Computational and Technology Research

Program Mission

The mission of the Computational and Technology Research (CTR) program, which consists of two distinct activities, is

- to foster and support fundamental research in advanced computing research — applied mathematics, computer science, and networking — and to operate supercomputer, networking, and related facilities to enable the analysis, modeling, simulation, and prediction of complex phenomena important to the Department of Energy, and
- 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 this program is continuously evaluated through the use of merit based peer review and scientific advisory committees.

Program Goals

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

Program Objectives

- *Foster and support fundamental, peer-reviewed research.* — Foster research to create new fundamental knowledge in areas of advanced computing research important to the Department.
- *Apply advanced computing knowledge to complex problems of importance to DOE.* — Promote the transfer of the results of advanced computing research to contribute 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.
- *Plan, fabricate, assemble, and operate premier supercomputer and networking facilities.* — Serve researchers at national laboratories, universities, and industry, thus enabling both new understanding

through analysis, modeling, and simulation of complex problems 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.

Performance Measures

The Computational and Technology Research program measures performance in various ways, depending on the objective. However, performance measures fall into four broad categories: (1) peer review; (2) indicators or metrics (i.e., things that can be counted); (3) customer evaluation and stakeholder input; and (4) qualitative assessments, which might include historical retrospectives and annual program highlights. The key process is peer review of all research activities. Facility performance measures include achievement of performance specifications, operating time, throughput, user satisfaction, and effective utilization of resources as determined by reports from external review panels, user steering committees, and internal Office of Science (SC) program manager committees. In addition, CTR supercomputer and network facilities have periodic external performance reviews. The Energy Science Network (ESnet) operations and management were reviewed in this manner in FY 1998.

In FY 2000, (1) facilities, including the National Energy Research Scientific Computing Center (NERSC) and ESnet, will be operated within budget to meet user and overall SC program requirements; (2) the operating time lost at scientific facilities due to unscheduled NERSC/ESnet downtime will be less than 10 percent of the total scheduled possible operating time, on average; (3) all new and continuing research projects will undergo regular peer review and merit evaluation based on the principles set down in 10 CFR 605 for grants or in cooperative agreements supported by the Office of Science; (4) work performed by investigators supported by CTR will continue to be recognized as outstanding through the receipt of major prizes and awards; and (5) initiate 7 Laboratory Technology Research projects that address the Department's top priorities for science and technology, through cost-shared research partnerships with industry.

Significant Accomplishments and Program Shifts

The CTR program builds on decades of leadership in advanced computing. Some of the pioneering accomplishments of this program are:

Mathematical, Information, and Computational Sciences

- **Remote, interactive access to supercomputers**. At the National Magnetic Fusion Energy Computing Center [the predecessor to the National Energy Research Scientific Computing Center (NERSC)], DOE pioneered the concept of remote, interactive access to supercomputers. Before this time, scientists using supercomputers had to travel to the location of the computer to make use of it. In addition, users were only able to use these computers by submitting jobs and waiting for hours or days to see the output. The Mathematical, Information, and Computational Sciences (MICS) subprogram developed the first interactive operating system for supercomputers, Cray Time Sharing System (CTSS), as well as a nationwide network to allow remote users to have effective access to the computers. This operating system revolutionized access to supercomputers by enabling users to see their jobs as they executed. When the National Science Foundation (NSF) initiated its Supercomputer Centers program in the 1970's, the CTSS operating system was adopted by the San Diego Supercomputing Center and the National Center for Supercomputing Applications to enable users to access NSF's first CRAY machines.

- **Numerical Linear Algebra Libraries**. Today's high performance computations rely on high performance, efficient libraries of numerical linear algebra software. These libraries, which are the core of numerical efforts in the solution of differential and integral equations — LINPACK, EISPAC, LAPACK, SCALAPACK — are the direct result of decades of DOE funding and basic research in this area. These libraries are used by thousands of researchers worldwide and are a critical part of the world's scientific computing infrastructure.

- **High Performance Parallel Interface (HiPPI)**. In order to develop a standard interface between supercomputers and other devices, such as disk arrays and archival tape systems, and visualization computers, DOE laboratories developed the high performance network interface (HiPPI) and led a consortium of vendors to make it the industry standard for the highest bandwidth interconnects between computers and peripheral devices. Many research issues in high speed signaling, data parallelism and high speed protocol design needed to be understood to enable this advance.

- **Parallel Virtual Machine (PVM) and Message Passing Interface (MPI).** DOE researchers developed PVM and MPI to enable scientists to make effective use of networks of workstations and massively parallel computers. Both of these software packages have become standards in the industry and are implemented by virtually all of the high performance computer manufacturers in the world. Both of these developments were enabled by over a decade of basic research in message passing and distributed computing supported by DOE along with many experiments to apply these techniques to real scientific problems.

- **Slow Start Algorithm for the Transmission Control Protocol (TCP)**. Transmission Control Protocol (TCP) part of TCP/IP (Internet Protocol) is responsible for ensuring that packets arrive at their destination. In 1987, as DOE and the other Federal agencies were interconnecting their networks to form the core of the Internet, critical parts of the infrastructure began to fail. There was concern that this represented a fundamental flaw in the TCP/IP architecture; however, a researcher at LBNL applied ideas from fluid flow research to understand the problem and develop a solution. This new TCP algorithm was incorporated in virtually every commercial version of Internet software within 6 months and enabled the Internet to scale from a small research network to today's worldwide infrastructure.

Building on this long history of accomplishments, principal investigators of the Computational and Technology Research program this year received recognition through numerous prizes, awards, and honors. A sample of the significant accomplishments produced by the program this year is given below.

- **1998 Gordon Bell Prize for Best Performance of a Supercomputing Application.** An international team of scientists including those from the Department's Oak Ridge and Lawrence Berkeley National Laboratories won the 1998 Gordon Bell prize for best performance of a supercomputing application for simulation of magnetism metal alloys. The team won the award for their modeling of 1,024 atoms of a metallic magnet. The calculation submitted by the team to the Gordon Bell Prize judges performed at 657 Gigaflops performance level (657 billion calculations per second); however, the team subsequently was able to raise the performance of their application to more than one teraflop (one trillion calculations per second.) In addition to supporting the winning entry, the CTR program partially supported the two other finalists.

- **1998 Fernbach Award**. Dr. Phillip Colella, a mathematician at DOE's Lawrence Berkeley National Laboratory, received the 1998 Sidney Fernbach Award at the SC98 conference in Orlando. Colella received the award for his "outstanding contribution in the application of high performance computers using innovative approaches." The Fernbach Award, created in memory of a computer scientist at DOE's Lawrence Livermore National Laboratory, is presented by the IEEE (Institute of Electrical and Electronics Engineers) Computer Society. Dr. Colella's research has focused on problems in Computational Fluid Dynamics and advanced techniques for the generation of the grids that form the basis for many scientific computations. His research has been applied to problems in the simulation of combustion devices as well as national security applications.

- **Simulation of Instabilities in Fluid Layers**. Many important physical systems can be described as layered fluids. For example, layers of oil float on salt water in geological structures. Even structures with layers of metal behave like layered fluids under high pressure and temperature. These types of systems develop very complex, unstable structures at the boundary between fluids. Computation of the characteristics of these boundaries is especially difficult, because their locations must be tracked accurately as part of the calculation. In addition, the most common formulations of the problem require many more grid points than are practical on even the largest computers. Applied mathematicians at the State University of New York, Stony Brook, in collaboration with physicists at Los Alamos National Laboratory (LANL) have discovered and implemented clever numerical schemes capable of following surfaces that can evolve into complicated shapes over time. They have used these new techniques to simulate this type of fluid instability under conditions where experiments are not possible, thus allowing design of devices for Defense Program's Advanced Strategic Computing Initiative and fusion energy applications. Modeling and simulation are absolutely vital since experimental data will never be available for conditions of importance to the designers.

- **Research in Optimization Impacts U.S. Industry**. Applied mathematicians at Rice University, working with engineers at Boeing, developed a software package for improving the manufacture of airplane components. The software combines new approaches to the optimization of systems having hundreds of thousands of parameters with research in the theory of control systems to enable engineers to optimize manufacturing processes. Previous design schemes for Boeing's production processes were based on simple "rules of thumb" that failed often in practice, and standard optimization packages were easily overwhelmed by the sheer size of the problems. This necessitated a complete rethinking of design and optimization strategies in order to accommodate industrial-size problems. The Rice researchers developed novel techniques in nonlinear optimization theory to correctly identify the underlying problems and provide feasible solutions. Boeing is using the software on its production line to lower costs and improve quality.

- **Law for Turbulent Stress Proved to be Invalid.** To design systems such as gas turbines, airplanes, or combustion devices where flow of gasses over physical structures is important, it is critical to be able to accurately describe the turbulent stress generated by the flow of the gas over the structure. Since the turn of the century, the standard methodology for calculating this was a simple mathematical model called the "law of the wall." This model, which is found in every engineering textbook, is based on relatively crude approximations that have been accepted without serious thought since the turn of the century. Mathematicians at Lawrence Berkeley National Laboratory (LBNL) have recently developed a rigorous mathematical basis for calculating turbulent wall effects and demonstrated that the "law of the wall" is not valid in general. They showed by careful mathematical analysis that the correct description of turbulent effects on solid objects requires a family of "scaling laws" rather than a single "law of the wall." This is difficult research since it must blend hard mathematical analysis with the proper physical insights from the fluid dynamics of turbulent flows. The discovery will have profound consequences in the engineering and design of airplanes, gas turbines, and other systems where controlling turbulence is critical to performance. The mathematical predictions of the Berkeley group have been verified recently by experiments at Princeton.

- **R&D 100 Award to Sandia Researchers.** Researchers at Sandia National Laboratories were awarded an R&D 100 award in FY 1998 for the Aztec software package. Aztec is a collection of very high performance software routines that run on the highest performance computers in the nation to solve important linear algebra problems such as solving systems of millions of linear equations. This type of linear algebra is at the core solving ordinary and partial differential equations on computers as well as many other types of scientific computations. Aztec grew out of the successful research programs at Sandia in numerical linear algebra and parallel programming techniques.

- **Advanced Computing Software Tools Enable Rapid Application Development.** One of the major challenges in modern high performance computing is to develop tools that enable scientists to quickly create computer software to solve scientific problems. Otherwise, chemists, materials scientists, and others would spend their entire effort creating software for computers that would be obsolete just as the applications were ready. The speed of change in the underlying computer architectures and the complexity of these computers and their operating systems makes this a major area of research. The Parallel Object-Oriented Methods and Applications (POOMA) Framework effort at LANL is one promising research approach to developing effective tools to help scientists in the disciplines develop software. In an early test with POOMA, a post-doc with no parallel programming experience developed computer software to solve a three dimensional fluid turbulence problem (including the tools to visualize the results while the program was running) in only six weeks rather than the 6-9 months required in similar efforts. POOMA is used extensively by two of the scientific applications pilots — the computational accelerator physics and numerical tokamak turbulence projects.

- **New Scientific Application Enabled By Interfacing Two Software Packages.** One research challenge facing advanced computing is to enable software developed by different teams to work together on massively parallel computers. Recently researchers at Argonne National Laboratory (ANL) and Lawrence Livermore National Laboratory (LLNL) have demonstrated that it is possible for well-designed components developed at different laboratories to be easily used together by providing each with a common interface. The latest generation of ordinary differential equation solvers for systems, whose behavior combines fine scale and large scale features, developed at LLNL has been interfaced with a large family of parallel algebraic solvers developed at ANL. This coupled software system has enabled several new applications. One of these is a collaboration of researchers at Louisiana Tech University and Oak Ridge National Laboratory

(ORNL) to develop a code for fully three-dimensional simulations of the dynamics of micro-structural interactions in materials. This code would not have been possible before the researchers had access to the coupled ANL-LLNL system.

- **Significant Speedups Achieved For Laue Crystallographic Analysis.** Advances in numerics coupled with work to enable software to run in parallel on many processors has enabled researchers in the "Supercomputer Solution of Massive Crystallographic and Microtomographic Structural Problems" scientific application pilot project to dramatically reduce the time required to analyze the data from Laue diffraction experiments. A typical illustration comes from progress made on a scientific applications project—diffraction patterns and intensities produced by x-rays passing through a crystal lattice are used to deduce information about molecular structures. The Laue diffraction technique is the most important tool in time-resolved crystallography, where structure data are captured rapidly to image the structure of a molecule at various stages of a reaction. Improved optimization and numerical techniques have been applied to the code that performs the complicated task of analyzing Laue diffraction data to obtain the structure of the molecules. The dramatic improvements in time-to-solution that have been made will significantly enhance experimental capabilities because runtimes are reduced from hours to minutes. These results are expected to be especially important for the 30 percent of users of DOE light sources who are involved in discovering protein structures.

- **ESnet Demonstrates Priority Service For Internet Traffic.** Scientists at two national laboratories successfully selected marked Internet traffic for priority service over unmarked traffic in a cross-country demonstration. This demonstration is a key milestone in the development of a broad set of capabilities called "differentiated services," which are required for the Internet to be able to give different levels of service on demand to network customers. The demonstration of such capabilities for production-mode scientific research between Lawrence Berkeley National Laboratory and Argonne National Laboratory across the ESnet paves the way for more reliable and constant connectivity via priority bandwidth on the Internet. Achieving this improved level of service is essential to the work of the Department, which is pioneering the use of various technologies to allow scientists at more than 30 DOE national labs to share access to some of the Nation's most advanced research facilities. The complex interactions between software on computers, network hardware such as routers, and telecommunications equipment operated by commercial carriers make this a difficult research problem. In addition, all of these components must be capable of efficiently scaling up to operate across the worldwide Internet which processes tens of billions of packets a month.

- **Collaboratory Tool Attracts Users.** The electronic notebook collaboratory tool project has been so well received that over eighty groups across the country have adopted the prototype electronic notebook. Some are DOE projects, but many are from outside the Department including pharmaceuticals, chemical processing and medicine. The electronic notebook is valuable to researchers because: it can be shared by a group of researchers; it can be accessed remotely; it cannot be misplaced, lost, or accidentally destroyed (if backed up); it is easy to incorporate computer files, plots, etc.; notarization and authentication are possible; it can easily be searched for information; it can include multimedia; and it can include hyperlinks to other information. In order to enable these capabilities, the electronic notebook project has had to overcome a number of challenges including the development of new technologies for describing types of data such as experimental protocols and experimental devices, which are not well treated by traditional techniques.

Laboratory Technology Research

- In FY 1999, the SC single-purpose laboratories (Fermi National Accelerator Laboratory, Thomas Jefferson National Accelerator Facility, Princeton Plasma Physics Laboratory, and Stanford Linear Accelerator Center) and Ames Laboratory were reinstated into the Laboratory Technology Research (LTR) subprogram, thus restoring the subprogram to its original participants. LTR now provides coverage to more regions of the country where small businesses, in particular, can take advantage of the resources at SC laboratories. LTR capabilities have been enhanced for cost-shared partnerships in fusion energy sciences, high energy physics, nuclear physics, materials sciences, chemical sciences, structural biology, and other disciplines.
- The LTR subprogram received two R&D-100 awards in 1997 and three R&D-100 awards in 1998. The 1998 awards were given to:
  - Oak Ridge National Laboratory, in collaboration with the Society of Exploration Geophysicists, for “Advanced Computational Tools for 3-D Seismic Analysis.” This research was cosponsored by DOE’s Office of Fossil Energy.
  - Argonne National Laboratory in collaboration with Front Edge Technology of Baldwin Park, CA, Stirling Motors of Ann Arbor, MI, and Diesel Technology of Wyoming, MI, for “Near-Frictionless Carbon Coatings.”
  - Argonne National Laboratory, in collaboration with Commonwealth Edison of Chicago, IL, for “Combined Expert System/Neural Network for Process Fault Diagnostics.”

Advanced Energy Projects

- The Advanced Energy Projects subprogram will be terminated in FY 2000.

Scientific Facilities Utilization

The CTR program request includes $27,500,000 in FY 2000 to support the NERSC Center. This investment will provide research time for about 3,500 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 the Nation’s largest, premier unclassified computing center, 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.

Funding of Contractor Security Clearances

In FY 1999, the Department divided the responsibility for obtaining and maintaining security clearances. The Office of Security Affairs, which was responsible for funding all Federal and contractor employee clearances, now pays only for clearances of Federal employees, both at headquarters and the field. Program organizations are now responsible for contractor clearances, using program funds. This change in policy enables program managers to make the decisions as to how many and what level clearances are necessary for effective program execution. In this way, it is hoped that any backlog of essential clearances which are impeding program success can be cleared up by those managers most directly involved. The Office of Science is budgeting $46,000 and $47,000 for estimated contractor security clearances in FY 1999 and FY 2000, respectively, within this decision unit.

Scientific Simulation Initiative

The CTR program also leads DOE’s Scientific Simulation Initiative (SSI), an integrated effort bringing together computational and communication resources, focused research in scientific disciplines, and research in computer science and other enabling technologies to solve the complex problems that characterize DOE’s scientific research needs. The SSI is a significant component of the broader President’s Information Technology Research Initiative (ITRI) which responds to the recommendations of the President’s Information Technology Advisory Committee (PITAC). The SSI couples research in advanced scientific applications in the programs of the Office of Science (SC) with research in computer science and enabling technologies and advanced computing and communications facilities. It is a joint program between CTR and the other program offices in SC. The overview of the integrated program is given here; however, the specific contributions of the other programs (Basic Energy Sciences and Biological and Environmental Research) are included in their budgets. This initiative will build on and benefit from the demonstrated capabilities brought about by the Department’s Accelerated Strategic Computing Initiative (ASCI) that have made it possible to obtain computational capabilities 100 times faster than currently in common use. The mission of this proposed effort is to develop further and to employ the emerging generation of very high performance computers as major tools for scientific enquiry. These resources will revolutionize our approach to solving complex problems in energy, environment, fundamental research, and technology development as well as stimulate our national system of innovation. The goal of the SSI is to:

> *“Revolutionize our ability to solve scientific problems of extraordinary complexity and to apply these resources to scientific problems relevant to the Department’s mission through the exploitation of the emerging power of exceptional computational capabilities.”*

As has been discussed earlier in this budget, scientific research has long been characterized as the interplay of theory and experimentation. Over the last half century the emergence of computers for solving complex mathematical problems and for analyzing large sets of data has introduced a third activity to complement both theory and experimentation. Defined as simulation and modeling, but also encompassing a broad range of data analyses, the application of computational resources to complex scientific problems has made increasingly important contributions to scientific discovery and understanding as those resources have grown more powerful. The SSI builds on DOE’s 50 year history of transforming advances in information technology into tools for scientific discovery to accelerate this process through a partnership with DOE laboratories, the academic community and industry.

This program began with workshops focused on the scientific opportunities that could result from access to multi-teraflop computing. All of the programs in the Office of Science made convincing cases that access to computing at this level would open new areas of research and enable solutions to new classes of problems. These workshops culminated in a workshop on advanced scientific and engineering computing jointly sponsored by DOE and the National Science Foundation at the National Academy of Sciences in July, 1998. This workshop validated the opportunity for scientific discovery and advanced engineering afforded by terascale computing. Throughout this document, the term teraflop is used in discussing the computational resources under consideration. One teraflop is one trillion (10^{12}) floating point operations per second. For comparison, desktop personal computers are generally capable of one hundred thousand (10^{5}) floating point operations per second. Current commercially available supercomputers are capable of nearly half teraflop performance.

Based on the input from these workshops, a plan for the SSI was developed to realize the goals of the SSI. It is a balanced plan of research which includes advanced computing and communications facilities, such as terascale computers and very high performance networks, computer science and enabling technology research to make these facilities useful, and investments in scientific disciplines to enable development of the advanced applications that will be required. The strategy for SSI applications has two components. The first was a selection of two mission critical application areas: global systems and combustion systems. In both of these cases, the initial SSI planning workshops identified significant opportunities for advanced computation to dramatically advance the state of the art with important impact on the nation's ability to respond effectively to issues such as the effect of greenhouse gasses on global climate and development of internal combustion engines able to meet societal goals for efficiency and pollution control. These two initial applications play a significant role in defining the requirements for computer and communications facilities as well as the computer science and enabling technology which will be required. These two applications will also play an important role in testing, integrating, and debugging both hardware and software components. The second component of the SSI applications strategy is an open, peer reviewed competition among basic science disciplines to select a small number for initial inclusion in the initiative and access to SSI computer and communications facilities.

The Scientific Simulation Initiative Builds on ASCI

One reason the Department is prepared to undertake this initiative at this time is that the proposed effort builds upon the DOE Defense Program (DP) Accelerated Strategic Computing Initiative (ASCI) which was launched in 1996 with a focus on multi-teraflop scale computing to meet the imposing challenges posed by Science Based Stockpile Stewardship.

Meeting the goals set for ASCI requires computers with capabilities exceeding those available today by a factor of one thousand and the ASCI program is following an aggressive hardware and software technology plan that will achieve the development and use of 100 teraflop computers by 2004. As a result of the ASCI effort, much more powerful computer systems, designed for full simulation of all scientific aspects of nuclear stockpile stewardship, have been developed. ASCI continues to develop the computational infrastructure needed to make effective use of these systems, including large-scale scientific data storage capabilities, computational grids for high-speed communication, software development systems for massively parallel computer systems, high-performance visualization systems, etc. These advances, although focused on DOE's stockpile stewardship responsibility, presage the development of a comprehensive simulation capability for a whole range of scientific problems.

A Critical Element in the President’s Initiative in Information Technology

Although the SSI is focused on revolutionary uses of computing as a tool for science, many of the investments in computer science and enabling technology will contribute to and benefit from the broader ITRI initiative. PITAC recommended significant increases in support for basic research in: Software, Scalable Information Infrastructure, High End Computing, and Socio-Economic and Workforce Impacts, as well as support of research projects of broader scope and visionary “Expeditions to the 21st century” to explore new ways that computing could benefit our world.

The SSI Computer Science and Enabling Technology (CSET) program has four components: Algorithms, Models, Methods and Libraries; Problem Solving Environments and Tools; Distributed Computing and Collaboration Technology; and Visualization and Data Management. The underlying strategy in all of these areas is based on taking advantage of the most recent work in software components and extending it to enable it to function at very high performance. In addition, the realities of the hardware platforms that will be available at SSI performance levels requires that the components be fault tolerant and that techniques be developed to enable the assembly of components into reliable fault tolerant assemblies. Thus, with respect to the PITAC recommendation on Software development, we expect SSI to be a significant program for fundamental research in software development methods and component technologies and to begin the development of a national repository of software components.

In addition, challenges facing SSI in the area of data management and visualization will require significant new research in the areas of human computer interfaces and interaction to enable researchers to navigate and extract scientific knowledge from petabyte-scale data archives resulting from SSI scale simulations and the coming generation of large experimental facilities. Both the SSI planning workshops as well as the ASCI Visualization Corridors workshops identified this as a critical area that would perhaps be an “Expedition to the 21st Century.”

However, there are areas of software research included in the PITAC report that SSI will not cover, particularly software for systems such as air traffic control.

With regard to the PITAC recommendation on a Scalable Information Infrastructure, SSI will have a more modest impact because the problem for The SSI is primarily providing very high performance access for a modest number of people rather than providing access to billions of users worldwide.

With regard to the PITAC recommendation on High End Computing, SSI will have a major impact on R&D to improve the performance of high end computers, acquiring high end systems to support scientific research and moving towards petaop (1000 teraflop) systems by 2010. However, while SSI is a part of the solution to the problem of access to high end computing for science and engineering research and will develop many of the technologies needed to make this possible, the computing systems proposed are scaled to support a modest number of focused scientific projects, which will be collaborative efforts of large teams of disciplinary scientists, computer and computational scientists and mathematicians.

With regard to the PITAC recommendation on Socio-Economic and Workforce Impacts, The SSI will increase research funding in relevant areas of IT and will develop strategies to retrain existing scientists and IT workers as well as train new undergraduate and graduate students in these disciplines.

In addition to the research contributions that SSI will make to the broader ITRI, SSI facilities will provide crucial testbeds for testing many of the ideas that emerge from basic research in information technology.

As stated at the outset the SSI is an integrated initiative from the Office of Science. The descriptions of the various components, corresponding to the goals are included in the budgets of the responsible offices. The SSI has five principal objectives:

- **Revolutionize scientific research by the application of teraflop computational resources.** Whereas the scientific accomplishments of this century have resulted in seeking and understanding the fundamental laws that govern our physical universe, the science of the coming century will be characterized by synthesis of this knowledge into predictive capabilities for understanding and solving a wide range of scientific problems, many with practical consequences. In this endeavor, the computer will be a primary instrument of scientific discovery. Many areas of scientific inquiry critical to the Department's mission will be advanced dramatically with access to teraflop scale computing including but not limited to materials sciences, structural genomics, high energy and nuclear physics, subsurface flow, and fusion energy research.
- **Discover, develop and deploy crosscutting computer science and applied mathematics technology.** The practical and intellectual challenges to making effective use of terascale computers require the development of a terascale technology base in software, networking, data management and visualization, communications, and operating environments to enable scientists to make effective use of the simulation infrastructure.
- **Establish a national terascale distributed scientific simulation infrastructure.** This network of terascale computers, ultra-high speed communications, science centers, and support centers in academia, the national laboratories, and industry will provide the advanced computing testbeds which enable the accomplishment of the first three objectives.
- **Understand, model, and predict the effects on the earth's global environment of atmospheric greenhouse gas emissions, with an emphasis on carbon dioxide.** Through the use of teraflop scale computers, accelerate progress in general circulation model development and application to reduce substantially the uncertainties in decade-to-century model-based projections of global environmental change and to increase the utility of such projections to the broad research community. Current models of global systems cannot presently achieve regional specification in global environmental change projections with the requisite accuracy and reliability needed to support national and international energy and environmental policies. Work towards this objective will also play a significant role in defining, testing, and integrating the SSI facilities and SSI crosscutting computer science and applied mathematics technologies into tools for scientific dsicovery.
- **Understand, model, and predict the behavior and properties of combustion processes and devices.** With teraflop scale computing resources and a concerted research program in combustion modeling, develop a new generation of combustion modeling tools for accelerated design of combustion devices meeting national goals of emission reduction and energy conservation. Through provision of a comprehensive understanding of the details of combustion, design engineers will have the computational tools to predict the chemical outcome of combustion processes with practical reliability, thus avoiding a time consuming, trial and error approach to the

design of combustion devices (gas turbines, internal combustion engines, etc.) Work towards this objective will also play a significant role in defining, testing, and integrating the SSI facilities and SSI crosscutting computer science and applied mathematics technologies into tools for scientific discovery.

This initiative is a part of the broader President's ITRI. The Department's primary focus is advancing science through terascale computing. To accomplish these goals, DOE will partner with other agencies, particularly the National Science Foundation, to leverage the strengths of both agencies. In the case of research in global systems the research is already part of a broad interagency effort coordinated under the U.S. Global Change Research Program (USGCRP). Computer science and enabling technologies activities have been coordinated between agencies through the Computing, Information, and Communications R&D Subcommittee of the NSTC for a number of years. As a part of this initiative, closer ties will be established between DOE-funded activities and activities funded by other agencies. The development of enabling technologies for SSI will also be coordinated with the development of related technologies for ASCI through a joint CTR - ASCI research management committee. Where appropriate, joint interagency programs in other scientific disciplines may be established.

The CTR budget includes descriptions of the SSI computational science and enabling technology, as well as advanced computing and communications facilities elements of the SSI. In addition, the CTR budget includes funding for the competitive, peer reviewed selection of a small number of basic science applications to complement the two larger integrated applications efforts. The description of the Global Systems element of the SSI is included in the Biological and Environmental Research (BER) budget and the description of the Combustion Systems element is included in the Basic Energy Sciences (BES) budget. For reference, a high level summary of the proposed budget for the entire SSI initiative is given in the table below. These amounts will be reduced for the required SBIR/STTR assessments of 2.65%.

Objective	Program	FY 2000 Request
Earth's Global Environment	Biological and Environmental Research	$10 million
Combustion Systems	Basic Energy Sciences	$7 million
Basic Science Applications	Computational and Technology Research	$6 million
Computer Science and Enabling Technology	Computational and Technology Research	$16 million
SSI Facility Operations	Computational and Technology Research	$30 million
Staffing Resources	Science Program Direction	$1 million
Total SSI		$70 million

Funding Profile

(dollars in thousands)

	FY 1998 Current Appropriation	FY 1999 Original Appropriation	FY 1999 Adjustments	FY 1999 Current Appropriation	FY 2000 Request
Computational and Technology Research					
Mathematical, Information, and Computational Sciences	124,026	139,300	-466	138,834	184,575
Laboratory Technology Research . . .	15,379	16,200	-58	16,142	14,300
Advanced Energy Projects	7,374	2,500	-5	2,495	0
Subtotal, Computational and Technology Research	146,779	158,000	-529	157,471	198,875
Use of Prior Year Balances	-1,714[a]	-1,573[a]	0	-1,573[a]	0
General Reduction	0	-529	+529	0	0
Total, Computational and Technology Research	145,065[b]	155,898	0	155,898	198,875

Public Law Authorization:

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

[a] Share of Science general reduction for use of prior year balances assigned to this program. The total general reduction is applied at the appropriation level.

[b] Excludes $3,582,000 which has been transferred to the SBIR program and $215,000 which has been transferred to the STTR program.

Funding by Site

(dollars in thousands)

	FY 1998	FY 1999	FY 2000	$ Change	%Change
Albuquerque Operations Office					
Los Alamos National Laboratory	14,614	13,034	10,894	-2,140	-16.4%
National Renewable Energy Laboratory	498	127	0	-127	-100.0%
Sandia National Laboratories	5,232	5,293	3,779	-1,514	-28.6%
Total, Albuquerque Operations Office	20,344	18,454	14,673	-3,781	-20.5%
Chicago Operations Office					
Ames Laboratory	2,290	1,939	1,490	-449	-23.2%
Argonne National Laboratory	16,869	15,430	13,176	-2,254	-14.6%
Fermi National Accelerator Laboratory	100	50	332	+282	+564.0%
Brookhaven National Laboratory	2,843	1,457	2,589	+1,132	+77.7%
Princeton Plasma Physics Laboratory	90	121	332	+211	+174.4%
Total, Chicago Operations Office	22,192	18,997	17,919	-1,078	-5.7%
Oakland Operations Office					
Lawrence Berkeley National Laboratory	57,916	53,938	49,377	-4,561	-8.4%
Lawrence Livermore National Laboratory . . .	2,755	2,940	640	-2,300	-78.2%
Stanford Linear Accelerator Center	980	357	782	+425	+119.0%
Total, Oakland Operations Office	61,651	57,235	50,799	-6,436	-11.2%
Oak Ridge Operations Office					
Oak Ridge Inst. For Science & Education. . .	335	0	244	+244	+100.0%
Oak Ridge National Laboratory	19,434	10,415	6,876	-3,539	-34.0%
Thomas Jefferson National Accelerator Facility	190	100	283	+183	+183.0%
Total, Oak Ridge Operations Office	19,959	10,515	7,403	-3,112	-29.6%
Richland Operations Office					
Pacific Northwest National Laboratory	4,188	3,238	3,584	+346	+10.7%
All Other Sites[a]	18,445	49,032	104,497	+55,465	+113.1%
Subtotal, Computational and Technology Research	146,779	157,471	198,875	+41,404	+26.3%
Use of Prior Year Balances	-1,714[b]	-1,573 [b]	0	+1,573[b]	+100.0%
Total, Computational and Technology Research	145,065[c]	155,898	198,875	+42,977	+27.6%

[a] Funding provided to laboratories, universities, industry, other Federal agencies and other miscellaneous contractors.

b Share of Science general reduction for use of prior year balances assigned to this program. The total general reduction is applied at the appropriation level.

c Excludes $3,582,000 which has been transferred to the SBIR program and $215,000 which has been transferred to the STTR program.

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 also 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 advanced computing research facility (ACRF) at ANL focuses on advanced computers in the IBM-SP family of technologies as well as the interaction of those architectures with advanced visualization hardware. The LTR subprogram at ANL also 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 facility. The advanced computing research facility (ACRF) 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 also 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 a 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 facility. The Advanced Computing Research Facility (ACRF) 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 CTR 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 also conducts research motivated by practical energy payoffs through cost-shared collaborations with industry. The areas emphasized in the research portfolio are high temperature superconducting wire, 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 also conducts research motivated by practical energy payoffs through cost-shared collaborations with industry. The areas emphasized in the research portfolio are mathematical simulation of glass production, interactions of biological polymers with model surfaces, and characterization of micro-organisms 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 Tonapah, 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 CTR program funds research at 71 colleges/universities located in 24 states. This line also includes funding of research awaiting distribution pending completion of peer review results.

Mathematical, Information, and Computational Sciences

Mission Supporting Goals and Objectives

The Mathematical, Information, and Computational Sciences (MICS) subprogram supports advanced computing research — applied mathematics, high performance computing, and networking — and operates supercomputer and associated facilities that are available to researchers 24 hours a day, 365 days a year. The combination of support for fundamental research, computational and networking tools development, and high-performance computing facilities provides scientists with the capabilities to analyze, model, simulate, and — most importantly — predict complex phenomena of importance to the Office of Science and to the Department of Energy.

Examples of the complex problems addressed by the Office of Science and the Department of Energy are: climate modeling, including the effects of greenhouse gases on global climate change; the combustion process, including the simulation of combustion in devices such as diesel engines; the subsurface transport of pollutants; the rational design of new materials to produce, for example, new alloys, superconductors, polymers, and catalysts; the effects of aging on the nuclear stockpile; and the analysis of vast amounts of real data from experiments at the Office of Science facilities for high energy physics, nuclear physics, materials sciences, chemical sciences, structural biology, and other disciplines.

For many such problems, traditional theoretical and experimental approaches may not be suitable; theory may be inadequate to handle the complexity, and experiments may not be feasible, because they are too dangerous, too expensive, or simply impossible owing to the length and time scales involved. Furthermore, experiments at the Office of Science facilities may generate millions of gigabytes (petabytes) of data per year (which would fill the disk drives of millions of today's personal computers) presenting significant computational and communications challenges in analyzing and extracting understanding from the data.

The increases in computational and communications capabilities achieved over the past two decades have made simulation a third branch of science that now complements theory and experiment and allows a new approach to previously intractable problems. However, as computational and communications capabilities have increased, so too have the challenges associated with effectively using these capabilities. The rate of increase in these capabilities itself poses formidable challenges. Significant hardware changes can occur as often as every 18 months; these, in turn, may require completely new approaches to computing software. This has already changed the way in which computation is performed in scientific disciplines such as materials science, biology, and fusion energy. Teams involving scientists from the disciplines as well as mathematicians; computer scientists; and experts in computer graphics, data management, and advanced computer networks are now required to address scientific problems that make use of the most modern computer and communications capabilities. This situation is quite different from that of even ten years ago when individual scientists could code and perform calculations with little or no support from others.

These same increases in computational and communications capabilities are having a second important impact on the way science is conducted. It is now possible for large geographically distributed teams to easily and effectively collaborate using major experimental facilities, computational resources, and data resources. The name given to such optimized linkages among geographically distributed resources is "collaboratories." In the coming years, they are expected to play an increasingly important role in the Nation's scientific enterprise.

In order to enable scientists in the Office of Science to make effective use of these capabilities, the MICS subprogram supports research in three areas:

- Applied Mathematics. Research on the underlying mathematical understanding and 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. 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 which 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. 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. The development of high speed communications and collaboration technologies will allow scientists to view, compare, and integrate data from multiple sources remotely.

In all of these areas — applied mathematics, computer science, and networking — the requirements far exceed the current state-of-the-art; furthermore, the requirements far exceed the tools that the commercial marketplace will deliver. In addition to these fundamental research efforts, the MICS subprogram takes the results of these efforts and forms partnerships with users in scientific disciplines to validate the usefulness of the ideas and to develop them into tools. MICS also provides the advanced computing and communications facilities that enable scientists to use these tools.

MICS provides two types of advanced computing and communications facilities. The first type of facility enables scientists to use the tools developed. Examples are NERSC and ESnet. The second type of facility is itself a research project. The principal current examples of this type of facility are the Advanced Computing Research Facilities (ACRFs) at Argonne National Laboratory, Lawrence Berkeley National Laboratory, and Los Alamos National Laboratory that represent the evolution of the High Performance Computing Research Centers and were established as part of the High Performance Computing and Communications initiative. The ACRFs combine research in computer software and hardware with targeted applications. One of the major issues that these facilities attempt to address is how different choices in computer architecture affect the ability of a system to scale to very large numbers of processors and very high performance. In order to address these issues, computers at a scale that push the state-of-the-art must be sited at the ACRFs. These computers enable research in computers and computing, but they are not sufficiently mature or robust for production computing by large numbers of users.

Partnerships with the scientific disciplines are an important management philosophy in the MICS subprogram. Partnerships with the scientific disciplines are critical, because they test 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. Finally, to develop future generations of scientists with the breadth of skills required to be effective both in advanced computing research and in interacting with disciplinary sciences, MICS supports the Computational Science Graduate Fellowship program.

The MICS subprogram includes the Department's participation in the President's Next Generation Internet (NGI) Initiative. This initiative will create the foundation for more powerful and versatile networks of the 21st century, just as previous federal investments in information technology R&D created the foundation for today's Internet. DOE's participation in this initiative is focused on network requirements that will enable data-intensive scientific research not now possible because of network limitations. It is anticipated that the results and “spinoffs” of this research, after testing and prototyping by the scientific community, will impact broad commercial use of networks.

DOE's NGI research program is focused on discovering, understanding, developing, testing and validating the networking technologies needed to enable wide area, data intensive and collaborative computing that is not currently possible. This program will integrate scientists working on fundamental research in applied mathematics, computer science, and networking with scientists working on DOE applications to develop new ways to link scientists with DOE's major scientific user facilities and computational centers. Such research is needed to enable effective use of petabyte/year High Energy and Nuclear Physics facilities such as the Relativistic Heavy Ion Collider (RHIC); to provide remote visualization of terabyte to petabyte data sets from computational simulation; to develop advanced collaboratories; and to enable effective remote access to tomorrow's advanced scientific computers. These applications share two important characteristics. They all involve extremely large data sets, and they all require that scientists be able to interact with the data in (nearly) real time. Current network technology limitations significantly limit our ability to address either of these characteristics.

The NGI activities are critical for DOE's fundamental science research. For example, using the current Internet, it would take about 2,500 hours to transmit one day's data from RHIC to one remote site for analysis. Typical RHIC experimental collaborations involve thousands of scientists and hundreds of institutions spread across the country and the world. Management of scientific data is further complicated because the data must be managed in large units. This is very different than data management in the commercial sector. For example, a standard telecommunications billing record is only 180 bytes long; a large web page is 500 kilobytes; but the data from a global climate model at any given time during the simulation (a single time step) may be greater than a gigabyte or 2,000 times as large as that of a large web page!

The technologies developed for commercial network traffic are simply inadequate for scientific network traffic. Significant *research* is needed to enable today's commercial networks to be used for scientific data. This research must include advanced protocols, special operating system services for very high speed, and very advanced network control. For example, a coast to coast gigabit/second network may have as much as a gigabit of data in transit at any one time. If this is scientific data this might be one packet. Current protocols would require that all of the switches along the way be able to buffer a gigabit of data at very high rate, which is not possible with today's switches. Therefore research is required. The issues related to managing a single gigabit packet versus 2,000 500 kilobit packets are very different.

In addition to the complications posed by the transmission of large data sets, there are additional complications that result from the many different kinds of network devices, network-attached devices, and services that need to be integrated together. 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. This type of integration, as well as the issues of improving the performance of the individual components, all require significant research because the issues are currently not well understood. Indeed, the first identification of many of these issues is the result of previous work in collaboratories and visualization supported by this subprogram. Thus DOE's participation in the NGI builds on previous results of the MICS subprogram to address critical issues in the network. Furthermore, the differences between the requirements of commercial networks and networks for scientific research require DOE to conduct this research because these tools and technologies will not be developed by commercial R&D. However, there will be tremendous “spinoff” benefits to the Internet in general after the research is completed and the results enhance commercial networks.

Scientific Simulation Initiative

This budget also includes the MICS subprogram’s contribution to DOE’s Scientific Simulation Initiative (SSI), an integrated effort bringing together computational and communication resources, focused research in scientific disciplines, and research in computer science and other enabling technologies to solve the complex problems that characterize DOE’s scientific research needs. The SSI couples research in advanced scientific applications in the programs of the Office of Science with research in computer science and enabling technologies and advanced computing and communications facilities. It is a joint program between CTR and the other program offices in SC.

It is a requirement for the success of this program to use the expertise that exists in national laboratories, universities, and industry and, where appropriate, formation of multidisciplinary teams of researchers. Participants from various institutions will be encouraged through open peer reviewed competitions. Also, strong partnerships with the DOE ASCI program and complementary programs in other agencies will be essential. The overview of the integrated program is given in the Program Mission statement of the CTR budget; however, the specific contributions of the MICS subprogram are described below.

MICS contributes to the SSI in three ways:

Management of the Computer Science and Enabling Technology (CSET) component of the SSI;

Management of the SSI Advanced Computing and Communications Facilities; and

Management of the peer reviewed selection process for the basic science application efforts to be initiated in FY 2000.

The goal of the CSET component of the SSI is to develop needed software systems, to deploy these systems into the DOE computing infrastructure, to support users, and to conduct the needed research to address future problems, all with close involvement of the applications scientists and the computer systems providers. CSET teams will also work closely with the vendor community to ensure that needed capabilities are incorporated into products, systems, and services. Working with the university research community, CSET teams will ensure a steady flow of young researchers into the field and will form long-term collaborations in the areas needed by scientific simulation. CSET teams will also establish natural partnerships with the DOE ASCI program and with the NSF Partnerships for Advanced Computational Infrastructure (PACI) program, which are addressing complementary goals. These partnerships could result in the joint funding of research and development activities and joint deployment of infrastructure capabilities. CSET is a critical crosscutting activity for the Scientific Simulation Initiative — one that will enable the program to accomplish far more than if applications teams were required to develop all the necessary software and infrastructure themselves. Through the sharing of common tools and technologies, the CSET activities will dramatically improve the cost-effectiveness of the SSI. In order to accomplish these goals, CSET must support a vigorous research program as well as significant development and software deployment activities to address the following critical issues.

First, despite considerable progress during the past ten years in making massively parallel computer systems usable for applications, much remains to be done. Computer systems targeted for SSI will scale from approximately 1,000 nodes today to 5,000 or 10,000 nodes for systems in 2004. This scaling will require substantial improvements in parallel computing tools, parallel I/O (input/output) systems, data management, algorithms, and program libraries. Not only will the number of compute nodes increase, but the nodes themselves will become more complex, as systems designers are forced to introduce more layers of memory hierarchy to maintain performance and develop new hardware features to support rapid communication and synchronization. The end result five years from now will be hardware systems that, while showing their roots in today's systems, will be substantially different and substantially more complex and therefore more challenging to exploit for high performance.

Second, because applications software systems typically outlive hardware by an order of magnitude (i.e., computer hardware typically has a useful lifetime of three years, whereas large-scale scientific codes last one to two decades), new software must be designed for machines not only for the near future but for the next decade. This effort will require close working relationships between code developers from the applications areas and computer scientists who are involved in next-generation systems and architecture design.

Third, the applications themselves are getting more complex. They are incorporating more sophisticated physical models, are using advanced numerical techniques (e.g., adaptive mesh refinement, unstructured and implicit solvers, and hierarchical methods), and are beginning to be combined into large-scale "simulation systems" that include the linkage of two or more previously stand alone models (e.g., ocean-atmosphere-biosphere or fluid-structures-chemistry). Furthermore, in some cases these models are being driven by an optimization environment to enable them to address policy questions.

Finally, in addition to rapidly changing hardware and applications targets, the user and development community is also evolving. No longer can groups afford to work only with local collaborators or to use only locally available hardware and software resources. The development and user environments of the near future will need to enable ubiquitous collaboration and distributed computing capabilities to address these complexities. As the applications codes themselves become more interdisciplinary, so too will the teams that write them. As the need to incorporate deep knowledge of computer science techniques into future codes increases, so will the need to form long-term collaborations among applications scientists and computer scientists and mathematicians. These unavoidable trends mean that the problem-solving environments of the future will need to support human-to-human interactions as easily as we support computer to computer interactions today.

The management of SSI Computing and Communications facilities also poses challenges. Specifying, developing procurement specifications for, and finally delivering installed SSI facilities require great skill and experience. All of the facilities must work together because the performance of the system is determined by the slowest component. The SSI facilities must provide the linkage between the computing environment developed for the user and the systems management and systems integration required to support these environments. The SSI requires three distinct types of computing facilities: computing, networking and associated supporting hardware.

DOE will establish an open solicitation process that seeks the widest participation in establishing its terascale computing infrastructure, including competition among national laboratories, universities, and industry, based on their qualifications. The sites for the major teraflop computers will be selected through peer-reviewed competition. A number of considerations are important in selecting organizations to manage and locations at which to site SSI facilities. One of the most important is an expertise to perform the necessary computing systems integration into the existing nationwide DOE Office of Science computing infrastructure. These large-scale systems have requirements for scalable systems management that will enable relatively small systems administration teams to manage systems with 1000's of nodes and 10,000's of processors. These systems also need to be tightly integrated with data storage environments, mass stores, visualization environments, and distributed computing frameworks. Other considerations include incremental site preparation costs, the cost of connecting at very high speed to the networking infrastructure, and site financial leverage in providing operational support to the facility.

Finally, MICS will manage the open peer reviewed competition among basic science disciplines to select a small number for initial inclusion in the initiative. These applications will be chosen from throughout the portfolio of the Office of Science based on the scientific and technical merit of the proposals; the importance to DOE missions; and the readiness of the area and the associated scientific communities to move quickly to terascale computers.

Funding Schedule

(dollars in thousands)

	FY 1998	FY 1999	FY 2000	$ Change	% Change
Mathematical, Computational, and Computer Sciences Research	47,857	51,560	71,407	+19,847	+38.5%
Advanced Computation, Communications Research and Associated Activities	76,169	83,800	108,682	+24,882	+29.7%
SBIR/STTR	0	3,474	4,486	+1,012	+29.1%
Total, Mathematical, Information, and Computational Sciences	124,026	138,834	184,575	+45,741	+32.9%

Detailed Program Justification

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000

Mathematical, Computational, and Computer Sciences Research

- **Applied Mathematics:** Research on the underlying mathematical understanding and numerical algorithms to enable effective description and prediction of physical systems. This activity supports research at DOE laboratories, universities, and private companies to provide the DOE laboratory community and the wider national scientific and engineering communities with the most powerful and effective mathematical and computational tools for modeling, analyzing, and simulating complex phenomena in the core disciplinary and technology areas of DOE. Laboratory, academic, and industry researchers supported by the program work at laboratory sites with DOE mentors. To accomplish its goals, the program supports research in a number of areas including: **Mathematical Physics** including string theory, superstring theory, geometry of space-time, and quantum effects; **Ordinary and Partial Differential Equations** including numerical methods, high performance algorithms, massively parallel algorithms, distributed computing, novel gridding schemes, numerical linear algebra, iterative methods, sparse solvers, and dense solvers; **Control Theory** including differential-algebraic systems, order reduction, queuing theory; **Shock Wave Theory** including hyperbolic systems, multipole expansions, mixed elliptic-hyperbolic problems, and wavelet transforms; **Fluid Dynamics** including compressible, incompressible, and reacting flows, turbulence modeling, and multiphase flows; **Dynamical Systems** including chaos-theory and control, and bifurcation theory; **Programming and Optimization** including linear and nonlinear programming, interior-point methods, and discrete and integer programming; and **Geometric and Symbolic Computing** including minimal surfaces and automated theorem proving. The FY 2000 budget includes the continuation of work initiated in FY 1999

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
to develop the mathematical basis for modeling and simulating complex stochastic phenomena of the type that arise in vital DOE areas such as global climate modeling, environmental remediation, and stockpile stewardship. CTR will provide opportunities for college faculty and students to spend time at DOE laboratories, to participate in world-class research projects. Faculty/Student Science Teams will visit DOE labs during the academic/summer semesters, be involved in conducting research, writing proposals, utilizing technology and pursuing technical or scientific careers. Primary goals of the Science Teams are to build long-term partnerships among DOE laboratories and provide faculty/students with a deeper understanding of DOE science associated needs for research and development. Funds will be provided to pay for faculty/student stipends, travel, housing, and subsidizing laboratory scientists' time for this activity ($1,947,000).	23,576	25,232	27,179
■ **Computer Science:** Research in computer science to enable large scientific applications. 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. The first area includes research in protocols for message passing and parallel input/output (IO) as well as tools to monitor the performance of scientific applications. The second area includes research in effective techniques for retrieving data with complex internal structure from massive data archives that may be geographically distributed as well as advanced techniques for visualizing very large scale scientific data.	14,000	14,000	14,000

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
■ **Advanced Computing Software Tools:** R&D to develop the results of fundamental research in applied mathematics and computer science into an integrated set of tools that can be used by scientists in various disciplines to develop high performance scientific applications (e.g., to simulate the behavior of materials) that will have a useful life spanning many generations of computer hardware. These tools will include capabilities for representing complex geometries, solving diverse numerical equations, simplifying multi-language parallel execution, evaluating and enhancing code performance, and dynamically steering calculations during execution. This effort began as a part of the DOE2000 initiative.	5,000	5,000	5,000
■ **Scientific Applications Pilot Projects:** R&D to apply computational techniques and tools developed in the Advanced Computing Software Tools effort to basic research problems in order to test the usefulness of current advanced computing research, transfer the results of this research to the scientific disciplines, and help define promising areas for future research. Examples of pilot projects include: research in simulations of the earth's climate; research in the fundamental structure and properties of magnetic materials; creation of advanced tools to understand the chemistry of actinides; and partnerships with experimental disciplines such as high energy and nuclear physics, human genomics, and crystallography to improve the ability of these disciplines to manage and analyze the petabytes of data (that would fill the hard disks of millions of today's PC's) produced by their experiments and simulations. These efforts represent the evolution of the Grand Challenge projects that were initiated as part of DOE's component of the Federal High Performance Computing and Communications program, which started in FY 1991. These projects will now be phased out in an orderly manner	5,281	7,328	3,785

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
■ **Scientific Simulation Initiative (SSI) Algorithms, Models, Methods and Libraries:** New classes of algorithms will need to be developed to effectively utilize SSI systems. SSI computers will be characterized by large-scale distributed shared memory, with deep memory hierarchies. To achieve maximum performance, new hierarchical algorithms and methods need to be developed that can take advantage of these architecture features. In addition, mathematical techniques that are already under development may yield fundamentally new ways to conduct individual simulations and ensemble calculations, improving predictability and our understanding of errors in computations. New strategies for building mathematical software based on components are also likely to make a large contribution to programmer effectiveness for the SSI applications.	0	0	2,947
■ **SSI Problem Solving Environments and Tools (PSET):** A major challenge for CSET will be the development of an integrated problem solving environment that supports the rapid construction and testing of applications codes. These environments need to support the most advanced parallel tools and libraries and enable users to collaborate on the development of codes, planning runs, debugging and performance analysis. PSET systems will also allow the user to monitor jobs visually and to steer computations from their desktop and to share these results with others.	0	0	2,920
■ **SSI Distributed Computing and Collaboration Technology (DCCT):** Many SSI applications teams will be distributed around the country. They need the capability to work together not only to develop codes, but also to share data, design experiments, plan work and jointly conduct long running jobs. DCCT will support the creation of a software technology base to support SSI simulations that will enable users to run jobs at remote sites and migrate data from one site to another. Distributed computing technology will enable users to begin to see a single logical view of the DOE computing infrastructure and request, coordinate, and manage resources from multiple sites to attack a problem.	0	0	2,920

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
■ **SSI Visualization and Data Management Systems:** Simulations run on 10 teraflop or 40 teraflop class computers will generate 10's or 100's of Terabytes of output per job. Advanced scientific visualization and analysis systems will be needed to translate this mass of output to human understanding. Some data will best be analyzed visually while other data will need to be post-processed to extract features and statistics. Both forms of analysis will need sophisticated data management systems to enable the rapid and facile manipulation of datasets. Large-scale databases will be used to manage data from simulations and advanced data organization techniques will support rapid and hierarchical traversals of data providing the user the possibility of flying through data in real-time. Associated capital requirements for large scale data management and visualization testbeds are also supported . . .	0	0	6,815
■ **SSI Basic Science Applications:** Support two basic science SSI application efforts. These efforts will be selected through a competitive process that evaluates the scientific and technical merit of the proposed project, the potential of the proposed work to make significant new science possible, and the readiness of the scientific discipline and associated scientific community to make early use of terascale computers.	0	0	5,841
Total Mathematical, Computational, and Computer Sciences Research	47,857	51,560	71,407
Advanced Computation, Communications Research, and Associated Activities			
■ **Networking:** Research in high performance computer networks and information surety required to support high performance computer 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 both local area and wide area networks. In addition, this activity supports research in network protocols to enable applications to request, and be guaranteed, certain levels of network capability.	5,987	4,500	4,500

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
■ **Collaboratory Tools:** R&D to develop the results of fundamental research in computer science and networking into an integrated set of tools to enable scientists to remotely access and control facilities and share data in real time. In order to accomplish this goal a number of issues are under investigation including: definition and demonstration of a general and modular security architecture that can protect open network applications such as control of experimental devices; development of a modular electronic notebook prototype that can be used in a number of desktop computer environments to enable the sharing of scientific results, data from scientific instruments, and design of scientific procedures; development of tools to manage distributed collaborations such as tools for managing multipoint videoconferences ranging from the current "whoever is speaking has the floor" to more formal meetings where a meeting leader controls who has the "floor"; development of advanced techniques for managing and returning to the electronic record of the collaboration; and exploration of techniques such as virtual reality to enable large groups to work together effectively at a distance. This effort began as a part of the DOE2000 initiative.	3,000	3,000	3,000
■ **National Collaboratory Pilot Projects:** R&D to test, validate, and apply collaboratory tools in partnership with other DOE programs. The two pilot projects are: (1) the Materials MicroCharacterization Collaboratory, a partnership with Basic Energy Sciences and Energy Efficiency and Renewable Energy to provide remote access to facilities located at Oak Ridge National Laboratory, Lawrence Berkeley National Laboratory, Argonne National Laboratory, and the National Institute of Standards and Technology, and the University of Illinois for electron beam microcharacterization of materials; and (2) the Diesel Combustion Collaboratory, a partnership with Basic Energy Sciences, Energy Efficiency and Renewable Energy, and three U.S. manufacturers of diesel engines, to link the research and researchers at Sandia National Laboratory, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, and the University of Wisconsin with efforts and researchers at industrial			

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
laboratories in Indiana and Michigan to develop the next generation of clean diesel engines. This effort began as a part of the DOE2000 initiative. These partnerships with the scientific disciplines and the technology programs enable MICS to be particularly successful in bringing advances in advanced computing research to bear on important problems faced by scientists in other disciplines. In addition, they provide important feedback to the researchers on what problems are most important.	3,000	3,000	3,000
■ **National Energy Research Scientific Computing Center (NERSC):** NERSC, located at LBNL, provides high performance computing for investigators supported by the Office of Science. The Center serves 3,500 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. These computational resources will be integrated by a common high performance file storage system that facilitates interdisciplinary collaborations. This file storage system, the Archival Systems Upgrade, is a major item of equipment in FY 2000 with a total estimated cost of $2,000,000.	26,500	26,500	27,500

■ **Advanced Computing Research Facilities (ACRFs):** ACRFs support advanced computational hardware testbeds for scientific application pilot projects and fundamental research in applied mathematics and computer science. 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 facilities must procure and develop software to manage and make useful the largest scale systems that can be afforded. In addition, the ACRFs, taken together, must have a full range of different 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.

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
Therefore, all of these systems are managed as research programs and not as information technology investments. ACRFs are located at Los Alamos National Laboratory (Nirvana Blue partnership with the DOE Biological and Environmental Research (BER) program and Defense Programs (DP), based on SGI/Cray Technology); Argonne National Laboratory (IBM-SP); and Lawrence Berkeley National Laboratory (SGI/Cray T3E and Next Generation procurement). Related capital equipment needs such as high speed disk storage systems, archival data storage systems and high performance visualization hardware are also supported. The ACRFs represent the evolution of the High Performance Computing Research Centers that DOE initiated as a part of the Federal High Performance Computing and Communications initiative.	22,895	17,411	11,876
■ **Energy Sciences Network (ESnet):** ESnet provides worldwide access to the Office of Science facilities, including: advanced light sources; neutron sources; particle accelerators; fusion reactors; spectrometers; ACRFs; and other leading-edge science instruments and facilities. ESnet provides the communications fabric that links 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, ESnet management at LBNL contracts with commercial vendors for advanced communications services including Asynchronous Transfer Mode (ATM) and Wave Division Multiplexing (WDM). ESnet management provides system integration to provide a uniform interface to these services for DOE laboratories. In addition, ESnet management is responsible for the interfaces between the network fabric it provides and the worldwide Internet including the National Science Foundations's very high performance backbone network service 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			

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
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. Related capital equipment needs are also supported such as high speed network routers, ATM switches, and network management and testing equipment.	14,787	14,787	14,787

- **Next Generation Internet (NGI):** DOE's program will focus on developing, testing and validating the networking technologies needed to enable wide area, data intensive and collaborative computing. The program, which began in FY 1999, has three subcomponents. First, research in basic underlying technologies such as: protocols and techniques for coordinating multiple, heterogeneous network-attached devices; congestion and flow control techniques; multi-gigabit end system interfaces, analyzers, and switches along with mechanisms to reduce operating system overhead for data transfers; mechanisms to provide application controlled Class of Service and Quality of Service; and middleware to provide Internet Protocol (IP), ATM, and WDM resource and admission control, scheduling, management, prioritization, accounting, and debugging. Second, Application-Network Technology-Network Testbed Partnerships to: integrate and test advanced network R&D and testbeds with DOE mission applications such as HENP Data, remote visualization of simulation results, advanced collaboratories; define what network & middleware services are required to permit these applications to effectively run over wide area networks; define the features and the API's necessary to allow the application and middleware to communicate; integrate local and wide-area network technologies to create distributed collaboratories; and integrate Differentiated Services, or other Quality of Service functions, into wide area networks and production network testbeds without compromising the existing production network services. Third, DOE-University Technology Testbeds focused on: R&D to implement advanced network services across multiple, interconnected networks; deployment of advanced differentiated services technology

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
across autonomous networks when priority flow represents a significant fraction of the available capability; development and testing of advanced tools to manage "peering" of networks with advanced services; cross-domain implementations of security and authentication technologies; development and testing of network performance monitoring and characterization software which applications can use in this environment to optimize their performance; and development of policy frameworks and specification languages to facilitate the negotiation of capabilities across autonomous system boundaries.	0	14,602	14,602

- **SSI Research Computing Facilities and Network Infrastructure:** Initiation of a procurement for one 5 Teraflop-class computer system and associated data archive facilities in the third quarter of FY 2000. This activity includes not only procurement of the computing and storage systems but also the associated operational costs and facility improvements which will be required. This process will have two distinct stages. DOE will establish an open solicitation process that seeks the widest participation in establishing its terascale computing infrastructure, including competition among national laboratories, universities, and industry, based on their qualifications. The sites for the major teraflop computers will be selected through peer-reviewed competition. A number of considerations are important in selecting organizations to manage and locations at which to site SSI facilities. One of the most important is an expertise to perform the necessary computing systems integration into the existing nationwide DOE Office of Science computing infrastructure. These large-scale systems have requirements for scalable systems management that will enable relatively small systems administration teams to manage systems with 1000's of nodes and 10,000's of processors. These systems also need to be tightly integrated with data storage environments, mass stores, visualization environments, and distributed computing frameworks. Other considerations include incremental site preparation costs, the cost of connecting at very high speed to the networking infrastructure, and site

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
financial leverage in providing operational support to the facility. After the site is selected, the site must conduct an open competition to procure the 5 teraflop-class computer system as well as the other supporting computing and data storage systems. This activity also includes significant enhancements to the ESnet infrastructure to support access to the SSI computer facility by SSI scientists. SSI computing facilities and remote sites require very high performance network facilities to be effective. Planning estimates indicate that aggregate network data flow into and out of SSI facilities must be near 100 gigabits/second, almost 200 times faster than the fastest links on today's ESnet. In addition, remote sites with major participation in SSI science areas and CSET require significantly greater network capabilities than ESnet can provide within its current funding profile. Enhanced network services will be provided to remote sites on the basis of need and the resources available. The balance of funding between the computer facility and enhanced networking will be determined by detailed review after the selection of the computer facility site. However, preliminary estimates indicate that the computer facility will require between 80 percent and 85 percent of the funding in this activity. Associated requirements for capital equipment and GPP funding are included here..	0	0	29,417
Total, Advanced Computation, Communications Research, and Associated Activities	76,169	83,800	108,682
SBIR/STTR			
■ In FY 1998, $3,004,000 and $180,000 were transferred to the SBIR and STTR programs, respectively. The FY 1999 and FY 2000 amounts are the estimated requirement for the continuation of the SBIR and STTR programs.	0	3,474	4,486
Total, Mathematical, Information, and Computational Sciences	124,026	138,834	184,575

Explanation of Funding Changes from FY 1999 to FY 2000

	FY 2000 vs. FY 1999 ($000)
Mathematical, Computational, and Computer Sciences Research	
■ Increase in funding will provide opportunities for college faculty and students to spend time at DOE laboratories to participate in world class research projects.	+1,947
■ Close out all scientific application pilot projects in an orderly fashion in mid-FY 2000.	-3,543
■ Initiate support for SSI's Computational Sciences and Enabling Technology and Basic Science Applications..	+21,443
Total, Mathematical, Computational, and Computer Sciences Research	+19,847
Advanced Computation, Communications Research, and Associated Activities	
■ Cost-of-living increase for operations of NERSC.	+1,000
■ Reduced support for ACRF's including termination of the ACRF at ANL in mid-FY 2000.	-5,535
■ Support for SSI's Research Computing Facilities and Network Infrastructure.	+29,417
Total, Advanced Computation, Communications Research, and Associated Activities	+24,882
SBIR/STTR	
■ Increase in SBIR/STTR due to increase in operating expenses.	+1,012
Total Funding Change, Mathematical, Information, and Computational Sciences	+45,741

Laboratory Technology Research

Mission Supporting Goals and Objectives

The mission of the Laboratory Technology Research (LTR) subprogram is to support high risk, energy related 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 operate 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 program component provides rapid access by small business 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, intelligent processes and controls, and energy-related applications of biotechnology.

Performance Measures

- Initiate about 7 Laboratory Technology Research projects that address the Department’s top priorities for science and technology, through cost-shared research partnerships with industry.

Funding Schedule

	(dollars in thousands)				
	FY 1998	FY 1999	FY 2000	$ Change	% Change
Laboratory Technology Research	14,983	12,751	13,921	+1,170	+9.2%
SBIR/STTR	0	423	379	-44	-10.4%
Congressional Direction	396	2,968	0	-2,968	-100.0%
Total, Laboratory Technology Research	15,379	16,142	14,300	-1,842	-11.4%

Detailed Program Justification

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000

Laboratory Technology Research

- This activity supports research to advance the fundamental science and technology at the Office of Science 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. In FY 2000, about 7 new Laboratory Technology Research projects will be initiated. The research portfolio consists of approximately 100 projects and emphasizes the following topics: advanced materials processing and utilization, intelligent processes and controls, and energy-related applications of biotechnology. Efforts underway include the exploration of (1) a process to produce ultra-smooth diamond coatings on rotating and sliding mechanical parts in order to reduce energy consumption, improve product reliability, and reduce toxic emissions to the environment; (2) a new family of wireless, single-chip luminescent-sensing devices for use in monitoring of environmental pollutants and in high throughput screening of new therapeutic drugs; and (3) an x-ray imaging module, consisting of a cadmium zinc telluride detector and an application-specific integrated circuit, to determine bone density in the diagnosis of osteoporosis by performing whole

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
body scans with lower dose rate and higher resolution. 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 an industry. Recent projects supported the development of (1) a low-temperature oxygen plasma technology for remediation of soils contaminated by pesticides and other chemicals; (2) a new class of fluids that can transfer heat more efficiently than conventional fluids and offer several energy-related benefits, including decreased pumping power needs and reduction in required heat exchanger size, for use in the transportation, electronics, and textile industries; and (3) improved catalytic materials, which can lead to substantial energy savings in petroleum refining and chemical manufacturing.	14,983	12,751	13,921
SBIR/STTR			
■ In FY 1998, $391,000 and $24,000 were transferred to the SBIR and STTR programs, respectively. The FY 1999 and FY 2000 amounts are the estimated requirement for the continuation of the SBIR and STTR programs.	0	423	379
Congressional Direction			
■ Funds the University of Southwestern Louisiana (per FY 1997 Congressional Direction).	396	2,968	0
Total, Laboratory Technology Research	15,379	16,142	14,300

Explanation of Funding Changes from FY 1999 to FY 2000

	FY 2000 vs. FY 1999 ($000)
Laboratory Technology Research	
■ Increase in multiyear technology research partnership projects.	+1,170
SBIR/STTR	
■ Decrease in SBIR/STTR due to decrease in operating expenses.	-44
Congressional Direction	
■ Reduction is due to completion of Congressionally directed projects	-2,968
Total Funding Change, Laboratory Technology Research.	-1,842

Advanced Energy Projects

Mission Supporting Goals and Objectives:

The Advanced Energy Projects (AEP) subprogram funded research that established the feasibility of novel, energy-related concepts that span the Department's energy mission and goals. Funded projects were based on innovative ideas that spanned multiple scientific and technical disciplines and did not fit into any other DOE program area. A common theme for each project was the initial linkage of new research results to an energy application with a potentially significant payoff. Typically, AEP supported projects up to a level of about $250,000 per year for a period of about 3 years. Projects were selected from proposals submitted by universities and national laboratories. Funding criteria emphasized scientific merit as judged by external peer review.

Funding Schedule

(dollars in thousands)

	FY 1998	FY 1999	FY 2000	$ Change	% Change
Advanced Energy Projects	7,374	2,429	0	-2,429	-100.0%
SBIR/STTR	0	66	0	-66	-100.0%
Total, Advanced Energy Projects	7,374	2,495	0	-2,495	-100.0%

Detailed Program Justification

(dollars in thousands)

	FY 1998	FY 1999	FY 2000
Advanced Energy Projects			
■ Support for high-risk, high-payoff research at universities and national laboratories established the feasibility of novel energy related concepts that were at an early stage of scientific definition. Final funds for these projects were provided in FY 1999.	7,374	2,429	0

	(dollars in thousands)		
	FY 1998	FY 1999	FY 2000
SBIR/STTR			
■ In FY 1998, $187,000 and $11,000 were transferred to the SBIR and STTR programs, respectively. The FY 1999 amount is the estimated requirement for the continuation of the SBIR and STTR programs.	0	66	0
Total, Advanced Energy Projects	7,374	2,495	0

Explanation of Funding Changes From FY 1999 to FY 2000

	FY 2000 vs. FY 1999 ($000)
Advanced Energy Projects	
■ Termination of existing AEP program. The decision to terminate the AEP subprogram resulted from a change in CTR program priorities	-2,429
SBIR/STTR	
■ Decrease in SBIR/STTR due to termination of the program.	-66
Total Funding Change, Advanced Energy Projects	-2,495

Capital Operating Expenses & Construction Summary

Capital Operating Expenses

(dollars in thousands)

	FY 1998	FY 1999	FY 2000	$ Change	% Change
General Plant Projects	397	70	4,000	+3,930	+5,614.3%
Capital Equipment (total)	6,912	6,275	11,275	+5,000	+79.7%
Total, Capital Operating Expense	7,309	6,345	15,275	+8,930	+140.7%

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

(dollars in thousands)

	Total Estimated Cost (TEC)	Prior Year Approp- riations	FY 1998	FY 1999	FY 2000	Accept- ance Date
Archival Systems Upgrade - LBNL	2,000	0	0	0	2,000	FY 2002
Total, Major Items of Equipment		0	0	0	2,000	