academic researchers are working on this project closely with engineers from Caterpillar and Cummins.

- *Solving Optimization Problems Easily and Inexpensively*. Optimization applications range from designing circuits, to estimating the value at risk of financial institutions, to determining routing patterns on the Internet, to finding energy functions for molecular structures. Argonne researchers (together with Northwestern University) completed a project to attack such problems successfully. The project involves development of a novel environment, called the Network-Enabled Optimization System (NEOS). NEOS allows users to solve optimization problems over the Internet with state-of-the-art software without downloading and linking code. Given the definition of a nonlinear optimization problem, NEOS determines an appropriate solver, uses tools to compute derivatives and sparsity patterns, compiles all subroutines, links with the appropriate libraries, and executes the solver. The user is given a solution in a matter of hours instead of days or weeks. The NEOS project has recently gained considerable visibility with the release of a new portable version that can be run on various computers, Web servers, and email servers. Over the past year, the number of users has risen to an average of 2,600 problem submissions per month. Moreover, NEOS is now being used as an educational tool at universities worldwide. Students and professors alike find it easy to use and have rated it as one of the more efficient ways to access state-of-the-art optimization software.

- *Bringing Powerful Scientific Visualizations to the Desktop*. As part of their work in DOE's Combustion Corridor research project, visualization researchers at Lawrence Berkeley National Laboratory created a new application that will enable distributed scientific visualization of large data volumes on remote workstations. Such visualizations typically represent complex, three-dimensional scientific problems varying over time, such as how two gases mix in a turbulent environment. To visualize these models, researchers previously required access to very powerful computers — moving such large files onto local workstations is either impossible or impractical. The fundamental idea behind Image-Based-Rendering-Assisted Volume Rendering (IBRAVR) is that large data are partially pre-rendered on a large "computational engine" close to the data, with the final image rendering performed on a workstation. The benefits and advantages of the IBRAVR technique include using a new type of visualization and rendering technology which "decomposes" nicely for parallel processing and sharing the workload between a remote multiprocessor machine and a local workstation.

- *Parallel Multigrid Methods*. LLNL researchers, in close collaboration with university colleagues, developed parallel multigrid methods for the solution of large, sparse systems of linear equations. The solution of these large systems (having upwards of one billion unknowns) is often the key computational bottleneck in scientific and engineering application codes. Multigrid is a so-called scalable algorithm, and LLNL's work in this area is already reducing overall simulation times of important LLNL application codes by up to ten-fold (linear solve time is up to 100 times faster than previous solution methods). The largest problem solved by LLNL researchers was a one billion unknown anisotropic diffusion problem, which took just 54 seconds on 3150 processors of the Advanced Strategic Computing Initiative (ASCI) Red platform. These methods are being developed for a variety of different applications, including inertial confinement fusion, structural dynamics, and flow in porous media. The problems of interest are defined on a variety of grids, including structured grids, block-structured grids, adaptive mesh refinement composite grids, overset grids, and unstructured grids. Both geometric and algebraic multigrid techniques are being investigated.

- *Adaptive Laser-Plasma Simulation*. LLNL computational scientists have developed new numerical algorithms and software to enable the use of parallel adaptive mesh refinement in the simulation of laser plasma interaction. The ability to predict and control the interaction of intense laser light with plasmas is critical in the design of laser-driven fusion energy experiments such as those to be conducted at the National Ignition Facility currently under construction at LLNL. These new algorithms solve a system of plasma fluid equations coupled with a laser light propagation model, incorporating parallel adaptive mesh refinement to use fine meshes only where they are needed. This approach has been shown to reduce overall computational requirements ten-fold on certain problems. These novel algorithms have been implemented in a research code called ALPS (Adaptive Laser Plasma Simulator). ALPS solves problems in two or three spatial dimensions and runs on a variety of high performance computing platforms, including massively parallel computers. Substantial speedups relative to conventional uniform grid algorithms has been achieved on problems involving single mode beams and beams smoothed with random phase plates.

FACILITY ACCOMPLISHMENTS

Mathematical, Information and Computational Sciences

- *NERSC Accepts New IBM Supercomputer*. DOE's National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory accepted the first phase of its new IBM RS/6000 SP system. As part of the acceptance testing, a group of computational scientists at national laboratories and universities were given early access to the machine to thoroughly test the entire system. Those researchers noted the high performance and ability to scale problems on the IBM SP, which is providing useful scientific results in such areas as climate modeling, materials science, and physics research. The first phase of NERSC's IBM system consists of an RS/6000 SP with 304 POWER3 SMP nodes, each with two processors per node. In all, Phase I has 512 processors for computing, 256 gigabytes of memory and 10 terabytes of disk storage for scientific computing. The system has a peak performance of 410 gigaflops, or 410 billion calculations per second. Phase II, delivered in December 2000, will have 2,048 processors dedicated to large-scale scientific computing and another 384 processors devoted to system tasks. The system will have a peak performance capability of more than 3 teraflops, or 3 trillion calculations per second.
- *ESnet Capability Expanded*. DOE's Energy Sciences Network, known as ESnet, selected Qwest Communications to provide advanced data communications services for the scientific research network. The awarded contract is up to seven years in length, including option years, and is valued at approximately $50,000,000. Under the contract, Qwest will provide a terabit (one million megabits/second) network by the year 2005, offering 500 times the highest speed available on the highest speed networks today. This network will also offer a staggering amount of capacity — more than that of all today's carriers combined.
- *Scientists Install New Testbed for Open Source Software*. A 512-CPU Linux cluster was installed at Argonne's Mathematics and Computer Science Division. The cluster provides a flexible development environment for scalable open source software in four key categories: cluster management, high-performance systems software (file systems, schedulers and libraries), scientific visualization, and distributed computing. Its modular design makes it easily reconfigurable for systems management experiments, and its availability for testing open-source code and algorithms ensures broad use by researchers both within the Laboratory and externally.

AWARDS

Mathematical, Information and Computational Sciences

- *Supercomputing (SC'99) Awards*. Special awards were given to scientists from Argonne National Laboratory and the University of Chicago for their achievements in simulating incompressible flows and another special award was given to a team of scientists from NASA and DOE laboratories for their achievements in fluid dynamics simulations.
- *The Maxwell Prize*. A new international prize in applied mathematics, the Maxwell Prize, has been awarded for MICS supported research at the Applied Mathematics Research program at Lawrence Berkeley National Laboratory. The research involved analysis of problems dominated by complexity, such as turbulence, failure and cracks in solids, flow in porous and inhomogeneous media, and combustion. The work on crack formation provided some of the basic tools used today in failure analysis, especially failure due to fatigue. The prize was awarded by the Committee for International Conferences on Industrial and Applied Mathematics (CICIAM) that is made up of the major applied mathematics organizations in 14 countries.

- *Herbrand 2000 Award*. A senior scientist at Argonne National Laboratory was named recipient of the Herbrand Award for the year 2000. The award is given by the Conference on Automated Deduction Inc. to honor exceptional contributions to the field of automated reasoning. Automated reasoning involves the use of powerful computers to solve the logical (as distinct from the numerical) aspects of problems. Applications range from the design and validation of electronic circuits to assisting in research in mathematics and logic. Only six Herbrand awards have been given in the past decade.
- *An Algorithm for the Ages*. Among the top 10 "Algorithms of the Century" announced by Computing in Science and Engineering magazine is the integer-relation algorithm dubbed PSLQ. PSLQ was discovered by a mathematician at the University of Maryland and implemented in practical computer software by a LBNL researcher. PSLQ has unearthed many surprising relations in mathematics and physics, although its most startling result may well be a simple formula for calculating any binary digit of pi without calculating the digits preceding it. Before PSLQ, mathematicians had not thought that such a digit-extraction algorithm for pi was possible. Using the remarkably simple formula, even a personal computer can calculate pi's millionth binary digit in about 60 seconds. As a tool of experimental mathematics, PSLQ's purpose is to discover new mathematical relations among numbers — for example, constants that occur in various groups of mathematical formulas. Very high precision arithmetic is needed by PSLQ, or else nonsense results are obtained.
- *Global Arrays Power Award Winning Applications*. Supported by the DOE Advanced Computational Testing and Simulation program, Pacific Northwest National Laboratory's Global Array (GA) toolkit provides a unique and portable "shared memory" interface that enables codes to efficiently access massive data structures distributed across thousands of processors. GA also provides easy-to-use interfaces to key numerical libraries, and interoperability with message passing standards like MPI, forming a powerful and complete environment for parallel software development. GA is a critical component of a series of award-winning applications, including Molecular Science Software Suite (MS^3), which was awarded both an R&D-100 award in 1999 and a Federal Laboratory Consortium Technology Transfer award in 2000; and COLUMBUS, featured in the SC'98 Conference Best Overall Paper.
- *Supercomputing (SC 2000) Awards*. Scientists at LBNL won the "Fastest and Fattest" category in the High-Bandwidth Applications Competition for achieving a peak performance of over 1.48 gigabits/second in a prototype application for the visualization of terascale datasets. A team that included members from ANL, LLNL and LBNL won the "Hottest Infrastructure" category for demonstrating secure, high-performance data transfer and replication for large-scale climate modeling data sets. These awards recognize DOE's unique requirements for network services that far exceed usual commercial offerings and DOE's success in solving the significant research problems in networking and applications design to satisfy critical mission requirements.
- *1999 R&D 100 Awards*. Awards were made: to PNNL for the Molecular Science Software Suite (MS3) as the first package to offer advanced computational chemistry components optimized for high-performance massively parallel computers; to ORNL, the University of California at San Diego and the University of Tennessee for Netsolve 1.2, which unifies disparate high-performance computers and software libraries and delivers them as a powerful, easy to use computational service; and to the University of Tennessee and ORNL for ATLAS which automatically generates optimized basic linear algebra subprograms resulting in significant improvements in performance and portability.

The Laboratory Technology Research (LTR) subprogram received one R&D-100 Awards in 2000 for the following research:

- *Carbon Monoxide Monitoring Device.* Lawrence Berkeley National Laboratory (LBNL), in collaboration with Quantum Group, has developed an inexpensive, passive carbon monoxide (CO) occupational dosimeter, which has achieved control of sensor reversibility, humidity dependence, and sensor variability. CO is one of the most deadly environmental pollutants encountered in indoor and outdoor occupational settings. An inexpensive but sensitive and accurate CO monitoring device should help to identify and mitigate high CO environments.

The LTR subprogram received five Federal Laboratory Consortium (FLC) Awards for Excellence in Technology Transfer in 2000 for the following research:

- *Quick, Cost-effective Filler for Potholes.* Argonne National Laboratory, in collaboration with Rostoker, Inc., has developed a versatile, unique, room-temperature technology that uses a nontoxic, nonflammable binder to form products that are impermeable to groundwater and have twice the strength of cement. The technology is filling an important need in the road repair industry as a quick, cost-effective filler for potholes. It also benefits the environment by reducing landfill requirements by encapsulating waste more efficiently than traditional methods.

- *Innovative Purification Method.* LBNL, in collaboration with WaterHealth International, has developed a device that uses readily available, energy-efficient, low-maintenance technologies and materials to disinfect water. For less than 2 cents per metric ton of water, the ultraviolet light from a single 40-watt compact UV bulb disrupts the DNA of contaminating bacteria and viruses within 12 seconds. As of late 1999, about 100,000 people used daily drinking water disinfected by the LBNL device in several developing countries. The children in these countries will now have better survival rates, and more of them will grow up without being stunted from repeated diarrheal episodes.

- *Proton Therapy Used to Treat Cancer.* LBNL, in collaboration with General Atomics, has developed an accelerator-based proton therapy center at Massachusetts General Hospital in Boston to safely destroy cancerous tumors. The recent addition of proton therapy to the arsenal of cancer treatments promises new hope in the battle against the disease. In the U.S. each year, about 375,000 cancer patients use conventional radiation therapy for curative treatments. Of these, about 130,000 patients will benefit if treated with 3D conformal therapy, which is best delivered using proton beams.

- *Superplastic Forming Process for Automotive Components.* Pacific Northwest National Laboratory (PNNL), in collaboration with General Motors and MARC Analysis Research, has significantly reduced the technical and economic impediments to using superplastic forming processes for automotive and other manufacturing. This technology is a metal-forming process that can reduce the weight and cost of manufactured devices such as automotive structural components. Cars with lightweight automotive components that are strong enough to meet design requirements will be safer, more fuel efficient, and have lower emissions.

- *Software to Solve Complex Environmental Problems.* PNNL, in collaboration with DuPont and Amoco, has developed the first general-purpose software that provides chemists with access to high-performance, massively-parallel computers for a wide range of applications. The software is now used by more than 37 universities and supercomputer centers, 14 national laboratories or Federal agencies, and 15 industries. The software will enable the scientific community to quickly and cost-effectively solve complex environmental problems in the atmosphere, aquatic systems, and the subterranean environment.

In addition to the R&D-100 and FLC awards, two scientists supported by the LTR subprogram were recipients of the following distinguished awards in 2000:

- The 1999 Tennessee Industrial Scientist of the Year Award to a group leader in the Metals and Ceramics Division of Oak Ridge National Laboratory.

- The 1999 Battelle Inventor of the Year Award to an analytical biochemist at PNNL.

In FY 2000, the Laboratory Technology Research subprogram initiated a portfolio of Rapid Access Projects that addresses research problems of small businesses by utilizing the unique facilities of the Office of Science laboratories. These projects were selected on the basis of scientific/technical merit and commercial potential, using competitive external peer review.

PROGRAM SHIFTS

The FY 2002 ASCR budget continues the research portfolio enhancements, initiated in FY 2001, to create the next generation of high performance computing and communications tools to support the missions of the Office of Science and the Department of Energy in the next century.

A Federally-chartered advisory committee was established for the Advanced Scientific Computing Research program in FY 2000 and is charged with providing advice on: promising future directions for advanced scientific computing research; strategies to couple advanced scientific computing research to other disciplines; and the relationship of the DOE program to other Federal investments in information technology research. This advisory committee will play a key role in evaluating future planning efforts for research and facilities.

Taking the Next Steps in Scientific Computing, Networking, and Collaboration: Scientific Discovery through Advanced Computing.

In FY 2002, the MICS subprogram of ASCR will continue its components of the collaborative program across the Office of Science to produce the scientific computing, networking and collaboration tools that DOE researchers will require to address the scientific challenges of the next decade. This program was described in the report to Congress entitled, “Scientific Discovery through Advanced Computing,” (SciDAC). These enhancements build on the historic strength of the Department of Energy’s Office of Science in computational science, computer science, applied mathematics, and high-performance computing and in the design, development, and management of large scientific and engineering projects and scientific user facilities. They also take full advantage of the dramatic increases in computing capabilities being fostered by the *Accelerated Strategic Computing Initiative* (*ASCI*) in the National Nuclear Security Agency (NNSA).

The ASCR contributions to this effort, which are described in detail in the MICS subprogram description, are briefly described below:

1. Continue the competitively selected partnerships focused on discovering, developing, and deploying to scientists key enabling technologies that were initiated in FY 2001. These partnerships, which are

called Integrated Software Infrastructure Centers, play a critical role in providing the software infrastructure that will be used by the SciDAC applications research teams.

2. Continue the integrated program of fundamental research in networking and collaboratory tools, partnerships with key scientific disciplines, and advanced network testbeds to provide the electronic collaboration tools that SciDAC teams, and DOE's next generation of experimental facilities, need to accomplish their goals. These projects were competitively selected in FY 2001.

Interagency Environment

The research and development activities supported by the MICS subprogram are coordinated with other Federal efforts through the Interagency Principals Group, chaired by the President's Science Advisor, and the Information Technology Working Group (ITWG). The ITWG represents the evolution of an interagency coordination process that began under the 1991 High Performance Computing Act as the High Performance Computing, Communications, and Information Technology (HPCCIT) Committee. DOE has been a key participant in these coordination bodies from the outset and will continue to coordinate its R&D efforts closely through this process.

In FY 1999, the President's Information Technology Advisory Committee (PITAC) recommended significant increases in support of basic research in: Software; Scalable Information Infrastructure; High End Computing; Socio-Economic and Workforce Impacts; support of research projects of broader scope; and visionary "Expeditions to the 21st Century" to explore new ways that computing could benefit our world.

Although the focus of the enhanced DOE program is on solving mission critical problems in scientific computing, this program will make significant contributions to the Nation's Information Technology Basic Research effort just as previous DOE mission-related research efforts have led to DOE's leadership in this field. In particular, the enhanced MICS subprogram will place emphasis on software research to improve the performance of high-end computing as well as research on the human-computer interface and on information management and analysis techniques needed to enable scientists to manage, analyze and visualize data from their simulations, and develop effective collaboratories. DOE's program, that focuses on the information technology research needed to enable scientists to solve problems in their disciplines, differs from the National Science Foundation's portfolio, that covers all of information technology. In addition, DOE's focus on large teams with responsibility for delivering software that other researchers can rely on differs from NSF's single investigator focus.

Scientific Facilities Utilization

The ASCR program request includes $28,244,000 in FY 2002 to support the National Energy Research Scientific Computing (NERSC) Center, which is ASCR's component of the SC-wide Scientific Facilities Initiative that started in FY 1996. This investment will provide computer resources for about 2,000 scientists in universities, federal agencies, and U.S. companies. It will also leverage both federally and privately sponsored research, consistent with the Administration's strategy for enhancing the U.S. National science investment. The proposed funding will enable NERSC to maintain its role as one of the Nation's largest, premier unclassified computing centers, which is a critical element in the success of many SC research programs. Research communities that benefit from NERSC include structural biology; superconductor technology; medical research and technology development; materials, chemical, and plasma sciences; high energy and nuclear physics; and environmental and atmospheric research.

Workforce Development

The R&D Workforce Development mission is to ensure the supply of computational and computer science and Ph.D. level scientists for the Department and the Nation through graduate student and post doctoral research support. In FY 2002, this program will support approximately 800 graduate students and post doctoral investigators, of which 500 will be supported at Office of Science user facilities.

ASCR will continue the Computational Science Graduate Fellowship Program with the successful appointment of 20 new students to support the next generation of leaders in computational science.

Funding Profile

(dollars in thousands)

	FY 2000 Comparable Appropriation	FY 2001 Original Appropriation	FY 2001 Adjustments	FY 2001 Comparable Appropriation	FY 2002 Request
Advanced Scientific Computing Research					
Mathematical, Information, and Computational Sciences	113,914	160,000	-3,830	156,170	156,170
Laboratory Technology Research......	8,424	10,000	-420	9,580	6,880
Subtotal, Advanced Scientific Computing Research........................	122,338	170,000	-4,250	165,750	163,050
General Reduction	0	-1,732	1,732	0	0
General Reduction for Safeguards and Security.................................	0	-2,153	2,153	0	0
Omnibus Rescission.........................	0	-365	365	0	0
Subtotal, Advanced Scientific Computing Research	122,338 [a] [b]	165,750	0	165,750	163,050
Pending Budget Amendment	0	0	0	0	2,700 [c]
Total, Advanced Scientific Computing Research	122,338	165,750	0	165,750	165,750

Public Law Authorization:

Public Law 95-91, "Department of Energy Organization Act"

Public Law 103-62, "Government Performance and Results Act of 1993"

[a] Excludes $3,041,000 which has been transferred to the SBIR program and $182,000 which has been transferred to the STTR program.

[b] Excludes $2,322,000 for Safeguards and Security activities transferred to consolidated Safeguards and Security program in FY 2001.

[c] A Budget Amendment transferring $2,700,000 from this program will be submitted shortly. The narrative description for this program has already been adjusted to reflect the revised levels.

Funding by Site

	(dollars in thousands)				
	FY 2000	FY 2001	FY 2002	$ Change	% Change
Albuquerque Operations Office					
Los Alamos National Laboratory	11,637	5,020	5,020	0	--
Sandia National Laboratories	4,961	3,889	3,889	0	--
Total, Albuquerque Operations Office	16,598	8,909	8,909	0	--
Chicago Operations Office					
Ames Laboratory	1,957	1,668	1,668	0	--
Argonne National Laboratory	12,861	10,447	10,047	-400	-3.8%
Brookhaven National Laboratory	1,847	1,566	1,266	-300	-19.2%
Fermi National Accelerator Laboratory	59	60	60	0	--
Princeton Plasma Physics Laboratory	38	0	0	0	--
Chicago Operations Office	8,760	7,240	7,240	0	--
Total, Chicago Operations Office	25,522	20,981	20,281	-700	-3.3%
Oakland Operations Office					
Lawrence Berkeley National Laboratory	57,069	54,501	54,151	-350	-0.6%
Lawrence Livermore National Laboratory	2,884	3,068	3,068	0	--
Stanford Linear Accelerator Center	590	234	234	0	--
Oakland Operations Office	3,160	960	960	0	--
Total, Oakland Operations Office	63,703	58,763	58,413	-350	-0.6%
Oak Ridge Operations Office					
Oak Ridge Inst. For Science and Education	147	99	99	0	--
Oak Ridge National Laboratory	12,016	10,563	10,223	-340	-3.2%
Thomas Jefferson National Accelerator Facility	49	0	0	0	--
Total, Oak Ridge Operations Office	12,212	10,662	10,322	-340	-3.2%
Richland Operations Office					
Pacific Northwest National Laboratory	2,844	2,038	1,738	-300	-14.7%
Washington Headquarters	1,459	64,397	63,387	-1,010	-1.6%
Subtotal, Advanced Scientific Computing Research	122,338 [a] [b]	165,750	163,050	-2,700	-1.6%
Pending Budget Amendment	0	0	2,700 [c]	+2,700	--
Total, Advanced Scientific Computing Research	122,338	165,750	165,750	0	--

[a] Excludes $3,041,000 which has been transferred to the SBIR program and $182,000 which has been transferred to the STTR program.

[b] Excludes $2,322,000 for Safeguards and Security activities transferred to consolidated Safeguards and Security program in FY 2001.

[c] A Budget Amendment transferring $2,700,000 from this program will be submitted shortly. The narrative description for this program has already been adjusted to reflect the revised levels.

Site Description

Ames Laboratory

Ames Laboratory is a Multiprogram Laboratory located on 10 acres in Ames, Iowa. The MICS subprogram at Ames Laboratory conducts research in the materials scientific application pilot project, which focuses on applying advanced computing to problems in microstructural defects, alloys, and magnetic materials, and in computer science. The LTR subprogram at Ames conducts research in the physical, chemical, materials, mathematical, engineering, and environmental sciences through cost-shared collaborations with industry.

Argonne National Laboratory

Argonne National Laboratory (ANL) in Argonne, Illinois, is a Multiprogram Laboratory located on a 1,700 acre site in suburban Chicago. ANL has a satellite site located in Idaho Falls, Idaho. The MICS subprogram at ANL conducts basic research in mathematics and computer science, as well as research in advanced computing software tools and collaboratory tools. ANL also participates in several scientific application and collaboratory pilot projects as well as supporting an advanced computing research facility. The testbed at ANL focuses on a large cluster of Intel-based compute nodes with an open source operating system based on LINUX, this cluster has been given the name of "Chiba City." The LTR subprogram at ANL conducts research motivated by practical energy payoffs through cost-shared collaborations with industry. The areas emphasized in the research portfolio are chemistry of ceramic membranes, separations technology, near-frictionless carbon coatings, and advanced methods for magnesium production.

Brookhaven National Laboratory

Brookhaven National Laboratory (BNL) is a Multiprogram Laboratory located on a 5,200 acre site in Upton, New York. The LTR subprogram at BNL conducts research motivated by practical energy payoffs through cost-shared collaborations with industry. The areas emphasized in the research portfolio are materials for rechargeable lithium batteries, sensors for portable data collection, catalytic production of organic chemicals, and DNA damage responses in human cells.

Fermi National Accelerator Laboratory (Fermilab)

Fermilab is located on a 6,800-acre site about 35 miles west of Chicago, Illinois. The LTR subprogram at Fermilab conducts research in areas such as superconducting magnet research, design and development, detector development and high-performance computing through cost-shared collaborations with industry.

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory (LBNL) is a Multiprogram Laboratory located in Berkeley, California. The Laboratory is on a 200 acre site adjacent to the Berkeley campus of the University of California. The MICS subprogram at LBNL conducts basic research in the mathematics and computer science, as well as research in advanced computing software tools and collaboratory tools. LBNL participates in several scientific application and collaboratory pilot projects as well as supporting an advanced computing research testbed. The testbed at LBNL currently focuses on very large scale computing on hardware in the T3E architecture from SGI-Cray including issues of distributing jobs over all the processors efficiently and the associated system management issues. LBNL manages the Energy

Sciences Network (ESnet). ESnet is one of the world's most effective and progressive science-related computer networks that provides worldwide access and communications to Office of Science (SC) facilities. In 1996, the National Energy Research Scientific Computing Center (NERSC) was moved from the Lawrence Livermore National Laboratory to LBNL. NERSC provides a range of high-performance, state-of-the-art computing resources that are a critical element in the success of many SC research programs. The LTR subprogram at LBNL conducts research motivated by practical energy payoffs through cost-shared collaborations with industry. The areas emphasized in the research portfolio are molecular lubricants for computers, advanced material deposition systems, screening novel anti-cancer compounds, and innovative membranes for oxygen separation.

Lawrence Livermore National Laboratory

Lawrence Livermore National Laboratory (LLNL) is a Multiprogram Laboratory located on an 821 acre site in Livermore, California. The MICS subprogram at LLNL involves significant participation in the advanced computing software tools program as well as basic research in applied mathematics.

Los Alamos National Laboratory

Los Alamos National Laboratory (LANL) is a Multiprogram Laboratory located on a 27,000 acre site in Los Alamos, New Mexico. The Mathematical Information and Computational Sciences (MICS) subprogram at LANL conducts basic research in the mathematics and computer science and in advanced computing software tools. LANL also participates in several scientific application and collaboratory pilot projects as well as supporting an advanced computing research testbed. The testbed at LANL focuses on a progression of technologies from SGI – Cray involving Origin 2000 Symmetric Multiprocessor Computers linked with HiPPI crossbar switches. This series of research computers has been given the name "Nirvana Blue."

Oak Ridge Institute for Science and Education

Oak Ridge Institute for Science and Education (ORISE) is located on 150 acres in Oak Ridge, Tennessee. ORISE provides support for education activities funded within the ASCR program.

Oak Ridge National Laboratory

Oak Ridge National Laboratory (ORNL) is a Multiprogram Laboratory located on a 24,000 acre site in Oak Ridge, Tennessee. The MICS subprogram at ORNL conducts basic research in the mathematics and computer science, as well as research in advanced computing software tools and collaboratory tools.

ORNL also participates in several scientific application and collaboratory pilot projects. The LTR subprogram at ORNL conducts research motivated by practical energy payoffs through cost-shared collaborations with industry. The areas emphasized in the research portfolio are high temperature superconducting wires, microfabricated instrumentation for chemical sensing, and radioactive stents to prevent reformation of arterial blockage.

Pacific Northwest National Laboratory

Pacific Northwest National Laboratory (PNNL) is a Multiprogram Laboratory located on 640 acres at the Department's Hanford site in Richland, Washington. The MICS subprogram at PNNL conducts basic research in the mathematics and computer science, as well as research in advanced computing software tools and collaboratory tools. PNNL also participates in several scientific application pilot projects. The LTR subprogram at PNNL conducts research motivated by practical energy payoffs through cost-shared collaborations with industry. The areas emphasized in the research portfolio are mathematical simulations of glass production, interactions of biological polymers with model surfaces, and characterization of microorganisms in environmental samples.

Princeton Plasma Physics Laboratory

The Princeton Plasma Physics Laboratory (PPPL), a laboratory located in Plainsboro, New Jersey, is dedicated to the development of magnetic fusion energy. The LTR subprogram at PPPL conducts research in areas that include the plasma processing of semiconductor devices and the study of beam-surface interactions through cost-shared collaborations with industry.

Sandia National Laboratories

Sandia National Laboratories (SNL) is a Multiprogram Laboratory, with a total of 3,700 acres, located in Albuquerque, New Mexico, with sites in Livermore, California, and Tonopah, Nevada. The MICS subprogram at SNL conducts basic research in mathematics and computer science, as well as research in advanced computing software tools and collaboratory tools. SNL also participates in several scientific application and collaboratory pilot projects.

Stanford Linear Accelerator Center

The Stanford Linear Accelerator Center (SLAC) is located at the edge of Silicon Valley in California about halfway between San Francisco and San Jose on 426 acres of Stanford University land. The LTR subprogram at SLAC conducts research in areas such as advanced electronics, large-scale ultra-high vacuum systems, radiation physics and monitoring, polarized and high-brightness electron sources, magnet design and measurement, and controls systems through cost-shared collaborations with industry.

Thomas Jefferson National Accelerator Facility

The Thomas Jefferson National Accelerator Facility (TJNAF) is a basic research laboratory located on a 200 acre site in Newport News, Virginia. The LTR subprogram at the TJNAF conducts research in such areas as accelerator and detector engineering, superconducting radiofrequency technology, speed data acquisition, and liquid helium cryogenics through cost-shared collaborations with industry.

All Other Sites

The ASCR program funds research at 71 colleges/universities located in 24 states supporting approximately 117 principal investigators. Also included are funds for research awaiting distribution pending completion of peer review results.

A number of Integrated Software Infrastructure Centers will be established at laboratories and/or universities. Specific site locations will be determined as a result of competitive selection. These centers will focus on specific software challenges confronting users of terascale computers.

Mathematical, Information, and Computational Sciences

Mission Supporting Goals and Objectives

The Mathematical, Information, and Computational Sciences (MICS) subprogram is responsible for carrying out the primary mission of the ASCR program: discovering, developing, and deploying advanced scientific computing and communications tools and operating the high performance computing and network facilities that researchers need to analyze, model, simulate, and — most importantly — predict the behavior of complex natural and engineered systems of importance to the Office of Science and to the Department of Energy. The MICS subprogram supports fundamental research and research facilities in all of the areas in which MICS supports research:

- *Applied Mathematics*. This includes research on the underlying mathematical understanding as well as the numerical algorithms to enable effective description and prediction of physical systems such as fluids, magnetized plasmas, or protein molecules. This includes, for example, methods for solving large systems of partial differential equations on parallel computers, techniques for choosing optimal values for parameters in large systems with hundreds to hundreds of thousands of parameters, improving our understanding of fluid turbulence, and developing techniques for reliably estimating the errors in simulations of complex physical phenomena.
- *Computer Science*. This includes research in computer science to enable large scientific applications through advances in massively parallel computing, such as very lightweight operating systems for parallel computers, distributed computing such as development of the Parallel Virtual Machine (PVM) software package that has become an industry standard, and large scale data management and visualization. The development of new computer and computational science techniques will allow scientists to use the most advanced computers without being overwhelmed by the complexity of rewriting their codes every 18 months.
- *Networking*. This includes research in high performance networks and information surety required to support high performance applications – protocols for high performance networks, methods for measuring the performance of high performance networks, and software to enable high speed connections between high performance computers and networks, and scalable methods for providing scientific users of networks the security services they need. The information security and assurance research supported by MICS is focused on the requirements of scientists engaged in unclassified research for applications such as remote control of experimental devices and research collaborations. The development of high-speed communications and collaboration technologies will allow scientists to view, compare, and integrate data from multiple sources remotely.

MICS also operates supercomputer and network facilities that are available to researchers 24 hours a day, 365 days a year. The computing and networking requirements of the Office of Science far exceed the current state-of-the-art; furthermore, the requirements far exceed the tools that the commercial marketplace will deliver. For this reason, the MICS subprogram must not only support basic research in the areas listed above, but also the development of the results from this basic research into software usable by scientists in other disciplines and partnerships with users to test the usefulness of the research. These partnerships with the scientific disciplines are critical because they provide rigorous tests of the usefulness of current advanced computing research, enable MICS to transfer the results of this research to scientists in the disciplines, and help define promising areas for future research. This integrated approach is critical for MICS to succeed in providing the extraordinary computational and communications tools that DOE's civilian programs need to carry out their missions. It is important to note that these tools have applications beyond the Office of Science in the NNSA and in the private sector after they have been initially discovered and developed by MICS.

As noted earlier, in FY 2002, the MICS subprogram will continue its components of the collaborative SciDAC program across the Office of Science to produce the scientific computing, networking and collaboration tools that DOE researchers will require to address the scientific challenges of the next decade. The MICS components include investments in scientific computing research and networking and collaboration research that are complemented by investments in computing and networking facilities.

Performance will be measured by reporting accomplishments on the common performance measures on leadership, excellence and relevance; quality; and safety and health. In addition, **performance will be measured by** serving researchers at national laboratories, universities, and industry, thus, enabling new understanding through analysis, modeling and simulation of complex natural and engineered systems and effective integration of geographically distributed teams through national collaboratories.

Scientific Computing Research Investments

As noted earlier, advances in computing technologies during the past decade have set the stage for a major step forward in modeling and simulation. The key measure of success in translating peak computing power into science is the percent of peak performance that is delivered to an application over the *entire* calculation. In the early to mid-1990's on computers such as the Cray Research C-90, many scientific codes realized 40% to 50% of the peak performance of the supercomputer. In contrast, on today's parallel supercomputers, scientific computing codes often realize only 5% to 10% of "peak" performance, and this fraction could decrease as the number of processors in the computers grows.

This challenge is a direct result of the fact that the speed of memory systems and the speed of interconnects between processors is increasing much more slowly than processor speed. For many scientific applications these factors dominate the performance of the application. Two types of solutions are available to the computer hardware designer in addressing the mismatch of speed between the components: (1) clever hierarchical arrangements of memory with varying speeds and software to find data before it is needed and move it into faster memory, closer to the processor that will need it; and (2) techniques to increase parallelism, for example, by using threads in the processor workloads or by combining parallel data streams from memory or disks. Current technology forecasts indicate a doubling or quadrupling in the numbers of layers in the memory hierarchy, and a 100- to 1000-fold increase in the amount of parallelism in disk and tape systems to accommodate the relative increase in the mismatch between processor speed and memory, disk and tape speeds in the next five years.

One result of this increasing complexity of high-performance computer systems is the importance of the underlying systems software. Operating systems, compilers, runtime environments, mathematical libraries, and end-user applications must all work together efficiently to extract the desired high performance from these systems.

In addition to the challenges inherent to managing the required level of parallelism, technology trends and business forces in the U.S. computer system industry have resulted in radically reduced development and production of high-end systems necessary for meeting the most demanding requirements of scientific research. In essence, the U.S. computer industry has become focused on the computer hardware and software needs of business applications, and little attention is paid to the special computational needs of the scientific community. Therefore, to achieve the performance levels required for agency missions and world leadership in computational science, large numbers of smaller commercial systems must be combined and integrated to produce terascale computers. Unfortunately, the operating systems software and tools required for effective use of these large systems are significantly different from the technology offered for the individual smaller components. Therefore, new enabling software must be developed if scientists are to take advantage of these new computers in the next five years.

The following are specific examples of *computer science* research challenges:

- Efficient, high-performance operating systems, compilers, and communications libraries for high-end computers.
- Software to enable scientists to store, manage, analyze, visualize, and extract scientific understanding from the enormous (terabyte to petabyte) data archives that these computers will generate.
- Software frameworks that enable scientists to reuse most of their intellectual investment when moving from one computer to another and make use of lower-level components, such as runtime services and mathematical libraries, that have been optimized for the particular architecture.
- Scalable resource management and scheduling software for computers with thousands of processors.
- Performance monitoring tools to enable scientists to understand how to achieve high performance with their codes.

In addition to these computer science challenges, significant enhancements to the MICS applied mathematical research activity are required for the Department to satisfy its mission requirements for computational science. Over the history of computing, improvements in algorithms have yielded at least as much increase in performance as has hardware speedup. Large proportions of these advances are the products of the MICS applied mathematics research activity. In addition to improving the speed of the calculations, many of these advances have dramatically increased the amount of scientific understanding produced by each computer operation. For example, a class of mathematical algorithms called "fast multipole algorithms," were discovered for a number of important mathematical operations required to process 1,000 datapoints by a factor of 1,000; 10,000 datapoints by a factor of 10,000; and so on. Another example of how powerful these methods can be is that they enable a scientist to process 10,000 datapoints in the time that it would have taken to process 100 using earlier techniques, or 1,000,000 datapoints in the time older techniques would have needed to process 1,000. The requirements of scientific domains for new algorithms that can scale to work effectively across thousands of processors and produce the most science in the fewest number of computer operations drives the need for improved mathematical algorithms and the supporting software libraries that must be made available for ready use by domain scientists. In this area of research the MICS applied mathematics activity is the core of the nationwide effort.

The MICS subprogram will address these challenges by continuing the competitively selected partnerships (based on a solicitation notice to labs and universities) focused on discovering, developing, and deploying to scientists key enabling technologies that were initiated in FY 2001. These partnerships, which are called Integrated Software Infrastructure Centers, must support the full range of activities from basic research through deployment and training because the commercial market for software to support terascale scientific computers is too small to be interesting to commercial software providers. These centers play a critical role in providing the software infrastructure that will be used by the SciDAC applications research teams. The management of these centers will build on the successful experience of the MICS subprogram in managing the DOE2000 initiative, as well as on the lessons learned in important programs supported by DARPA such as Project Athena at MIT, the Berkeley Unix Project, and the initial development of the Internet software and the Internet Activities Board (IAB). These Integrated Software Infrastructure Centers will have close ties to key scientific applications projects to ensure their success.

The efforts initiated in FY 2001 address the important issues of understanding and developing the tools that applications developers need to make effective use of machines that will be available in the next several years.

The MICS activities that respond to these challenges are described later in the Detailed Program Justification section of the MICS subprogram budget under the headings:

Applied Mathematics,

Computer Science, and

Advanced Computing Software Tools.

Networking and Collaboration Research Investments

Advances in network capabilities and network-based technologies now make it possible for large geographically distributed teams to effectively collaborate on the solution of complex problems. This is especially important for the teams using the major experimental facilities, computational resources, and data resources supported by DOE.

- Significant research is needed to enable today's Internet to be effectively used for scientific data retrieval and analysis and collaboratories. The requirements this places on the network are very different than the requirements of the commercial sector where millions of users are moving to small web pages. The MICS subprogram includes research on advanced protocols, special operating system services to support very high-speed transfers, and advanced network control.
- Research is also needed to understand how to integrate the large number of network devices, network-attached devices, and services that collaboratories require. Examples of the components and services that need to be integrated include network resources, data archives on tape, high

performance disk caches, visualization and data analysis servers, authentication and security services, and the computer on a scientist's desk. Common software framework building blocks or "middleware" to enable the collaboratories of the future to succeed must tie all of these physical and software services together.

The MICS subprogram will address these challenges through an integrated program of fundamental research in networking and collaboratory tools, partnerships with key scientific disciplines, and advanced network testbeds.

The MICS activities that respond to these challenges are described in the Detailed Program Justification section of the MICS subprogram budget under the headings:

Networking,

Collaboratory Tools, and

National Collaboratory Pilot Projects.

Enhancements to Computing and Networking Facilities

To realize the scientific opportunities offered by advanced computing, enhancements to the Office of Science's computing and networking facilities are also required. The MICS subprogram supports a suite of high-end computing resources and networking resources for the Office of Science:

- **Production High Performance Computing Facilities.** The National Energy Research Scientific Computing Center (NERSC) provides high performance computing for investigators supported by the Office of Science. NERSC provides a spectrum of supercomputers offering a range of high performance computing resources and associated software support.
- **Energy Sciences Network (ESnet).** ESnet provides worldwide access to Office of Science facilities, including light sources, neutron sources, particle accelerators, fusion reactors, spectrometers, high-end computing facilities and other leading-edge instruments and facilities.
- **Advanced Computing Research Testbeds.** These testbeds provide advanced computational hardware for testing and evaluating new computing hardware and software. In addition, these testbeds will provide specialized computational resources to support SciDAC applications teams in FY 2002.

Current production resources provide less than half of the computer resources that were requested last year. The pressure on production facilities will only increase in future years as more applications become ready to move from testing the software to using the software to generate new science. In addition, as the speed of computers increases, the amount of data they produce also increases. Therefore, focused enhancements to the Office of Science's network infrastructure are required to enable scientists to access and understand the data generated by their software. These network enhancements are also required to allow researchers to have effective remote access to the experimental facilities that the Office of Science provides for the Nation.

The MICS activities that respond to these challenges are described later in the Detailed Program Justification section of the MICS subprogram budget under the headings:

National Energy Research Scientific Computing Center (NERSC),
Advanced Computing Research Testbeds, and
Energy Sciences Network (ESnet).

Funding Schedule

(dollars in thousands)

	FY 2000	FY 2001	FY 2002	$ Change	% Change
Mathematical, Computational, and Computer Sciences Research	42,713	70,654	70,681	+27	--
Advanced Computation, Communications Research and Associated Activities	71,201	81,543	81,543	0	--
SBIR/STTR	0	3,973	3,946	-27	-0.7%
Total, Mathematical, Information, and Computational Sciences	113,914	156,170	156,170	0	--

Detailed Program Justification

(dollars in thousands)

	FY 2000	FY 2001	FY 2002
Mathematical, Computational, and Computer Sciences Research	**42,713**	**70,654**	**70,681**
Applied Mathematics	**20,391**	**32,339**	**32,366**

Research is conducted on the underlying mathematical understanding and numerical algorithms to enable effective description and prediction of physical systems. Research in applied mathematics is critical to the DOE because of the potential of improved mathematical techniques to enable large computational simulations. As discussed earlier in the ASCR overview, improvements in mathematical algorithms are responsible for greater improvement in scientific computing capabilities than the increases in hardware performance. This activity supports research at DOE laboratories, universities, and private companies at a level similar to previous years. Many of the projects supported by this activity are partnerships between researchers at universities and DOE laboratories. To accomplish its goals, the program supports research in a number of areas including: ordinary and partial differential equations, including numerical linear algebra, iterative methods, sparse solvers, and dense solvers; fluid dynamics, including compressible, incompressible and reacting flows, turbulence modeling, and multiphase flows; optimization, including linear and nonlinear programming, interior-point methods, and discrete and integer programming; mathematical physics, including string theory, superstring theory, geometry of space-time, and quantum effects; control theory, including differential-algebraic systems, order

	(dollars in thousands)	
FY 2000	FY 2001	FY 2002

reduction, queuing theory; shock wave theory systems, multipole expansions, mixed elliptic-hyperbolic problems, including hyperbolic and wavelet transforms; dynamical systems, including chaos-theory and control, and bifurcation theory; programming; and geometric and symbolic computing, including minimal surfaces and automated theorem proving.

The FY 2002 budget continues the FY 2001 increased level of funding for the Computational Sciences Graduate Fellowship program and the competitively selected Integrated Software Infrastructure Centers (ISICs) partnerships focused on algorithms and mathematical libraries for critical DOE applications on terascale computers that are a significant component of the SciDAC program.

Performance will be measured in a number of ways. Efforts in applied mathematics will be continuously evaluated for their leadership and significant contributions to the worldwide applied mathematics effort using a number of measures including awards, significant advances, and invited participation and membership on organizing and program committees of major national and international conferences. Progress reviews of the Integrated Software Infrastructure Centers (ISICs) will be conducted to ensure that these partnerships are moving forward to accomplish their missions. The Computational Science Graduate Fellowship Program will appoint 20 new students to develop the next generation of leaders in computational science for DOE and the Nation.

	FY 2000	FY 2001	FY 2002
Computer Science	**13,747**	**21,051**	**21,051**

Research in computer science to enable large scientific applications is critical to DOE because its unique requirements for high performance computing significantly exceed the capabilities of computer vendors' standard products. Therefore, much of the computer science to support this scale of computation must be developed by DOE. This activity supports research in two general areas: the underlying software to enable applications to make effective use of computers with hundreds or thousands of processors as well as computers that are located at different sites; and large scale data management and visualization under circumstances where the underlying resources and users are geographically distributed. The first area includes research in protocols and tools for interprocessor communication and parallel input/output (I/O) as well as tools to monitor the performance of scientific applications and advanced techniques for visualizing very large-scale scientific data. Researchers at DOE laboratories and universities, often working together in partnerships, carry out this research. In FY 2002, support for the competitively selected Integrated Software Infrastructure Centers to address critical computer science and systems software issues for terascale computers will be continued. The teams in these Centers focus on critical issues including: tools for analyzing and debugging scientific simulation software that uses thousands of processors; and the development of data management and visualization software capable of handling terabyte scale data sets extracted from petabyte scale data archives. These Integrated Software Infrastructure Centers are a critical component in DOE's strategy for SciDAC.

	(dollars in thousands)		
	FY 2000	FY 2001	FY 2002

Performance in computer science will be measured through peer review, periodic external expert review of ongoing projects, production of significant research results, and adoption of these technologies by other researchers supported by the Office of Science. In addition, the ISICs initiated in FY 2001 will undergo a progress review to ensure effective coupling both between the ISICs, and between the ISICs and application teams in the MICS Scientific Applications Pilot Projects efforts and with the SciDAC teams funded by the other Programs in the Office of Science.

Advanced Computing Software Tools	**4,910**	**8,473**	**8,473**

This research uses the results of fundamental research in applied mathematics and computer science to develop an integrated set of software tools that scientists in various disciplines can use to develop high performance applications (such as simulating the behavior of materials). These tools, that provide improved performance on high-end systems, are critical to the ability of scientists to attack the complex scientific and engineering problems that can only be solved with high-end computing systems. The initial goal of this program element was to develop foundational tools (math libraries, runtime systems, etc.) that will have a useful life spanning many generations of computer hardware. From the experience gained with end user application scientists applying these tools, it has become clear that to promote wide usage across the scientific community, the tools must also be reliable, documented, and easy to use. In addition, users of the tools need the tools to be maintained so that the tools continue to be available, have bugs fixed, etc. Since many of the tools needed in the high performance arena have no commercial market, the Integrated Software Infrastructure Centers initiated in FY 2001 will provide a means for focused investment to deploy these tools to the scientific community. These competitively selected centers focus research in several areas that include software frameworks, problem-solving environments, distributed computing and collaboration technologies, as well as visualization and data management.

Performance will be measured through peer review, periodic external expert review of ongoing projects, production of significant research results, and adoption of these technologies by other researchers in the Office of Science. In addition, the ISICs initiated in FY 2001 will undergo a progress review to ensure effective coupling between the ISICs and between the ISICs and application teams in the MICS Scientific Applications Pilot Projects efforts and in the SciDAC teams funded by the other Programs in the Office of Science.

Scientific Applications Pilot Projects	**3,665**	**8,791**	**8,791**

This research is a collaborative effort with disciplinary computational scientists to apply the computational techniques and tools developed by MICS supported research to basic research problems relevant to the mission of SC. This effort tests the usefulness of current advanced computing research, transfers the results of this research to the scientific disciplines, and helps define promising areas for future research. The FY 2002 funding for this activity will allow the continuation of the pilot projects that were competitively selected in FY 2001. These pilot projects are tightly coupled to the Integrated Software Infrastructure Centers (described above in applied

	FY 2000	FY 2001	FY 2002
(dollars in thousands)			

mathematics, computer science and advanced computing software tools) to ensure that these activities are an integrated approach to the challenges of terascale simulation and modeling that DOE faces to accomplish its missions.

	FY 2000	FY 2001	FY 2002
Advanced Computation, Communications Research, and Associated Activities	**71,201**	**81,543**	**81,543**
Networking	**5,892**	**7,066**	**7,066**

Research is needed to develop high-performance networks that are capable of supporting distributed high-end computing and secure large-scale scientific collaboration. High performance networks enable scientists to collaborate effectively and to have efficient access to distributed computing resources such as tera-scale computers, and experimental scientific instruments, and large scientific data archives. This research is carried out at national laboratories and universities. It focuses in areas such as high-performance transport protocols for high-speed networks; scalable techniques for measuring, analyzing, and controlling traffic in high performance networks; network security research to support large-scale scientific collaboration; advanced network components to enable high-speed connections between terascale computers, large scientific data archives, and high-speed networks; and research on high-performance "middleware." Middleware is a collection of network-aware software components that scientific applications need in order to couple efficiently to advanced network services and make effective use of experimental devices, data archives, and terascale computers at different locations. In all of these cases, the network and middleware requirements of DOE significantly exceed those of the commercial market.

	FY 2000	FY 2001	FY 2002
Collaboratory Tools	**2,946**	**5,527**	**5,527**

This research uses the results of fundamental research in computer science and networking to develop an integrated set of software tools to support scientific collaborations. This includes enabling scientists to remotely access and control facilities and share data in real time, and to effectively share data with colleagues throughout the life of a project. These tools provide a new way of organizing and performing scientific work that offers the potential for increased productivity and efficiency. These tools will also enable broader access to important DOE facilities and data resources by scientists and educators across the country. This research includes, for example, developing and demonstrating an open, scalable approach to application-level security in widely distributed, open network environments that can be used by all the collaboratory tools as well as by the advanced computing software tools whenever access control and authentication are issues. Having demonstrated feasibility of the security architecture on a small scale, an additional investment is needed to support the integration of collaboratory tools with advanced networking services in a research setting. In this way, security features can be integrated into more end user applications or collaboratory tools, and demonstrated on a large user base. An example of research in collaboratory tools is the development of a modular electronic notebook that extends the capabilities of its paper counterpart by allowing scientists located across the country to share a common record of ideas, data and events of their joint experiments and research programs. These are designed to be valid as a long term, legally defensible record of research, invention, and records management. Tools are being developed for managing distributed collaborations where videoconferencing, whiteboards, and other shared applications are important. Software to enable geographically distributed teams to collaboratively control visualization of data is also being investigated.

National Collaboratory Pilot Projects	**4,910**	**10,857**	**10,857**

This program is intended to test, validate, and apply collaboratory tools in real-world situations in partnership with other DOE programs. The competitively selected partnerships involve national laboratories, universities, and U.S. industry. It is important to continue to demonstrate and test the benefits of collaboratory tools technology in order to promote its widespread use and enable more effective access to the wide range of resources within the Department, from light sources to terascale computers to petabyte data storage systems. The partnerships that were initiated in FY 2001 focus on developing user environments where collaboration is ubiquitous and distributed computing is seamless and transparent for DOE mission applications such as High Energy and Nuclear Physics data as well as remote analysis and visualization of experimental and simulation results. This level of funding will permit the continuation of the efforts funded in FY 2001.

National Energy Research Scientific Computing Center (NERSC)	**26,252**	**28,244**	**28,244**

NERSC, located at LBNL, provides high performance computing for investigators supported by the Office of Science. The Center serves 2,000 users working on about 700 projects; 35 percent of users are university based, 60 percent are in National Laboratories, and 5 percent are in industry. NERSC provides a spectrum of supercomputers offering a range of high performance computing resources and associated software support. The two major computational resources at NERSC are a 512 processor Cray T3E computer and a large IBM SP computer whose initial installation was completed in early FY 2000 in a fully competitive procurement process. The FY 2002 funding will support the operation of the IBM-SP computer at about 3.5 teraflops "peak" performance. These computational resources will be integrated by a common high performance file storage system that facilitates interdisciplinary collaborations. Related requirements for capital equipment, such as high-speed disk storage systems, archival data storage systems, and high performance visualization hardware, and general plant projects (GPP) funding are also supported. The MIE for the distributed visualization server was completed in FY 2001. FY 2002 capital equipment requirements continue at the same level as in FY 2001; however, no individual anticipated capital purchase exceeds the MIE threshold.

Performance will be measured in a number of ways. The operating time lost due to unscheduled NERSC downtime, which will be less than 10 percent of the total scheduled possible operating time. In addition, user surveys will continue to show a high degree of satisfaction with the services at NERSC and annual reports will continue to demonstrate production of world-class

	(dollars in thousands)	
FY 2000	FY 2001	FY 2002

science from the facility. In FY 2000, the measured operating time lost to unscheduled downtime on systems at NERSC ranged from .08% to .52% of total scheduled possible operating time. NERSC will be operated within budget while meeting user needs and satisfying overall SC program requirements.

	FY 2000	FY 2001	FY 2002
Advanced Computing Research Testbeds	**15,739**	**13,061**	**13,061**

This activity supports the advanced computational hardware testbeds that play a critical role in testing and evaluating new computing hardware and software, especially with regard to their applicabilities to scientific problems. Current testbeds are located at Argonne National Laboratory (IBM/ Intel Cluster); and ORNL (Compaq-Alpha technology). These testbeds represent the evolution of Advanced Computing Research Facilities that supported the computational requirements of the scientific application partnerships that were completed in FY 2000. Support for the Nirvana Blue Computer Testbed at LANL was phased out in FY 2001. This activity also supports the distributed high performance storage system (HPSS) testbed collaboration between ORNL and LBNL. Because many of the issues to be investigated only appear in the computer systems at significantly larger scale than the computer manufacturers' commercial design point, these testbeds must procure the largest scale systems that can be afforded and develop software to manage and make them useful. In addition, the ACRTs, taken together, must have a full range of computer architectures to enable comparison and reduce overall program risk. These all involve significant research efforts, often in partnership with the vendors to resolve issues including operating system stability and performance, system manageability and scheduling, fault tolerance and recovery, and details of the interprocessor communications network. Therefore, these systems are managed as research programs and not as information technology investments. In addition, these testbeds will provide specialized computational resources to support SciDAC applications teams in FY 2002.

Performance will be measured by the importance of the research that results from these testbeds as viewed by publications in the scientific literature, the ASCR Advisory Committee and external reviews and the demand for access to these facilities by the nationwide computer and computational science communities.

	FY 2000	FY 2001	FY 2002
Energy Sciences Network (ESnet).......................................	**15,462**	**16,788**	**16,788**

ESnet provides worldwide access to the Office of Science facilities, including: advanced light sources; neutron sources; particle accelerators; fusion reactors; spectrometers; supercomputers; Advanced Computing Research Testbeds (ACRTs); and other leading-edge science instruments and facilities. ESnet provides the communications fabric that links worldwide DOE researchers to one another and forms the basis for fundamental research in networking, enabling R&D in collaboratory tools, and applications testbeds such as the national collaboratory pilot projects. To provide these facilities, DOE employs ESnet management at LBNL, who contracts with commercial vendors for advanced communications services including Asynchronous Transfer Mode (ATM) and Wave Division Multiplexing (WDM). LBNL ESnet management provides system integration to provide a uniform interface to these services for DOE laboratories. In

		(dollars in thousands)	
	FY 2000	FY 2001	FY 2002

addition, LBNL ESnet management is responsible for the interfaces between the network fabric it provides and the worldwide Internet including the University Corporation for Advanced Internet Development (UCAID) Abilene network that provides high performance connections to many research universities. One reason that ESnet, in the words of the 1998 external review committee, is able to provide the capabilities and services to its users "at significantly lower budgets than other agencies" is its management structure with strong user and site coordination committees. This management structure is built on DOE's experience in operating large user facilities. The funding in FY 2002 will continue support for an advanced network testbed established in FY 2001 to enable research in collaboratory tools and pilots. Related capital equipment needs are also supported such as high-speed network routers, ATM switches, and network management and testing equipment.

Performance will be measured in a number of ways. The operating time lost due to unscheduled ESnet downtime will be less than 10 percent of the total scheduled possible operating time. In addition, the ESnet program will be reviewed by an external committee during FY 2001 to ensure its effectiveness in meeting its goals. In FY 2000, the measured operating time lost to unscheduled downtime on ESnet was 4 percent of total scheduled possible operating time. ESnet will operate within budget while meeting user needs and satisfying overall SC program requirements. Network enhancements improve researchers access to high performance computing and software support, and enhance scientific opportunities by enabling scientists to access and understand greater amounts of scientific data.

	FY 2000	FY 2001	FY 2002
SBIR/STTR	**0**	**3,973**	**3,946**

In FY 2000, $2,821,000 and $169,000 were transferred to the SBIR and STTR programs, respectively. The FY 2001 and FY 2002 amounts are the estimated requirement for the continuation of the SBIR and STTR programs.

	FY 2000	FY 2001	FY 2002
Total, Mathematical, Information, and Computational Sciences	**113,914**	**156,170**	**156,170**

Explanation of Funding Changes from FY 2001 to FY 2002

	FY 2002 vs. FY 2001 ($000)
Mathematical, Computational, and Computer Sciences Research	
Provides a slight increase for the competitively selected Integrated Software Infrastructure Centers (ISICs) partnerships focused on algorithms and mathematical libraries for critical DOE applications on terascale computers that are a significant component of the SciDAC program.	+27
SBIR/STTR	
Decrease in SBIR/STTR due to decrease in operating expenses.	-27
Total Funding Change, Mathematical, Information, and Computational Sciences	0

Laboratory Technology Research

Mission Supporting Goals and Objectives

The mission of the Laboratory Technology Research (LTR) subprogram is to support high risk research that advances science and technology to enable applications that could significantly impact the Nation's energy economy. LTR fosters the production of research results motivated by a practical energy payoff through cost-shared collaborations between the Office of Science (SC) laboratories and industry.

An important component of the Department's strategic goals is to ensure that the United States maintains its leadership in science and technology. LTR is the lead program in the Office of Science for leveraging science and technology to advance understanding and to promote our country's economic competitiveness through cost-shared partnerships with the private sector.

The National Laboratories under the stewardship of the Office of Science conduct research in a variety of scientific and technical fields and operates unique scientific facilities. Viewed as a system, these ten laboratories — Ames Laboratory, Argonne National Laboratory, Brookhaven National Laboratory, Fermi National Accelerator Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory, Pacific Northwest National Laboratory, Princeton Plasma Physics Laboratory, Stanford Linear Accelerator Center, and Thomas Jefferson National Accelerator Facility — offer a comprehensive resource for research collaborations. The major component of the LTR research portfolio consists of investments at these laboratories to conduct research that benefits all major stakeholders — the DOE, the industrial collaborators, and the Nation. These investments are further leveraged by the participation of an industry partner, using Cooperative Research and Development Agreements (CRADAs). Another LTR subprogram component provides funding to the Office of Science national laboratories to facilitate rapid access to the research capabilities at the SC laboratories through agile partnership mechanisms including personnel exchanges and technical consultations with small business. The LTR subprogram currently emphasizes three critical areas of DOE mission-related research: advanced materials processing and utilization, nanotechnology, intelligent processes and controls, and energy-related applications of biotechnology.

Funding Schedule

(dollars in thousands)

	FY 2000	FY 2001	FY 2002	$ Change	% Change
Laboratory Technology Research..........	8,424	9,326	6,698	-2,628	-28.2%
SBIR/STTR	0	254	182	-72	-28.3%
Total, Laboratory Technology Research..	8,424	9,580	6,880	-2,700	-28.2%

Detailed Program Justification

	(dollars in thousands)		
	FY 2000	FY 2001	FY 2002
Laboratory Technology Research	**8,424**	**9,326**	**6,698**

This activity supports research to advance the fundamental science at the Office of Science (SC) laboratories toward innovative energy applications. Through CRADAs, the SC laboratories enter into cost-shared research partnerships with industry, typically for a period of three years, to explore energy applications of research advances in areas of mission relevance to both parties. The research portfolio consists of 45 projects and emphasizes the following topics: advanced materials processing and utilization, nanotechnology, intelligent processes and controls, and energy-related applications of biotechnology. Efforts underway include the exploration of: (1) a new mask-less photoelectrochemical method for depositing conductive metal patterns with nanometer-scale precision for use in fabricating miniaturized and rugged electrical interconnects and biomolecular electronic devices on any surface or in solution; (2) an innovative injector design for a compact, low-cost, high brightness electron gun to be used in linear colliders, short wavelength Free Electron Lasers, and as a test bed for advanced accelerator concepts; and (3) a new design for a compact scintillation camera for medical imaging for use in the detection of thyroid disease and for pre-surgical imaging of breast cancer and nodal metastases. A small but important component of this activity provides industry, particularly small businesses, with rapid access to the unique research capabilities and resources at the SC laboratories. These research efforts are usually supported for a few months to quantify the energy benefit of a specific problem posed by industry. Recent projects supported the development of: (1) Metal Plasma Immersion Ion Implantation and Deposition to produce various layers of reduced size in copper metallization for the next generation of integrated circuits; (2) the optimization of the composition and mechanical properties of new steels for elevated temperature applications used in the power generation, chemical, and petrochemical industries; and (3) infrared thermography and computer modeling of electrical surge arresters which should lead to increased reliability of electrical energy transmission and reduced inconvenience and expense of power outages for end users. The FY 2002 budget will allow the LTR subprogram to continue technology research partnership projects in emerging areas of interest to DOE; however, the number of research projects supported by CRADAs will be reduced by approximately 30 percent from FY 2001. The Rapid Access portion of the LTR subprogram will be preserved.

Performance in this activity will be measured through merit-based peer and on-site reviews.

SBIR/STTR	**0**	**254**	**182**

In FY 2000, $220,000 and $13,000 were transferred to the SBIR and STTR programs, respectively. The FY 2001 and FY 2002 amounts are the estimated requirement for the continuation of the SBIR and STTR programs.

Total, Laboratory Technology Research................................	**8,424**	**9,580**	**6,880**

Explanation of Funding Changes from FY 2001 to FY 2002

	FY 2002 vs. FY 2001 ($000)
Laboratory Technology Research	
Decrease the number of supported collaborative research projects by approximately 30 percent.	-2,628
SBIR/STTR	
SBIR/STTR decreases due to decrease in operating expenses.	-72
Total Funding Change, Laboratory Technology Research	-2,700

Capital Operating Expenses & Construction Summary

Capital Operating Expenses

(dollars in thousands)

	FY 2000	FY 2001	FY 2002	$ Change	% Change
General Plant Projects......................	0	0	1,000	+1,000	+100.0%
Capital Equipment (total)..................	3,797	6,250	6,250	0	0.0%
Total, Capital Operating Expenses	3,797	6,250	7,250	+1,000	+16.0%

Major Items of Equipment (TEC $2 million or greater)

(dollars in thousands)

	Total Estimated Cost (TEC)	Prior Year Approp-riations	FY 2000	FY 2001	FY 2002	Accept-ance Date
Archival Systems Upgrade – LBNL........	2,000	0	2,000	0	0	FY 2002
Distributed Visualization Server – LBNL..	2,500	0	0	2,500	0	FY 2001
Total, Major Items of Equipment............		0	2,000	2,500	0	