

INTRODUCTION

The Office of Technology Assessment (OTA) recently completed a report entitled *Information Technology R&D: Critical Trends and Issues*. This report explored the structure and orientation of selected foreign programs, issues of manpower, institutional change, new research organizations developing out of Bell Laboratories, and trends in science and technology policy. Four specific areas of research: advanced computer architecture, artificial intelligence, fiber optics, and software engineering were also examined. To supplement this earlier work, the House Committee on Science and Technology requested that OTA examine issues of resource management, networking, and the role of supercomputers in basic research. This background paper will explore issues raised in the earlier R&D assessment and examine new and ongoing Federal programs in large-scale computer research.

Supercomputer is the term applied to the class of the most powerful computers available at any particular time. The cost/performance ratio of all classes of computers, from the largest to the smallest, continues to decrease rapidly, and today's desk-top computer has the power that years ago was available only in mainframes. Speed is gained both by improving the logical design of the computer and by making electronic components of the machine operate faster. Hence, each generation of supercomputers has tested many new design ideas and component technologies that were later introduced in smaller, less expensive machines.

Since the 1950s most large computers have shared an architecture named for John von Neumann, a prominent mathematician who played a major role in the invention and development of the digital computer. In the von Neumann architecture, *data* and *program instructions* both reside in memory, and instructions are acted on one by one, sequentially by the "processor" (other parts are the "control" and the "memory"). Many of the supercomputers now popular, such as the Cray 1 and the Cyber 205, are still based on variations of

the von Neumann design. Called "vector" machines, they gain their speed by breaking up computational tasks (such as addition and multiplication) into separate "pipelines," which allows certain problems to be executed far faster. (See figure 1.)


Most computer scientists have concluded that the sequential, von Neumann design can no longer sustain the rapid growth to which we have become accustomed (though component speeds will continue to improve).¹ They are looking elsewhere for new design ideas, and their interest has turned to parallelism. In a parallel machine, rather than one processor working sequentially on the steps of solving a problem, many processors work simultaneously on the computation. This interest in parallel design is based on three propositions:

1. the parallel computer will theoretically be far more powerful than the current von Neumann design;
2. the parallel multiprocessor could be less costly for a given task, especially when utilizing mass production technologies; and
3. parallel architectures will achieve higher computational speeds.

As the Federal Government sponsors more and more research in parallel computation, it is important to recognize this new design direction as a key component of the government's computer research effort. At the same time, it must be recognized that computer scientists and mathematicians are only beginning to understand how to use optimally the types of highly parallel designs that computer architects are exploring. Because of the growing importance of parallel computation, the terms "largescale computing" and "advanced scientific computing" refer in this background paper to both current vector supercomputers that employ von Neumann architecture and systems based on multiprocessor technologies.

¹U.S. Congress, office of Technology Assessment, *Microelectronics Research and Development—A Background Paper*, OTA-BP-CIT-40 (Washington, DC: J. S. Government Printing office, March 1986).

Figure I.—Current and Projected Supercomputers, 1960-90



SOURCE Sidney Fernbach

Federal interest in largescale computing devices is based on many concerns including:

- the ability of selected Federal agencies to fulfill their mission requirements in national defense, space technologies, energy technologies, and other areas;
- the viability of the U.S. supercomputer industry, particularly in light of increasing foreign competition;
- the research and development that is performed in hopes of increasing computational speed and the capabilities of these machines; and

- the availability of these machines to members of the scientific, research, and industrial communities to perform new research in a variety of fields.

These and other needs have led the Federal Government to expand its program in ways designed to give scientists and researchers greater access to large-scale computing facilities. This access will also foster the development of new architectures, and will lead to new generations of information technologies and the design and development of new software.

PRINCIPAL FINDINGS

- Several interagency panels were established in 1983 by the Federal Coordinating Council on Science, Engineering, and Technology (FCCSET) Panel as a forum for discussion on specific supercomputer issues. These panels have succeeded in this role, and they remain as a forum where national interests, goals, and programs can be fully considered. At the same time, FCCSET panels hold limited authority to alter or implement government policy based on their findings.
- No single government agency holds lead authority in advanced computer research and access. Each agency's programs represent their own mission requirements. Though it is clear that there is a need for this diversity of government programs, there is also a need for enhanced coordination of these efforts to ensure that national goals are realized. This may be especially true as greater fiscal constraints are placed on the funding agencies.
- Federal efforts have grown substantially in the past 2 years in response to a series of reports that noted shortcomings in the national supercomputer program. With the diversity of programs, program goals, and mission requirements now underway throughout the government, it may be advisable to assess the Federal efforts to ensure that the original concerns noted in the reports (e.g., the need for more research in computational mathematics, software, and algorithms) are still valid or have been replaced by new more pressing concerns. If it is to be effective, such a reexamination should:
 - involve scientific and research users, members of the private sector, and pertinent agency administrators; and
 - include a broader examination of the role of new information technologies and the conduct of scientific research.
- It is difficult to accurately assess the Federal investment in large-scale computing programs as the agencies employ a variety of terms to describe comparable or similar efforts.
- At least over the short term, limited human resources will be a critical factor in the success of the supercomputer programs. The opening of seven new centers by the National Science Foundation (NSF), the Department of Energy (DOE) and Florida State University, and the National Aeronautics and Space Administration (NASA) in the next fiscal year will generate a large demand for expert personnel to manage and operate the centers, but relatively few are available. Demand will be particularly heavy in the areas of applications software design and development, and this can only worsen as significantly different architectures proliferate.
- Software is an important determinant of the efficiency of the machines and the types of problems that can be tackled on them. It also influences the design of the next generation of machine. Therefore an investment in algorithm and software development is essential and integral to any large-scale computation program.
- Research in parallel computation has become a key component of the Federal Government's computer research effort and one result has been a proliferation of significantly different architectures. Most scientists and researchers consider these experimental designs necessary and fundamental to the efforts of advancing computational speed.
- Our current understanding of software and algorithm development is inadequate to fully realize the benefits of the new architectures. Resources need to be directed to:
 - develop an understanding of the new architectures,
 - define the research necessary to move software and algorithms from current generation supercomputers to other supercomputers and architectures, and
 - develop software and algorithms for the new architectures.
- Advanced data communication networks are important to the conduct of research science and technology, because they pro-

vide nationwide (and sometimes international) access to resources and information. Networks can:

- expand interconnections between research communities,
 - encourage greater joint or collaborative efforts, and
 - broaden access to a variety and number of resources.
- Government and private networks are proliferating, many employing a variety of technologies, standards, and protocols. This diversity may merit concern in the near future if it makes use of the networks difficult for users.
 - NSF is establishing NSFnet, a network that will link researchers with the large-scale computing resources. It is also NSF's intention that NSFnet will be the basis of a

national research network for scientists, researchers, and interested members of the industrial community. The coupling of NSFnet and a national research network could have far-reaching implications, and merits an in-depth and detailed study by an organization such as the National Academy of Sciences. In the interim, there are key issues regarding technology development and management operations of NSFnet that need to be considered by NSF:

- NSFnet is developing quickly and choices made today, of pilot projects for example, may affect the future configuration of a national research network; and
- the industrial community has not been included in plans for NSFnet, which may restrict private researchers' access to resources and NSFnet users access to industrial resources.

NATIONAL POLICY

Over the past few years, a combination of events has broadened awareness of and interest in Federal programs for the design and use of largescale computing facilities. Prompted by a decline in the computational facilities and services at American universities and colleges, particularly scientific researchers' lack of adequate access to large-scale computing facilities, NSF convened a panel of scientists to review the situation in 1982. The Department of Defense (DOD), DOE, and NASA joined NSF in sponsoring the panel, which was chaired by Peter Lax of New York University. Unlike previous studies, which explored the needs of specific segments of the research community, this panel (referred to as the "Lax panel") examined the large-scale computing needs of the entire U.S. research community. The panel noted two key problems: access to supercomputer facilities was limited and R&D on new architectures was insufficient to meet the perceived need for more sophisticated computers. The panel recommended four actions:

1. provide the scientific research community with increased access to supercomputing facilities and experimental computers through high bandwidth networks;
2. increase research in computational mathematics, software, and algorithms;
3. train personnel to use these facilities; and
4. conduct R&D of large-scale computing systems.²

A second report sponsored by NSF, the Bardon-Curtis report, outlined how NSF could respond to the problems noted by the Lax panel. The Bardon-Curtis report laid the groundwork for the new NSF supercomputing centers. The report recommended that NSF take six steps:

1. enhance coordination between Federal and private programs and supercomputer research projects;
2. increase support for local scientific computing facilities;

3. elicit proposals for supercomputer research centers, and support up to 10 centers within 3 years;
4. support networks to link universities and laboratories with each other and with supercomputer centers, thus providing access to facilities, file transfer capability, and scientific communication;
5. create an advisory committee to assist and oversee NSF's decisions concerning computer services and networks; and
6. support academic research and training programs in the areas of advanced computer systems design, computational mathematics, software, and algorithms.³

While the Lax panel was studying large-scale computing needs in the United States, the Japanese Government was working intensively to develop two programs, the National Super Speed Computer Project and the Fifth-Generation Computer Project. These programs, both designed to meet Japan's domestic supercomputer needs, also give entry into the international marketplace. The National Super Speed Computer Project is a lo-year program that seeks to develop a machine one-thousand times faster than a current supercomputer. The Fifth-Generation Computer Project is focusing on development of a machine with artificial intelligence applications. Both projects are supported by Ministry of International Trade and Industry, and private companies. Recently, three Japanese companies, Fujitsu, Nippon Electric Corp. and Hitachi, announced supercomputers that appeared to be faster than U.S. machines.

In 1983, the British Government also announced a research effort in this area based on the recommendations of the Alvey Committee, the committee that formulated the British response to the Japanese Fifth-Generation Program. The British effort is focused on artificial intelligence and largescale integrated circuits, software engineering, and man/machine

²National Science Foundation, *Report of the Panel on Large Scale Computing in Science and Engineering, 1982.*

³M. Bardon and K. Curtis, *A National Computing Environment for Academic Research*, National Science Foundation, July 1983.

interfaces. The European Economic Community has initiated the ESPRIT project, which will fund research in advanced microelec-

tronics, software, office automation, advanced information processing, and computer integrated manufacturing.

NATIONAL GOALS

Together, foreign competition and pressure from the academic community have heightened concern over what the U.S. Government role should be in the development and use of advanced scientific computers and large-scale facilities. Several themes have emerged from the various reports that describe and elucidate the national role and goals with respect to large-scale computing machines.

The advancement of *science* is one of the most commonly cited goals.

Perhaps the most significant applications of scientific computing lie not in the solution of old problems but in the discovery of new phenomena through numerical experimentation. They [supercomputers], permit the solution of previously intractable problems, and motivate scientists and engineers to explore and formulate new areas of investigation."⁴

Integral to achieving this goal are *education* and *access*. Because computers and computing have become essential tools in scientific and engineering research many trained and knowledgeable personnel are needed to operate and to use them. With access to these machines and facilities, researchers can be trained in large-scale computing and also conduct research using high-performance machines.

The reports also stress that the United State's *economic strength* and ability to *compete internationally* both now and in the future are dependent on the continuing development of and access to large-scale computing machines. Supercomputers are now integral in the design of aerospace, automotive, chemical, pharmaceutical, and microelectronic products. Over the last two decades, the United States has led the industrialized world in computer technology; each succeeding generation of supercomputer has led to new and innovative applications and designs. The relationship between a large-scale computing program to advance the state of the art in numerous fields

and the U.S. position in the international marketplace is quite clear and is discussed in depth in two OTA reports: *Information Technology R&D: Critical Trends and Issues* and *International Competitiveness in Electronics*.⁵ As the number of supercomputers available to the scientific and research communities increases, more and more applications will be employed expanding the commercial and economic benefits to the United States. Clearly, the potential benefits from employing large-scale computing machines within the defense, industrial, and scientific communities are enormous.

The development of supercomputers in support of *national defense and national security programs* are critical goals of a national program. High-performance computing is needed to verify arms control, analyze intelligence data and information, and protect national security secrets. Supercomputers have always been essential in defense programs for military preparedness, and the design of weapon systems and ballistic trajectories.

Largescale computing facilities are also central to R&D in such fields as atmospheric sciences, aeronautics, nuclear reactor theory and design, and geophysics. "Research at the forefront of contemporary and future science and technology will demand adequate access to supercomputer power."⁶ With the help of supercomputers, scientists can now tackle problems and investigations in areas not possible before. Problems such as simulating the circulation of the oceans and the atmosphere, and the effects of carbon dioxide buildup in the atmosphere are examples of questions previously too large and too time-consuming to attempt solution on computing facilities. "The

⁴U.S. Congress, Office of Technology Assessment, *Information Technology R&D: Critical Trends and Issues*, OTA-CIT-268 (Washington, DC: U.S. Government Printing Office, February 1985); and U.S. Congress, Office of Technology Assessment, *International Competitiveness in Electronics*, OTA-ISC-200 (Washington, DC: U.S. Government Printing Office, November 1983).

⁵U.S. Department of Energy, *The Role of Supercomputers in Energy Research Programs*, February 1985, p. 4.

⁴National Science Foundation, *Report of the Panel on Large Scale Computing in Science and Engineering*, 1982, p. 6.

Table 1.—Partial List of Problems/Applications That Will Benefit From Use of Large-Scale Facilities

Use of supercomputers In engineering:	Design of membranes
Automobile design	Design of new materials and chemicals
Aircraft design	Use of supercomputers in energy research and development:
Digital simulation	Petroleum exploration
Nuclear power system safety	Reservoir modeling
Stress/strain analysis of structures	Power grid modeling
Use of supercomputers In manufacturing:	Fusion engineering
Robotics	Exploration support
Assembly	Use of supercomputers in support of new computer design and development:
System control	Simulation of new computer architectures
Plant design	Supercomputer design
Computer-aided design	Computer graphics
Use of supercomputers In support of national security missions:	Examples of scientific research employing supercomputers:
Command and control	Cell growth in biological systems
Embedded systems	Lattice quantum chromodynamics
Weapons systems design	Quantum calculations of molecular energy surfaces
Mapping	Molecular dynamics simulations on genetically engineered proteins
Intelligence collection and analysis	General circulation models of the atmosphere to study the climatic effects of aerosols
Use of supercomputers In exploring environmental questions:	Simulation and condensed matter physics
Weather modeling and climate	Free oscillations of the earth
Satellite imaging	Determination of the structure of an animal virus. human rhinovirus-14
Chemical flow models	Simulations of DNA molecules in an aqueous environment
Use of supercomputers in medicine and health-related issues and problems ¹	Simulations of the origins of the universe
Diagnosis	Neural modelling of the frog retina
Biochemical processes	Simulations of interactions between enzymes and inhibitor molecules
Design of drugs	
Genetic research	
Use of supercomputers in economics research:	
Real-time models of world economic relationships	
Prediction effects of phenomena	
Use of supercomputers In chemistry:	
Design of catalysts	

SOURCE Office of Technology Assessment


story of numerical weather prediction and climate modeling is one of a never-ending fight for faster computers, because they raise the level of realism we can put into our models and hence the level of realism we can expect in our results. (See table 1.) As increased availability of large-scale computing facilities has extended the range of problems to be investigated by the scientific and research communities in the United States, it has also sharpened interest in facilities on either side of the supercomputing spectrum; those not nearly as

powerful as the Cray 2 or Cyber 205, dubbed the minisupercomputers, and those machines not yet available, the next generation.

The minisupercomputers, such as the Intel Personal Supercomputer, the IPSC, or that of Convex Computer Corp., the C-1, present the scientist with a cost-effective alternative to a supercomputer with hands-on availability and large-number crunching capabilities. But there are problems still intractable on current generation supercomputers in fields such as hydrodynamics, fusion, plasma physics, and others that drive scientists to design new architectures with increased computational capabilities. (See figure 2.)

¹Merry Maisel, "Questions for Supercomputers." *Gather Scatter*, San Diego Supercomputer Center, January 1986, p. 5.

Figure 2.—Range of Federal Policies Possible With a Supercomputer 200 Times the Current Capabilities



KEY.
 [White box] = New science
 [Light gray box] = More science
 [Dark gray box] = Wider access/educational use

Table 1.— Partial List of Problems/Applications That Will Benefit From Use of Large-Scale Facilities

Use of supercomputers in engineering:	Design of membranes
Automobile design	Design of new materials and chemicals
Aircraft design	Use of supercomputers in energy research and development:
Digital simulation	Petroleum exploration
Nuclear power system safety	Reservoir modeling
Stress/strain analysis of structures	Power grid modeling
Use of supercomputers in manufacturing:	Fusion engineering
Robotics	Exploration support
Assembly	Use of supercomputers in support of new computer design and development:
System control	Simulation of new computer architectures
Plant design	Supercomputer design
Computer-aided design	Computer graphics
Use of supercomputers in support of national security missions:	Examples of scientific research employing supercomputers:
Command and control	Cell growth in biological systems
Embedded systems	Lattice quantum chromodynamics
Weapons systems design	Quantum calculations of molecular energy surfaces
Mapping	Molecular dynamics simulations on genetically engineered proteins
Intelligence collection and analysis	General circulation models of the atmosphere to study the climatic effects of aerosols
Use of supercomputers in exploring environmental questions:	Simulation and condensed matter physics
Weather modeling and climate	Free oscillations of the earth
Satellite imaging	Determination of the structure of an animal virus, human rhinovirus-14
Chemical flow models	Simulations of DNA molecules in an aqueous environment
Use of supercomputers in medicine and health-related issues and problems ⁷	Simulations of the origins of the universe
Diagnosis	Neural modeling of the frog retina
Biochemical processes	Simulations of interactions between enzymes and inhibitor molecules
Design of drugs	
Genetic research	
Use of supercomputers in economics research:	
Real-time models of world economic relationships	
Prediction effects of phenomena	
Use of supercomputers in chemistry:	
Design of catalysts	

SOURCE office Technology Assessment


story of numerical weather prediction and climate modeling is one of a never-ending fight for faster computers, because they raise the level of realism we can put into our models and hence the level of realism we can expect in our results. (See table 1.) As increased availability of large-scale computing facilities has extended the range of problems to be investigated by the scientific and research communities in the United States, it has also sharpened interest in facilities on either side of the supercomputing spectrum; those not nearly as

powerful as the Cray 2 or Cyber 205, dubbed the minisupercomputers, and those machines not yet available, the next generation.

The minisupercomputers, such as the Intel Personal Supercomputer, the IPSC, or that of Convex Computer Corp., the C-I, present the scientist with a cost-effective alternative to a supercomputer with hands-on availability and large-number crunching capabilities. But there are problems still intractable on current generation supercomputers in fields such as hydrodynamics, fusion, plasma physics, and others that drive scientists to design new architectures with increased computational capabilities. (See figure 2.)

⁷ Merry Maisel, "Questions for Supercomputers," *Gather Scatter*. San Diego Super-computer Center, January 1986. p. 5.

Figure 2.—Range of Federal Policies Possible With a Supercomputer 200 Times the Current Capabilities



KEY.
 = New science
 = More science
 = Wider access/educational use

CURRENT FEDERAL PROGRAMS

This section will briefly review the Federal large-scale computing programs in the pertinent agencies to illustrate the scope and nature of the [J. S. Government's investment in large-scale computing programs and facilities. The OTA review will describe five major programs at: NSF, NASA, DOE, the Supercomputing Research Center, and the Defense Advanced Research Projects Agency (DARPA). The programs vary in nature, ranging from those of DOE and NSF, which combine R&D funding and access to facilities, to DARPA, which is focused on R&L) of advanced computer research.

National Science Foundation

Within the past 2 years, NSF has broadened its large-scale computing program efforts in direct response to the Lax and Bardou-Curtis reports and the pressures and problems they cited. To rectify two commonly cited problems, NSF established an Office of Advanced Scientific Computing (OASC). The new office is designed to provide U.S. researchers with access to supercomputers or advanced computing services; and encourage the growth and development of advanced scientific computing in the United States.

NSF is providing researchers with increased access to advanced computing services in several ways. In July 1984, NSF funded three operating computer centers (phase 1), at Purdue University, University of Minnesota, and Boeing Computer Services in Seattle. This action presented qualified researchers with immediate opportunities to conduct research at these computing facilities and at the same time become familiar with four supercomputers: the Cyber 205, the Cray 2, the Cray 1A, and the Cray 1S. The majority of users of the three centers are current NSF grantees, who also comprise the bulk of new proposals soliciting computer center time. Since those original grants were made, three other facilities were funded. Colorado State houses a Cyber 205; Digital/Vector Productions has a Cray X-MP; and AT&T Bell Laboratories now

has a Cray X-MP. The OASC allocated 22,000 hours of supercomputer time to NSF researchers in fiscal year 1985; 5,000 hours had been used by the end of fiscal year 1985. At the beginning of fiscal year 1986, time is being used at a rate of over 1,000 hours per month.

The new supercomputer centers represent a major new Federal investment. Over the next 5 years, NSF will invest approximately \$200 million in five more centers (phase II) at Cornell University, the Consortium for Scientific Computing near Princeton (a consortium of 13 universities), the University of California at San Diego (a consortium of 19 universities and research institutions), the University of Illinois, and the Pittsburgh Center (a consortium of 10 to 20 universities). (See table 2.) NSF funds serve also as "seed" money, and have already generated interest and support in other sectors. In addition to Federal funding, the centers receive some money from State governments and industry. For example, the State of New Jersey, private industry, and consortium members have already committed \$39.3 million to the John von Neumann Center for Scientific Computing near Princeton, New Jersey. Private industry is also expected to fund specific research projects at these five centers.

Each of the five new supercomputer centers are expected to develop a different research

Table 2.—NSF/OASC Budget (in millions)

	Fiscal year 1985	Fiscal year 1986	Fiscal year 1987
Centers:			
Phase I	9.7	33	2.3
Phase II	19.3	25.1	34.5
Training	0.5	1.0	1.0
Networks:			
NSFnet	3.7	5.0	6.5
Local access	2.2	1.8	.5
New Technologies:			
Cornell Center	4.9	5.2	6.2
Other experimental access	.0	2.3	1.0
Software	.9	1.5	1.6
Total	\$41.4	\$45.2	\$53.6

SOURCE: John Connolly, Director, Office of Advanced Scientific Computing, National Science Foundation

emphasis. The Center for Theory and Simulation in Science and Engineering, located at Cornell University, has been designated an experimental center where research will focus on parallel processing and software productivity. Researchers there use an IBM 3084 QX mainframe computer attached to FPS 164 and 264 scientific processors. IBM has donated both equipment and personnel to support this center. An important aspect of the Cornell program is the plan to bring in interdisciplinary teams of scientists to develop new algorithms. Unlike the other center programs, this program focuses on experimental equipment, and this configuration means that it will serve a few users with large needs, rather than a large number of users in need of computing cycles.

The John von Neumann Center, located near Princeton, will be managed by the Consortium for Scientific Computing, which represents 13 universities.⁸ At first, the von Neumann Center will use a Cyber 205, then later will upgrade the facilities to include an ETA-10. The center was established to provide researchers with access to the facilities for scientific research and to develop new architectures and algorithms.

The San Diego Supercomputer Center, located at the University of California at San Diego, is managed by GA Technologies and supported by a consortium of 19 universities and research institutions.⁹ The State of California is committing \$1 million per year to this center. The San Diego Supercomputer Center will use a Cray X-MP/48 and plans to use a Cray-compatible minisupercomputer pledged

⁸Members are: University of Arizona; Brown University; Columbia University; University of Colorado; Harvard University; Institute for Advanced Study-Princeton, iNJ; Massachusetts Institute of Technology; New York University; University of Pennsylvania; Pennsylvania State University; Princeton University; University of Rochester; and Rutgers University.

⁹The members are: Agouron Institute, San Diego, CA; California Institute of Technology; National Optical Astronomy Observatories; Research Institute of Scripps Clinic; Salk Institute for Biological Studies; San Diego State University; Southwest Fisheries Center; Stanford University; University of California at Berkeley; University of California at Los Angeles; University of California at San Diego; Scripps Institution of Oceanography; University of California at San Francisco; University of Hawaii; University of Maryland; University of Michigan; University of Utah; University of Washington; and the University of Wisconsin.

by Scientific Computer Systems. At this center, the focus will be on providing research time on the supercomputer facilities. Members of the consortium believe that the center will develop strengths in particular disciplines, such as microbiology.

The fourth center, the National Center for Supercomputing Applications is located at the University of Illinois at Urbana-Champaign. Like the San Diego Center, it will use a Cray X-MP/24 and upgrade to a Cray X-MP/48. The Illinois center will be closely affiliated with the Center for Supercomputer Research and Development, a program supported by DOE and NSF. The Illinois center will provide computing cycles to the research community, and through a visiting scholar program, it will also focus on the development of new architectures and algorithms. The Illinois center has received extensive support from the State of Illinois.

The Pittsburgh Center for Advanced Computing will not be funded at the same level as the other phase II centers, although NSF is committed to its long-term operation. A Cray 1S donated by NASA prompted the establishment of this new center, which will be dedicated to providing time on the Cray facilities. A Cray X-MP/48 and SSP will be delivered in April 1986 to update the center's facilities. The University of Pittsburgh Center will be managed by Carnegie-Mellon University with participation by Westinghouse Electric Corp.

The funding cycles of these centers vary. The phase I centers will be funded for 2 years, through 1986, after which phase II centers will begin full operation. ¹⁰ Funding for each phase II center will be approximately \$40 million per center over a 5-year period. Prototype centers (e.g., Cornell University), will be funded for 3 years at \$20 million. NSF projects that the program will require between \$300 million and \$500 million within 5 years. This will cover the costs of network development and of establishing 11 to 13 supercomputer centers nationwide, with two systems per center. This estimate is based on an analysis of projected

¹⁰Phase I centers are now going through year two funding cycles.

CURRENT FEDERAL PROGRAMS

This section will briefly review the Federal large-scale computing programs in the pertinent agencies to illustrate the scope and nature of the U.S. Government's investment in large-scale computing programs and facilities. The OTA review will describe five major programs at: NSF, NASA, DOE, the Supercomputing Research Center, and the Defense Advanced Research Projects Agency (DARPA). The programs vary in nature, ranging from those of DOE and NSF, which combine R&D funding and access to facilities, to DARPA, which is focused on R&D of advanced computer research.

National Science Foundation

Within the past 2 years, NSF has broadened its large-scale computing program efforts in direct response to the Lax and 13ardon-Curtis reports and the pressures and problems they cited. To rectify two commonly cited problems, NSF established an Office of Advanced Scientific Computing (OASC). The new office is designed to provide U.S. researchers with access to supercomputers or advanced computing services; and encourage the growth and development of advanced scientific computing in the United States.

NSF is providing researchers with increased access to advanced computing services in several ways. In July 1984, NSF funded three operating computer centers (phase 1), at Purdue University, University of Minnesota, and Boeing Computer Services in Seattle. This action presented qualified researchers with immediate opportunities to conduct research at these computing facilities and at the same time become familiar with four supercomputers: the Cyber 205, the Cray 2, the Cray 1A, and the Cray 1 S. The majority of users of the three centers are current NSF grantees, who also comprise the bulk of new proposals soliciting computer center time. Since those original grants were made, three other facilities were funded. Colorado State houses a Cyber 205; Digital/Vector Productions has a Cray X-MP; and AT&T Bell Laboratories now

has a Cray X-MP. The OASC allocated 22,000 hours of supercomputer time to NSF researchers in fiscal year 1985; 5,000 hours had been used by the end of fiscal year 1985. At the beginning of fiscal year 1986, time is being used at a rate of over 1,000 hours per month.

The new supercomputer centers represent a major new Federal investment. Over the next 5 years, NSF will invest approximately \$200 million in five more centers (phase II) at Cornell University, the Consortium for Scientific Computing near Princeton (a consortium of 13 universities), the University of California at San Diego (a consortium of 19 universities and research institutions), the University of Illinois, and the Pittsburgh Center (a consortium of 10 to 20 universities). (See table 2.) NSF funds serve also as "seed" money, and have already generated interest and support in other sectors. In addition to Federal funding, the centers receive some money from State governments and industry. For example, the State of New Jersey, private industry, and consortium members have already committed \$39.3 million to the John von Neumann Center for Scientific Computing near Princeton, New Jersey. Private industry is also expected to fund specific research projects at these five centers.

Each of the five new supercomputer centers are expected to develop a different research

Table 2.—NSF/OASC Budget (in millions)

	Fiscal year 1985	Fiscal year 1986	Fiscal year 1987
Centers:			
Phase I	\$ 9.7	\$ 3.3	\$ 2.3
Phase II	19.3	25.1	34.5
Training	0.5	1.0	1.0
Networks:			
NSFnet	3.7	5.0	6.5
Local access	2.2	1.8	.5
New Technologies			
Cornell Center	4.9	5.2	6.2
Other experimental access0	2.3	1.0
Software9	1.5	1.6
Total	\$44.4	\$45.2	\$53.6

SOURCE: John Connolly, Director, Office of Advanced Scientific Computing, National Science Foundation

emphasis. The Center for Theory and Simulation in Science and Engineering, located at Cornell University, has been designated an experimental center where research will focus on parallel processing and software productivity. Researchers there use an IBM 3084 QX mainframe computer attached to FPS 164 and 264 scientific processors. IBM has donated both equipment and personnel to support this center. An important aspect of the Cornell program is the plan to bring in interdisciplinary teams of scientists to develop new algorithms. Unlike the other center programs, this program focuses on experimental equipment, and this configuration means that it will serve a few users with large needs, rather than a large number of users in need of computing cycles.

The John von Neumann Center, located near Princeton, will be managed by the Consortium for Scientific Computing, which represents 13 universities.⁸ At first, the von Neumann Center will use a Cyber 205, then later will upgrade the facilities to include an ETA-10. The center was established to provide researchers with access to the facilities for scientific research and to develop new architectures and algorithms.

The San Diego Supercomputer Center, located at the University of California at San Diego, is managed by GA Technologies and supported by a consortium of 19 universities and research institutions.⁹ The State of California is committing \$1 million per year to this center. The San Diego Supercomputer Center will use a Cray X-MP/48 and plans to use a Cray-compatible minisupercomputer pledged

⁸Members are: University of Arizona; Brown University; Columbia University; University of Colorado; Harvard University; Institute for Advanced Study-Princeton, NJ; Massachusetts Institute of Technology; New York University; University of Pennsylvania; Pennsylvania State University; Princeton University; University of Rochester; and Rutgers University.

⁹The members are: Agouron Institute, San Diego, CA; California Institute of Technology; National Optical Astronomy Observatories; Research Institute of Scripps Clinic; Salk Institute for Biological Studies; San Diego State University; Southwest Fisheries Center; Stanford University; University of California at Berkeley; University of California at Los Angeles; University of California at San Diego; Scripps Institution of Oceanography; University of California at San Francisco; University of Hawaii; University of Maryland; University of Michigan; University of Utah; University of Washington; and the University of Wisconsin.

by Scientific Computer Systems. At this center, the focus will be on providing research time on the supercomputer facilities. Members of the consortium believe that the center will develop strengths in particular disciplines, such as microbiology.

The fourth center, the National Center for Supercomputing Applications is located at the University of Illinois at Urbana-Champaign. Like the San Diego Center, it will use a Cray X-MP/24 and upgrade to a Cray X-MP/48. The Illinois center will be closely affiliated with the Center for Supercomputer Research and Development, a program supported by DOE and NSF. The Illinois center will provide computing cycles to the research community, and through a visiting scholar program, it will also focus on the development of new architectures and algorithms. The Illinois center has received extensive support from the State of Illinois.

The Pittsburgh Center for Advanced Computing will not be funded at the same level as the other phase II centers, although NSF is committed to its long-term operation. A Cray 1S donated by NASA prompted the establishment of this new center, which will be dedicated to providing time on the Cray facilities. A Cray X-MP/48 and SSP will be delivered in April 1986 to update the center's facilities. The University of Pittsburgh Center will be managed by Carnegie-Mellon University with participation by Westinghouse Electric Corp.

The funding cycles of these centers vary. The phase I centers will be funded for 2 years, through 1986, after which phase II centers will begin full operation. Funding for each phase II center will be approximately \$40 million per center over a 5-year period. Prototype centers (e.g., Cornell University), will be funded for 3 years at \$20 million. NSF projects that the program will require between \$300 million and \$500 million within 5 years. This will cover the costs of network development and of establishing 11 to 13 supercomputer centers nationwide, with two systems per center. This estimate is based on an analysis of projected

¹⁰Phase I centers are now going through year two funding cycles.

needs of the 20 disciplines funded by NSF. By 1990, one-third of these disciplines will require two or more large-scale computing facilities; one-third will require one facility; and the remaining one-third less than one-half of a facility. This totals 22 to 24 state-of-the-art large-scale facilities and minisupercomputers." These projections were made prior to the passage of the Gramm-Rudman-Hollings legislation. In a recent briefing on the NSF program, program staff stated that no more centers would be established.

An important facet of the establishment of these centers is the creation of a new network to allow users to communicate with one other, both at the centers and around the country. The stated goal of the NSF networking program is to provide the research community with universal access to the large-scale research facilities. This network is intended to be the basis for a national research network for the academic community and eventually will connect with international networks. The NSF strategy appears to be twofold: 1) immediate access through existing networks, such as ARPANET (the first long-haul computer network developed under contract by DARPA); 2) followed by development of high-speed networks that will be the backbone of the new network, connecting all of the data centers. To carry out this strategy, the program will soon begin funding pilot networking projects.

OASC also intends to fund projects in five areas of software productivity and computational mathematics: computer science research on programming environments, development of software tools, numerical analysis, algorithm development, and increasing research effectiveness in using advanced computers.¹⁴ The fiscal year 1985 budget for this program was almost \$1 million, but only \$500,000 was committed to new grants, some funded jointly with the Divisions of Mathematical Sciences and Computer Research. The funded proposals

¹⁴Larry Lee, National Science Foundation, personal communication. June 7, 1985.

¹⁵Ibid., June 14, 1985.

have focused on algorithm development and numerical techniques. In fiscal year 1986, \$1.5 million is earmarked for software proposals.

Other divisions of NSF also support research on large-scale facilities, networking, software engineering, and related areas. Projects funded by the Division of Computer Research over the past 20 years are now the basis for many information technologies in use or have led to prototype development elsewhere (many of the DARPA projects described originated within the Division of Computer Research). Projects may be directly related to the work planned at the new centers and in OASC, but the focus is varied and distinct from that of OASC. Unlike OASC, the other divisions place a greater emphasis on funding projects that use multiprocessor technologies. For example, two NSF projects with this emphasis, now in the prototype stage, are the Texas Reconfigurable Array Computer project at the University of Texas, Austin, and the Cedar project at the University of Illinois.

National Aeronautics and Space Administration

NASA has used supercomputers at various locations around the country for several years in support of a number of mission programs. Until recently, however, NASA has funded very little advanced scientific supercomputer research designed to create or develop new architectures or machines, although the agency did fund the development of the Massively Parallel Processor (MPP), now being used for image processing. I¹⁵ (See table 3.)

NASA's three research centers have supercomputers: Ames, Langley, and Lewis. The Goddard Space Flight Center also has supercomputer facilities (Cyber 205) though it is not designated as a research center."

¹⁵MPP, a limited application computer, has over 16,000 processors operating simultaneously and a custom integrated circuit containing eight complete processors.

¹⁶Ames, Cyber 205, X-hi P; Langley, modified Cyber 205; Lewis, Cray 1 S and awaiting an X-hi P; and Marshall Space Flight center has an RFP out to purchase a supercomputer in the near future.

Table 3.—NAS Development Budget (in millions)

	Fiscal year 1986	Fiscal year 1987
Hardware acquisition and lease ^a	\$17.3	\$20.8
Software research and development	3.4	4.0
System engineering, testing and integration	7.5	5.2
Total	\$28.2	\$30.0
^a High speed processors	\$ 8.3	\$13.9

SOURCE: Randy Graves, National Aeronautics and Space Administration.

One of NASA's major supercomputer programs is the Numerical Aerodynamic Simulation (NAS) Program. NAS is designed to solve problems of aerodynamic and fluid dynamics, but it is also intended to:

... act as the pathfinder in advanced large-scale computer system capability through systematic incorporation of state-of-the-art improvements in computer hardware and software technologies.¹⁵

When NAS becomes operational in fiscal year 1986, it will be available to interested individuals from NASA, DOD, other government agencies, industry, and universities. The NAS processing system will employ state-of-the-art high-speed processors (designated HSP 1 and 2), a mass storage system, a support-processing subsystem, a workstation subsystem, a graphics subsystem, and a long-haul communications subsystem. HSP-1 will be a Cray 2 supercomputer with four processors configuration and 256 million words of memory. The HSP-2, to be developed as NAS becomes operational, is expected to achieve four times the computational capabilities of a Cray 2 and will include upgraded subsystems and graphics capabilities and expanded wideband communications. ¹⁶

(continued from previous page)

Ames Research Center utilizes the Cray IS—upgraded to X-MP for computational aerodynamic research; Lewis Research Center uses the Cray 1S for internal computational fluid mechanics, and for thermal and structural performance analysis of propulsion system components; Goddard Space Flight Center utilizes their Cyber 205 for atmospheric science investigations.

¹⁵NASA testimony, House Science and Technology Committee, June 25, 1985.

¹⁶Lewis Peach and Randy Graves, NASA, personal communication, February 1985.

In fiscal year 1984, anticipating the arrival of the Cray 2, NASA began software development projects. This work has been carried out both at Cray Research and on a Cray X-MP at Ames. Early in the process, NASA chose a UNIX operating system.

In conjunction with the new facilities, NASA plans to establish a network to link all of NASA communications including computer facilities. High-speed communications between the four large-scale computing centers, an integral part of the network, will be facilitated through satellite and terrestrial links. The NAS center will be included in this network and in ARPANET and MILNET (the Defense Data Network) for access by the private sector and university researchers.

Department of Energy

DOE has a long history of using supercomputers and supporting architecture development for them. Since the 1950s, a DOE laboratory has acquired the first or one of the first manufactured units of nearly every largescale computer. DOE's National Laboratories still hold the greatest concentration of users of supercomputers; approximately 35 percent of the supercomputers in use in the United States are located in these laboratories.

DOE uses supercomputers to support a variety of its missions. The nuclear weapons program relies on large-scale computers to aid in highly complex computations in the design process. The Magnetic Fusion Energy and the Inertial Confinement Fusion Programs are heavily dependent on supercomputers as well. The machines are required to model the complex behavior of hot plasmas, including the effects of electric and magnetic fields, atomic physics, the interaction with intense radiation, and various boundary conditions. To this end, DOE hopes by fiscal year 1986 to have an installed base of 26 supercomputers throughout the agency and its laboratories.¹⁷

¹⁷Briefing, DOE, January 1985; 18 in support of defense programs, 3 in support of magnetic fusion energy, 2 in support of the energy sciences, 2 in support of the naval reactors program, and finally, 1 in support of the uranium enrichment program.

More recently, DOE has started using large-scale computers to support other programs in the Office of Energy Research (OER) in addition to Fusion Energy Research, known also as the Supercomputer Access Program. In February 1983, noting that various disciplines in OER needed supercomputing time, DOE set aside 5 percent of the National Magnetic Fusion Energy Computer Center (NMFCECC) facility at the Lawrence Livermore Laboratory for the energy research community through the Energy Sciences Advanced Computation Program. This supercomputer time, first available in June 1983, was immediately filled, and the hours requested exceeded the availability by an order of magnitude.⁸

In addition to providing access to largescale computational facilities for research purposes, DOE also funds supercomputer research to accelerate the development of systems designed and built in the United States. Historically, DOE laboratories needed the fastest computers available to fulfill their missions. As a result, the laboratories have traditionally played a key role in applying each succeeding generation of supercomputers, and developing software, since initially most machines have arrived without usable software. The DOE Applied Mathematical Sciences Research Program within OER supports new computer architectures and also mathematical and computer science research. Recently, program emphasis has shifted to new parallel multiprocessors. DOE has supported numerous projects and prototypes, such as the Hypercube machine, a California Institute of Technology project; the Ultracomputer at NYU; and the Cedar project at the University of Illinois.

The Magnetic Fusion Energy Computer Network, established in 1974, provides access to these supercomputer facilities, while also satisfying other communication needs. Lo-

⁸U.S. Department of Energy, *The Role of Supercomputers in Energy Research Programs*, February 1985, p. 1. Two Cray 1 computers at the National Magnetic Fusion Energy Computer Center (NMFCECC) and a Cray X-MP will be used in fiscal year 1985 as well to help meet the Class VI needs of the energy research community.

cated at Lawrence Livermore National Laboratory, the network includes one Cray 2, one Cray X-MP, and two Cray 1 computers that provide large-scale computing capability for the program. The four largest fusion contractors are linked by satellite communications, and others have access through telephone lines.

DOE and Florida State University, in a cooperative agreement, established the Florida State University Supercomputer Computational Research Institute in 1985. The new institute was established to explore aspects of energy-related computational problems, algorithms, and architectural research. DOE and the State of Florida will both contribute to the development of the institute. The State will provide 10 faculty positions and the computer facility; DOE will provide 69.5 percent of the funds to establish the institute; and Control Data Corp. will contribute equipment discounts and some personnel. (See table 4.)

Supercomputing Research Center, National Security Agency

Recognizing a need for research in supercomputing, the Institute for Defense Analyses has established the Supercomputing Research Center (SRC) for the National Security Agency. The new "center of excellence" in parallel processing will focus its development efforts on algorithms and systems, and also conduct research on national security programs. Still relatively new, the research agenda for the center has not been set, although basic re-

Table 4.—DOE Budget (in millions)

	Fiscal year 1985	Fiscal year 1986	Fiscal year 1987
Energy research	\$17.8	\$20.55	\$22.1
Florida State University SCRI	7.0	8.5	0
Additional funds in support of supercomputer operations at: Oak Ridge, Sandia(2), Los Alamos Scientific Laboratory(2), Livermore, MFE Center:			170.0

SOURCE James Decker Deputy Director Department of Energy

search in parallel algorithms, operating systems, languages and compilers, and computer architectures (including in-house construction) are the areas that will be investigated.¹⁸

The center will perform both classified and unclassified research. It is recognized that participation by the academic and industrial communities is essential and some exchange of information will be allowed.

Budget figures are not available though staffing levels and facility data are available: 100 professionals; 70 support staff; 100,000 square feet/permanent facility including current supercomputer and other equipment.²⁰

Defense Advanced Research Projects Agency

DARPA supports a large number of research projects that seek to advance the state of the art in multiprocessor system architectures. Unlike other programs, DARPA does not use supercomputers to fulfill agency mission requirements, but rather funds promising research projects that may advance the knowledge and design of current computer architectures. Support is directed towards fulfilling the goals of the Strategic Computing Program, which seeks to create a new generation of "machine intelligence technology."²¹ The program is partly focused on symbolic processing for artificial intelligence applications.

¹⁸ "Parallel Processor Programs in the Federal Government," draft report, 1985, p. 1.

¹⁹ Paul Schneck, Director, Supercomputing Research Center, personal communication, December 1985.

²⁰ Defense Advanced Research Projects Agency, *Strategic Computing, 1983*, p. 1.

DARPA funds several stages of R&D, from simulation to prototype construction and finally to benchmarking, a procedure using a set of programs and files designed to evaluate the performance of the hardware and software of a computer in a given configuration. In fiscal year 1985, DARPA expanded the multiprocessor program and funded a greater number of research projects. DARPA's efforts represent the bulk of government R&D in multiprocessor research, surpassing programs of DOE and NSF. Multiprocessor projects supported by DOD include the Butterfly Multiprocessor, which can accommodate up to 256 commercially available microprocessors, the Connection Machine, contains 64,000 processors, and the NONVON machine, with up to 8,000 large and small processors. DARPA has also participated in joint funding of projects with other agencies, such as the Cosmic Cube project at California Institute of Technology. (See table 5.)

Table 5.— Defense Advanced Research Projects Agency Budget (in millions)

	Fiscal year 1985	Fiscal year 1986	Fiscal year 1987
13Basic Research and Exploratory Development for Advanced Computer Research (this figure includes funding figures for other areas such as machine Intelligence and robotics)	\$1241	\$125	\$125
Machine Architecture	204	40	40
Distributed Computing and Software Systems	17.8	18	18
Network and Research Facilities	26.7	27	27
Total	\$1890	\$210	\$210

SOURCES FCCSET Panel on Advanced Computer Research in the Federal Government, summer 1985, and Craig Fields, personal communication February 1985

NETWORKS

A number of networks are operating or being created to support government computational programs by linking the users and the computational resources. Networks allow large, diverse, and geographically dispersed research communities to share resources in this case, large-scale computing facilities, exchange information, and share software.

As most commonly used, the term “network” refers to a communications system designed to provide links between two or more of a large collection of users. As used in the computer community, the term applies more specifically to a system that provides data transmission links among a set of computers—called “hosts”—and among users of those computers, who link to the network by means of a “terminal.” The term “network” commonly is applied to the entire assemblage of computers, communication lines, and user terminals.

A wide variety of networks already exist. They extend in scale from so-called “local area networks” that connect desk-top computers within an organization, to national and international networks, such as ARPANET, that connect thousands of users with several very powerful machines. Networks use many different data transmission speeds, encoding techniques, and protocols (basic sequences of message that tell the network what is to be done) that attune them to particular types of data communication and use. Because of these variations, different networks can be very difficult to interconnect, even when they use the same fundamental communications technology—e.g., satellites or fiber optics—and the same type of computers. Because such incompatibility can be a barrier to development and use of networks, some standards organizations are developing common descriptive models and interconnection standards at both the national and international level. The combination of activities by the standard organizations and the numerous networking activities will eventually lead to common standards and protocols that will satisfy the differing needs of the individual networks.

The scientific and research communities are already using several networks. ARPANET, developed and operated by DOD, hosts over 200 computers at nearly 100 universities, government laboratories and private sector research companies. CSNET is a data communications network linking computer scientists and engineers at over 120 university, government, and commercial sites throughout the United States and Canada, with gateways to Europe and the Far East. BITNET is a network of more than 350 computers at over 100 higher education and research institutions, with direct links to counterparts in Canada and Europe. Commercial value-added networks, such as TELENET, TYMNET, and UNINET, provide users with low-speed terminal access and moderate speed host-to-host access.

The concept of a network as it is used by most government agencies, goes beyond the notion of supercomputer access for remote users. Rather, networks are viewed as effective means to make the fullest possible use of the resources. Three networking programs are described below: NSF, DOE, and NASA.

The first recommendation of the Lax panel called for “increased access for the scientific and engineering research community through high bandwidth networks to adequate and regularly updated supercomputer facilities and experimental computers.”²² The development of NSFnet is designed to meet this need, as is the new NASA network in part. DOE already has a highly successful and efficient network, MFENET. With the substantial activity in this area it is important to be cognizant of the plans of the Federal programs.

National Science Foundation

A major new effort within OASC is the creation of a network to link researchers with the large-scale computing facilities. The new net-

²²National Science Foundation, *Report of the Panel on Large Scale Computing in Science and Engineering, 1982*, p. 10.

work, called NSFnet, is intended to be the basis for a national research network. The NSF concept for the new network is to "leverage" existing resources and networks with a new national network that is limited in both funding and authority. NSFnet would then try to take advantage of existing and new campus, community, State, and consortium networks—a network of networks following the DARPA internet model. To achieve this "internet" environment, NSF has adopted interim and long-term standards. (Initially, DOD internet protocol suite—TCP/IP plus existing applications—and eventual migration to IS0/OS1 protocols.) Future plans for NSFnet are uncertain because the network design and structure will be based on knowledge gained during phase I. Conceptually, the network will be designed to link end users with end resources.

Early efforts of NSFnet have focused on providing interim services—i.e., linking the researcher to the resource. To this end, phase I efforts provided funds to universities to buy equipment such as workstations; to arrange links between local networks, phase I centers, and other networks; and to fund consortium networks and pilot projects. The goal of phase I is to provide the top 50 to 60 campuses with shared or dedicated 56,000 bits per second circuits. NSF allocated \$5.9 million to the network program in fiscal year 1985, \$2.2 million of which was for local-access projects in fiscal year 1985 only.

NSF is also trying to link three existing networks, ARPANET, BITNET, and CSNET. OASC and DOD have signed a Memorandum of Understanding to expand the current ARPANET by one-third, with NSF funding 25 percent of the total costs of the expanded network. It is anticipated that ARPANET will be expanded by the time the phase II centers become operational.

There has also been some discussion of enhancing BITNET with NSF standard protocols, because BITNET has extensive links with the American academic community and international connections to European Academic Research Network and Japan. The addi-

tion of CSNET to NSFnet would greatly expand the access base for users.

NSFnet will also be linked with each of the four centers described above, which together encompass a large percentage of the U.S. academic community. Of the four centers, only the San Diego Center will be MFENET based, with migration to the NSF internet standards planned. Discussions are underway to include regional networks, such as MERIT and the Colorado State Network, in NSFnet.

Pilot projects proposed for phase I of NSFnet's development take advantage of available technologies in an effort to enhance the communications between the user and the resources. OASC is now considering three projects: Vitalink Translan,²³ DARPA wideband, and workstation projects. Funding for local-access projects has been allocated for fiscal year 1985, and will be used to connect campus networks and the NSF internet and for local computing facilities for supercomputer users. The NSF strategy is to favor service organizations on campus for handling the users concerns, rather than placing that burden on the future administrators and managers of NSFnet.²⁴

Department of Energy

The Magnetic Fusion Energy Network (MFENET) is an integral part of NMFEEC. This network links NMFEEC, located at the Lawrence Livermore National Laboratory, with computer centers at major fusion laboratories and other facilities nationwide. Over 4,000 users in 100 separate locations use MFENET. MFENET interconnects all computers at the center, and links remote user service centers that support local computing, experimental data acquisition, printers, terminals, remote user service stations, ARPANET (via the user service center at NMFEEC), and

²³Vitalink Translan is designed to interconnect several Ethernet local area networks in order for those networks to appear as a single large network. This would employ satellite and terrestrial lines.

²⁴Based on discussions and briefing by D. Jennings, National Science Foundation, June, September, and November 1985.

dial-up terminals (via TYMNET and commercial telephone lines) .25

The DOE fusion laboratories are linked by dual 56,000 bits per second (bps) satellite links to NMFEC. Many of the other users are connected to these centers by 4,800 or 9,600 bps leased telephone lines. Other users gain access to the centers through TYMNET, ARPANET, direct commercial dial, or Federal Telephone System. over 125 user service stations located at national laboratories, universities, or elsewhere provide local computing capabilities and, through the connection of the NMFEC, can function as remote output and job-entry stations. Those users not within the local dialing area of a user service station may dial access to the main computers and other network hosts.²⁶

DOE and the scientific communities have discussed expanding the resources available through MFECC and MFENET. However, budgetary restraints preclude this expansion in the near future.²⁷ (See table 6.)

National Aeronautics and Space Administration

Boeing Computer Services is currently developing the Program Support Communications Network (PSCN) for NASA. This network will serve NASA and selected user sites by providing wideband and other transmission services. The Numerical Aerodynamic Simulation (NAS) Program (see above) will not develop a separate network to support its pro-

²⁵D. Fuss and C. Tull, "Centralized Supercomputer Support for Magnetic Energy Fusion Research," *Proceedings of the IEEE*, January 1984, p. 36.

²⁶Ibid.

²⁷Don Austin, Department of Energy, personal communication, Aug. 14, 1985; and John Cavalini, Department of Energy, personal communication, June 5, 1985, and December 1985.

Table 6.—DOE Network Budget (in millions)

	Fiscal year 1985	Fiscal year 1986	Fiscal year 1987
Energy research	\$2.8	\$3.0	\$3.3

SOURCE James Decker, Deputy Director Department of Energy

gram and users, but instead will employ PSCN. PSCN will be operating by April 1986, with completion expected in the fall of 1986. Interim service for NAS and a limited number of users will be provided by a Long Haul Communication Prototype, which will examine and evaluate communication requirements of the NAS network.

PSCN is composed of several levels of service. Eight NASA centers and facilities will be linked by satellite: Ames, Dryden, Johnson, Kennedy, Marshall, Langley, Lewis, and Goddard. Eleven centers and facilities will be linked via 1.5 Mbps (megabits per second) terrestrial backbone links: Ames, Western Launch Operations, Jet Propulsion Laboratory, Dryden, Johnson, Michoud Assembly, Slidell Computer Complex, National Space Technology Laboratories, Kennedy, Langley, NASA Headquarters, Wallops, and Lewis. Four centers will be connected by 6.3 Mbps: Ames, Langley, Goddard, and Lewis Research Center. NASA's three research centers (Ames, Langley, and Lewis) will be linked through the Computer Network Subsystem (CNS).

CNS, a subsystem of PSCN, is a high-speed computer network linking selected NASA centers. Like PSCN, CNS will send jobs and files between these centers, while also maintaining updated information on the files and jobs within the network. This subsystem is designed for high-volume data transfer; it will use a T1 link for inter-site transfers and a 9,600 bps link for control information. (See table 7).

Table 7.—NASA Network NAS Budget (in millions)

	Fiscal year 1986	Fiscal year 1987
PSCN/NAS contribution	\$1.6	\$2.0

SOURCE Randy Graves, National Aeronautics and Space Administration

SOFTWARE DESIGN AND DEVELOPMENT

A wide variety of software is necessary to run any large computer system. Some broad categories of software are as follows:

- *Operating System:* The operating system software manages the flow of work through the machine. It has responsibilities such as assigning and controlling physical devices attached to the computer, cataloging and keeping track of data in the computer memory, and controlling the input of programs and data and the output of results. It also provides a user with a set of basic tools for accomplishing certain basic tasks common to most computer applications.
- *Programming Tools and Languages:* Users developing programs can make use of several types of automated aids. Higher level languages allow the user to express a program in a form that is simpler, more readable and understandable by a human, and more closely related to the technical language in which users communicate their problems, than is the basic computer "machine language." Programs written in these languages are easier to develop, more easily understandable by others, and often are more easily transferred between different machines. In addition, the compiler, which translates the higher level language program into machine language, can assist a user in taking advantage of particular characteristics of a specific computer by restructuring the program during translation. This latter advantage is particularly important for supercomputers, which have unusual features that must be employed efficiently to obtain the very high computation speeds.

Software engineers are developing many other types of tools to help users develop software, both in the programming and in the diagnostic and testing phase. The OTA report, *Information Technology R&D*, has additional discussion of the state of software engineering.

- *Applications Programs:* Applications programs are the software developed to solve specific problems. They range in scale and purpose from small, relatively simple programs designed to solve a specific problem once, to large programs usable by many researchers to solve a variety of related problems.

Underlying these levels of software is a rapidly developing body of computational theory flowing from mathematics and computer science. Computational mathematics, for example, by examining how basic methods for calculating solutions to equations behave in terms of efficiency and accuracy, helps to develop improved methods. In particular, computer scientists and mathematicians are only beginning to understand how to use optimally the types of highly parallel designs computer architects are exploring. As such knowledge is developed, it will lead to more efficient use of existing supercomputers, and will help computer architects design even better future machines.

Theoretical computer scientists are developing an understanding of how large, complex computational systems behave. For example, they study techniques for scheduling and allocating resources, they study the structure of higher order languages, and they explore the theoretical basis for determining whether a program is correct.

When OTA examined the field of software engineering in its report *Information Technology R&D*, it found that: "The lack of applications software for supercomputers has been a significant barrier to their adoption and use."²⁸ Concern over the availability and use of software of all types, not only applications software, for large-scale computing continues to be a critical issue, and "will grow worse in the near future because of the proliferation of

²⁸U.S. Congress, Office of Technology Assessment, *Information Technology R&D: Critical Trends and Issues*, OTA-CIT 268 (Washington, DC: U.S. Government Printing Office, February 1985), p. 64.

significantly different architectures.²⁹ Applications and systems software is available for the current generation vector machines, although in some cases, "it does not fully use the capabilities of the machine."³⁰ Also, the amount and breadth of systems and applications software appears insufficient to satisfy both current and projected demands generated by the expanded Federal programs." The software programs of NSF, NASA, SRC, and DOE are described below.

National Science Foundation

NSF funds academic and some corporate research in software engineering, software development, computational mathematics, and related areas. Together, the Division of Computer Research and the OASC provide over \$3 million for software development, although not all of this research is directly applicable to the needs of advanced scientific computing. Within the Division of Computer Research, the Software Engineering Program, the Software Systems Science, Computer Systems Design, the Theoretical Computer Science and Special Projects programs each fund a variety of research projects.

The OASC will be funding projects in direct support of the supercomputing centers described above. These efforts will focus on software productivity and computational mathematics. Research will be concentrated in the following areas: computer science research on programming environments, development of software tools, numerical analysis and algorithm development, and increasing effectiveness of advanced computers in research.

National Aeronautics and Space Administration

Since the mid-1970s, NASA research centers have supported algorithm development for supercomputers (Illiack and Star-100). Agency officials noted that earlier advanced scientific

computing efforts were hindered by insufficient or inadequate software, and so designed the NAS program schedule to avoid this problem. At both Cray and the Ames Research Center, development work has been underway since 1984 in the areas of operating systems modifications, network communication, distributed file systems, batch queuing systems, and common graphics services. The NAS Project developed a common user environment, based on a UNIX operating system, spanning a network of computers from multiple vendors. Except for specific user service projects, these software projects are for the NAS facility, not "research.

Supercomputing Research Center

The SRC software efforts concentrate on systems software development, which can be divided into three research areas: operating systems and compilers, language, and performance measurement. The research in these fields will be conducted at SRC.

Department of Energy

To support its varied research programs DOE is involved in numerous software development efforts that address systems, applications, and tools. These efforts, however, are closely tied to other DOE programs so their budget figures cannot be broken out. For example, at least some of the \$7.8 million allocated in fiscal year 1985 for the Department analytical and numerical methods program was devoted to computational software. Similarly, some of the \$10.4 million budgeted for advanced computer concepts was allocated to software engineering technologies. This past year, OER requested that \$2 million be set aside specifically for software tool development, but the Office of Management and Budget did not approve this request.³²

²⁹Office of Technology Assessment Workshop, Apr. 29, 1985.

³⁰Ibid.

³¹Ibid.

³²Don Austin, Department of Energy, personal communication, Aug. 14, 1985.

ISSUES: MANAGEMENT AND INSTITUTIONAL QUESTIONS

As the Federal programs in large-scale computing continue to grow in fiscal year 1986, several management and institutional questions arise concerning coordination, network development, and software. These questions may merit greater attention than they now receive given their importance and the enormity of the task facing the agencies. Such issues are particularly important because no single agency dominates largescale computer research or policy. Topics discussed below are: coordination between agencies, including issues of center management, resource allocation, and manpower and training; the Federal role in software design and development; and network development.

Coordination

ISSUE:

With the growing Federal efforts in large-scale computer research and access, **coordination among agencies is necessary to the success of the overall national program. Are** the current coordinating mechanisms, the FCCSET panels, succeeding at promoting "efficiency" in national programs while also allowing individual agencies to accomplish their missions?

The Federal investment in large-scale computing programs is already large and is growing. Coordination among programs and government agencies is one way to realize the best return on this investment. Coordination among agencies is also one way to ensure that national goals and interests are achieved, despite agencies disparate goals. To date, FCCSET has been the coordinating mechanism for all Federal programs, and has involved the private and academic communities in its deliberations and actions. The FCCSET panels successfully brought about discussions between agencies, and established a forum for discussion.

In 1983, FCCSET established the Panel on Supercomputers, which was charged with exploring what the U.S. Government could do to advance the development and use of large-scale computers. Subsequently, three sub-

panels were formed focusing on issues of procurement, access, and research coordination. Representatives from DOE, NSF, NASA, NSA, DOD, the Department of Commerce, and the Central Intelligence Agency are members of pertinent FCCSET subpanels.

Deliberations of the procurement subpanel, which was chaired by a DOE representative, focused on the question, "what should the government do to ensure that the United States retains its lead in supercomputers?" Specific recommendations included "guaranteed" buying of supercomputers by the government, increasing government purchasing of machines, and designing a supercomputer 200 times the current capabilities. After issuing a report in late 1983, this group merged with the access panel."

The second subpanel, on access, also chaired by a representative of DOE, published a report that called for upgrading current programs in a variety of ways and also recommended steps to meet the long-term objectives of providing scientists and engineers with access to supercomputers.³⁴

The research coordination subpanel, chaired by a representative of DOD, seeks to coordinate government agency programs that fund research that contributes to the U.S. technology base. A recent report of this subpanel outlines present and proposed federally sponsored research in very-high-performance computer research, summarizes agency funding in this area, including budgets for fiscal year 1983 to 1985, and presents findings and recommendations for future action.³⁵

³⁴Report to the Federal Coordinating Council on Science, Engineering, and Technology Supercomputing Panel on Recommended Government Actions To Retain U.S. Leadership in Supercomputers, n.d.

³⁴Report of the Federal Coordinating Council on Science, Engineering, and Technology Supercomputer Panel on Recommended Government Actions To Provide Access to Supercomputers, n.d.

³⁵Report of the Federal Coordinating Council on Science, Engineering and Technology Panel on Advanced Computer Research in the Federal Government, draft report, August 1985.

More recently, an ad hoc Office of Science and Technology Policy (OSTP) panel was established to explore avenues to improve cooperation between Federal agencies and the private and academic sectors. Although the committee headed by Jim Browne of the University of Texas, Austin, has not yet published its findings, Browne in a recent briefing, recommended a higher level of inter-agency coordination, either through FCCSET or a new mechanism. Browne also suggested the formation of a "continuing policy review committee."³⁶

The call for a higher level of interagency coordination reflects the limited abilities of FCCSET committees to substantially alter or establish government policy. Although members of the FCCSET committees are themselves intimately familiar with their agency's programs, they lack authority to implement recommendations made through the FCCSET subpanels. Implementation and policy directives must come from either the OSTP, the administrators of each agency, or both. This remains a potentially serious gap in the overall coordination of government programs. OSTP may wish to consider creating a panel whose membership would include pertinent agency administrators, scientific and research users, and private sector participants. Such a panel could serve several purposes:

- *Evaluate the Federal efforts underway to ensure that they are addressing the original concerns noted in the various reports (e.g. Lax, Bardon, Press, and others).* Such a reexamination could be a conduit for recommending "mid-course" corrections that may be necessary within the Federal program.
- *Provide a high level of authority to implement suggestions and recommendations made by the current FCCSET panels.*
- *Review of the Federal efforts in large-scale computing periodically.* Such a panel could review the overall health and direction of the Federal programs; identifying

new scientific opportunities made possible by the government programs; targeting specific areas of concern for FCCSET subpanels to address; recommending and authorizing actions that cut across more than one agency program.

The government emphasis on funding research, rather than prototype development, is one such issue that crosscuts most agency boundaries, has been the focus of a FCCSET subpanel, and remains a concern of scientists, researchers, and government officials. Some steps are being taken toward prototype development. DARPA's Strategic Computing Program is actively pushing technologies from the research stage into the prototype phase. A much smaller effort is underway in NSF Division of Computer Research and in the DOE program. In this case, DARPA's efforts are significantly greater than the others, which is cause for concern, since the research is focused on a specific mission and cannot have as broad a perspective in experimental technologies as can NSF. Also, DARPA's funding in this area greatly exceeds that of other agencies. At present, a FCCSET subpanel is examining current Federal efforts in the funding of the R&D of very-high-performance computer research. In this instance, an OSTP panel could assist by:

- determining if the current Federal efforts are sufficient to develop and "support" new experimental technologies,
- identifying the "next" steps that would facilitate the transfer of research results to utilization, and
- examining the balance of agency programs in the funding of experimental technologies.

Center Management

ISSUE:

The FCCSET subpanels on access, procurement, and research coordination have acted as forums for information exchange between Federal programs in large-scale computer research. Are comparable FCCSET bodies needed to discuss center management issues, such as resource policies and funding, and manpower and training?

³⁶Briefing of the NSF Advisory Committee, June 1985.

As the number of large-scale computing sites proliferate, need continues for FCCSET to coordinate supercomputer issues. FCCSET can remain a single forum to consider in full national interests, goals, and programs; the centers are also viewed here as national resources whose significance extends beyond the mission needs of the funding agency.

The DOE supercomputing programs, for example, are designed to achieve specific missions. Nevertheless, careful planning, an appreciation of user needs, and the establishment of a network have allowed these supercomputing programs to become more national in scope. NSF has considered a comparable effort—a strong program plan and well-defined research agenda to provide the framework for success in the NSF centers' program. A program plan will provide the basis for decisions on how time should be allocated at the facilities, and in turn will define the type of research performed at the sites. Soon after the centers begin operation, research trends of users will be evident allowing one to see both ongoing research and also new research opportunities that have been created. This is especially important since the NSF centers' program has no single constituency, but relies on the support of 20 disciplines within the agency.

FCCSET also operates as a forum for agencies and supercomputer center managers to exchange information about experiences in operating the centers, software, network development, and general lessons learned. For instance, the experiences gleaned by DOE in establishing its facilities and network will be invaluable to NSF and NASA, despite the agencies' very different missions. DOE experience may be particularly germane over the next few years, as NSF judges the success of its funding formula or "seed money philosophy." Unlike other mission agencies that fully fund their largescale facilities programs, NSF has required the individual centers to raise additional funds to supplement the NSF grants. Some scientists have claimed that this policy has hindered the long-term program by streamlining center staff.

Resource Policies

Federal supercomputer programs operate under a variety of resource or allocation policies. The NSF policies for allocation of time on computing facilities have undergone some changes since the phase I centers were established. Originally, NSF allocated time in batches of 10 hours or less, with no peer review. Currently, in addition to those who request 10 hours or less, one can submit a request for time to NSF with no maximum time limit set, and program directors there distribute hours at their discretion. If the requestor is not an NSF grantee, an application for time will undergo normal NSF peer review. Under the new policy, 60 percent of the centers' service units will be distributed by NSF, with the remaining 40 percent allocated by the individual centers. Up to 25 percent of each centers' allocation may be available for proprietary research. Also, if use of facilities is sold to for-profit institutions, the organization will be charged the full cost of using the service. Each center is expected to create a review panel of scientists from multiple disciplines to evaluate each research proposal and to adopt standard NSF peer review principles.

At DOE, time allocation in the Supercomputer Access Program is based on scientific merit and need. Investigators make requests to the Office of Energy Research program directors at the beginning of each year, and the program directors rank the proposals. These determinations are not always final; program administrators shift allocations during the year to allow for immediate program needs. This "flexible" program policy highlights the distinction between the allocation plan of a research agency like NSF and a mission-directed allocation plan, such as DOE'S.

NASA is in the process of establishing NAS usage policy, and guidelines for user allocation. Allocation of time on the NAS system will be made on the basis of uniqueness and suitability of the proposal to the NAS facilities. An "announcement of opportunities" was released recently by NASA, alerting interested researchers that the NAS system will

be available for use in fiscal year 1986. Once NASA receives responses to this announcement, the agency will finalize its allocation policies. Generally, program managers foresee a rough breakdown of 55 percent NASA-related research, 5 percent university research (non-NASA sponsored), 25 percent DOD and other government agency sponsored research, and the remaining 15 percent, proprietary research.

Because each agency has a different mission, each requires its own allocation policies. FCCSET has established a mechanism for allocating time on facilities for scientists who are funded by more than one agency. After the NSF centers and the NAS have been operating for awhile, it may be advantageous to examine allocation policies to ensure that Federal goals, such as the advancement of science, are attained. At present, though, a FCCSET panel could contribute little else to the issue of allocation.

Manpower and Training

Training and manpower concerns are twofold: how to use the available computer resources effectively, and how to make the best use of available personnel at each site. To promote the understanding of the facilities that leads to effective use as well as foster new use of the large-scale computing resources, NSF and DOD sponsored three summer institutes to train researchers in the development of codes, vector techniques, and networks for remote access. These institutes also gave researchers a forum to present their work and discuss the role of supercomputers in performing it.

DOE regularly sponsors conferences and tutorials on the supercomputer operating system (DOE's Compatible Time-Sharing System) and on network access, both at the Lawrence Livermore Laboratory and elsewhere around the country. Special conferences are also held that explore specific problems and issues pertinent to members of DOE research communities.

In anticipation of the NAS starting operation, NASA's training efforts include extensive "onboard" support, a user interface group consisting of 40 organizations across the country, and workshops to familiarize users with the system.

These combined efforts will help educate professionals at universities who are not computer scientists on how to use the facilities. Generally the pool of qualified individuals available to support the supercomputing facilities is limited. With the opening of five new NSF centers, the Florida State University Center, and one NASA center in fiscal year 1987, the demand for talented personnel will grow, both for people to run the facilities and to work through problems with research scientists. And, too few personnel are available to meet the current and projected software demands, since too few have had access to experimental computing facilities and the number of graduate students in computational science and mathematics is small. Two key concerns of scientists are: the training of new personnel to help operate the new centers and make them effective and well utilized; and preventing "raids" of experienced personnel from an existing center to a planned center, to avoid hurting an ongoing program. The human resource questions are of great importance in determining the success of the supercomputer programs. It may be worthwhile for a FCCSET panel to consider a review of agency training efforts and mechanisms to ensure that greater training efforts be undertaken, particularly in light of the development of new architectures.

Problems and Prospects for Software Development

ISSUE:

The availability of software for large-scale computing continues to be a critical concern. What Federal efforts are necessary to tackle this software problem?

The greater availability of time on supercomputers as a result of the Federal investments and programs described above, can

both exacerbate and ease the current software demands. As more and more users turn to supercomputers to solve their scientific problems, new software will be generated that will increase the usefulness and capabilities of the machines. Once a network is in place, those who use it will share software, as well. At the same time, until more software is developed, users will inevitably be frustrated by their inability to access the machines easily and the lack of standard "libraries" of applications.³⁷ Until more scientists use large-scale computers for their research, the lack of applications software will remain a problem. Traditionally, manufacturers have not developed a substantial base of software (including applications software) for their machines. In fact, it has been stated that:

... it has taken over 5 years to develop software to efficiently support and utilize vector technology .38

Also, the ability to move software from machine to machine and site to site is needed to encourage users to experiment with the advanced scientific computing facilities. Transportability of software will be very important in building a base of new users: "We should be in a position where we can make it very easy to move from one machine to another."³⁹ Missing a comfortable bridge to move code/software from machines in use (e.g., a Vax) to a supercomputer or parallel processor machine, the scientific community may withhold its support, because of delays in performing research and the complexity of undertaking the research on a supercomputer. As noted in the FCCSET report:

... Researchers may not use this important scientific tool if its use proves difficult to learn, frustrating and generally inconvenient."

The nurturing of new architectures will also bring a host of software problems. The new architectures demand "a rather fundamental

rethinking of problems and a rather fundamental rethinking of strategies,"⁴⁰ in the creation and development of algorithms. At a meeting sponsored by NSF in December 1984, researchers concluded that a new type of research activity was necessary.⁴² They recommended interdisciplinary teams be formed, consisting of mathematicians, computer scientists, and scientists from disciplines with problems to be solved, to tackle the software needed for the new highly parallel architectures. Recently, NSF convened an ad hoc panel to consider the establishment of a National Supercomputer Software Institute. Although the panel recommended against such a move, it did encourage NSF to:

mount a program to encourage research on parallel software by increasing the level of university access to experimental parallel computers and by funding promising research projects at sufficient scale to allow major software developments to be undertaken either by university groups or consortia or by joint university-industry teams.⁴³

Software is an important **determinant** of the efficiency of these machines and the types of problems that can be tackled with them. It is also an influence on the design of the next generation of machine. For these reasons, an investment in algorithms and software development is integral to any large-scale computation program.

Finally, as noted earlier, the shortage of personnel needed to meet the current and projected software demands will grow worse when coupled with the introduction of radically different and new architectures.

The issues of supercomputer software require better definition and resolution. To this end, several ideas were proposed at the OTA Workshop to address these issues:

- *A follow-on study to the Lax report that specifically addresses software issues, defines the problems, and explores possible*

³⁷Office of Technology Assessment Workshop, Apr. 29, 1985.

³⁸D. Fuss and C. Tull, "Supercomputer Support for Magnetic Fusion Research," *Proceedings of the IEEE*, January 1984, p. 41.

³⁹Office of Technology Assessment Workshop, Apr. 29, 1985.

⁴⁰FCCSET, op. cit., Access, Section III.

⁴⁰Office of Technology Assessment Workshop, Apr. 29, 1985.

⁴¹Ibid.

⁴²Jacob Schwartz, Chairman, Ad Hoc Panel on a National Supercomputer Institute, letter to Erick Block, Director, National Science Foundation, Nov. 4, 1985.

solutions or actions. This study should be sponsored by one or more agencies, such as NSF, DOE, NASA, or DOD. The Lax report recommended: "Increased research in computational mathematics, software and algorithms necessary to the effective and efficient use of supercomputer systems." ⁴

- *Establishment of a software development/engineering institute separate from each agency but supported by funding from each of the agencies.* The institute's charge would be to develop software and algorithms for use on all types of large-scale computing machines and to advance the state of the art.
- *The formation of a new subpanel within FCCSET to address issues of software and algorithm development.* This group would act as a coordinating body within the government to keep track of agency efforts and suggest needed improvements and directions.

Network Design and Development

ISSUE:

The National Science Foundation has stated that NSFnet will be the basis for a national research network. Should NSFnet be the basis for a national research network? And until this network is in place, how should NSFnet be administered and managed?

Networks permit users easier access to resources and facilitate the exchange of information between and among them. "The ability of a network to knit together the members of a sprawling community has proved to be the most powerful way of fostering scientific advancement yet discovered." ⁴ Networks make information and resources available to the researcher regardless of geography, thus expanding scientists' research base. Networks may also provide uniform access to their resources and users without requiring the user

⁴National Science Foundation, *Report of the Panel on Large Scale Computing in Science and Engineering*, 1982.

⁵Peter Denning, "The Science of Computing, Computer Networks," *American Scientist*, March-April 1985, p. 127.

to know the physical location of resources or other users.

As an increasing number of networks are designed, developed, and created to support Federal programs in largescale computation, their importance will become more evident. Networks are essential to the success of these programs because they provide access to resources and information on a national scale.

The proliferation of computer networks, combined with the expansion of Federal large scale computing programs and the identified need for access to these computational resources, has generated interest in the creation of a *national research network*. The research needs and goals of the NSF program require a network that will serve a diverse, geographically dispersed group of scientists and engineers. This broad-based approach also suits the development of national research network, a long-term goal of certain segments of the research community. The national research network would be a broad-based telecommunications network designed to serve the complex and diverse requirements of the national research community and address the broader issue of providing scientific information to this community. Such a network could provide researchers with information, and a means to exchange it. Services such as file transfer, computer-conferencing, electronic mail, and bulletin boards could be available. It could also stimulate the formation of major new databases and new scientific opportunities; and facilitate access to remote resources, including large-scale computing resources.

The proposed linking of NSFnet with a national research network raises several questions and issues.

- What are the goals of a national research network?
- Is NSFnet the appropriate base for a national research network?
- Are the goals of NSFnet and a national research network compatible?
- Is the proposed design of a national research network the most feasible approach?

- What is the Federal role in the development of this national network?
- Is NSF the appropriate lead agency for this endeavor?

Some of these questions are more easily answered than others, but all merit and need thorough discussion. What these questions illuminate is the need, in any discussion of a national network, to keep sight of the issues of advanced scientific computing and access to these resources. NSFnet is developing quickly and choices made today for pilot projects and other elements, will affect the future configuration of a national research network.

As NSFnet develops and is used more and more by the user communities, NSF will be called on to be a network manager. But, with one exception, the agency has not managed the day-to-day operations of a large-scale project following the project's early stages. In the near future the agency will need to decide what role it will play as NSFnet evolves. NSF-OASC must soon decide which of several courses of action it should take in the management of NSFnet. Three are outlined below:

1. *Retain all management operations within the OASC, including day-to-day operations (network operations), user services, and financial services, etc.* Under this option, OASC will be responsible for all management aspects of the network including interactions with all users and other networks.
2. *Identify and select a private firm with previous experience in the management of a net work to manage daily operations of NSFnet, including interactions with users and other networks.* Under this arrangement, NSF would retain overall policy development responsibilities.
3. *Create a Board of Users that includes representatives from government agencies with participating networks such as DOD-DA RPA-ARPANET, individuals*

from other participating networks such as CSNET, and users including the center consortiums of NSFnet to assist NSF with policy development/direction. NSF would retain lead authority and be the "chairman" of this Board of Users. A qualified outside contractor could also be made responsible for network management, user services, etc.

Regardless of the style of management chosen by NSF for NSFnet, several other issues need to be addressed in the near future by the agency. For NSFnet to be successful as it is currently designed, networks wishing to join NSFnet must either now employ interim internet working standards, or agree to change to these standards in the future. This shift to compatible protocols and standards must be tightly managed and centrally coordinated for the strategy to succeed.

The proliferation of government and private networks, many of which use a variety of technologies, standards, and protocols, may also prompt concern in the near future. To address the issue, the FCCSET has recently formed a networking subcommittee. It is charged with examining both individual agency networks and community networks that are necessary in order to fulfill agency mission requirements and goals. The committee will also examine the feasibility of a single computer network infrastructure for supercomputer access and research collaboration in the United States. The creation of NSFnet is seen as a unique opportunity for coordinating with other Federal agency networks.

Finally, the industrial community as such has not been included in the NSFnet plans. This may cause two problems: first, users within this community may lack easy access to facilities, including the network; and second, NSFnet communities may not be able to tap private sector institutions with large-scale resources.