Advanced Scientific Computing Research

Funding Profile by Subprogram

(dollars in thousands)

	FY 2003 Comparable Appropriation	FY 2004 Original Appropriation	FY 2004 Adjustments	FY 2004 Comparable Appropriation	FY 2005 Request
Advanced Scientific Computing Research					
Mathematical, Information, and Computational Sciences	160,367	200,490	-1,198[a]	199,292	204,340
Laboratory Technology Research	2,818	3,000	0	3,000	0
Subtotal, Advanced Scientific Computing Research	163,185	203,490	-1,198[a]	202,292	204,340
Use of Prior Year Balances	0	-481	0	-481	0
Total, Advanced Scientific Computing Research	163,185[bcd]	203,009	-1,198[a]	201,811	204,340

Public Law Authorizations:
Public Law 95-91, "Department of Energy Organization Act"
Public Law 103-62, "Government Performance and Results Act of 1993"

Mission

In the past two decades leadership in scientific computation has become a cornerstone of the Department's strategy to ensure the security of the nation and succeed in its science, energy, environmental quality, and national security missions. According to a number of authorities, ranging from the President's Science Advisor and the President's Council of Advisors on Science and Technology to the National Research Council and the Council on Competitiveness, this scientific leadership is critical to the economic health of the nation. The mission of the Advanced Scientific Computing Research (ASCR) program is to underpin DOE's world leadership in scientific computation by supporting research in applied mathematics, computer science and high-performance networks and providing the high-performance computational and networking resources that are required for world leadership in science.

Benefits

ASCR supports DOE's mission to provide world-class scientific research capacity through peer-reviewed scientific results in mathematics, high performance computing and advanced networks and

[a] Excludes $1,197,753 for a rescission in accordance with the Consolidated Appropriations Act, 2004, as reported in conference report H.Rpt. 108-401, dated November 25, 2003.
[b] Excludes $4,017,000 which was transferred to the SBIR program and $241,000 which was transferred to the STTR program.
[c] Excludes $1,115,315 for a rescission in accordance with the Consolidated Appropriations Resolution, FY 2003.
[d] Excludes $3,029,000 transferred for Department of Homeland Security activities in FY 2003.

through the application of terascale computing to advanced scientific applications. High-performance computing provides a new window for researchers to observe the natural world at a fidelity that could only be imagined a few years ago. Research investments in advanced scientific computing equip researchers with premier computational tools to advance knowledge and to solve the most challenging scientific problems facing the Nation.

Strategic and Program Goals

The Department's Strategic Plan identifies four strategic goals (one each for defense, energy, science, and environmental aspects of the mission plus seven general goals that tie to the strategic goals. The ASCR program supports the following goals:

Science Strategic Goal
General Goal 5, World-Class Scientific Research Capacity: Provide world-class scientific research capacity needed to ensure the success of Department missions in national and energy security, to advance the frontiers of knowledge in physical sciences and areas of biological, medical, environmental, and computational sciences, and to provide world-class research facilities for the Nation's science enterprise.

The ASCR program has one program goal which contributes to General Goal 5 in the "goal cascade":

Program Goal 05.23.00.00: Deliver forefront computational and networking capabilities - Deliver forefront computational and networking capabilities to scientists nationwide that enable them to extend the frontiers of science, answering critical questions that range from the function of living cells to the power of fusion energy.

Contribution to Program Goal 05.23.00.00 (Deliver forefront computational and networking capabilities)

Within the ASCR program, the Mathematical, Information and Computational Sciences subprogram contributes to Program Goal 05.23.00.00 by: delivering the fundamental mathematical and computer science research that enables the simulation and prediction of complex physical and biological systems; providing the advanced computing capabilities needed by researchers to take advantage of this understanding; and delivering the fundamental networking research and facilities that link scientists across the nation to the computing and experimental facilities and their colleagues to enable scientific discovery. This subprogram supports fundamental research in applied mathematics, computer science, computer networks, and tools for electronic collaboration; integrates the results of these basic research efforts into tools and software that can be used by scientists in other disciplines, especially through efforts such as Scientific Disovery through Advanced Computing (SciDAC); and provides the advanced computing and network resources that enable scientists to use these tools to deliver extraordinary science. Applied Mathematics enables scientists to translate the world into a computer with extraordinary fidelity, and provides the algorithms the computer requires to manipulate that representation of the world effectively, exposing the underlying structure. Computer science research provides the link between the mathematics and the actual computer systems. Finally, scientific discovery results from simulations conducted on the advanced computers themselves, including experimental computers with hardware designs optimized to enable particular types of scientific applications, and the largest computing capabilities available to the general scientific community. All of these elements are critical to advance the frontiers of simulation. Shrinking the distance between scientists and the resources they need is also critical to the Office of Science (SC). The challenges that

SC faces require teams of scientists distributed across the country as well as the full national portfolio of experimental and computational tools. High performance networks and network research provide the capability to move the millions of gigabytes that these resources generate to the scientists' desktops. The collaboratory activity, which is one component of the integrated ASCR Network Environment, provides the tools that enable scientists to discover, coordinate, and safely use the resources on the network.

Therefore, the ASCR program contributes to General Goal 5 by enabling research programs across SC, as well as other elements of the Department, to succeed. The following indicators establish specific long term (10 years) goals in Scientific Advancement that the ASCR program is committed to, and progress can be measured against.

- Develop multiscale mathematics, numerical algorithms, and software that enable effective models of systems such as the Earth's climate, the behavior of materials, or the behavior of living cells that involve the interaction of complex processes taking place on vastly different time and/or length scales.
- Develop, through the Genomics: GTL partnership with the Biological and Environmental Research (BER) program, the computational science capability to model a complete microbe and a simple microbial community. This capability will provide the science base to enable the development of novel clean-up technologies, bio-energy sources, and technologies for carbon sequestration.

Annual Performance Results and Targets

FY 2000 Results	FY 2001 Results	FY 2002 Results	FY 2003 Results	FY 2004 Targets	FY 2005 Targets
05.23.00.00 Deliver forefront computational and networking capabilities					
Mathematical, Information and Computational Sciences					
			Completed the definitive analysis of the advantages and issues associated with lightweight kernel operating systems rather than full kernels for the compute nodes of extreme-scale scientific computers, resolving a critical issue for the future of high performance computers in the U.S. [Met]		
			Began installation of next generation National Energy Research Scientific Computing Center (NERSC) computer, NERSC-4, that will at least double the capability available in FY 2002 to solve leading edge scientific problems. [Not Met]	Maintain Procurement Baselines. Percentages within (1) original baseline cost for completed procurements of major computer systems or network services, and (2) original performance baseline versus integrated performance over the life of the contracts - FY04 – <10%	Maintain Procurement Baselines. Percentages within (1) original baseline cost for completed procurements of major computer systems or network services, and (2) original performance baseline versus integrated performance over the life of the contracts - FY05 – <10%
			Initiated at least 5 competitively selected interdisciplinary research teams to provide computational science and applied mathematics advances that will accelerate biological discovery in microbial systems and develop the next generation of computational tools required for nanoscale science based on peer review, in partnership with the Biological and	Improve Computational Science Capabilities. Average annual percentage increase in the computational effectiveness (either by simulating the same problem in less time or simulating a larger problem in the same time) of a subset of application codes within the Scientific Discovery through Advanced Computing (SciDAC) effort. FY04 – >50%	Improve Computational Science Capabilities. Average annual percentage increase in the computational effectiveness (either by simulating the same problem in less time or simulating a larger problem in the same time) of a subset of application codes within the Scientific Discovery through Advanced Computing (SciDAC) effort. FY05 – >50%

FY 2000 Results	FY 2001 Results	FY 2002 Results	FY 2003 Results	FY 2004 Targets	FY 2005 Targets
			Environmental Research and Basic Energy Sciences programs, respectively, of submitted proposals. [Met]		
			Maintained and operated facilities, including NERSC and ESnet, so the unscheduled downtime on average is less than 10 percent of the total scheduled operating time. [Met]	Focus usage of the primary supercomputer at the National Energy Research Scientific Computing Center (NERSC) on capability computing. Percentage of the computing time used that is accounted for by computations that require at least 1/8 of the total resource - FY04 – 50%	Focus usage of the primary supercomputer at the National Energy Research Scientific Computing Center (NERSC) on capability computing. Percentage of the computing time used that is accounted for by computations that require at least 1/8 of the total resource - FY05 – 50%
	Initiated project to understand the advantages and issues associated with lightweight kernel operating systems rather than full kernels for the compute nodes of extreme-scale scientific computers. [Met]	Completed the development of the Cougar lightweight kernel for clusters of Alpha processor-based computers and began the assessment of scalability and performance for selected applications. [Met]			
Developed advanced computing capabilities, computational algorithms, models, methods, libraries, and advanced visualization and data management systems that enabled new computing applications to science. [Met]	Continued to fabricate, assemble, and operate premier supercomputer and networking facilities that served researchers at national laboratories, universities and within industry, enabling understanding of complex problems and effective integration of geographically distributed teams in national collaborations. [Met]	Achieved operation of the IBM-SP computer at 5.0 teraflop "peak" performance. These computational resources were integrated by a common high performance file storage system that facilitates interdisciplinary collaborations. Transferred the users with largest data processing and storage needs to the IBM-SP from the previous generation Cray T3E. [Met]			

Means and Strategies

The ASCR program will use various means and strategies to achieve its goals. However, various external factors may impact the ability to achieve these goals.

The ASCR program will support fundamental, innovative, peer-reviewed research to create new knowledge in areas of advanced computing research that are important to DOE. In addition, the ASCR program will plan, fabricate, assemble, and operate premier supercomputer and networking facilities that serve researchers at national laboratories, universities, and industry, thus enabling new understanding through analysis, modeling, and simulation for complex problems, and effective integration of geographically distributed teams through national laboratories. Finally, the program will continue its leadership of the SC-wide Scientific Discovery through Advanced Computing (SciDAC) initiative with Basic Energy Sciences (BES) and Biological and Environmental Research (BER) in the areas of nanotechnology and Genomics: GTL. All research projects undergo regular peer review and merit evaluation based on procedures set down in 10 CFR 605 for the extramural grant program, and under a similar process for the laboratory programs and scientific user facilities. All new projects are selected through peer review and merit evaluation.

External factors that affect the programs and performance include: (1) mission needs as described by the DOE and SC mission statements and strategic plans; (2) evolving scientific opportunities, which sometimes emerge in a way that revolutionizes disciplines; (3) results of external program reviews and international benchmarking activities of entire fields or subfields, such as those performed by the National Academy of Sciences; (4) unanticipated failures, for example, in the evaluation of new computer architectures for science, that cannot be mitigated in a timely manner; (5) strategic and programmatic decisions made by other (non-DOE) Federal agencies and by international entities; and (6) the evolution of the commercial market for high performance computing and networking hardware and software.

The fundamental research program and facilities are closely coordinated with the information technology research activities of other Federal Agencies (DARPA, EPA, NASA, NIH, NSA, and NSF) through the Computing Information and Communications R&D subcommittee of the National Science and Technology Council (NSTC), under the auspices of the Office of Science and Technology Policy. This coordination is periodically reviewed by the President's Information Technology Advisory Committee (PITAC). In addition to this interagency coordination, ASCR has a number of partnerships with other programs in SC and other parts of the Department, focused on advanced application testbeds to apply the results of ASCR research. Finally, ASCR has a significant ongoing coordination effort with the National Nuclear Security Administration's (NNSA) Advanced Science Computing (ASC) Campaign to ensure maximum effectiveness of both computational science research efforts.

Validation and Verification

Progress against established plans is evaluated by periodic internal and external performance reviews. These reviews provide an opportunity to verify and validate performance. Quarterly, semiannual, and annual reviews consistent with specific program management plans are held to ensure technical progress, cost and schedule adherence, and responsiveness to program requirements.

Program Assessment Rating Tool (PART) Assessment

In the PART review, OMB gave the Advanced Scientific Computing Research (ASCR) program a relatively high score of 84% overall which corresponds to a rating of "Moderately Effective." OMB found that the program supports world-class scientific user facilities, has demonstrated an improved level of interagency communication and cooperation, is in the process of drafting a long-term strategic vision, and has been very successful with a major effort in interdisciplinary software. Although ASCR is establishing a Committee of Visitors (COV), to provide outside expert validation of the program's merit-based review processes for impact on quality, relevance, and performance, this committee has not yet met. Once the COV issues a report, ASCR will develop an action plan to respond to the findings and recommendations within 30 days. The assessment found that ASCR has developed a limited number of adequate performance measures. However, OMB noted concerns regarding the collection and reporting of performance data. To address these concerns, ASCR will work with its Advisory Committee to develop research milestones for the long-term performance goals, will include the long term research goals in grant solicitations, will work to improve performance reporting by grantees and contractors, and will work with the CFO to improve ASCR sections of the Department's performance documents. OMB also found that the ASCR Advisory Committee is underutilized. ASCR will meaningfully engage the Advisory Committee in thorough assessments of research performance and in regularly revisiting the strategic priorities for the program. ASCR's role in providing scientific research facilities is strongly supported by the Administration. Funding is provided in FY 2005 to operate the program's facilities at maximum capacity.

Funding by General and Program Goal

(dollars in thousands)

	FY 2003	FY 2004	FY 2005	$ Change	% Change
General Goal 5, World-Class Scientific Research Capacity					
Program Goal 05.23.00.00, Deliver forefront computational and networking capabilities					
Mathematical, Information and Computational Sciences	160,367	199,292	204,340	+5,048	+2.5%
Laboratory Technology Research	2,818	3,000	0	-3,000	-100.0%
Total, Program Goal 05.23.00.00, Deliver forefront computational and networking capabilities	163,185	202,292	204,340	+2,048	+1.0%
Use of Prior-Year Balances	0	-481	0	+481	+100.0%
Total, Advanced Scientific Computing Research	163,185	201,811	204,340	+2,529	+1.3%

Overview

Computational modeling and simulation are among the most significant developments in the practice of scientific inquiry in the 20th Century. Scientific computing is particularly important for the solution of research problems that are insoluble by traditional theoretical and experimental approaches, hazardous to study in the laboratory, or time-consuming or expensive to solve by traditional means. All of the research programs in DOE's Office of Science—Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, and High Energy and Nuclear Physics—have identified major scientific challenges that can only be addressed through advances in scientific computing.

ASCR research underpins the efforts of the other programs in the Office of Science. The applied mathematics research activity produces the fundamental mathematical methods to model complex physical and biological systems. The computer science research efforts enable scientists to efficiently perform scientific computations on the highest performance computers available and to store, manage, analyze, and visualize the massive amounts of data that result. The networking research activity provides the techniques to link the data producers, e.g., supercomputers and large experimental facilities with scientists who need access to the data.

ASCR's other principal responsibility is to provide the high-performance computational and networking resources that are required for world leadership in science. Recent dramatic advances in scientific computation by researchers and computer companies underscore the importance of strengthening our position in computational sciences in strategic areas. In March 2002, Japan's NEC Earth Simulator became operational. With a peak speed of 40 teraflops and a demonstrated sustained capability of over 25 teraflops, it is faster by approximately a factor of 50 than the most advanced supercomputer for civilian science in the United States. The potential long-term implications of the Earth Simulator on DOE's computational sciences capability was the principal message of the report on this subject delivered to the Director of the Office of Science by the Advanced Scientific Computing Advisory Committee. To strengthen the program's position in this area, the ASCR program is proposing a new effort in Next Generation Computer Architecture (NGA) to identify and address major bottlenecks in the performance of existing and planned DOE science applications.

How We Work

The ASCR program uses a variety of mechanisms for conducting, coordinating, and funding research in applied mathematics, network and computer sciences, and in advanced computing software tools. The program is responsible for planning and prioritizing all aspects of supported research, conducting ongoing assessments to ensure a comprehensive and balanced portfolio, regularly seeking advice from stakeholders, supporting core university and national laboratory programs, and maintaining a strong infrastructure to support research in applied mathematics, network and computer science, and advanced computing software tools. The quality of the research supported by the ASCR program is continuously evaluated through the use of merit-based peer review, scientific advisory committees, and interagency coordinating bodies.

Advisory and Consultative Activities

The ***Advanced Scientific Computing Advisory Committee (ASCAC),*** established on August 12, 1999, provides valuable, independent advice to the Department of Energy on a variety of complex scientific and technical issues related to the ASCR program. The ASCAC is charged with providing advice on: promising future directions for advanced scientific computing research; strategies to couple advanced scientific computing research to other disciplines; and the relationship of the DOE program to other Federal investments in information technology research. ASCAC's recommendations include advice on long-range plans, priorities, and strategies to address more effectively the scientific aspects of advanced scientific computing including the relationship of advanced scientific computing to other scientific disciplines, and maintaining appropriate balance among elements of the program. This advisory committee plays a key role in assessing the scientific and programmatic merit of presently funded activities and in evaluating plans for the future. The Committee formally reports to the Director, Office of Science and includes representatives from universities, national laboratories, and industries who are involved in advanced computing research. Particular attention is paid to obtaining a diverse membership with a balance among scientific disciplines, institutions, and geographic regions. ASCAC operates in accordance with the Federal Advisory Committee Act (FACA, Public Law 92-463; 92nd Congress, H.R. 4383; October 6, 1972) and all applicable FACA Amendments, Federal Regulations and Executive Orders.

The activities funded by the ASCR program are coordinated with other Federal efforts through the *Interagency Principals Group*, chaired by the President's Science Advisor, and the *Information Technology Working Group (ITWG).* The ITWG evolved through an interagency coordination process that began under the 1991 High Performance Computing Act as the High Performance Computing, Communications, and Information Technology (HPCCIT) Committee. The Federal IT R&D agencies have established a 10-year record of highly successful collaborative accomplishments in multiagency projects and in partnerships with industry and academic researchers. The multiagency approach leverages the expertise and perspectives of scientists and technology users from many agencies who are working on a broad range of IT research questions across the spectrum of human uses of information technology. DOE has been an active participant in these coordination groups and committees since their inception and the ASCR program will continue to coordinate its activities through these mechanisms including an active role in implementing the Federal IT R&D FY 2002-2006 Strategic Plan under the auspices of the National Science and Technology Council and the President's Science Advisor.

ASCR is a participant in the Interagency Committee for Extramural Mathematics Programs (ICEMAP), a coordinating committee with representatives from Federal agencies that manage programs in mathematical research, including the National Science Foundation, DOE (through ASCR), the National Aeronautics and Space Administration, the National Institute for Standards and Technology, the Air Force Office of Scientific Research, the Army Research Office, and the Office of Naval Research. Meetings are held to coordinate activities across mathematical research programs, ensuring that the Federal agencies coordinate their investments in basic mathematical research. The ASCR program regards ICEMAP as an important component in their efforts to maintain coordination with other Federal agencies.

In addition, ASCR, both through ASCAC and independently, supported a number of workshops to support its planning. These include:

- Blueprint for Future Science Middleware and Grid Research and Infrastructure, August 2002 (http://www.nsf-middleware.org/MAGIC/default.htm);
- DOE Science Network Meeting, June 2003 (http://gate.hep.anl.gov/may/ScienceNetworkingWorkshop/);
- DOE Science Computing Conference, June 2003 (http://www.doe-sci-comp.info);
- Science Case for Large Scale Simulation, June 2003 (http://www.pnl.gov/scales/);
- Workshop on the Road Map for the Revitalization of High End Computing (http://www.cra.org/Activities/workshops/nitrd/);
- Cyber infrastructure Report (http://www.cise.nsf.gov/evnt/reports/toc.htm); and
- ASCR Strategic Planning Workshop (http://www.fp-mcs.anl.gov/ascr-july03spw).

Facility Operations Reviews

The ASCR program has undertaken a series of operations reviews of the National Energy Research Scientific Computing Center (NERSC), the Energy Sciences Network (ESnet), and the Advanced Computing Research Testbeds (ACRTs).

NERSC, operated by the Lawrence Berkeley National Laboratory, annually serves about 2,000 scientists throughout the United States. These researchers work at DOE laboratories, universities, industrial laboratories and other Federal agencies. Allocations of computer time and archival storage at NERSC are awarded to research groups based on a review of submitted proposals. As proposals are submitted, they are subjected to peer review to evaluate the quality of science, the relevance of the proposed research to Office of Science goals and objectives and the readiness of the proposed application to fully utilize the computing resources being requested.

The ESnet, managed and operated by the Lawrence Berkeley National Laboratory, is a high-speed network serving thousands of Department of Energy scientists and collaborators worldwide. A pioneer in providing DOE mission oriented high-bandwidth, reliable connections, ESnet enables researchers at national laboratories, universities and other institutions to communicate with each other using the leading edge collaborative capabilities, not available in the commercial world that are needed to address some of the world's most important scientific challenges. The ESnet Steering Committee (ESSC) was established in 1985 to ensure that ESnet meets the needs of the Office of Science programs. All program offices in the Office of Science appoint members, who represent their scientific communities, to serve on the ESSC. The ESSC is responsible for reviewing and prioritizing network requirements, for establishing performance objectives, and for proposing innovative techniques for enhancing ESnet capabilities. In addition to the ongoing oversight from the ESSC, ASCR conducts external peer reviews of ESnet performance on a three year interval. The last such review was chaired by a member of ASCAC and took place in September 2001.

Advanced Computing Research Testbeds (ACRTs) play a critical role in testing and evaluating new computing hardware and software. Current testbeds are located at Oak Ridge National Laboratory (IBM Power-4 Technology and CRAY X1 technology). In FY 2002, ASCAC conducted a review of NERSC and the ACRTs. The charge to ASCAC, posed the following questions:

- What is the overall quality of these activities relative to the best-in-class in the U.S. and internationally?
- How do these activities relate and contribute to Departmental mission needs?
- How might the roles of these activities evolve to serve the missions of the Office of Science over the next three to five years?

The essential finding of the Subcommittee was that NERSC and the ACRTs are among the best worldwide in their respective categories. It was the opinion of the Subcommittee that these ASCR activities and the related spin-off research efforts contribute significantly to the mission needs of the DOE, and profoundly and positively impact high performance computing activities worldwide. The complete report is available on the web. ***(http://www.science.doe.gov/ascr/ASCAC-sub.doc)***

In FY 2001, ASCR conducted a peer review of the Center for Computational Sciences (CCS) at the Oak Ridge National Laboratory. The findings from this review validated the contributions that the CCS made to the Advanced Computing Research Testbed activity within the ASCR program.

Program Reviews

The ASCR program conducts frequent and comprehensive evaluations of every component of the program. Results of these evaluations are used to modify program management as appropriate.

In FY 2003, ASCR conducted a peer review of the Numerical Linear Algebra, Optimization, and Predictability Analysis areas within the Applied Mathematics activity. These areas represent 33 percent of this activity. In FY 2004, ASCR will conduct a peer review of the Differential Equations and Advanced Numerical Methods for High Performance Computing areas within the Applied Mathematics activity, representing an additional 33 percent of this activity. In FY 2005, ASCR will conduct a peer review of the remaining 34 percent of the Applied Mathematics activity, which consists of Computational Fluid Dynamics and Meshing Techniques. Also in FY 2003, ASCR completed reviews of all of the SciDAC Integrated Software Infrastructure Centers (ISICs). There are a total of seven such centers (three with a mathematics focus and four with a computer science focus), and this represents over 50 percent of the ASCR SciDAC budget. In FY 2004, ASCR will initiate a comprehensive review of the Computer Science base research activity.

In FY 2003, ASCR also conducted peer reviews of all the SciDAC Collaboratory Pilot and Middleware Projects. These reviews focused on accessing progress and the possible need for mid-course corrections.

In FY 2002, following a comprehensive peer review, the ASCR program approved a proposal from the Lawrence Berkeley National Laboratory (LBNL) to manage and operate the National Energy Research Scientific Computing Center for FY 2002 – FY 2006.

Planning and Priority Setting

The ASCR program must coordinate and prioritize a large number of goals from agency and interagency strategic plans. One of the most important activities of ASCAC is the development of a framework for the coordinated advancement and application of network and computer science and applied mathematics. This framework must be sufficiently flexible to rapidly respond to developments in a fast paced area of research. The key planning elements for this program are:

- The Department and Office of Science Strategic Plan, as updated through program collaborations and joint advisory committee meetings. (***http://www.science.doe.gov/production/bes/strat_pln.htm)***;
- Scientific Discovery through Advanced Computing (SciDAC) plan delivered to Congress in March 2000. (***http://www.science.doe.gov/scidac/)***;
- The Interagency Working group for Information Technology Five Year Plan – FY 2002-FY 2006 (with key appendixes); and
- ASCAC report on the Japanese Earth Simulator. (***http://www.sc.doe.gov/ascr/ascac.reports.htm)***

How We Spend Our Budget

The ASCR program budget has two subprograms: Mathematical, Information and Computational Sciences (MICS) and Laboratory Technology Research (LTR). The MICS subprogram has two major components: research and facility testbed and network operations. The FY 2005 budget request continues the core and SciDAC research efforts and strengthens the research partnerships with other SC offices. The testbed and network operations expenditures account for 37 percent of the National Laboratory Research, or 24 percent of the total ASCR budget. The LTR subprogram will be brought to a successful completion in FY 2004.

**Advanced Scientific Computing Research Budget Allocation
FY 2005**

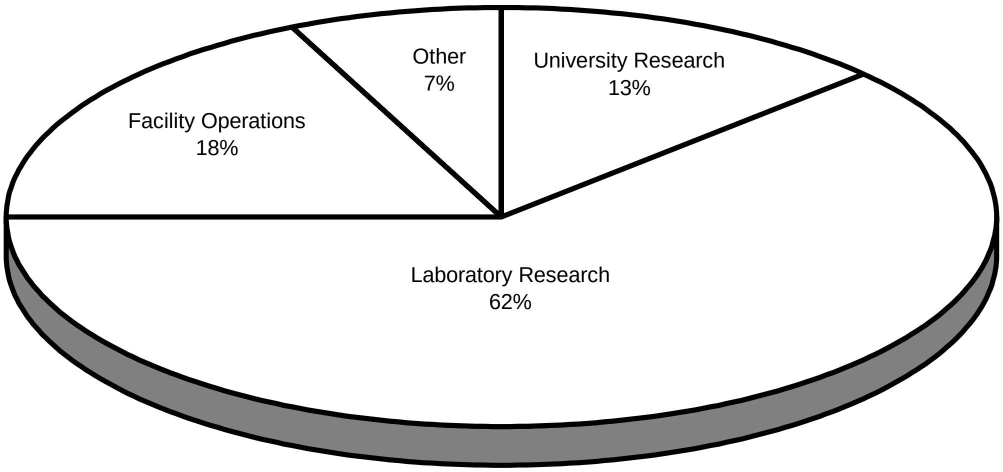

Research

Over 74 percent of the ASCR program's FY 2005 funding will be provided to scientists at universities and laboratories to conceive and carry out the research or to fund advanced computing testbeds and network operations. National laboratory research scientists work together with the other programs of the Office of Science to develop the tools and techniques that allow those programs to take advantage of terascale computing for scientific research. The laboratories provide state-of-the-art resources for testbeds and novel applications. The division of support between national laboratories and universities is adjusted to maximize scientific productivity.

- ***University Research***: University researchers play a critical role in the nation's research effort and in the training of graduate students. During FY 2002, the ASCR program supported over 150 grants to the nation's university researchers and graduate students engaged in civilian applied mathematics, large-scale network and computer science research. In addition, ASCR supports a Computational Science Graduate Fellowship and an Early Career Principal Investigator activity in Applied Mathematics, Computer Science and High-Performance Networks. In FY 2002, ASCR selected 24 new graduate fellows representing 17 universities and 13 states and expects to make up to forty awards to early career principal investigators. Approximately one-half of those who received Ph.D.'s in the Computational Sciences Graduate Fellowship program between 1991 and 2001 are pursuing careers outside universities or national labs. ASCR also provides support to other Office of Science research programs.

  The university grants program is proposal driven, similar to the computer science and applied mathematics programs at NSF. However, ASCR grant solicitation notices are focused on topics that have been identified as important for DOE missions. ASCR funds the best among the ideas submitted in response to grant solicitation notices (**http://www.science.doe.gov/production/grants/**). Proposals are reviewed by external scientific peers and competitively awarded according to the guidelines published in 10 CFR 605 (**http://www.science.doe.gov/production/grants/605index.html**).

- ***National Laboratory Research***: ASCR supports national laboratory-based research groups at Ames, Argonne, Brookhaven, Los Alamos, Lawrence Berkeley, Lawrence Livermore, Oak Ridge, Pacific Northwest, and Sandia National Laboratories. The directions of laboratory research programs are driven by the needs of the Department and the unique capabilities of the laboratories to support large scale, multidisciplinary, collaborative research activities. In addition, laboratory-based research groups are highly tailored to the major scientific programs at the individual laboratories and the computational research needs of the Office of Science. Laboratory researchers collaborate with laboratory and academic researchers, and are important for developing and maintaining testbeds and novel applications of high performance computing and networking in Office of Science research. At Los Alamos, Livermore and Sandia, ASCR funding plays an important role in supporting basic research that can improve the applied programs, such as the Accelerated Strategic Computing Initiative (ASCI) and the Science Stockpile Stewardship program.

ASCR funds field work proposals from the national laboratories. Proposals are reviewed by external scientific peers and awarded using procedures that are equivalent to the 10 CFR 605 guidelines used for the grants program. Performance of the laboratory groups is reviewed by ASCR staff annually to examine the quality of their research and identify needed changes, corrective actions or redirection of effort. Individual laboratory groups have special capabilities or access to laboratory resources that can be profitably utilized in the development of the research program.

Significant Program Shifts

The ASCR program advances mathematics and computer science, and develops the specialized algorithms, the scientific software tools, and the software libraries needed by DOE researchers to effectively use high-performance computing and networking hardware for scientific discovery. The ASCR program has been a leader in the computational sciences for several decades and has been acknowledged for pioneering accomplishments.

Research efforts initiated in FY 2001 in Scientific Discovery through Advanced Computing (SciDAC) will be continued, as planned. In FY 2005, ASCR will continue efforts initiated in FY 2004 to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures. The near term activities are represented by enhancements to NERSC while the longer term activities are a part of the Next Generation Computer Architecture (NGA). This effort will continue to increase the delivered computing capability available to address the Office of Science mission through optimization of computer architectures to meet the special requirements of scientific problems. This investment positions the nation to realize extraordinary scientific opportunities in computing for science and enable new classes of scientific problems to be addressed. The NGA effort complements SciDAC and integrates advanced computer architecture researchers and engineers, application scientists, computer scientists, and applied mathematicians. The NGA efforts, as well as the enhancement of NERSC are aligned with the plan developed by the High End Computing Revitalization Task Force (HECRTF) established by the Office of Science and Technology Policy. These efforts will play a critical role in enabling potential future Leadership Class Machines, which could lead to solutions for scientific and industrial problems beyond what would be attainable through a continued simple extrapolation of current computational capabilities.

The FY 2005 budget request includes $7,500,000 for continued support of the Genomics: GTL research program, in partnership with the Biological and Environmental Research program; and $2,600,000 for the Nanoscale Science, Engineering and Technology initiative led by the Basic Energy Sciences program. ASCR's contributions to these partnerships will consist of advancing the mathematics and developing new mathematical algorithms to simulate biological systems and physical systems at the nanoscale. In addition to this continued partnership support, the FY 2005 request includes $1,350,000 to expand SciDAC partnerships with the Fusion Energy Sciences Program to lay the groundwork for the Fusion Simulation Project (FSP). The FSP will be a focused, interdisciplinary effort, whose objective is to develop the capability to predict reliably the behavior of fusion plasmas.

The FY 2005 budget also includes $8,500,000 for the new "Atomic to Macroscopic Mathematics" (AMM) research support in applied mathematics needed to break through the current barriers in our understanding of complex physics processes that occur on a wide range of interacting length- and time-scales. Achieving this basic mathematical understanding will provide enabling technology to virtually every challenging computational problem faced by the Office of Science.

In FY 2005, the Mathematical, Information and Computational Sciences subprogram will continue to support core research activities in applied mathematics, computer science, network research, collaboratory tools and collaboratory pilot projects at current levels.

The Laboratory Technology Research subprogram was brought to a successful conclusion in FY 2004 as planned with orderly completion of all existing CRADAs. This does not mean that technology transfer activities have ended; rather, these activities are now institutionalized as a part of the process of doing research at DOE sites.

Interagency Environment

The activities funded by the MICS subprogram are coordinated with other Federal efforts through the Interagency Principals Group, chaired by the President's Science Advisor, and the Information Technology Working Group (ITWG). The ITWG evolved through an interagency coordination process that began under the 1991 High Performance Computing Act as the High Performance Computing, Communications, and Information Technology (HPCCIT) Committee. DOE has been an active participant in these coordination groups and committees since their inception. The MICS subprogram will continue to coordinate its activities through these mechanisms and will lead the development of new coordinating mechanisms as needs arise. The DOE program solves mission critical problems in scientific computing. In addition, results from the DOE program benefit the Nation's Information Technology Basic Research effort. The FY 2005 program positions DOE to make additional contributions to this effort. In the area of high performance computing and computation, ASCR has extensive partnerships with other Federal agencies and the NNSA. Examples include: participating in the program review team for the DARPA High Productivity Computing Systems program; serving on the planning group for the Congressionally mandated DOD plan for high performance computing to serve the national security mission; serving on the OSTP High End Computing Revitalization Task Force; and extensive collaboration with NNSA-Advanced Simulation Computing.

Scientific Discovery through Advanced Computing

The Scientific Discovery through Advanced Computing (SciDAC) program is a set of coordinated investments across all Office of Science mission areas with the goal of achieving breakthrough scientific advances via computer simulation that were impossible using theoretical or laboratory studies alone. The power of computers and networks is increasing exponentially. Advances in high-end computing technology, together with innovative algorithms and software, are being exploited as intrinsic tools for scientific discovery. SciDAC has also pioneered an effective new model of multidisciplinary collaboration among discipline-specific scientists, computer scientists, computational scientists, and mathematicians. The product of this collaborative approach is a new generation of scientific simulation codes that can productively exploit terascale computing and networking resources. The program is bringing computation and simulation to parity with experiment and theory in the scientific research enterprise as demonstrated by major advances in climate modeling and prediction, plasma physics, particle physics, accelerator design, astrophysics, chemically reacting flows, and computational nanoscience.

The research focus of ASCR SciDAC activities includes Integrated Software Infrastructure Centers (ISICs) and collaboratories. ISICs are partnerships between DOE national laboratories and universities focused on research, development, and deployment of software to accelerate the development of SciDAC application codes. Progress to date includes significant improvements in performance modeling and analysis capabilities that have led to doubling the performance on 64 processors of the Community Atmosphere Model component of the SciDAC climate modeling activity. Collaboratories address complex science projects undertaken in the SciDAC program that involve geographically distributed computing resources, research teams, and science. Progress includes development of GridFTP, a high-performance data transport program that has become the de facto standard for data transport on the Grid. In high-energy physics, collaboratory technology enabled a single job to generate 1.5 million simulated events for the Compact Muon Solenoid. The three Mathematics ISICs, now 1.5 years into their 3-5 year life, are bringing a new level of mathematical sophistication to computational problems throughout the Office of Science. One of these, the Terascale Optimal Partial Differential Equations (PDE) Simulations (TOPS) Center, is combining the Hyper and Portable Extensible Tool kit for Scientific Computation (PETSc) libraries, together with newly developed algebraic multigrid solvers, to create fast algorithms for a variety of tough and important problems, including biochemical reaction diffusion equations, advection equations for combustion simulation, and so forth. The Terascale Simulation Tools and Technologies Center is working to develop a framework for coupling different types of grids together in a single application. For example, in a simulation of engine combustion, one might want an unstructured grid for the complex geometry around the valves, but a regular grid in the rest of the cylinder. Finally, the Applied Partial Differential Equations Center is focused on using structured adaptive grids for simulation in a variety of application domains, including ground water flow, combustion chemistry, and magneto-hydrodynamics. Given the difficulty of magneto-hydrodynamic simulation, this center is having a strong impact on the design of new particle accelerators.

Next Generation Computer Architecture for Science and Industry

The Next Generation Computer Architecture for Science and Industry (NGA) research activity is an integral part of an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures. Total funding for the NGA was $38,268,000 in FY 2004 and $38,212,000 in FY 2005. The goal of the NGA is to identify and address major architectural bottlenecks, such as internal data movement in very large systems, in the performance of existing and planned DOE science applications. Emphasis will be placed on understanding the impact of alternative computer architectures on application performance with particular attention paid to data movement from memory to processor and between processors in highly parallel systems. Software research will be initiated to improve application performance and system reliability through innovative approaches to next generation operating systems. Emphasis will also be placed on hardware evaluation testbeds of sufficient size to understand key issues impacting application performance scalability and portability. These testbeds will also enable significant scientific progress by delivering significant increases in performance to critical DOE mission applications. These testbeds will also enable industrial researchers to find opportunities for virtual prototypes and simulation of industrial processes that result in enhanced competitive position because of sharply reduced 'time to market.' The NGA activity is coordinated with other Federal agencies to gain additional insight into research directions, optimize the utilization of resources, and establish the framework for a national effort. This effort is aligned with the HECRTF plan.

Scientific Facilities Utilization

The ASCR program request includes support to the National Energy Research Scientific Computing Center (NERSC), a component of the Office of Science-wide Facilities Optimization effort. This investment will provide computer resources for about 2,000 scientists in universities, DOE laboratories, Federal agencies, and U.S. companies. The proposed funding will enable NERSC to maintain its role as one of the Nation's premier unclassified computing centers, a critical element for success of many Office of Science research programs.

	FY 2001	FY 2002	FY 2003	FY 2004 Est.	FY 2005 Est.
Maximum Hours – NERSC	8,760	8,760	8,760	8,760	8,760
Scheduled Hours – NERSC....................	8,497	8,585	8,585	8,585	8,585
Unscheduled Downtime – NERSC	1%	1%	1%	1%	1%

Workforce Development

The R&D Workforce Development mission is to ensure the supply of computational and computer science and Ph.D. level scientists for the Department and the Nation through graduate student and postdoctoral research support. In FY 2005, this program will support approximately 800 graduate students and post doctoral investigators, of which 500 will be supported at Office of Science user facilities.

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

	FY 2001	FY 2002	FY 2003	FY 2004 est.	FY 2005 est.
# University Grants...........	170	163	144	140	142
Size, Duration...................	$157,000/yr-3yrs	$157,000/yr-3yrs	$197,000/yr-3yrs	$197,000/yr-3yrs	$197,000/yr-3yrs
# Lab Groups....................	226	209	165	165	165
# Grad Students	370	354	354	354	354
# PhD's Awarded.............	660	604	675	675	675

Mathematical, Information, and Computational Sciences

Funding Schedule by Activity

	(dollars in thousands)				
	FY 2003	FY 2004	FY 2005	$ Change	% Change
Mathematical, Information, and Computational Sciences					
Mathematical, Computational, and Computer Sciences Research	68,748	83,301	86,405	+3,104	+3.7%
Advanced Computation, Communications Research and Associated Activities	91,619	110,553	112,389	+1,836	+1.7%
SBIR/STTR	0	5,438	5,546	+108	+2.0%
Total, Mathematical, Information, and Computational Sciences	160,367	199,292	204,340	+5,048	+2.5%

Description

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

Benefits

MICS supports ASCR's contribution to DOE's mission to provide world-class scientific research capacity by providing world-class, peer-reviewed scientific results in mathematics, high performance computing and advanced networks and applying the potential of Terascale computing to advanced scientific applications. High-performance computing provides a new window for researchers to observe the natural world at a fidelity that could only be imagined a few years ago. Research investments in advanced scientific computing equip researchers with premier computational tools to advance knowledge and to solve the most challenging scientific problems facing the Nation.

Supporting Information

The computing and the networking required to meet Office of Science needs exceed the state-of-the-art by a wide margin. Furthermore, the algorithms, software tools, the software libraries and the software environments needed to accelerate scientific discovery through modeling and simulation are beyond the realm of commercial interest. To establish and maintain DOE's modeling and simulation leadership in scientific areas that are important to its mission, the MICS subprogram employs a broad, but integrated research strategy. The MICS subprogram's basic research portfolio in applied mathematics and computer science provides the foundation for enabling research activities, which includes efforts to advance networking, to develop software tools, software libraries and software environments. Results from enabling research supported by the MICS subprogram are used by computational scientists supported by other Office of Science and other DOE programs. This link to other DOE programs provides a tangible assessment of the value of the MICS subprogram for advancing scientific discovery and technology development through simulations.

In addition to its research activities, the MICS subprogram plans, develops, and operates supercomputer and network facilities that are available to researchers working on problems relevant to DOE's scientific missions 24 hours a day, 365 days a year.

The Early Career Principal Investigator (ECPI) activity was initiated in FY 2002 for scientists and engineers in tenure track positions at U.S. universities. Seventeen (17) awards were made in FY 2002. Additional awards will be made in FY 2003 for this activity pending peer reviews of applications. The goal of the ECPI activity is to support Office of Science mission-related research in applied Mathematics, computer science and high-performance networks performed by exceptionally talented university investigators, who are at an early stage in their professional careers.

Accomplishments

- *New Robust WAN File Replication and Movement Available for Large Scientific Data* - Terascale computing and large scientific experiments produce enormous quantities of data that require effective and efficient management for large scientific collaborations dispersed across wide-area networks. Using storage resource management (SRM) technology developed at Lawrence Berkeley Laboratory, it is now possible to achieve continuous replication of hundreds of files with a single request with no further human intervention. The SciDAC funded Earth System Grid is using this technology for a generalized data access for climatologists. The technology also provides an important feature--interoperability of archival systems at DOE laboratories and the National Center for Atmospheric Research. The replication process relies on GridFTP, developed at Argonne National Laboratory as part of the Globus Toolkit® for the reliable, secure and policy-aware large-scale data movement. GridFTP servers are used to stage input data and move results to mass storage systems. This software has become the de facto standard worldwide for the movement of large data.
- *Commodity Grid Kits Make Grids Easy to Program and Use* - Many scientific applications, including climate modeling, astrophysics, high energy physics, structural biology, and chemistry need numerous distributed resources to make advances in multidisciplinary scientific research. The Grid provides an infrastructure that can be used to accomplish this. Work at Argonne National Laboratory and Lawrence Berkeley National Laboratory has built on commodity technologies (Java and Python) already in use by thousands of scientists to access the Grid from higher level programming frameworks. These frameworks form the basis for scientific portals promoting collaboration between large scientific teams and allow the harnessing of distributed resources that are needed as part of the scientific problem solving process. The new Commodity Grid (CoG) Kit technology has been used while building a number of scientific portals to the Grid. Indeed, it already has become the de facto community standard for developing Grid portal applications. Examples of applications relying on this technology include the SciDAC funded Particle Physics Data Grid and Earth System Grid projects, portals for astrophysical black-hole simulations, and portals for structural biology.
- *Increased Scientific Productivity through Advanced Collaborative Environments* - The SciDAC funded National Fusion Collaboratory Project has deployed a production computational and data grid that is accelerating scientific understanding and innovation towards the design of an attractive new fusion energy source. The design of this persistent infrastructure to enable scientific

collaboration is being put forth by the U.S. as the design template for ITER, the International Thermonuclear Experimental Reactor that is planned as the one next generation experimental device for the worldwide magnetic fusion community.

- *ParamBench demonstrates the significant impact of concurrent memory accesses* - Computer scientists at the Lawrence Livermore National Laboratory, in collaboration with researchers at the University of Utah and North Carolina State University have implemented ParamBench, low-level benchmarks of memory performance in symmetric multiprocessors (SMPs). These benchmarks measure the raw memory performance of SMPs, including the effect of multiple processors accessing the memory system concurrently. Results with this benchmark suite demonstrate that standard latency-hiding techniques, such as hardware prefetching, are less effective in SMPs, even with a crossbar-based memory interconnection.
- *New Analysis Tools for Innovative Materials* - Mathematicians at the Oak Ridge National Laboratory have extended the class of materials science problems that can be solved by a powerful technique known as the "Boundary Element Method." This numerical method significantly reduces the number of operations that are needed to solve materials science problems, but has traditionally been applicable only to homogeneous materials. The researchers derived the fundamental solution to a set of integral equations for "Functionally Graded Materials," an important class of materials that are not homogeneous, but whose properties vary smoothly. These materials already play an important role in many applications, including coatings for protecting turbine blades, special optical materials, and dental implants and other bio-materials.
- *Scientific Data Objects: A Common Language for Exchanging Parallel Data* - Arrays, or matrices, are one of the basic data structures of scientific computing. In large-scale simulations, arrays are often so large that they must be distributed across many processors. In order for different software modules to work together on a distributed array, a method must exist to precisely describe the distribution of the data. As part of the SciDAC Center for Component Technology for Simulation Software, researchers at the Oak Ridge National Laboratory developed such a description, thus greatly simplifying the development of components that need to exchange distributed array data objects with other components. The new interface specification is capable of describing the layouts used by a wide range of distributed array tools, including CUMULVS (ORNL), Global Arrays (PNNL), High Performance Fortran, A++/P++ (LLNL), and others.
- *New Scientific Data Index Performs 100 Times Faster Than Commercial Database Systems* - Terascale computing and large scientific experiments produce enormous quantities of data that require effective and efficient management. The task of managing scientific data is overwhelming. Researchers at the Lawrence Berkeley National Laboratory have developed a specialized index for accessing very large datasets that contain a large number of attributes that may be queried. This new index performs 12 times faster than the previous best-known method, and 100 times faster than conventional indexing methods in commercial database systems. The prototype index is being used by researchers in high energy physics and combustion modeling.
- *Faster Reconstruction Methods are Making Waves* - Mathematicians at the Lawrence Berkeley National Laboratory have developed efficient and fast techniques for solving the problem of multiple arrivals; that is, detecting and separating the arrival of waves that have taken differing paths through a medium. Example applications include geophysical analysis, which is important for oil exploration, and antenna design. The methods are fast enough that they can be embedded inside

"inverse solvers," computer codes that use information about the arriving waves to deduce the characteristics of an unknown body between the source and detector. This will result in new computational tools to examine hidden objects, accurately reconstruct inaccessible regions, and rapidly test proposed models.

- *Increasing Scientific Productivity through Automated Optimization* - Many complex problems in science, engineering and business require the solution of optimization problems, but the conventional approach to solving such problems can be extremely time-consuming and difficult to apply. Researchers at Argonne National Laboratory have developed the Network-Enabled Optimization System (NEOS) that allows users to solve optimization problems over the Internet with state-of-the-art software and without tedious downloading and linking of specialized optimization code. Because of its ease of use, the NEOS server has gained widespread popularity with more than 5,000 job requests each month from users around the world. Recent NEOS applications include circuit simulation, protein folding, circuit design, brain modeling, airport crew scheduling, and modeling of electricity markets.

- *Open Source Cluster Application Resources (OSCAR) Cluster Software Distribution A Big "Hit" Worldwide* - The Open Source Cluster Application Resources package, OSCAR, is a collection of software tools for managing Linux-based computer clusters developed by a consortium of academic, research, and industry members led by scientists at the Oak Ridge National Laboratory. According to the Top 500 Clusters web site, OSCAR has become the most used cluster computing distribution available today. OSCAR is also used as the core cluster base package in the MacNeil Schwindler (MSC) Linux commercial cluster distribution as well as the NCSA "in-a-box" series of cluster computing solutions. OSCAR has a "market share" of over 30% according to the poll – more than twice its nearest competitor. OSCAR has been downloaded over 53,000 times and has received over 140,000 web page hits during the past year.

- *Tiled Displays: Automatic Calibration of Scalable Display Systems* - Today's scientific simulations and rich multimedia collaborative environments can easily produce tens of millions of pixels for display. Tiled display systems built by combining the images from arrays of projectors can provide massive numbers of pixel elements to visually represent large amounts of information. Multiprojector tiled arrays can be a cost-effective way to create these displays, and they may be the only practical way to create large information dense displays. But, it is difficult to create the illusion of a unified seamless display for a variety of reasons, including projector-to-projector color and luminosity differences, variation of luminosity across the image from a single projector, and optical distortion of the individual projector images caused by imperfections in the lenses and misalignment of projectors. Researchers at Argonne National Laboratory have developed methods to attack these fundamental issues providing an efficient and optimized measurement process using inexpensive components that is tolerant of a wide range of imperfections in components and measurement setup such as lighting conditions and camera optics.

- *Center for Computational Sciences (CCS) Deploys CrayX1 Computer System* - The Center for Computational Sciences at ORNL has acquired and begun deployment of a CrayX1 system to test its effectiveness in solving scientific problems of national scale in climate, biology, nanoscale materials, fusion, and astrophysics. "This partnership with Cray is one of the first steps in the initiative to explore computational architecture essential to 21st century scientific leadership," said the Director of the Office of Science. The Cray X1 is the first U.S. computer to offer vector processing and massively parallel processing capabilities in a single architecture. The system has been specifically designed for scientific applications. Preliminary results on climate applications show the potential for significant improvements in performance over current generation computers.

- *NERSC Improves Supercomputer Performance* **-** The Department of Energy's National Energy Research Scientific Computing Center (NERSC) at the Lawrence Berkeley National Laboratory has doubled the peak performance of its IBM RS/6000 SP supercomputer. The 10 teraflop/s (10 trillion calculations per second) NERSC-3E (enhanced) is now the most powerful computer for unclassified research in the United States. The supercomputer named *Seaborg* has 6,656 processors and has the largest aggregate memory of any unclassified computer in the U.S. – 7.8 terabytes (trillion bytes) – with 44 terabytes of disk storage.
- *Supernova Factory Makes Rapid Discoveries* - At the January 2003 meeting of the American Astronomical Society in Seattle, the Nearby Supernova Factory (SNfactory) based at Lawrence Berkeley National Laboratory announced that it had discovered 34 supernovae during its first year of operation — all but two of them in the last four months alone. This discovery rate of eight per month had been achieved by other supernova search projects only after years of work. The SNfactory processed a quarter-million images in its first year and archived 6 terabytes (trillion bytes) of compressed data at the National Energy Research Scientific Computing Center (NERSC) at Berkeley Lab — one of the few centers with an archive large enough to store this much data.
- *Computational Simulation Finds Correct Theoretical Model* - After three decades of uncertainty, the origins of at least some gamma-ray bursts (GRBs) are being revealed, thanks to a new generation of orbiting detectors, fast responses from ground-based robotic telescopes, and a new generation of computers and astrophysics software. A GRB detected on March 29, 2003 has provided enough information to eliminate all but one of the theoretical explanations of its origin. Computational simulations based on that model were already being developed at the National Energy Research Scientific Computing (NERSC) Center at Lawrence Berkeley National Laboratory when the discovery was made.
- *Improved Algorithm Speeds Up Fusion Code by a Factor of 5* - The NIMROD project, funded by the DOE Office of Fusion Energy Sciences and the SciDAC Center for Extended Magneto-hydrodynamic Modeling, is developing a modern computer code suitable for the study of long-wavelength, low-frequency, nonlinear phenomena in fusion reactor plasmas. The project's primary high-end computing resource is the 10 teraflop/s IBM SP (Seaborg) at the National Energy Research Scientific Computing (NERSC) Center at Lawrence Berkeley National Laboratory. Through a collaboration with members of the SciDAC Terascale Optimal PDE Simulations (TOPS) Center to implement the SuperLU linear solver software within NIMROD it now runs four to five times faster for cutting-edge simulations of tokamak plasmas, with a corresponding increase in scientific productivity.
- *ESnet deploys next general protocol* – ESnet has deployed Internet Protocol Version 6 (IPv6) on its production network. Enabling IPv6 on the network brings a new level of security (e.g. packet encryption and source authentication) and supports real-time traffic, such as video conferencing. IPv6 is expected to become the protocol of choice throughout the Internet.
- *ESnet deploys global Public Key Infrastructure (PKI)* – ESnet has played a key role in coordination and deployment of a Public Key Infrastructure for use by the new computational grids being developed around the world. ESnet is actively working with the Global Grid Forum, the European

Data Grid and Cross Grid Certificate Authority to ensure that the service has the widest possible acceptance.

Awards

Lovelace Medal Awarded to ANL Scientist - An Argonne National Laboratory scientist and colleague of the Information Sciences Institute at the University of Southern California were named as recipients of the prestigious Lovelace Medal, given by the British Computer Society (BCS). BCS cited "their work with the Globus Project and Grid computing," in giving the award for contributions with major significance in the advancement or development of information systems. This is the first time that this award has been given to a DOE-funded researcher. Previous recipients include the developer of the computer mouse and graphic interface; and the developer of the Linux operating system.

Sidney Fernbach Award goes to ORNL Scientist **-** An internationally recognized quantum chemist from ORNL, who is the principal architect of the Northwest Computational Chemistry Software (NWChem), was named the 2002 recipient of the IEEE Computer Society's Sidney Fernbach Award. The Sidney Fernbach Award was established by the IEEE Computer Society in 1992 and is awarded for outstanding contributions in the application of high performance computers using innovative approaches.

Scientific Computing Research Investments

High-performance computing hardware is important for meeting DOE's modeling and simulation needs. However, computer hardware can only enable scientific advances when the appropriate algorithms, scientific software tools, libraries, software environments, and the networking infrastructure are easy to use and are readily available to the users. The MICS subprogram differs from high performance computing efforts in other Federal agencies because of its management focus to integrate research investments to enable new science. Desktop systems realize advances in computing power primarily through increases in the processor's clock speed. High performance computers employ a different strategy for achieving performance, complicating the architecture and placing stringent requirements on software. The MICS subprogram supports software research over a broad range, but that research is tailored to DOE's science needs. Research is underway to improve the performance of simulations on high-end computers, to remove constraints on the human-computer interface and to discover the specialized information management and analysis techniques that scientists need to manage, analyze and visualize extremely large data sets.

Technology trends and business forces in the U.S. computer industry over the past decade caused most domestic vendors to curtail or abandon the development of high-end systems designed to meet the most demanding requirements of scientific research. Instead, large numbers of smaller commercial systems were combined and integrated into terascale computers to achieve the peak performance levels required for agency missions in computational science. The hardware is complicated, unwieldy and not balanced for scientific applications. Enabling software has been developed for scientists to take advantage of these new computers. However, this software is extraordinarily complex and can be a barrier to scientific progress. Consequently, the DOE, primarily through the MICS subprogram, and other Federal agencies whose missions depend on high-performance computing, must make basic research investments to adapt high-performance computing and networking hardware into tools for scientific discovery.

The NGA represents the first step in the adjustment to our strategy that is required to enable future progress in computational science. Continued emphasis on developing software-based solutions to enable scientific simulations on large clusters of computers designed for mid-range applications is no longer the basis for a sustainable strategy for many high-end applications. Rather, our emphasis needs to broaden to include computer hardware technology, architecture, and design trends motivated from a scientific and industrial user perspective. This can be accomplished by making research investments that couple computational scientists, computer scientists, and industrial researchers with U.S. computer vendors to orient future computer architectures towards the needs of science and industry. Additional research investments must be made to ensure the availability of software takes full advantage of these future computer architectures. The status of the technology, the conditions of the current business market for computing, and the success of the Earth Simulator supercomputer in Japan are strong indicators that this strategy can provide tangible near-term benefits for scientific simulation. While NGA will be instrumental in removing some architectural bottlenecks to performance on actual scientific simulations, others will remain and possibly become persistent obstacles in the future.

To illustrate the complexities involved, think of a high-performance computer as a large number of conference rooms distributed around a region. Each conference room is connected through the region's transportation and communications infrastructure. Now, the task of a successful scientific application is analogous to getting everyone in the region to a pre-assigned conference room on time. Instructions are given to each participant (systems software). Results from each conference (calculations) will be documented (stored in memory) for distribution. New conferences are convened, new instructions are given and new decisions are made. Now repeat this process trillions of times, as occurs in a scientific simulation! As one can appreciate, this process can only work if the region's infrastructure is properly configured and operating efficiently. That is, the buses, subways, taxicabs, roads, elevators and telephones can efficiently handle the demand. Most of the systems available from computer vendors are analogous to small regions, a limited number of conference rooms and an inefficient infrastructure. Computers for scientific simulation on the other hand, must be analogous to large cities, large numbers of conference rooms, and an efficient infrastructure, with alternative modes of transportation and communication.

Advances in *computer science* research can enable scientists to overcome these remaining barriers. For example:

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

Research advances in computer science do not provide the full range of capabilities that computational scientists need, especially for the complex problems faced by the Office of Science. Significant efforts in the applied mathematical research activity will be required for the Department to satisfy its mission requirements for computational science. Historically, improvements in mathematical algorithms have yielded at least as much increase in performance as have improvements in hardware. A large proportion of these advances resulted from the MICS subprogram applied mathematics research activity. The requirements of scientific domains for new algorithms that can scale to work effectively across thousands of processors and produce the most science in the fewest number of computer operations drives the need for improved mathematical algorithms and the supporting software libraries that must be made available for ready use by scientists. In this area of research, the MICS applied mathematics activity is the core of the nationwide effort.

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

- Applied Mathematics
- Computer Science
- Advanced Computing Software Tools

High Performance Networking, Middleware and Collaboratory Research Investments

Advances in network capabilities and network-enabled technologies now make it possible for large geographically distributed teams to effectively collaborate on solutions to complex problems. It is now becoming possible to harness and integrate the collective capabilities of large geographically distributed data archives, research teams, and computational resources. This collective capability is especially important for the teams using the major experimental facilities, computational resources, and data resources supported by DOE because all of the necessary resources are not available at one location. To successfully realize the potential of this collective research capability, additional research is needed to bring network, data, and computational resources to the members of a distributed team in a manner that is easy to use and guarantees end-to-end performance. For example:

- Significant research is needed to augment the capability of the Internet to support distributed high-end data-intensive applications and to secure large-scale scientific collaborations. The requirements of high-performance networks that support distributed data-intensive computing and scientific collaborations on a national and international scale are very different than the requirements of the current commercial networks where millions of users are moving small web pages. The MICS-supported research on high-performance networks includes research on high-performance protocols, network-aware operating system services, advanced network coprocessors, network measurement and analysis.
- Research is also needed for the development and testing of high-performance middleware needed to seamlessly couple scientific applications to the underlying transport networks. These include high-performance middleware such as advanced security services for grid computing, ultra-high-speed data transfer services, services to guarantee Quality of Service (QoS) for delay sensitive applications, and grid resources discovery. This high-performance middleware provides the scalable software components needed to integrate data, visualization, computation and high-speed networks into a scalable and secure scientific collaborative environment.

The MICS subprogram will address these challenges through fundamental research in networking; software tools that integrate networking and computer science to enable scientific collaboration (collaboratory tools); partnerships with key scientific disciplines; and advanced network testbeds.

Specific responses to these challenges are described in the Detailed Program Justification section of the MICS subprogram budget under the headings:

- Networking
- Collaboratory Tools
- National Collaboratory Pilot Projects

Enhancements to High Performance Computing and Networking Facilities

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

- *Production High Performance Computing Facilities.* The National Energy Research Scientific Computing Center (NERSC) provides high performance computing for investigators supported by the Office of Science. NERSC provides a spectrum of supercomputers offering a range of high performance computing resources and associated software support.
- *Energy Sciences Network (ESnet).* ESnet provides worldwide access to Office of Science facilities, including light sources, neutron sources, particle accelerators, fusion reactors, spectrometers, high-end computing facilities, massive data resources and other leading-edge instruments and facilities.
- *Advanced Computing Research Testbeds.* The Advanced Computing Research Testbeds (ACRTs) consist of high performance, advanced architecture computing platforms for testing and evaluation to ascertain the prospects for meeting future general, or specialized, computational science needs of the Office of Science. In FY 2005, the ACRTs will provide hardware resources for the NGA activity. Two types of computing platforms will be evaluated - early systems from vendors, and experimental systems. Based on an analysis of vendor offerings and a peer review of the potential that such offerings can meet Office of Science computational needs, hardware will be acquired at sufficient scale to address key performance and software scaling issues. The evaluation process will include computer science studies and tests of leading-edge Office of Science computational science applications, such as those being developed under SciDAC. In addition, the ACRTs will provide computing resources to SciDAC teams.
- *Trends for Future Supercomputing and Networking Resources.* The need for high performance computational resources will increase in future years as applications transition from the software development and testing phase to using the software to generate new science. As the peak performance of the computers increase, the amount of data produced in a simulation increases as well. Therefore, focused enhancements to the Office of Science's network infrastructure are required to enable scientists to access and understand the data generated by their software and by large-scale science experiments.

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

- National Energy Research Scientific Computing Center (NERSC)
- Advanced Computing Research Testbeds
- Energy Sciences Network (ESnet)

Detailed Justification

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005
Mathematical, Computational, and Computer Sciences Research	**68,748**	**83,301**	**86,405**
▪ **Applied Mathematics**	**21,332**	**22,635**	**29,363**

This activity supports research on the underlying mathematical understanding of physical, chemical and biological systems, and on advanced numerical algorithms that enable effective description and prediction of such systems on terascale computing systems. Research in Applied Mathematics supported by MICS underpins computational science throughout DOE. Historically, the numerical algorithms developed under this activity have produced more scientific advances through simulation than improvements in computer hardware. This activity supports research at DOE laboratories, universities, and private companies. Many of the projects supported by this activity involve research partnerships between DOE's national laboratories and universities. The activity supports research in a wide variety of areas of mathematics, including: ordinary and partial differential equations and solutions methods, including techniques to convert equations into discrete elements and boundary integral methods, advanced treatment of interfaces and boundaries, (fast marching and level set methods, and front tracking); numerical linear algebra (advanced iterative methods, general and problem-specific preconditioners, sparse solvers, and dense solvers); fluid dynamics (compressible, incompressible and reacting flows, turbulence modeling, and multiphase flows); optimization (linear and nonlinear programming, interior-point methods, and discrete and integer programming); mathematical physics; control theory (differential-algebraic systems, order reduction, and queuing theory); accurate treatment of shock waves; "fast" methods (fast multipole and fast wavelet transforms); mixed elliptic-hyperbolic systems; dynamical systems (chaos theory, optimal control theory, and bifurcation theory); and automated reasoning systems.

The FY 2005 budget continues the Computational Sciences Graduate Fellowship program at the current level of $3,500,000.

The FY 2005 budget also includes $8,500,000, for the new "Atomic to Macroscopic Mathematics" (AMM) research effort to provide the research support in applied mathematics needed to break through the current barriers in our understanding of complex physical processes that occur on a wide range of interacting length- and time-scales. The current state-of-the-art in the theory and modeling of complex physical systems generally requires that the physical phenomena being modeled either occur at a single scale, or widely separated scales with little or no interaction. Complex physical systems frequently involve highly nonlinear interactions among many

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005

phenomena at many different scales. Increases in computational power over the last decade have enabled scientists to begin the process of creating sophisticated models with fewer simplifying assumptions. These new models cannot succeed without a deeper understanding of the mathematics of phenomena at multiple scales and how they interact, from the atomic scale through the mesoscopic to the macroscopic. Achieving this basic mathematical understanding will provide enabling technology to virtually every challenging computational problem faced by the Office of Science.

Progress in AMM will best be achieved through a combination of investments, including (1) funds for innovative approaches to multiscale mathematics at universities throughout the country, (2) investments in partnerships between university researchers and investigators at the national laboratories, and (3) additional investments in multidisciplinary teams at the national laboratories. Category (1) represents investment in relatively high-risk / high-payoff approaches. Categories (2) and (3) follow the SciDAC model of building teams that involve national laboratory researchers in various critical applications. AMM research will support the development of new high-fidelity simulations that are crucial to our improved understanding of important problems across the Office of Science, including fuel cell design, understanding of microbial cells and communities, accelerator design and optimization, combustion processes including clean and efficient engine design, fusion reactor design and optimization, design of materials atom-by-atom, and many more.

	FY 2003	FY 2004	FY 2005
▪ **Computer Science**	**17,489**	**23,680**	**23,680**

This activity supports research in computer science to enable researchers to effectively utilize high-performance computers to advance science in areas important to the DOE mission. DOE has unique requirements for high performance computing that significantly exceed the capability of software products from computer vendors. This activity supports computer science research in two general areas: the underlying software to enable applications to make effective use of computers with hundreds or thousands of processors as well as computers that are located at different sites; and large scale data management and visualization under circumstances where the underlying resources and users are geographically distributed. The first area includes research in protocols and tools for interprocessor communication and parallel input/output (I/O) as well as tools to monitor the performance of scientific applications and advanced techniques for visualizing very large-scale scientific data. Researchers at DOE laboratories and universities, often working together in partnerships, propose and conduct this research.

Beginning in FY 2004, this activity incorporates the software research component of NGA to improve application performance and system reliability through innovative approaches to next generation operating systems. In FY 2005, NGA effort in this activity also includes $2,000,000 transferred from Scientific Applications Partnerships for applications teams working in close partnership with systems evaluation teams. Total funding for the NGA software research component research is $6,659,000.

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005
▪ **Advanced Computing Software Tools**....................	**19,362**	**20,256**	**19,362**

This activity supports research that builds on the results from research in applied mathematics and computer science to develop integrated software tools that computational scientists can use to develop high performance applications (such as characterizing and predicting phase changes in materials). These tools, which enable improved performance on high-end systems, are critical to the ability of scientists to attack the complex scientific and engineering problems that can only be solved with high-end computing systems.

In FY 2005, this activity will continue to support the Integrated Software Infrastructure Centers (ISICs), a SciDAC activity, competitively selected in FY 2001. The ISICs funded under this activity focus on: structured and unstructured mesh generation for large simulations and high performance tools for solving partial differential equations on parallel computers; tools for analyzing the performance of scientific simulation software that uses thousands of processors; the development of data management and visualization software capable of handling terabyte scale data sets extracted from petabyte scale data archives, and software for managing computers with thousands of processors; and software component technology to enable rapid development of efficient, portable, high performance parallel simulation software.

These Integrated Software Infrastructure Centers are a critical component in DOE's SciDAC strategy. The ISICs are responsible for the entire lifecycle of the software that they develop. These software tools must be reliable, understandable and well documented. Also, the scientific user community needs these tools to be maintained, bug-free and upgraded, as necessary. Software tools for high performance scientific simulations have no commercial market. The Integrated Software Infrastructure Centers initiated in FY 2001 provide the only means for developing and deploying these tools to the scientific community.

There is a decrease of $894,000 in the last year of the SciDAC program for ISICs resulting from accelerated rampdown of selected efforts. This decrease is required to support an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures.

▪ **Scientific Applications Partnerships**	**10,565**	**16,730**	**14,000**

This activity, formerly titled Scientific Application Pilot Projects, supports collaborative research with computational scientists in other disciplines to apply the computational techniques and tools developed by other MICS activities to address problems relevant to the mission of SC. This effort tests the usefulness of advances in computing research, transfers the results of this research to the scientific disciplines, and helps define opportunities for future research. The FY 2005 funding for this activity will allow the continuation of the multidisciplinary partnerships that were competitively selected in FY 2001. These projects are part of the SciDAC activity and are coupled to the Integrated Software Infrastructure Centers. Areas under investigation include design of particle accelerators with the High Energy Physics (HEP) and Nuclear Physics (NP) programs; plasma turbulence in tokamaks with the Fusion Energy Sciences (FES) program; global climate

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005

change with the Biological and Environmental Research (BER) program; and combustion chemistry with the Basic Energy Sciences (BES) program.

The FY 2005 request includes funds to continue at a reduced level the partnership with the Biological and Environmental Research Genomics: GTL and the partnership with the Basic Energy Sciences program in nanoscale science. The FY 2005 request also includes $1,350,000 to expand SciDAC partnerships with the Fusion Energy Sciences program to lay the groundwork for the Fusion Simulation Project (FSP). The FSP will be a focused, interdisciplinary effort, whose objective is to develop the capability to predict reliably the behavior of fusion plasma. The NGA effort ($2,000,000) for applications teams working in close partnership with systems evaluation teams is shifted to the Computer Science activity FY 2005.

	FY 2003	FY 2004	FY 2005
Advanced Computation, Communications Research, and Associated Activities..	**91,619**	**110, 553**	**112,389**
▪ **Networking**	**8,736**	**7,066**	**5,784**

The DOE national laboratories, the unique instruments and facilities at those laboratories, and the university community contributing to the DOE missions create a complex, distributed system that is conducting scientific research on a wide range of problems that depend, increasingly, on high-performance network infrastructure. This program activity is one component of an integrated ASCR effort to develop and deploy a scalable, secure, integrated environment to support these network-intensive science applications, a Network Environment. The components of this effort are the MICS activities: Network Research, Collaboratory Tools, and Collaboratory Pilots. Together, these activities support research and development in high performance networking and middleware along with collaboratory pilots and testbeds that are largely partnerships with other program offices to test, demonstrate, and validate the technologies that derive from that research.

This activity advances the Network Environment vision by supporting research and development in high-performance networks needed to develop and deploy advanced networking capabilities to address challenging issues such as ultra-high-speed data transfers, remote visualization, real-time remote instrumentation, and large-scale scientific collaboration. Networking research is carried out at national laboratories and universities and consists of two major elements:

Network R&D – to address the fundamental issues of high-performance networks to support access to the next generation of scientific facilities, terascale computing resources and distributed petabyte-scale data archives. Network R&D focuses on leading-edge networking technologies such as ultra optical transport protocols and services for ultra high-speed data transfers; techniques and tools for ultra high-speed network measurement and analysis; advanced network tools and services to enable network-aware, high-end scientific applications; and scalable cyber-security technologies for open science environment.

Advanced experimental networking – to accelerate the adoption of emerging networking technologies and to transfer networking R&D results into production networks that support science applications. It includes activities such as experimental networking testbeds, advanced deployment and evaluation of new networking technologies, and exploration of advanced

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005

networking concepts. A rapid adoption of emerging network capabilities into production networks will enable scientists pushing the limits of today's networks capabilities to use networking technologies to conduct far-reaching experiments.

There is a decrease of $1,282,000 in the level of support for network research activities in FY 2005. This will reduce research activities at universities and laboratories in high performance network protocols and optical networks. This decrease is required to support an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures.

▪ **Collaboratory Tools**	**7,403**	**5,527**	**5,527**

The DOE national laboratories, the unique instruments and facilities at those laboratories, and the university community contributing to the DOE missions create a complex, distributed system that is conducting scientific research on a wide range of problems that depend, increasingly, on high-performance network infrastructure. This program activity is one component of an integrated ASCR effort to develop and deploy a scalable, secure, integrated environment to support these network-intensive science applications, a Network Environment. The components of this effort are the MICS activities: Network Research, Collaboratory Tools, and Collaboratory Pilots. Together, these activities support research and development in high performance networking and middleware along with collaboratory pilots and testbeds that are largely partnerships with other program offices to test, demonstrate, and validate the technologies that derive from that research.

This activity advances the network environment vision by supporting research that builds on results of fundamental research in computer science and networking to develop an integrated set of software tools to support scientific collaborations. These tools provide a new way of organizing and performing scientific work that offers the potential for increased productivity and efficiency and will also enable broader access to important DOE facilities and data resources by scientists and educators throughout the country. It is particularly important to provide for efficient, high-performance, reliable, secure, and policy-aware management of large-scale data movement. This research includes an effort to develop a set of essential middleware services required to support large-scale data-intensive collaboratory applications. This research also includes an effort to research, develop, and integrate the tools required to support a flexible, secure, seamless collaboration environment that supports the entire continuum of interactions between collaborators. The goal is to seamlessly allow collaborators to locate each other; use asynchronous and synchronous messaging; share documents, progress, results, and applications; and hold videoconferences. There is also research for developing and demonstrating an open, scalable approach to application-level security in widely distributed, open network environments that can be used by all the collaboratory tools as well as by the advanced computing software tools whenever access control and authentication are issues.

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005
▪ **National Collaboratory Pilot Projects**....................	**9,380**	**10,857**	**8,013**

The DOE national laboratories, the unique instruments and facilities at those laboratories, and the university community contributing to the DOE missions create a complex, distributed system that is conducting scientific research on a wide range of problems that depend, increasingly, on high-performance network infrastructure. This program activity is one component of an integrated ASCR effort to develop and deploy a scalable, secure, integrated environment to support these network-intensive science applications, a Network Environment. The components of this effort are the MICS activities: Network Research, Collaboratory Tools, and Collaboratory Pilots. Together, these activities support research and development in high performance networking and middleware along with collaboratory pilots and testbeds that are largely partnerships with other program offices to test, demonstrate, and validate the technologies that derive from that research.

This activity advances the Network Environment vision by supporting research that tests, validates, and applies collaboratory tools in real-world situations in partnership with other DOE programs. The competitively selected partnerships involve national laboratories, universities, and U.S. industry. It is important to demonstrate and test the benefits of collaboratory tools' evolving technology in order to promote their widespread use and enable more effective access to the wide range of resources within the Department, from light sources to terascale computers to petabyte data storage systems. The partnerships that were initiated in FY 2001 focus on developing user environments where collaboration is ubiquitous and distributed computing is seamless and transparent for DOE mission applications. The Particle Physics Data Grid is developing middleware infrastructure to support High Energy Physics (HEP) and Nuclear Physics (NP) communities, and to enable grid-enabled data-management ("manipulation") and analysis capabilities "at the desk of every physicist." It is building one unified system that will be capable of handling the capture, storage, retrieval and analysis of particle physics experiments at the five most critical research facilities, a key collaboratory issue being the highly distributed access to, and processing of, the resulting data by a worldwide research community. In another scientific community, the Earth System Grid II is developing a virtual collaborative environment linking distributed centers, models, data, and users that will facilitate exchange among climatologists all over the world and provide a needed platform for the management of the massive amounts of global climate data that are being generated. Development of this and similar concepts is essential for rapid, precise, and convincing analysis of short- and long-term weather patterns, particularly in the period when increasing pollution introduces changes that may affect us for generations to come. The National Fusion Collaboratory is centered on the integration of collaborative technologies appropriate for widely dispersed experimental environments and includes elements of security, distributed systems, and visualization. All three of these pilot collaboratories will rely on the DOE Science Grid to provide the underpinnings for the software environment, the persistent grid services that make it possible to pursue innovative approaches to scientific computing through secure remote access to online facilities, distance collaboration, shared petabyte datasets and large-scale distributed computation.

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005

There is a decrease of $2,844,000 in the level of support for National Collaboratory Pilot projects due to the accelerated rampdown of selected efforts. This decrease is required to support an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures.

	FY 2003	FY 2004	FY 2005
▪ **National Energy Research Scientific Computing Center (NERSC)**...	**22,515**	**31,906**	**37,868**

Located in a single facility at the Lawrence Berkeley National Laboratory (LBNL), NERSC delivers high-end capability computing services and support to the entire DOE Office of Science (SC) research community. NERSC provides these services to the DOE community, to the other DOE laboratories, and to major universities performing work relevant to DOE missions. NERSC provides the majority of resources and services that are used to support the Office of Science SciDAC program. The Center serves 2,000 users working on about 700 projects; 35 percent of users are university based, 61 percent are in National Laboratories, 3 percent are in industry, and 1 percent in other government laboratories. The major computational resource at NERSC is an IBM SP computer. The initial installation of hardware, which was completed in FY 2001 following a fully competitive process, provided a peak performance of 5 trillion floating point operations per second (teraflops) to its users. The capability of this system was increased to 10 teraflops following the acquisition of additional computer hardware in FY 2003. The FY 2005 funding will support the continued operation of the IBM SP computer at 10 teraflops peak performance. These computational resources are integrated by a common high performance file storage system that facilitates interdisciplinary collaborations. Related requirements for capital equipment, such as high-speed disk storage systems, archival data storage systems, and high performance visualization hardware, and general plant projects (GPP) funding are also supported. FY 2005 capital equipment requirements for these types of capital equipment remain at the same level as in FY 2004.

In FY 2004 and FY 2005, the NERSC budget is increased as a part of an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures. This enhancement will enable NERSC to competitively procure a significant new high performance computer to support the missions of the Office of Science. The enhancement at NERSC will deliver an increase of approximately 30% in the high performance computing resources available to scientists as well as associated improvements in storage and network systems to enable scientists to most effectively use NERSC resources.

	FY 2003	FY 2004	FY 2005
▪ **Advanced Computing Research Testbeds (ACRTs)**	**25,934**	**36,909**	**36,909**

This activity supports the acquisition, testing and evaluation of advanced computer hardware testbeds to assess the prospects for meeting future computational needs of the Office of Science, such as SciDAC and special purpose applications. The ACRT activity will provide two types of

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005

computer testbeds for evaluation - early systems and experimental systems. Each testbed will involve significant research and architecture design activities. These research and evaluation (R&E) prototypes have been identified as a critical element in the HECRTF plan because they enable early partnership with vendors to tailor architectures to scientific requirements. The results from these partnerships also play a key role in the choice of both high performance production systems and potential leadership class systems government-wide.

The FY 2005 request continues an enhanced scope for the hardware evaluation testbed in the Next Generation Architecture (NGA) research activity as a part of an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures. A goal of these testbeds is to identify and address major architectural bottlenecks, such as: internal data movement in very large systems, to the performance of existing and planned DOE science applications. Total funding for the testbed activities in the NGA is $31,553,000. Emphasis will be placed on understanding the impact of alternative computer architectures on application performance with particular attention paid to data movement from memory to processor and between processors in highly parallel systems. The enhanced scope in the hardware evaluation testbed will improve our ability to understand key issues impacting application performance scalability. The NGA activity will be coordinated with other Federal agencies to gain additional insight into research directions, optimize the utilization of resources, and establish the framework for a national effort and is aligned with the goals of the HECRTF plan. Funding for these testbeds will be allocated through peer reviewed competition. These testbeds, coupled with the NGA software research, play a critical role in enabling potential future leadership class scientific computing facilities for open science.

▪ **Energy Sciences Network (ESnet)**	**17,651**	**18,288**	**18,288**

ESnet is a Wide Area Network (WAN) project that supports the scientific research mission of the Department of Energy. The ESnet project/investment supports the agency's mission and strategic goals and objectives by providing DOE with interoperable, effective, and reliable communications infrastructure and leading-edge network services in support of the agency's science research missions. ESnet supplies the DOE science community with capabilities not available through commercial networks or the commercial Internet. ESnet provides national and international high-speed access to DOE and Office of Science researchers and research facilities, including: light sources; neutron sources; particle accelerators; fusion reactors; spectrometers; supercomputers; and other high impact scientific instruments. ESnet provides the communications fabric that interconnects geographically distributed research facilities and large-scale scientific collaborations. ESnet supplies the critical infrastructure that links DOE researchers worldwide and forms the basis for advanced experimental research in networking, collaboratory tools, and distributed data-intensive scientific applications testbeds such as the national collaboratory pilot projects. ESnet provides network services through contracts with commercial vendors for advanced communications services including Asynchronous Transfer Mode (ATM), Synchronous Optical Networks (SONET) and Dense Wave Division Multiplexing (DWDM). ESnet provides interfaces between the network fabric it provides and peering arrangements with other Federal,

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005

education and commercial networks, international research network connections, and the University Corporation for Advanced Internet Development (UCAID) Abilene network that provides high performance connections to many research universities.

In FY 2005, funds will be used to operate ESnet and to continue support for upgrading the capability of the ESnet backbone to 10,000 million bits per sec (Mbps) from its current capability of 155 Mbps. Remaining funds will be used to upgrade networking hardware and services at high priority ESnet sites to exploit the enhanced performance capabilities of the backbone. FY 2005 capital equipment requirements remain at the same level as in FY 2004.

	FY 2003	FY 2004	FY 2005
SBIR/STTR	**0**	**5,438**	**5,546**

In FY 2003, $3,942,000 and $236,000 were transferred to the SBIR and STTR programs, respectively. The FY 2004 and FY 2005 amounts are the estimated requirement for the continuation of the SBIR and STTR program.

	FY 2003	FY 2004	FY 2005
Total, Mathematical, Information, and Computational Sciences	**160,367**	**199,292**	**204,340**

Explanation of Funding Changes

	FY 2005 vs. FY 2004 ($000)
Mathematical, Computational, and Computer Sciences Research	
▪ **Applied Mathematics**. Provides an increase to support initiation of Atomic to Macroscopic Mathematics research effort ($+8,500,000). The increase is offset by a decrease of $1,772,000 from the existing program.	+6,728
▪ **Computer Science.** Core research is decreased $2,000,000 to support an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures. This decrease is offset by NGA funding of $2,000,000 transferred from Scientific Application Partnerships.	0
▪ **Advanced Computing Software Tools.** Decrease in last year of SciDAC program for ISICs resulting from rampdown of selected efforts. This decrease is required to support an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures.	-894

	FY 2005 vs. FY 2004 ($000)
▪ **Scientific Application Partnerships.** The change in this activity includes reductions in partnerships with BER ($-367,000) and BES ($-472,000) and an increase of $1,350,000 to expand SciDAC partnerships with the Fusion Energy Sciences program to lay the groundwork for the Fusion Simulation Project (FSP). The NGA effort for applications teams working in close partnership with systems evaluations teams is shifted to the Computer Science activity in FY 2005 ($-2,000,000). A decrease in core research ($-1,241,000) is required to support an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures.	-2,730
Advanced Computation, Communications Research, and Associated Activities	
▪ **Network Research**. Decrease in level of support for network research activities. This will reduce research activities at universities and laboratories in high performance network protocols and optical networks. This decrease is required to support an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures.	-1,282
▪ **National Collaboratory Pilots.** Decrease in the level of support for National Collaboratory Pilot projects because of accelerated rampdown of selected efforts. This decrease is required to support an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures.	-2,844
▪ **National Energy Research Scientific Computing Center.** Provides an increase to enable installation of a major new resource for computational scientists with an architecture different from the current NERSC resource. This increase supports an Office of Science strategy to acquire additional advanced computing capability to support existing users in the near term and to initiate longer-term research and development on next generation computer architectures.	+5,962
SBIR/STTR	
▪ Increase in SBIR/STTR due to increase in operating expenses.	+108
Total Funding Change, Mathematical, Information, and Computational Sciences	**+5,048**

Laboratory Technology Research

Funding Schedule by Activity

	(dollars in thousands)				
	FY 2003	FY 2004	FY 2005	$ Change	% Change
Laboratory Technology Research	2,818	2,916	0	-2,916	-100%
SBIR/STTR	0	84	0	-84	-100%
Total, Laboratory Technology Research	2,818	3,000	0	-3,000	-100%

Description

The Laboratory Technology Research (LTR) subprogram will be brought to a successful conclusion in FY 2004 with orderly completion of all existing CRADAs. The mission of the Laboratory Technology Research subprogram was to support high-risk research that advances science and technology to enable applications that could significantly impact the Nation's energy economy. LTR fostered the production of research results motivated by a practical energy payoff through cost-shared collaborations between the Office of Science (SC) laboratories and industry. The termination of the LTR subprogram does not mean that technology transfer activities have ended; rather, due to the impact of this subprogram, these activities are now institutionalized as a part of the process of doing research at DOE sites.

Benefits

LTR supported ASCR's contribution to DOE's mission of world-class scientific research capacity by promoting the transfer of these research results to the private sector. The success of this program has institutionalized these processes in all of the programs within the Office of Science; therefore these processes are now integrated into the other programs and the LTR subprogram is no longer needed.

Detailed Justification

	(dollars in thousands)		
	FY 2003	FY 2004	FY 2005
Laboratory Technology Research	**2,818**	**2,916**	**0**

This activity supported research to advance the fundamental science at the Office of Science (SC) laboratories toward innovative energy applications. Through CRADAs, the SC laboratories entered into cost-shared research partnerships with industry, typically for a period of three years, to explore energy applications of research advances in areas of mission relevance to both parties. The existence of the LTR subprogram fostered the institutionalization of technology transfer activities at DOE sites. Now that these activities are institutionalized, a separate program to fund them is no longer necessary.

	FY 2003	FY 2004	FY 2005
(dollars in thousands)			
SBIR/STTR	**0**	**84**	**0**

In FY 2003, $75,000 and $5,000 were transferred to the SBIR and STTR programs, respectively. The FY 2004 amount is the estimated requirement for the continuation of the SBIR and STTR program.

	FY 2003	FY 2004	FY 2005
Total, Laboratory Technology Research	**2,818**	**3,000**	**0**

Explanation of Funding Changes

	FY 2005 vs. FY 2004 ($000)
Laboratory Technology Research	
▪ The Laboratory Technology Research subprogram will be brought to a successful completion in FY 2004, as planned.	-2,916
SBIR/STTR	
▪ Decrease in SBIR/STTR due to completion of the LTR subprogram.	-84
Total Funding Change, Laboratory Technology Research	**-3,000**

Capital Operating Expenses & Construction Summary

Capital Operating Expenses

	(dollars in thousands)				
	FY 2003	FY 2004	FY 2005	$ Change	% Change
Capital Equipment (total).........................	3,962	6,290	6,250	-40	-0.6%