Scientific Discovery through Advanced Computing

Office of Science

U.S. Department of Energy

March 24, 2000

Scientific Discovery through Advanced Computing

Office of Science

U.S. Department of Energy

March 24, 2000

"The computer literally is providing a new window through which we can observe the natural world in exquisite detail."

"National Workshop on Advanced Scientific Computing"
J. S. Langer, Workshop Chair, July 1998

Scientific Discovery through Advanced Computing

Executive Summary

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

Advances in computing technologies during the past decade have set the stage for a major step forward in modeling and simulation. Within the next five years, computers 1,000 times faster than those available to the scientific community today, *i.e.*, *terascale* computers, will be at hand. However, to deliver on this promise, these increases in "peak" computing power, *i.e.*, the maximum theoretical speed that a computer can attain, must be translated into corresponding increases in the capabilities of scientific codes. This is a daunting problem that will only be solved by increased investments in *computer software*—the scientific codes for simulating physical phenomena, the mathematical algorithms that underlie these codes, and the computing systems software that enables the use of high-end computer systems.

In the FY 2001 budget, the Office of Science (SC) has proposed a set of coordinated investments focused on the scientific and computational problems that must be solved to address the critical scientific challenges in all of SC's research programs. SC has a long history of use of and accomplishments in scientific computing and has often served as the proving ground for many new computer technologies. SC now intends to bring its experience and expertise to bear to realize the promise of terascale computers for its basic science programs. The SC-wide effort will focus on:

- *Scientific Challenge Codes*—research, development, and deployment of mathematical models, computational methods, and scientific codes to take full advantage of the capabilities of terascale computers as well as use of this new research capability to solve critical scientific problems in SC's research programs.
- *Computing Systems and Mathematical Software*—research, development, and deployment of software to accelerate the development of and protect long-term investments in scientific codes, to achieve maximum efficiency on high-end computers, and to enable a broad range scientists to use simulation in their research.
- *Collaboratory Software Infrastructure*—research on network technologies and research, development, and deployment of software to link geographically separated researchers, to facilitate movement of large (petabyte) data sets, and to ensure that academic scientists can fully participate in the activities described above.

These activities will be supported by upgrades to SC's existing *Scientific Computing Hardware Infrastructure*. This infrastructure has been designed to meet the needs of SC's research programs. It is *robust*—to provide computing resources for scientific research; *agile*—to respond to innovative advances in computer technology; and *flexible*—to ensure that the most effective and efficient resources are used to solve each class of problems.

The above investments by the Office of Science will produce a *Scientific Computing Software Infrastructure* that bridges the gap between advanced research in applied mathematics and computer science and computational science research in the physical, chemical, biological, and environmental sciences. The *Scientific Computing Software Infrastructure*, combined with the *Scientific Computing Hardware Infrastructure*, will allow researchers supported by the Office of Science to solve challenging scientific problems at a level of accuracy and detail never before achieved.

Computational modeling and simulation are among the most significant developments in the practice of scientific inquiry in the 20th Century. In the past century, scientific research has been extraordinarily successful in identifying the fundamental physical laws that govern our material world. At the same time, the advances promised by these discoveries have not been fully realized, because the real-world systems governed by these physical laws are extraordinarily complex. Computer-based simulation provides a means of solving the mathematical equations and predicting the behavior of complex systems that can only be described empirically at present. Since the development of digital computers in mid-century, scientific computing has greatly advanced our understanding of the fundamental processes of nature, *e.g.*, fluid flow and turbulence in physics, molecular structure and reactivity in chemistry, and drug-receptor interactions in biology. Computational simulation has even been used to explain, and sometimes predict, the behavior of such complex natural and engineered systems as weather patterns and aircraft performance.

Within the past two decades, scientific computing has become a contributor to essentially all scientific research programs. It is particularly important to the solution of research problems that are (*i*) insoluble by traditional theoretical and experimental approaches, *e.g.*, prediction of future climates or the fate of underground contaminants; (*ii*) hazardous to study in the laboratory, *e.g.*, characterization of the chemistry of radionuclides or other toxic chemicals; or (*iii*) time-consuming or expensive to solve by traditional means, *e.g.*, development of new materials, determination of the structure of proteins, understanding plasma instabilities, or exploring the limitations of the "Standard Model" of particle physics. In many cases, theoretical and experimental approaches do not provide sufficient information to understand and predict the behavior of the systems being studied. Computational modeling and simulation, which allows a description of the system to be constructed from basic theoretical principles and the available experimental data, are key to solving such problems.

Advances in the simulation of complex scientific and engineering systems provide an unparalleled opportunity for solving major problems that face the nation in the 21st Century.

Advances in computing technologies during the past decade have set the stage for a major step forward in modeling and simulation. Within the next five years, advances in microprocessor and networking technologies coupled with increased exploitation of parallel computing will lead to computers able to perform more than 100 trillion arithmetic operations per second (teraflops), *i.e.*, *terascale computers* that are 1,000 times faster than those generally available today. However, to exploit this opportunity, these increases in "peak" computing power, *i.e.*, the maximum theoretical speed that a computer can attain, must be translated into corresponding increases in the capabilities of scientific modeling and simulations codes. With the last generation of supercomputers, represented by the "vector" supercomputers produced by Cray Research, Inc., this was possible—many scientific codes realized 40% to 50% of the peak performance of the supercomputer. In contrast, on microprocessor-based parallel supercomputers, scientific computing codes often realize only 5% to 10% of "peak" performance, and this fraction will decrease with increasing use of parallelism.

At issue is *computer software*—the scientific codes for simulating physical phenomena, the mathematical algorithms that underlie these codes, and the computing systems software that enables the use of terascale computer systems. As noted in the report from the National Science Foundation (NSF)/DOE *National Workshop on Advanced Scientific Computing*, "[terascale computers] have far outstripped our ability to manage parallelism and to deliver large fractions of peak performance." The President's Information Technology Advisory Committee also noted that "Government organizations face the challenge of converting their applications to new high-end architectures without disrupting their missions. This transition has proved to be technically difficult…" Both reports recommended increased funding for software research and development to improve the performance of high-end computing and to provide access to the coming generation of terascale computers for the civilian research community.

Advanced scientific computing is key to accomplishing the missions of the U.S. Department of Energy (DOE). It is essential to the design of nuclear weapons, the development of new energy technologies, and the discovery of new scientific knowledge. All of the research programs in DOE's Office of Science—in Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, and High-Energy and Nuclear Physics—have identified major scientific questions that can only be addressed through advances in scientific computing (see sidebar). As the lead government funding agency for basic research in the physical sciences, the Office of Science (SC) has a special responsibility to ensure that its research programs continue to advance the frontiers of science. To do so requires significant enhancements to SC's scientific computing programs in Advanced Scientific Computing Research, Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, and High-Energy and Nuclear Physics. This plan outlines those enhancements.

SC RESEARCH PORTFOLIO:
CRITICAL QUESTIONS THAT NEED ANSWERS

- Can we predict the effects of cracking, aging and fatigue on materials?
- Can we improve the efficiency and specificity of the catalysts that produce the materials of the modern world?
- Can we predict the structure and function of proteins from a knowledge of the DNA sequence?
- Can we reliably predict the evolution of the earth's regional climates decades and centuries into the future?
- Can we control the instabilities that lead to the loss of power in fusion devices?
- Can we design more efficient heavy-ion accelerators for inertial fusion?
- Can we design more powerful particle accelerators for high-energy physics?
- Is the Standard Model of particle physics complete?

Computational Modeling and Simulation

Development of computational modeling and simulation software for terascale computers is a demanding task. The work flow is illustrated in Figure 1. Theoretical scientists must first lay the foundation for the mathematical description of the phenomena of interest—*basic theoretical research*. Computational scientists, in collaboration with applied mathematicians, who provide the basic mathematical algorithms, and computer scientists, who provide the basic computer systems software, cast these equations into algorithms and codes to make efficient use of high-performance computers—*scientific code development*. Computational and research scientists then benchmark the new software by comparing computed results with existing experimental data—*scientific code validation*. If the simulation does not reproduce the known experimental data, the work flow recycles to the theoretical scientists, if the problem appears to be with the basic mathematical model, or to the computational scientists, if the problem appears to be with the computational approach. If the simulations reproduce the experimental data, the code becomes a powerful tool for interpreting new experimental data, designing new experiments, and predicting new phenomena. As the code is applied to increasingly complex systems, computational scientists may have to extend the mathematical description of the phenomena of interest. This cycle, as well as the code performance cycle noted on the right in Figure 1, will be repeated many times, often over a span of years, as increasingly more realistic descriptions of physical phenomena are sought for minimal computational cost.

For large scientific codes, the above process requires many person-years of effort. For example, the GAUSSIAN code for which John Pople was awarded the 1998 Nobel Prize in Chemistry has been under continuous development since 1970. There are more than one-half million lines of code in GAUSSIAN, with an estimated investment of several hundred person-years. Because of this investment, algorithmic advances have contributed as much to increases in the performance of GAUSSIAN as have increases in the computing speed of computers and has resulted in a code that is widely used in the chemistry research community. Although GAUSSIAN does not run on massively parallel computers, a related program, NWCHEM, one of the projects in DOE's Grand Challenge program, was developed for massively parallel computers. NWCHEM was created by an international team of scientists. The "core" team consisted of 15 to 20 full-time theoretical and computational chemists, computer scientists, and applied mathematicians over the 5-year life of the project. This team was complemented by a somewhat larger number of collaborators from five U.S. universities, two U.S. national laboratories, three European universities, and two European laboratories. The total investment in NWCHEM thus far exceeds 100 person-years.

Figure 1. Work flow for the development of scientific modeling and simulation codes.

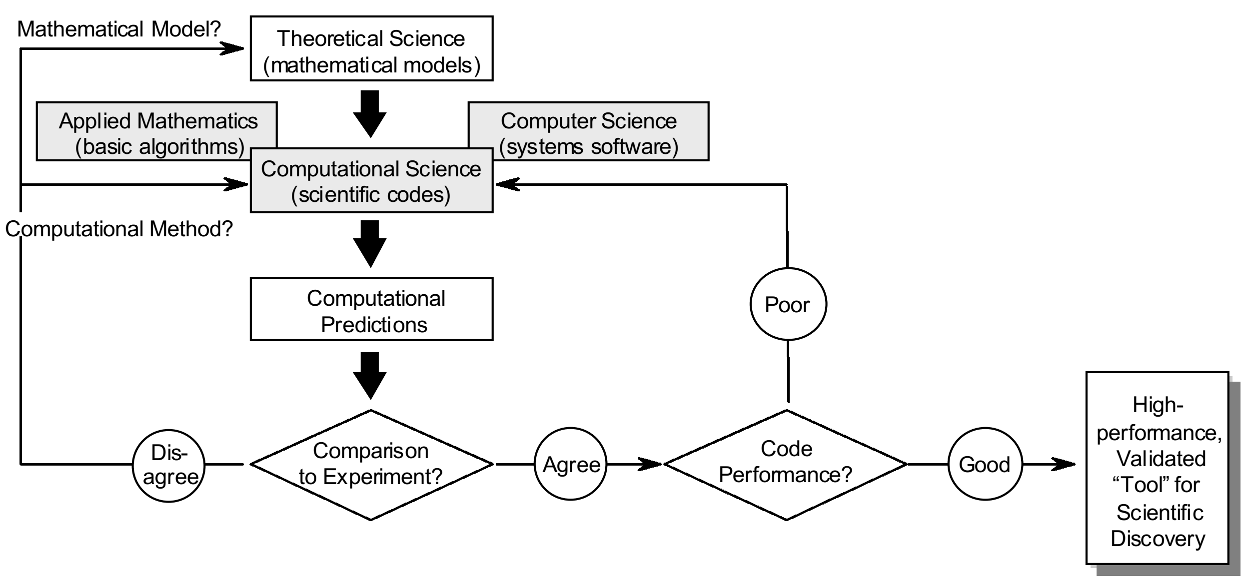

Three critical issues must be addressed in developing scientific modeling and simulation codes—performance, portability, and adaptability. Scientific codes are among the most demanding applications in computing. This point is illustrated in the sidebar, which lists the general characteristics of a calculation of the energy of the iso-octane molecule; such calculations will be repeated hundreds of times in research directed at understanding the chemistry of this important fuel hydrocarbon. As noted in the sidebar, calculation of the energy of iso-octane requires the iterative solution of 275 million non-linear equations, the exchange of 2½ petabytes of data between the processors and 15 terabytes of data between the processors and disks, and approximately 30 quadrillion arithmetic operations. This calculation clearly pushes the limits of what is possible—the computing power of the processors as well as the speed of the interprocessor and processor-disk communications channels. If the code does not achieve top performance from all of the computer hardware, the calculation quickly becomes impossible.

COMPLEXITY OF SCIENTIFIC PROBLEMS: COMPUTING THE ENERGY OF ISO-OCTANE

- Iterative solution of 275 million non-linear equations.
- Exchange of 2½ petabytes of data among the processors.
- Exchange of 15 terabytes of data between the processors and the disks.
- Execution of 30 quadrillion arithmetic operation (30 petaflops).

Portability and adaptability of scientific codes is important because the lifetimes of scientific codes (30 years in the case of GAUSSIAN!) are much longer than either the theories and methods on which they are based or the computers on which they are run. The architecture of NWCHEM was designed and implemented by a tightly integrated team of computational chemists and computer scientists with these goals in mind. The architecture, which is layered (to facilitate portability) and modular (to facilitate adaptation), is a prime example of the efficiencies that can be achieved through multidisciplinary collaborations. Only a small fraction of the code in one of the three layers of NWCHEM must be modified to run the software on a different computer—a process that typically takes less than a week. Further, many of the modules are "plug and play," allowing new theoretical and computational advances to be readily incorporated in the code. NWCHEM runs on a broad range of computers, from the parallel supercomputers available in the DOE and NSF national centers to the notebook computers carried by business travelers. The carefully crafted architecture of NWCHEM saves countless months porting the software from one computer to the next, or from one generation of computers to the next, or incorporating new theories or methods. NWCHEM has been installed in more than 200 research institutions worldwide.[1]

The work in scientific code development has led to tremendous progress in computational studies of scientific systems and processes. To measure progress, one need only consider the size of molecules that can now be modeled using GAUSSIAN—several hundred atoms now, *e.g.*, a zeolite catalyst, versus a few atoms in 1970—or the accuracy with which many important physical properties can now be calculated.

[1] NWCHEM, along with two companion software packages, the Extensible Computational Chemistry Environment (ECCE) and Parallel Software Development Tools (PARSOFT), won an R&D 100 Award in 1999 and a Federal Laboratory Consortium Excellence in Technology Transfer Award in 2000.

In Figure 2, the accuracy of predicted molecular atomization energies[2] is plotted from 1970-2000. Knowledge of atomization energies is critical for understanding many chemical processes, including combustion, atmospheric chemistry, and industrial chemical production. In 1970, the average error in calculated atomization energies was so large that the simulations were totally useless for real-world chemical applications. By 2000, however, through a combination of theoretical advances and increases in computing power, the error has been reduced by a factor of almost 100. Calculations are now more accurate than most experimental measurements! There are, of course, limitations on the size of molecules that can be simulated using computational techniques, but advances in theoretical methodology and computer technology continue to push this limit to larger and larger molecules.

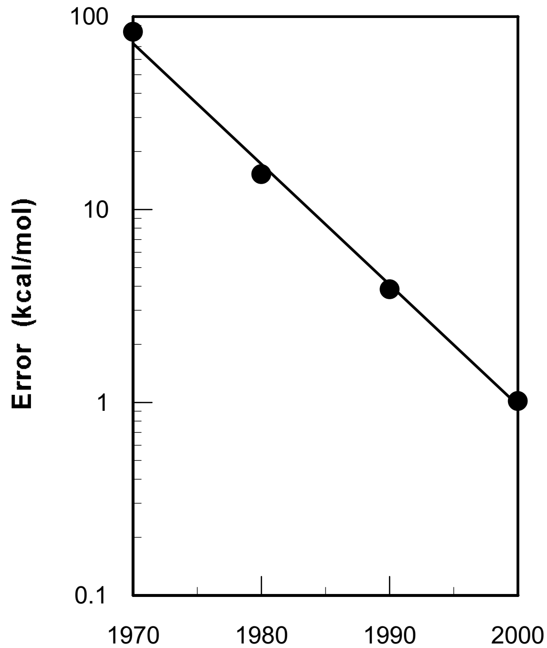

Figure 2. Accuracy of computational predictions of molecular atomization energies.

The above are typical of the issues addressed, results obtained, and conclusions drawn in DOE's most successful Grand Challenge projects, including projects in accelerator physics, materials science, ocean circulation, particle physics, and plasma physics.

COMPLEX PROBLEMS OF BASIC AND STRATEGIC IMPORTANCE

- Fundamental nature of matter
- Origin of the universe
- Behavior of plasmas
- Design of materials
- Weather prediction
- Climate prediction
- Analysis of impacts of earthquakes
- Improvement of human health
- Solution of environmental problems
- Design of combustion devices

From report of "National Workshop on Advanced Scientific Computing," July 30-31, 1998.

Opportunities and Challenges

The coming advances in computing performance, if they can be realized for scientific problems, herald a new era in scientific computing. With computers capable of 100 teraflops or more, it will not only be possible to dramatically extend our exploration of the fundamental processes of nature, it will be possible to predict the behavior of a broad range of complex systems. The opportunities and challenges offered by terascale computing have been the subject of two recent reports. The first report, from the "National Workshop on Advanced Scientific Computing," focused on the advances in science and technology that will be

[2] The atomization energy of a molecule is the energy required to dissociate the molecule into its constituent atoms. The larger this number, the more strongly the atoms are bound together in the molecule and the more difficult it is to break the molecule apart.

enabled by terascale computers.[3] The meeting's participants, which included the nation's leaders in computational science and engineering, computer science, and applied mathematics, noted that advances in scientific and engineering simulation would contribute to the solution of "complex problems of both basic and strategic importance" to the nation (see sidebar). The participants recommended that the U.S. launch a vigorous, 5-year program to make the most powerful computers available to the national scientific and engineering community and to invest concurrently in the development of the software, algorithms, and communications infrastructure necessary for effective use of these computers.

The second report, from the President's Advisory Committee on Information Technology (PITAC) in February 1999,[4] noted that high-end computing systems are essential for a broad range of activities of national importance, including scientific and engineering research. The Committee found that "new applications for high-end computing are ripe for exploration," but that investments in software research were inadequate to address these new opportunities. They recommended increased government funding for the development of innovative hardware and software approaches that overcome the limitations of today's high-end computers. Recognizing the nature of the challenges to be faced in the 21st Century, PITAC also recommended funding projects of larger scope and duration.

As noted in the above two reports, the promise of terascale computers will not be realized without additional investments by the U.S. Government. The advances in information technology are being driven by strong market forces in the commercial sector, not by scientific computing. Harnessing commercially based technology for scientific research poses problems unlike those encountered in previous generations of supercomputers, in magnitude as well as in kind. For example, to minimize costs several levels of memory (cache, local memory, remote memory) are used in microprocessor-based, parallel supercomputers, with more time being required to fetch the data from each succeeding level of memory. Because scientific calculations tend to be dominated by memory accesses (a feature first recognized by the legendary Seymour Cray), the scientific software developer must manage the location of data in the memory hierarchy to ensure that the data is as "close" to the processor as possible when it is needed. Otherwise, performance can decrease one or two orders of magnitude. Another example is the operating system software provided with large scientific computers. Commercial applications tend to use 100 processors or less, not the thousands required for scientific computations. As a result, the vendor operating systems are less complete and not as robust for scientific applications. Investments will be needed to enhance vendor-supplied operating systems.

SC Investment Plan

In the FY 2001 budget, SC has proposed a set of coordinated investments focused on the scientific and computational problems that must be addressed to overcome the barriers to progress in its research programs. SC has a long history of use of and accomplishments in scientific computing and has often

---

[3] This workshop was sponsored by the National Science Foundation and the Department of Energy and hosted by the National Academy of Sciences on July 30-31, 1998. Copies of the report may be obtained from: http://www.er.doe.gov/production/octr/mics/index.html

[4] Copies of the PITAC report may be obtained from: http://www.ccic.gov/ac/report/.

served as the proving ground for many new computer technologies—subjecting these technologies to the demands that only its most compute-intensive problems could provide. SC now intends to bring its experience and expertise to bear to realize the promise of terascale computers for its basic science programs. The SC-wide effort will address the following activities:

- *Scientific Challenge Codes*. This effort will be focused on research, development, and deployment of scientific codes, mathematical models, and computational methods to take full advantage of the extraordinary capabilities of terascale computers as well as application of this new capability to solve critical scientific problems in SC's research programs.
- *Computing Systems and Mathematical Software.* This activity focuses on research, development, and deployment of software to accelerate the development of and protect long-term investments in scientific codes, to achieve maximum efficiency on terascale computers, and to enable a broad range scientists to use computational modeling and simulation in their research.
- *Collaboratory Software Infrastructure.* This effort will be focused on research on network technologies and research, development, and deployment of software to link geographically separated researchers, to facilitate movement of large (petabyte) data sets, and to ensure that academic scientists can fully participate in the activities described above.

This set of investments by the Office of Science will produce a *Scientific Computing Software Infrastructure* that bridges the gap between advanced research in mathematics and computer science and computational science research in physics, chemistry, biology and environmental science. This gap has been identified as a significant barrier to effective use of terascale computers for scientific discovery.[3,4] These research efforts will be supported by:

- *Scientific Computing Hardware Infrastructure.* Upgrades to SC's existing computing infrastructure have been carefully designed to meet SC program needs. It is robust—to provide the computing resources needed for scientific research; agile—to respond to innovative advances in computer technology; and flexible—to ensure that the most effective and efficient resources are used to solve each class of problems.

The above investments build on DOE's established strengths in scientific modeling and simulation, computer science, applied mathematics, and high-performance computing as well as in the organization and management of large scientific projects and user facilities (see sidebar). The investments leverage the capabilities being developed in the existing SC programs as well as the leading-edge computing and simulation capabilities being developed in the

DOE ACCOMPLISHMENTS IN SCIENTIFIC COMPUTING

- Established the first national supercomputing center—for the fusion energy sciences community—in 1974.
- Developed the first time-share operating system for Cray supercomputers.
- Developed major mathematical libraries used by both vector and parallel supercomputers.
- Developed major communications libraries used by parallel supercomputers.
- Developed High Performance Parallel Interface (HiPPI)—the international standard for interconnecting supercomputers to other high-speed devices.
- "Computational Grand Challenges" led to dramatic increases in modeling and simulation capabilities in a number of scientific fields, including chemistry, materials science, plasma physics, ocean circulation, and accelerator physics.
- Developed the first scientific application to sustain over one teraflops.

Accelerated Strategic Computing Initiative (*ASCI*) in DOE's National Nuclear Security Agency (see section on "Relationship to ASCI").

The present effort also builds on SC's experience in its Grand Challenge program, which focused on the development of proof-of-principle modeling and simulation capabilities in a number of scientific disciplines. This program highlighted the importance of

- Integration of computational scientists with applied mathematicians and computer scientists in research and development of scientific codes for advanced computers.
- Multi-institutional teams from the national laboratories, universities, and government and industrial laboratories to provide the depth as well as breadth of scientific expertise needed to harness high-end computing for scientific applications.
- Complementary research and development efforts in applied mathematics and computer science to address problems generic to a number of scientific applications.

The Grand Challenge program demonstrated successfully what disciplinary computational scientists could achieve with access to advanced computer science, applied mathematics, and computer systems and set the stage for a major effort to change dramatically how scientists use computational modeling and simulation to solve the most challenging scientific problems. The goal of the current investments is to make this level of scientific simulation capabilities available to the broader scientific community—to theoretical and computational scientists as a tool for exploring fundamental scientific problems, to experimental scientists as a tool for guiding experimental studies, and to engineers and technologists as a tool for predicting the behavior of complex systems.

In the sections immediately below are descriptions of the three components of this plan—**Scientific Challenge Codes**; **Mathematical, Computing Systems, and Collaboratory Software**; and the **Scientific Computing Hardware**. Following these sections is an overview of the investment portfolio for the Division of Mathematical, Information, and Computing Sciences within the Office of Advanced Scientific Computing Research (ASCR) showing the FY 2000 base and the funding increments required to effect this plan. Incremental investments in disciplinary theoretical and computational science in other programs—in the Offices of Biological and Environmental Research (BER), Basic Energy Sciences (BES), Fusion Energy Sciences (FES), and High-Energy and Nuclear Physics (HENP)—and the relationship of these investments to the investments in ASCR are also shown. A section on future opportunities follows. Finally, the relationship of the elements of this plan to ASCI, NSF and other agencies, and to experimental efforts in other parts of DOE is discussed.

Scientific Challenge Codes

The availability of computers capable of 100 teraflops or more will have a profound impact on the ability to simulate the fundamental physical, chemical, and biological processes that underlie complex systems. Chemists will be able to model the diverse set of chemical processes involved in the combustion of hydrocarbon fuels and the catalytic production of materials. Materials scientists will be able to predict the properties of materials from knowledge of their structure, and then invert the process and design materials with a targeted set of properties. Physicists will be able to model a broad range of complex phenomena—from the fundamental interactions between elementary particles, to high-energy nuclear processes and particle-electromagnetic field interactions, to the interactions that govern the behavior of fusion plasmas. Biologists will be able to predict the structure and, eventually, the function of the 140,000 proteins coded in the human DNA. All of the SC programmatic offices—BES, BER, FES, and HENP—have identified major scientific challenges that only can be addressed through advances in computational modeling and simulation. The corresponding research goals of these offices are summarized in Table 1.

The "engines" that drive scientific discovery are the *Scientific Codes* and the physical and

Table 1. Research objectives in advanced computational modeling and simulation in the programs of Biological and Environmental Research, Basic Energy Sciences, Fusion Energy Sciences, and High-Energy and Nuclear Physics.

SC Office	Research Goals
BES	▪ Predicting and understanding aging, cracking, fatigue, and catastrophic failure in materials. ▪ Predicting and understanding the effects of substituents on the electronic, physical, and chemical properties of materials and designing new materials for desired properties. ▪ Predicting the interaction of chemical reactivity and fluid dynamics to understand the mixing of reactants and the removal of products for applications as diverse as combustion, subsurface transport, atmospheric chemistry, and chemical processing. ▪ Modeling chemical reactivity at the nanoscale to understand the behavior of such processes as catalysis and chemical vapor deposition. ▪ Modeling the transport of fluid and solid phases within the Earth's shallow crust to understand contaminant transport to enable waste disposal decisions, assess nuclear waste repositories, enable geophysical imaging and subsurface visualization, and determine the most efficient use of natural resources (coordinated with BER).
BER	▪ Predicting the earth's climate at both regional and global scales for decades to centuries, including levels of certainty and uncertainty. ▪ Understanding the fundamental processes that govern the workings of a cell, including the structure and function of the proteins and other molecular species involved. ▪ Predicting the fate of contaminants in the subsurface, including all relevant physical, chemical, and biological processes involved (coordinated with BES).
FES	▪ Understanding turbulence and transport, macroscopic equilibrium and stability, magnetic reconnection, electromagnetic wave/particle interactions, boundary layer effects in plasmas, and plasma/material interactions ▪ Predicting energy and particle confinement in plasmas, stability limits in magnetically confined plasmas, heating and current drive in plasmas using radio frequency waves, transport of heat and particles in the edge region of a fusion device, and plasma/wall interactions in a fusion device. ▪ Understanding electromagnetic fields and beam dynamics in heavy-ion accelerators to design an efficient, high-current accelerator for inertial fusion applications.
HENP	▪ Predicting electromagnetic fields and beam dynamics in particle accelerators, with particular attention to processes, such as beam halos, that impact the performance of current and proposed high-energy accelerators (coordinated with FES). ▪ Understanding the physical phenomena encompassed in the Standard Model of particle physics to determine whether additional theoretical concepts are needed to explain fundamental interactions at very high energies or short distances. ▪ Understanding the structure of nuclei as well as nuclear processes involved in energetic events such as stellar supernovae explosions.

mathematical models on which they are based. Development of scientific codes for terascale computers poses demanding technical challenges. The mathematical algorithms must achieve high-performance levels on each microprocessor and scale to thousands of processors (if not more). In contrast, many existing algorithms achieve only 5% to 10% of the "peak" performance of the processors, and few algorithms scale beyond a few hundred processors. In addition, the scientific codes must be carefully designed so that they can easily accommodate new mathematical models and computational methods as knowledge advances. Scientific codes must continually evolve as new knowledge is gained. It must also be possible to move the codes from one generation of computers to the next without undue difficulty as computer technology advances. The lifetime of scientific codes is measured in decades, the lifetime of computers in years.

As in the Grand Challenges, the scope and complexity of the scientific codes required to solve the problems listed in Table 1 require close collaboration among researchers from the disciplines of theoretical and computational sciences, computer science and applied mathematics. To obtain the needed depth and breadth of scientific expertise, close collaboration will also be required among researchers at the national laboratories, universities, and government and industrial laboratories. To achieve the goal—a new generation of terascale Scientific Challenge Codes—SC intends to fund multidisciplinary, multi-institutional teams for sufficient periods of time to allow the projects to reach completion:

- *Scientific Challenge Teams*—teams to research, develop, and deploy advanced computational modeling and simulation codes and new mathematical models and computational methods that take full advantage of the capabilities of terascale computers, and to use the new modeling and simulations codes to address the challenging scientific problems identified in Table 1.

The Scientific Challenge Teams will be selected in open, peer-reviewed competitions with selection based on the merit of the work proposed, the software engineering practices that will be used to create the codes, and the plan for deploying and supporting the codes. The teams will include, as needed, theoretical and computational scientists, possibly from a number of scientific and engineering disciplines; computer scientists; applied mathematicians; and software engineers. Collaboratory technologies, which has been pioneered in projects funded by ASCR, will help to knit these dispersed researchers into integrated, cohesive units. The teams will support the full life-cycle of the scientific codes—from research through software development and engineering to deployment and support. This is essential if computational modeling and simulation are to be adopted by the larger scientific community in pursuit of SC goals.

Mathematical, Computing Systems, and Collaboratory Software

Many of the computer science and applied mathematics issues in scientific computing are common to a number of scientific applications. These include libraries for basic mathematical operations, an environment and tools to facilitate the development of complex Scientific Challenge Codes, and the software infrastructure that underlies the management and analysis of massive (petabyte) data sets, distributed computing, and disciplinary problem-solving environments as well as enhancements to vendor operating systems. To eliminate duplication of effort, these issues will be addressed in a separate, but coordinated activity.

The overall goals of the activity in *Mathematical, Computing Systems, and Collaboratory Software* are to:

- Accelerate the development of the scientific applications codes needed to solve challenging scientific problems and protect these investments from evolution in computer hardware.
- Enable theoretical, computational, and experimental scientists as well as teams of scientists to *use* and *understand* the results of advanced modeling and simulation in their research.
- Enable the integration of multi-institutional, geographically dispersed research teams to address the most challenging research problems faced by the Office of Science.

Each of these goals represents a substantial commitment to a new research effort at the leading edge of computer science and applied mathematics. *Without these investments, however, it will not be possible to realize the full potential of terascale computers for scientific discovery.*

The SC effort in computer science and applied mathematics builds on existing research programs funded by ASCR, enhancing selected activities and building new programs in those areas required to meet the above goals. Scientific challenges that must overcome in the ASCR program include:

- Algorithms for basic mathematical operations that scale to thousands, and eventually millions, of processors for a variety of computer architectures.
- Software environments and tools that enable the development of scientific codes demanding the highest level of performance from all components of a computer system (processor, memory, communications channels, and disks).
- Scientific data management and analysis software systems to enable the extraction of knowledge from the massive (petabyte-scale) data sets produced by simulations on terascale computers.
- Collaboratory software infrastructure that enables the integration of researchers from different institutions into effective research teams, taking into account the differing performance levels of the local computing environments and networks.
- Operating systems for terascale computers with thousands of processors that provide all of the features needed for scientific computing.

The major research areas to be funded by ASCR are summarized in Table 2.

Multi-institutional teams will also be required to successfully pursue these problems. To this end, SC will fund:

- *Enabling Technology Centers*—teams of computer scientists, applied mathematicians, applications scientists, and software engineers responsible for researching, developing, and deploying the mathematics and computer systems software needed by a number of scientific applications to realize the full potential of terascale computers for scientific discovery.

The centers will be selected through open, peer-reviewed competitions with selection based on the merit of the proposed work as well as the plans proposed for ascertaining the critical needs of the scientific

Table 2. Research areas in computer science and applied mathematics that provide the basic computer science and applied mathematics software infrastructure needed by the Scientific Challenge Codes.

SC Office	Research Areas
ASCR	▪ Basic mathematical methods, algorithms and libraries that scale to thousands, and eventually a million, processors. ▪ Code development environments and tools to enable the development of complex, scientific challenge codes. ▪ Software infrastructure for disciplinary problem-solving environments to enable the use of computational modeling and simulation by a broad range of scientists. ▪ Scientific data management and analysis (visualization) systems to enable the extraction of knowledge from the massive (petabyte) data sets produced by scientific simulations. ▪ Collaboratory software to enable integration of multi-institutional, geographically dispersed researchers into effective, efficient teams. ▪ Operating system enhancements to provide the basic infrastructure needed to use massively parallel computer systems. ▪ Distributed computing and collaboration environments and tools to support use of remote computing. ▪ Computational Sciences Graduate Fellowships to provide additional fellowships to train the next generation of computational scientists.

research challenges and engineering and supporting the deployed software. The teams will support the full life-cycle of the computer systems software infrastructure. This is essential if this software is to serve as the basis for a new generation of scientific modeling and simulation codes. The Enabling Technology Centers will be closely coupled to the Scientific Challenge Teams through a variety of mechanisms, ranging from joint advisory groups to joint staff appointments. Only through such active coordination will the capabilities of terascale computers be fully realized.

Scientific Computing Hardware Infrastructure

Enhancements to SC's existing Scientific Computing Hardware Infrastructure will be required to solve the scientific problems identified in previous sections. New computing capabilities will be needed to provide resources for the computational simulations. New data storage and communications capabilities will be required to manage the massive data sets created by terascale simulations. However, within the next few years, many pathways to high-end computing are possible. It is critical that the computing infrastructure be sufficiently flexible that it can benefit from the continuing advances in computing.

The SC computing infrastructure has been tailored to meet the needs of its research programs. It is *robust*, to provide the computing resources needed by the scientific applications; *agile*, to respond to innovative advances in computer technology that impact scientific computing; and *flexible*, to ensure that the most appropriate and economical resources are used to solve each class of problems. The three major elements in the *Scientific Computing Hardware Infrastructure* are a Flagship Computing Facility, Topical Computing Facilities, and Computer Technology Testbeds (see sidebar and below).

Flagship Computing Facility

To provide the large-scale computing resources needed by the scientific applications, the ½-teraflops Cray computer system currently in production use at the National Energy Research Scientific Computing (NERSC) facility at Lawrence Berkeley National Laboratory will be upgraded to a 5-teraflops IBM system in FY 2001. This facility has a distinguished history of service to the SC research community, providing the community with general purpose, high-end computing capabilities not available elsewhere in the SC laboratory system. This upgrade can be achieved by making a modest additional investment in the NERSC budget.

SCIENTIFIC COMPUTING HARDWARE INFRASTRUCTURE

- *Flagship Computing Facility*—to provide high-end computing resources for all of the Office of Science.
- *Topical Computing Facilities*—to provide computing resources tailored for specific scientific applications and to serve as a focal point for an application community.
- *Experimental Computing Facilities*—to assess the promise of new computing technologies for scientific applications.

Topical Computing Facilities

The most appropriate, cost-effective computing resources for scientific simulations vary significantly from application to application. Climate modeling codes are extremely sensitive to the bandwidth between memory and processors, molecular modeling codes benefit substantially from large memories and local disk storage, and quantum chromodynamics codes run very efficiently on specially built computer systems. An excellent example of the potential of tailoring the computer to the application is illustrated by the IBM "Blue Gene" computer[5] that is targeted to become available in 2004-5. This computer, which is designed to operate at 1,000 teraflops (a petaflops) on protein-folding applications, is expected to cost less than 1% of a general purpose, 1-petaflops supercomputer.

The SC proposes to establish, beginning in FY 2002, two to three moderately sized Topical Computing Facilities focused on specific scientific applications. One or more closely related Scientific Challenge Teams will be associated with each facility. These Team(s) will work closely with the teams in the Enabling Technology Centers and facility management to select computers optimized for the application. In some instances, the computers in the Topical Facilities may be special configurations of more generally available computers. In other instances, the computers may be designed and constructed for a specific application area, *e.g.*, for quantum chromodynamics or protein folding.

The SC already operates two prototype Topical Computing Facilities—the Molecular Science Computing Facility at Pacific Northwest National Laboratory and the Columbia-BNL-Riken Quantum Chromodynamics (QCD) Computer System at Brookhaven National Laboratory. The first uses an IBM parallel computer configured for molecular science calculations. This not only allows the IBM computer to achieve record-breaking CPU efficiency rates, attesting to the synergy between the molecular science codes and the machine configuration, but also permits calculations that are not feasible with less optimal configurations. The NWCHEM code, developed at the MSCF, also serves as the basis of community

[5] "IBM Plans Computer To Work at Speed of Life," S. Lohr, *New York Times*, 6 December 1999. As noted in the article, "[t]he hardware design of 'Blue Gene' is innovative indeed, but the real challenge, as is so often the case in computing, will be the software."

development efforts for massively parallel computers. The second facility, which is based on a computer system built specifically for quantum chromodynamics calculations, won a Gordon Bell Award in 1998 as the best price/performance computer system currently available (the 0.6 teraflops computer cost less than $2 million when it was built and would cost even less at today's prices). These facilities clearly illustrate the benefits of focusing on a scientific application area.

To encourage the development of new computing approaches and capabilities, an appropriate fraction of the computing resources at the Topical Computing Facilities will be made available to scientific research efforts outside the area of specialization.

Experimental Computing Facilities

Although the high-performance computers of the future will rely on parallelism to achieve the highest performance levels, many implementations of this architecture are possible. Possible implementations range from the commodity-based architectures of the ASCI computers, to the high-performance processor-memory systems being developed and manufactured in Japan, to the application-specific design of the IBM "Blue Gene" computer system. As new computer technologies become available that promise to greatly decrease the cost or extend the range of scientific simulations, prototype computer systems will be installed in one or two Experimental Computing Facilities. These facilities provide critical experimental equipment for computer science and applied mathematics researchers as well as computational science pioneers to assess the potential of new computer systems and architectures for scientific computing. Production versions of successful new computer systems may be installed in either the Flagship or Topical Facilities as appropriate. The Experimental Computing Facilities represent a continuation of the Advanced Computing Research Facilities with increased emphasis on prototyping and evaluation studies aimed at long term gains in scientific computing capability.

High Speed Communications Network

Use of high-performance computers requires high-performance, high-speed networks. These requirements will be met through the strategic planning process for ESnet, which periodically assesses SC's communications requirements. ESnet is one of the nation's most effective and progressive science-related computer networks, providing worldwide access and communications to SC's computing and experimental facilities. The ESnet Strategic Plan is prepared by the ESnet Steering Committee and is available at http://www.es.net/.

Implementation of the entire SC Scientific Computing Hardware Infrastructure will be coordinated with other government agencies.

Technical Organization and Management

Development of the scientific modeling and simulation codes and computing systems software and mathematical algorithms, upgrades of SC's computers and networks, and exploration of complex scientific problems are interdependent activities. They must be closely coordinated if major advances in scientific modeling simulation capabilities are to be realized. As noted in previous sections, the SC

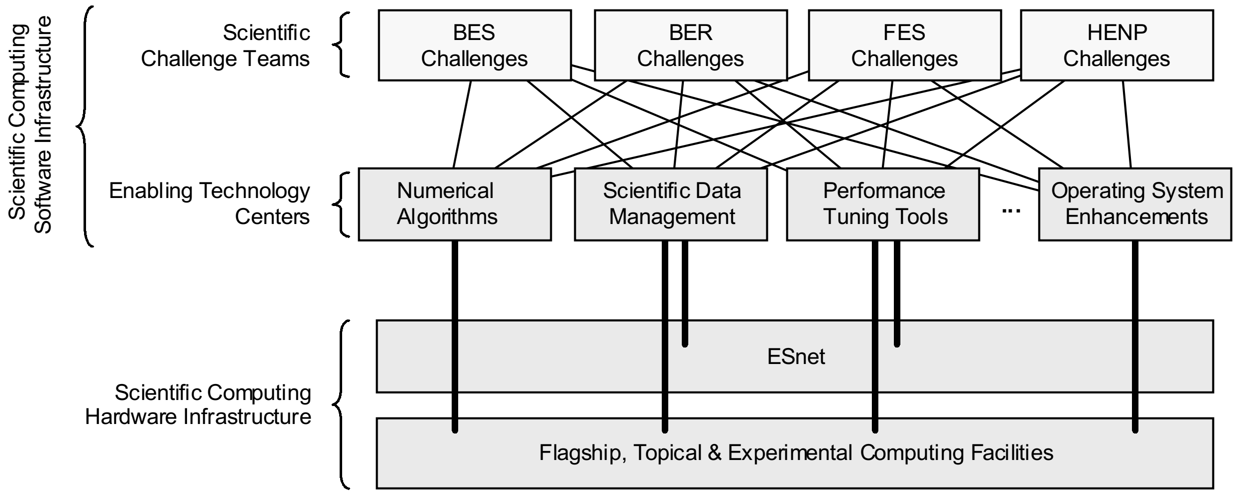

Figure 3. Relationship between the Scientific Challenge Teams, Enabling Technology Centers, Scientific Computing Software Infrastructure, and Scientific Computing Hardware Infrastructure (see also Figure 1).

investments will be organized around Scientific Challenge Teams, Enabling Technology Centers, and the Scientific Computing Hardware Infrastructure (Flagship Computing Facility, Topical Computing Facilities, Experimental Computing Facilities, and ESnet). The Scientific Challenge Teams and the Enabling Technology Centers will be focused on the development of a *Scientific Computing Software Infrastructure* that will bring the capabilities of terascale computers to bear on the scientific challenges faced by the Office of Science. This goal cannot be achieved without careful coordination of all of these activities. The relationships among the various efforts are illustrated in Figure 3. A management plan will be a required for each proposed Team, Center, and Facility. The management plan will address coordination of activities within the Team/Center/Facility, interaction and coordination of the Team/Center/Facility with other dependent Teams/Centers/Facilities, and development, deployment and support of the software created in the Teams, Centers, and Facilities.

DOE has a long history of managing and implementing complex scientific projects using distributed research teams. For example, the current U.S. Atlas Project for the Large Hadron Collider (LHC) at the European Laboratory for Particle Physics (CERN) is a $163,750,000 project involving three DOE laboratories and 29 universities. The Atlas Project also requires close interaction with European scientists associated with CERN. Such projects have been undertaken by all of the SC programmatic offices at one time or another. In the current effort, all of the SC programmatic offices will be involved. To provide high-level coordination of this effort, an *Assistant Director* has been appointed in the Office of Science. This individual is working closely with staff in the SC programmatic offices to coordinate the new efforts across the Office of Science, while simultaneously integrating those new efforts into the existing computational modeling and simulation programs. This individual will also be responsible for coordinating the activities in SC with those in ASCI and in other government agencies.

The FY 2001 Investment Portfolio

Figure 4 provides an overview of the ongoing activities in ASCR; the relationships between those activities and activities in the other SC programs; and the increases proposed in FY 2001 to effect the program described in this plan. More detailed descriptions of the new activities are given in the SC budget narratives.

The ASCR currently supports basic research in applied mathematics, computer science, and computer networks, and operates advanced computing and network facilities to provide scientific computing and collaboratory tools. The proposed FY 2001 investments build on this base, increasing the budget of the ASCR's Mathematics, Information, and Computational Sciences Division from $119,071,000 to $169,682,000. Investments fall into three categories: advancing science through computation, collaboratories, and advanced computing and network facilities.

In the category of "advancing science through computation," increases are requested in applied mathematics, computer science, advanced computing software tools, and scientific application pilot

Figure 4. Overview of *funding increases* requested in the research activities in the Office of Advanced Scientific Computing Research/Mathematical, Information and Computational Sciences Division as well as in the Offices of Basic Energy Sciences Biological and Environmental Research, Fusion Energy Sciences, and High-Energy and Nuclear Physics.

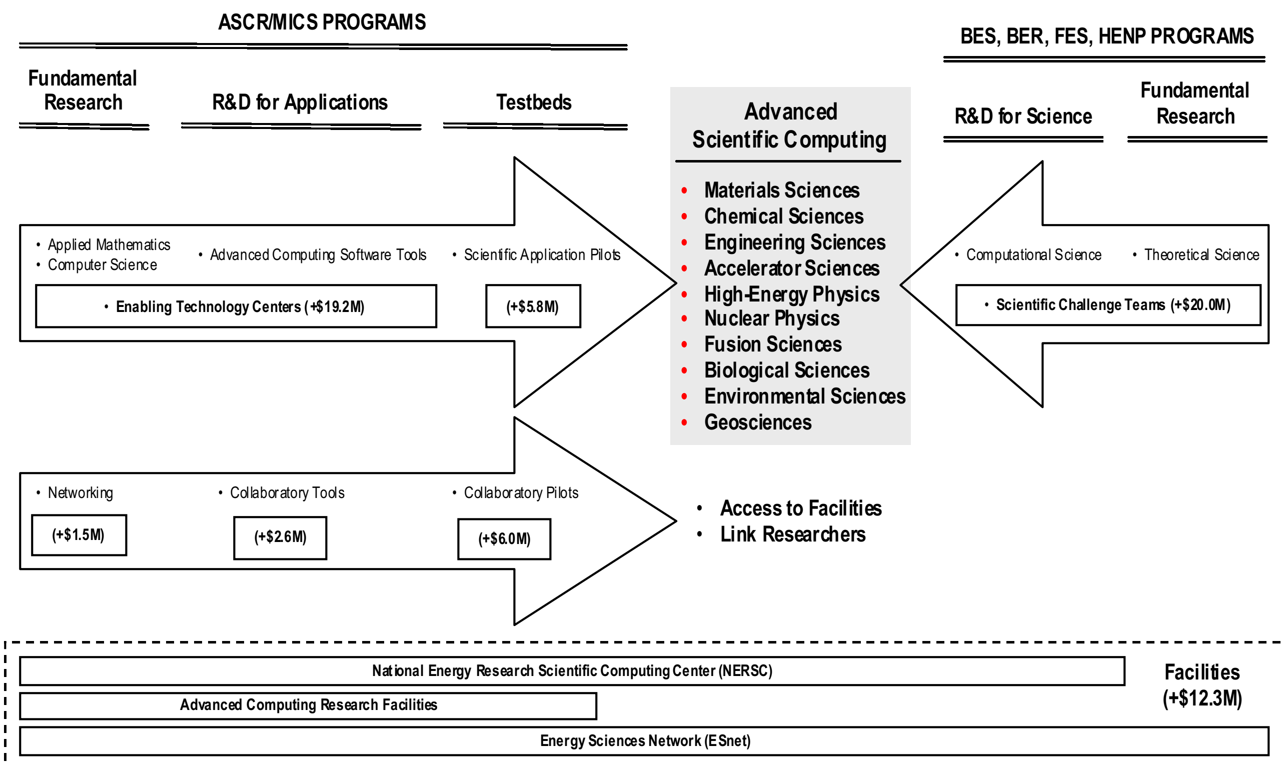

Note: SBIR/STTR for ASCR's Mathematical, Information, and Computational Sciences Division is increased by $1.2M in FY'01 due to increases in operating expenses. The above numbers also do not include the requested $2.0M increase in the Computational Sciences Graduate Fellowship Program.

projects. These increases are based on an existing $44,565,000 core program in this area. The increases in ASCR will double the support for the Computational Science Graduate Fellowship Program (FY 2000, $2,000,000; FY 2001, $4,000,000); begin new scientific application partnerships with the scientific disciplines (FY 2000, $3,732,000; FY 2001, $9,500,000); and establish enabling technology centers focused on mathematical and computer science technologies (FY 2000, $0; FY 2001, $19,176,000) needed by the scientific applications. These staff of these centers will include applied mathematicians, computer scientists, software engineers, and scientists from selected applications who conduct basic research in a particular technology and then move these advances to the scientific applications.

In the category of "national collaboratories," increases will support fundamental research in computer networks and collaboratory tools and will begin new collaboratory partnerships across SC. (FY 2000 $14,000,000; FY 2001 24,100,000). This program has been an international leader in the development of collaboratory tools and has founded two national collaboratories already. The ongoing work has identified shortcomings in underlying network technologies, in technologies for interconnecting SC networks with networks serving university researchers, and in handling large data streams and datasets.

In the category of "advanced computing and networking facilities," increases are requested to upgrade the NERSC III computer procurement from 3½ teraflops to 5 teraflops (FY 2000, $26,500,000; FY 2001, $32,278,000); to provide additional capabilities for ESnet (FY 2000, $15,747,000; FY 2001, $20,287,000); and to increase the funding for Advanced Computing Research Facilities (FY 2000, $13,749,000; FY 2001, $15,776,000).

The other SC programmatic offices have requested funding increases to support the activities focused on "advancing science through computation" through Scientific Challenge Teams. These teams will research, develop, deploy, and use the scientific modeling and simulation codes needed to tackle the highest priority projects listed in Table 1. The funding requests for FY 2001 are BES, $2,000,000; BER, $8,000,000; FES, $3,000,000; and HENP, $7,000,000.

The FY 2001 request will enable SC to deliver significant results for all of its programs with no further increase in funding through FY 2005. SC will significantly increase its capabilities for scientific discovery through advanced computing and the national collaboratory technologies that its scientific teams and facility users require to succeed. Examples of accomplishments by FY 2005 include:

- A 15+-teraflops computer will have been installed at NERSC to enable leading-edge computations on the most complex problems faced by SC.
- Eight to ten major research teams will have been formed across SC to develop the next generation of scientific modeling and simulation codes to address the critical science and engineering challenges facing the programs.
- New algorithms and mathematical libraries that provide the most advanced techniques for solving complex equations will have been developed and deployed to computational scientists.
- Effective techniques for managing, analyzing, and visualizing petabyte data sets will be developed and deployed to scientists.

- Collaboratory technology and the network infrastructure will be ready for the challenges of terascale computing as well as data-intensive experiments such as the Large Hadron Collider.

The Future

There are significant investment opportunities from FY 2002 through FY 2005 that would provide dramatically increased benefits for the nation. These opportunities are in the areas of:

- *Computing Systems Software for Terascale Computers*. Additional investments would accelerate the work on computing systems software and enable scientists to more effectively use high-end computers earlier in their life-cycle. These developments would have significant benefits to all of the government agencies that require high-performance scientific computing to achieve their mission objectives as well to the U.S. high-performance computing industry.
- *Topical Computing Facilities*. Additional investment is required to establish the Topical Computing Facilities. These facilities promise to substantially advance scientific computing by providing highly cost-effective increases in computing capabilities for one or more closely related scientific applications. They also provide a focal point for the development and use of modeling and simulation codes in the selected scientific research areas.
- *Scientific Challenge Codes*. The new scientific modeling and simulation capabilities realized by additional investments in scientific challenge codes would further extend the benefits of advances in computing technology to the scientific portfolio of SC. The new modeling and simulation capabilities created in these efforts will benefit the entire scientific research community, not just SC.

The timing of the establishment of the Topical Computing Facilities as well as decisions about further upgrades to NERSC will be based on technical evaluation of the effectiveness of the computational resources for scientific computing as well as the availability of funding.

Relationship to ASCI

The computational modeling and simulation efforts in both the National Nuclear Security Agency (NNSA), *i.e.*, the Accelerated Strategic Computing Initiative (ASCI), and SC are focused on large-scale scientific and engineering simulations—the former for complex nuclear weapons systems, the latter for complex scientific problems. There are, of course, significant differences in the detailed nature of the problems being addressed. In addition, the classified nature of many of the scientific and engineering applications in ASCI limits participation by the broader scientific community. Still, there is much to be gained by exploiting the natural synergy between the two programs. In fact, the SC plan outlined here builds on current and anticipated accomplishments of ASCI. Because of the investments in ASCI, SC will only require modest investments in some research areas to tailor the software environments and tools developed in ASCI to SC needs.

Many of the technical problems that arise in using terascale computing capabilities are generic to scientific simulation and apply equally to ASCI and SC applications. Areas of common interest include:

- Scientific data management.
- Analysis and visualization of petabyte data sets.
- Computer operating systems.

In these areas the SC Plan has been carefully coordinated with those of ASCI as well as those of the National Science Foundation.[6] However, there are needs unique to SC's mission. Examples of such needs include:

- Efficient access to petabyte data sets from both national laboratory and university computer and network environments.
- Collaboration technologies to support the integration of widely dispersed researchers.
- Mathematical and software libraries and technologies specific to the scientific mission of the SC, *e.g.*, global climate, molecular science, plasma physics, elementary particle physics, and accelerator physics.
- Robust, easy-to-use scientific applications codes for use by the larger scientific community ("community codes").

By building on the accomplishments of ASCI and working closely with ASCI on projects of common interest, SC will be able to focus more of its effort on the problems in computational science, computer science, and applied mathematics that are unique to SC's mission and research environment.

Relationship to Other Government Agencies

In the Implementation Plan for the Initiative in "Information Technology for the Twenty-first Century," the NSF and DOE were given leadership responsibility for the effort in "Advanced Computing for Science, Engineering and the Nation." Other agencies participating in this effort are Department of Defense, National Aeronautics and Space Administration, National Oceanic and Atmospheric Administration, and National Institutes of Health. NSF and DOE bring complementary strengths to this leadership role.

- DOE, a mission agency like DOD, NASA, NOAA, and NIH, is focused on scientific problems critical to its mission; NSF has a broad scientific portfolio. For example, in computer science and applied mathematics, DOE only funds activities critical to scientific computing, while NSF supports a wide range of activities important to the advancement of information technology as a whole.

---

[6] See, *e.g.*, "Data and Visualization Corridors: Report of the 1998 DVC Workshop Series," Ed. P. H. Smith and J. van Rosendale, California Institute of Technology, 1999.

- In fulfillment of its mission, DOE often establishes large, nationwide, integrated programs aimed at solving its most challenge scientific and engineering problems, problems not unlike those outlined here. NSF, on the other hand, seeks to stimulate the creativity of single investigators or small, local research groups.
- To achieve its goals, DOE, like the other mission agencies, must support the full life-cycle (research, development, and deployment and support) of the major software systems created in its research programs. This is a burden beyond the traditional scope of research, but is essential to realizing the full potential of advances in computing technology for scientific discovery.

Relationship to Other DOE Programs

The investments in modeling and simulation outlined here complement the investments in the experimental programs supported by DOE. The modeling and simulation effort will draw heavily on DOE's experimental programs to provide basic insights into and data on complex systems as well as the fundamental processes that govern their behavior. The experimental programs also provide the means of validating the computational models and simulations, and validation is essential to advancing the state-of-the-art in scientific simulation. In turn, the computational modeling and simulation will:

- Provide insights into fundamental physical, chemical, and biological processes that would otherwise be unattainable.
- Maximize the return on investments in experimental facilities by more quickly and efficiently harvesting scientific results.
- Provide the technical capability to design and construct innovative, new experimental facilities.

Computational modeling and simulation provide a strong link between the SC and the other DOE offices. By demonstrating how real-world devices and systems can be simulated from a knowledge of the fundamental physical, chemical, and biological processes involved, the investment in scientific computing connects the basic research programs in SC with the energy and environmental research and development programs in the Offices of Fossil Energy, Energy Efficiency, and Environmental Management.