

Originally Processed With FOIA(s):
2005-0336-F

FOIA Number:
2005-0336-F

FOIA MARKER

**This is not a textual record. This is used as an
administrative marker by the George Bush Presidential
Library Staff.**

Record Group/Collection: George H.W. Bush Presidential Records
Collection/Office of Origin: Science and Technology Policy, Office of (OSTP)
Series: Publications
Subseries: Reports and Publications Files

OA/ID Number: 62109
Folder ID Number: 62109-004

Folder Title:
Federal Coordinating Council for Science, Engineering, and Technology [2 of 2]

Stack:

Row:

Section:

Shelf:

Position:

0


0

0

0


Grand Challenges: High Performance Computing and Communications

The FY 1992 U.S. Research and Development Program



A Report by the Committee on
Physical, Mathematical, and Engineering Sciences

To Supplement the President's Fiscal Year 1992 Budget



On the Cover:

1. Numerically modelled thunderstorm.
2. Computational model of chemical carcinogen binding with DNA molecule.
3. Visualization of structure of superconducting material.
4. Simulation of acid rain pollutants over Ohio River basin.
5. Aerodynamic characteristics of space vehicle: computer simulated versus wind tunnel data.
6. Photo of wafer and multi-chip prototypes.
7. Numerical simulation of fuel jet.
8. Computer image of earth's biosphere components generated from satellite data.

The images used in this report were produced by ongoing scientific projects in areas of the planned HPCC program. They were selected to illustrate the breadth of subject matter of the HPCC program, and are elaborated upon in Chapter 4. The U.S. map suggests how the National Research and Education Network supports geographically distributed collaborative research activities.

Grand Challenges: High Performance Computing and Communications

The FY 1992 U.S. Research and Development Program

A Report by the Committee on
Physical, Mathematical, and Engineering Sciences

Federal Coordinating Council for Science,
Engineering, and Technology

Office of Science and Technology Policy

To Supplement the President's Fiscal Year 1992 Budget

**Office of Science and Technology Policy
Federal Coordinating Council for Science,
Engineering, and Technology
Committee on Physical, Mathematical, and
Engineering Sciences**

Acting Chairman

Charles Herzfeld, Department of Defense

Members

Barry Williamson, Department of the Interior
Charles E. Hess, Department of Agriculture
Robert White, Department of Commerce
William Raub, Department of Health and Human Services
James Decker, Department of Energy
John B. Childers, Department of Education
Norine Noonan, Office of Management and Budget
Eugene Wong, Office of Science and Technology Policy
Arnold Aldrich, National Aeronautics and Space Administration
Erich Bretthauer, Environmental Protection Agency
Fredrick M. Bernthal, National Science Foundation

Executive Secretary

Jane Stutsman, National Science Foundation

FCCSET Directorate

Maryanne C. Bach, Executive Director
Charles H. Dickens, Senior Staff Associate
Jean Grace, Executive Assistant

High Performance Computing and Communications Working Group

David Nelson, Department of Energy, Chairman
Barry Boehm, Defense Advanced Research Projects Agency, Co-Chairman
Charles N. Brownstein, National Science Foundation, Co-Chairman
Lee Holcomb, National Aeronautics and Space Administration, Co-Chairman
Lawrence E. Brandt, National Science Foundation
John S. Cavallini, Department of Energy
Melvyn Ciment, National Science Foundation
Jack D. Fellows, Office of Management and Budget
Stephen M. Griffin, National Science Foundation
Paul E. Hunter, National Aeronautics and Space Administration
Gary M. Johnson, Department of Energy
Thomas A. Kitchens, Department of Energy
Norman H. Kreisman, Department of Energy
Albert T. Landberg, Jr., National Institute of Standards and Technology
Fred Scoresby Long, National Oceanic and Atmospheric Administration
Daniel R. Masys, National Library of Medicine
Joan H. Novak, Environmental Protection Agency
William L. Scherlis, Defense Advanced Research Projects Agency
Paul H. Smith, National Aeronautics and Space Administration
K. Speierman, National Security Agency
Stephen L. Squires, Defense Advanced Research Projects Agency
Stephen Wolff, National Science Foundation

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF SCIENCE AND TECHNOLOGY POLICY
WASHINGTON, D.C. 20506

MEMBERS OF CONGRESS:

I am pleased to forward with this letter "Grand Challenges: High Performance Computing and Communications, The FY 1992 U. S. Research and Development Program," a report by the Committee on Physical, Mathematical, and Engineering Sciences of the Federal Coordinating Council for Science, Engineering, and Technology, a supplement to the President's Fiscal Year 1992 Budget.

The report presents an ambitious and well-coordinated research and development program designed to sustain and extend U.S. leadership in all advanced areas of computing and networking. The program not only provides a far-sighted vision for the underlying technologies but also gives recognition to the importance of both human resources and those applications that serve major national needs. This is a program of national investment that will bring both economic and social dividends.

The program is strategically related to other key components of the President's overall approach to challenges in science, technology, and education. It provides for the use of improved computational and communications technologies to contribute to more effective solutions of grand challenge problems.

The goal of the Federal High Performance Computing and Communications (HPCC) Program is to accelerate significantly the commercial availability and utilization of the next generation of high performance computers and networks. Recent advances offer the potential for a thousand-fold improvement in useful computing capability and a hundred-fold improvement in available computer communications capability by 1996. These advances will come through improvements in hardware and software. This increased capability will greatly expand the availability of these resources for research and education. It is my personal view, moreover, that the successful implementation of this program will lay the foundation for changes in education at all levels.

Several years of effort on the part of senior government, industry, and academic scientists and managers are reflected in this program. Acting Chairman Charles Herzfeld and his interagency committee members, associates, and staff are to be commended on the excellent work that is manifest in both the program and the report.



D. Allan Bromley
Director

To obtain a copy of this document — send request to:

Committee on Physical, Mathematical, and Engineering Sciences
c/o National Science Foundation
Computer and Information Science and Engineering
1800 G Street, N.W.
Washington, D.C. 20550

Table of Contents

	Page
Executive Summary	2
1. Program Goals and Overview	5
Introduction	
Needs and Benefits	
Program Description	
Goals	
Strategy	
Program Execution Strategy	
2. HPCC Program Components	12
High Performance Computing Systems	
Advanced Software Technology and Algorithms	
National Research and Education Network	
Basic Research and Human Resources	
3. Program Development and Agency Budgets	22
Program Planning	
Evaluation Criteria	
Agency Budgets	
Agency Program Descriptions	
Defense Advanced Research Projects Agency	
Department of Energy	
National Aeronautics and Space Administration	
National Science Foundation	
National Institute for Standards and Technology	
National Oceanic and Atmospheric Administration	
Environmental Protection Agency	
National Library of Medicine	
4. Grand Challenge and Supporting Technology Case Studies	37
Forecasting Severe Weather Events	
Cancer Genes	
Predicting New Superconductors	
Air Pollution	
Aerospace Vehicle Design	
Energy Conservation and Turbulent Combustion	
Microsystems Design and Packaging	
Earth's Biosphere	
High Speed Networks	
Education Using the NREN	
Glossary	56

List of Figures

	Page
Figure 1. The High Performance Computing and Communications Program	3
Figure 2. Performance Requirements for Grand Challenge Problems	7
Figure 3. Computer System Performance Trends for Grand Challenge Problems	13
Figure 4. Performance Improvement	15
Figure 5. NREN Applications	18
Figure 6. Evaluation Criteria for the HPCC Program	23
Figure 7. HPCC Budgets by Agency and Program Component	24
Figure 8. Agency Responsibilities	26

**Grand Challenges:
High Performance Computing
and Communications**

The FY 1992 U.S. Research and Development Program

Executive Summary

EXECUTIVE SUMMARY

- High performance computing and computer communications networks are becoming increasingly important to scientific advancement, economic competition, and national security. The technology is reaching the point of having a transforming effect on our society, industries, and educational institutions. The goal of the Federal High Performance Computing and Communications (HPCC) Program is to accelerate significantly the commercial availability and utilization of the next generation of high performance computers and networks in a manner consistent with the Strategic and Integrating Priorities shown in Figure 1.
- The HPCC Program is the result of several years of effort on the part of senior government, industry, and academic scientists and managers to design a research agenda to extend U.S. leadership in high performance computing and networking technologies.
- For FY 1992 the HPCC Program proposes to invest \$638 million in the four complementary and coordinated components shown in Figure 1. This investment represents a \$149 million, or 30%, increase over the FY 1991 enacted level.
- The HPCC Program is driven by the recognition that unprecedented computational power and capability is needed to investigate and understand a wide range of scientific and engineering "grand challenge" problems. These are fundamental problems whose solution is critical to national needs. Progress toward solution of these problems is essential to fulfilling many of the missions of the participating agencies. Examples of grand challenges addressed include: prediction of weather, climate, and global change; determination of molecular, atomic, and nuclear structure; understanding turbulence, pollution dispersion, and combustion systems; mapping the human genome and understanding the structure of biological macromolecules; improving research and education communications; understanding the nature of new materials; and problems applicable to national security needs.
- The HPCC Program nurtures the educational process at all levels by improving academic research and teaching capabilities. Advanced computing and computer communications technologies will accelerate the research process in all disciplines and enable educators to integrate new knowledge and methodologies directly into course curricula. Students at all levels will be drawn into learning and participating in a wide variety of research experiences in all components of this program.
- The FY 1992 Program and this document were developed by the HPCC Working Group under the direction of the Committee on Physical, Mathematical, and Engineering Sciences of the Federal Coordinating Council for Science, Engineering, and Technology.

Figure 1

The High Performance Computing and Communications Program

Goals: Strategic Priorities

Extend U.S. technological leadership in high performance computing and computer communications.

Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment.

Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.

Strategy: Integrating Priorities

Support solutions to important scientific and technical challenges through a vigorous R&D effort.

Reduce the uncertainties to industry for R&D and use of this technology through increased cooperation between government, industry, and universities and by the continued use of government and government-funded facilities as a prototype user for early commercial HPCC products.

Support the underlying research, network, and computational infrastructures on which U.S. high performance computing technology is based.

Support the U.S. human resource base to meet the needs of industry, universities, and government.

Program Components

High Performance Computing Systems

- Research for Future Generations of Computing Systems
- System Design Tools
- Advanced Prototype Systems
- Evaluation of Early Systems

Advanced Software Technology and Algorithms

- Software Support for Grand Challenges
- Software Components and Tools
- Computational Techniques
- High Performance Computing Research Centers

National Research and Education Network

- Interagency Interim NREN
- Gigabits Research and Development

Basic Research and Human Resources

- Basic Research
- Research Participation and Training
- Infrastructure
- Education, Training, and Curriculum

1. PROGRAM GOALS AND OVERVIEW

Introduction

High performance computing (HPC) is emerging as a powerful technology for industrial design and manufacturing, scientific research, communications, and information management. A robust U.S. high performance computing and computer communications capability contributes to leadership in critical technology and national security areas. Improved computational and communications technologies contribute to more effective approaches to problem solving, new products and services, and enhanced national competitiveness across broad sectors of the economy.

Recent advances offer the potential for a thousand-fold improvement in useful computing capability and a hundred-fold improvement in available computer communications capability by 1996. Based on several years of planning, under the auspices of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), Federal agencies and the technical community have developed the Federal High Performance Computing and Communications (HPCC) Program to realize this potential and to meet the challenges of advancing computing and associated communications technology and practices. Agencies have realigned and enhanced their high performance computing research and development programs, coordinated their activities with other agencies, and shared common resources to develop the program presented in this document.

Needs and Benefits

High performance computing has become a vital enabling force in the conduct of science and engineering research over the past three decades. Computational science and engineering has joined, and in some areas displaced, the traditional methods of theory and experiment. For example, in the design of commercial aircraft, many engineering issues are resolved through computer simulation rather than through costly wind tunnel experiments. This trend has been powered by computing hardware and software, computational methodologies and algorithms, availability and access to high performance computing systems and infrastructure, and the growth of a trained pool of scientists and engineers. This process has been nurtured by Federal investment in advanced research, agency supercomputer centers, and national networks through DARPA, DOE, NASA, NSA, and NSF. These facilities have contributed to national mission areas such as energy, space, health, defense, environment, weather, and basic science and

technology that could not be effectively addressed without the use of such advanced facilities.

High performance computing technology is knowledge and innovation intensive. Its development and use engages the entire scientific and engineering community. Building upon fundamental research of the early 1980's, a new computing technology of scalable parallel processing computers emerged. By the mid-1990's, this innovative approach to high performance computing systems promises to achieve sustained performance improvements of a thousand-fold compared to current systems.

In a growing number of science and technology fields, progress and productivity in modern research are increasingly dependent on the close interaction of people located in distant places, sharing and accessing computational resources across networks. Although the U.S. is the world leader in most of the critical aspects of computing technology, this lead is being challenged.

The Federal HPCC Program is a strategic Federal investment in the frontiers of computing and computer communications technologies and is formulated to satisfy national needs from a variety of perspectives including: technology, science applications, human resources, and technology transition. Needs are derived from the agency missions and based on the underlying science, engineering, and technology base required to carry out these missions. Many of these mission needs are related to solving very intensive large scale computing problems. These fundamental problems often cut across various agencies and missions and are called grand challenge problems (Figure 2).

The industrial and academic sectors provide major sources of innovation, cost effective development, and support of information technologies and their application to grand challenge problems. As these technologies are developed, the results support the Federal agency missions and become available nationally. The program provides for development of these revolutionary technologies within a framework of a partnership among government, industry, and academe and allows for rapid transition of laboratory results into new products that will then be applied within the program.

Figure 2

Performance Requirements for Grand Challenge Problems

Computer Performance
in Billions of Operations
per Second

1000

100

10

1

0.1

Grand Challenges

Climate Modeling
Fluid Turbulence
Pollution Dispersion
Human Genome
Ocean Circulation
Quantum Chromodynamics
Semiconductor Modeling
Superconductor Modeling
Combustion Systems
Vision and Cognition

Structural
Biology

Pharmaceutical
Design

Speech and
Natural Language

Vehicle Signature

ULSI Design

72 Hour
Weather

Estimate of Higgs
Boson Mass

48 Hour
Weather

Chemical
Dynamics

Airfoil Design

2D Plasma
Modeling

3D Plasma
Modeling

1980

1990

2000

Program Description

The Program consists of four integrated components representing the key areas of high performance computing and communications:

High Performance Computing Systems (HPCS)— the development of the underlying technology required for scalable parallel computing systems capable of sustaining trillions of operations per second on large problems.

Advanced Software Technology and Algorithms (ASTA)— the development of generic software technology and algorithms for grand challenge research applications to realize the performance potential of high performance computing systems in a networked environment.

National Research and Education Network (NREN)— the development of a national high speed network to provide distributed computing capability to research and educational institutions and to further advanced research on very high speed networks and applications.

Basic Research and Human Resources (BRHR)— support for individual investigator and multidisciplinary long term research drawn from diverse disciplines, including computer science, computer engineering, and computational science and engineering; initiation of activities to significantly increase the pool of trained personnel; and support for efforts leading to accelerated technology transition.

Advances in high performance computing enable advances in almost every other science and engineering discipline. There is a complex web of research interdependencies among the four components, and each area contributes to progress in other areas. Because of these dependencies, achieving and maintaining balance between the research components is a primary goal and the most important priority in the current context and environment. The HPCC Program is designed to provide balanced support both for technology areas including components, systems, software, and algorithms, and for applications, infrastructure, and human resources to achieve rapid overall research progress and productivity.

The component activities are planned to produce a succession of intermediate benefits on the way to meeting the long range programmatic goals. The HPCC Program builds on Federal programs already in place, providing additional resources in selected areas. Computational science and engineering grand challenges as illustrated in Figure 2 are the focal points for these efforts.

Goals

The goals of the High Performance Computing and Communications Program are to:

Extend U.S. technological leadership in high performance computing and computer communications.

Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment.

Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.

These goals will be realized by achieving: computational performance of one trillion operations per second (10^{12} ops, or teraops) on a wide range of important applications; development of associated system software, tools, and improved algorithms for a wide range of problems; a national research network capable of one billion bits per second (10^9 bits, or gigabits); sufficient production of Ph.D.'s and other trained professionals per year in computational science and engineering to enable effective use and application of these new technologies.

Strategy

The goals will be met through coordinated government, industry, and university collaboration to:

Support solutions to important scientific and technical challenges through a vigorous R&D effort.

Reduce the uncertainties to industry for R&D and use of this technology through increased cooperation between government, industry, and universities and by the continued use of government and government-funded facilities as a prototype user for early commercial HPCC products.

Support the underlying research, network, and computational infrastructures on which U.S. high performance computing technology is based.

Support the U.S. human resource base to meet the needs of industry, universities, and government.

At the program component level the strategy will: exploit and extend scalable parallelism and engage in intensive software development in HPCCS; use common requirements of the grand challenges to foster HPC

software progress in ASTA to strengthen HPC software development and coordination; evolve from the current Internet network to the NREN using a series of testbed systems; and strengthen academic activities in computer science and computational science and engineering as part of BRHR.

Program Execution Strategy

The strategy is based on the strengths of partnerships among the Federal agencies and other organizations. Major portions of the program will be cost-shared and leveraged by the participation of industry and universities. This general approach operates well today and provides strong evidence that the HPCC Program can be successful in the future. The specific elements of the approach are to:

Create a balanced, critical-mass program. The program must achieve sufficient scope and balance among the components. A technology program that created extremely fast processors without comparable memory, input-output, and mass storage systems would not succeed. Neither would a program that created powerful computers without adequate software, network access, and capable people. Similarly, a program that created only high performance networks would not satisfy the increased performance requirements needed for grand challenges. The HPCC Program must operate at a sufficient scale and coverage of technology areas that the new technologies can be effectively applied to grand challenge problems with acceptable levels of risk. The HPCC Program achieves balance by the extensive participation of experienced users, applications developers, and researchers in the HPCC disciplines throughout the design, development, and implementation process.

Build on agency strengths. The strategy builds on agency strengths by giving appropriate agencies the responsibility to coordinate activities for areas of demonstrated capability. It also ensures that the strengths of the other agencies are included by integrating their participation in various task areas. DOE, NASA, NSF, DOC (NOAA, NIST), and DOD (DARPA, NSA) have decades of experience in applying the world's most powerful computers, and thus provide a valuable perspective on high performance computing requirements and applications. NSF has the demonstrated technical and operational expertise needed for deploying high performance national networks in the research community, and is uniquely positioned to support basic research in computer science, computational science, and other scientific areas that can benefit from high performance computing in

interdisciplinary programs. DARPA, having pioneered technologies for high performance computing systems and microelectronic components, has a strong existing research program in teraops computing and gigabit networking, and offers rapid technology transition into commercial products in support of DOD science, technology, and applications base. EPA, NIH, and NOAA provide complementary HPCC requirements perspectives and application bases. NIST has extensive experience in HPC systems and networks instrumentation and evaluation, and provides a means for standards development.

Accelerate Technology Transfer. The transfer of technology from research to development and to application can be a very slow process due to a number of barriers to the use of new technology: high initial cost, inadequate and user-unfriendly software and systems, and lack of standards. The HPCC Program relies upon substantial industry participation to overcome these barriers and yield the benefits derived from moving new technologies to industry. The strategy accelerates technology transition by using a participative development process for each of the task areas, and by stimulating the growth of shared knowledge and capabilities. An example of this is the creation of network-accessible repositories of scalable HPC software and associated user groups to provide usage feedback and improvement of the HPC software base.

Overcome Barriers. To overcome high costs of creating successive generations of high performance computers, the program will emphasize scalable computer designs. Scalable computing and networking technologies enable exploratory use with small, lower cost prototype systems needed to eventually support the acquisition of larger systems. The development of user-friendly software and systems as part of the investment in the HPC software base is a major, integral part of the program. The central role of mission agencies, and the broad academic scientific research community involvement, ensures that the hardware, software, and networking technology developed will be responsive to user needs. A strategy of cooperative government, industry, and university activities is used to manage and coordinate the coupling of these sectors to achieve maximum synergy.

2. HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM COMPONENTS

The HPCC Program is composed of four integrated and coordinated components that are designed to enhance scientific productivity and support long term agency needs. The emerging scientific computing environment is that of advanced workstations with high resolution color displays connected to a high speed computer communications network and high performance computing resources. The regional and national networks provide a means to gain access to additional high performance computing systems and research resources. Realizing the full potential of these teraops computing and gigabits networking systems will require advanced software technology and algorithms and people educated and trained to use these tools and resources in this dynamic field.

High Performance Computing Systems (HPCS)

The HPCS component produces scalable parallel computing systems through the development of prototype systems. The program is designed to attack computational science problems by developing innovative systems that will provide a one-hundred to one-thousand-fold increase of sustained computing capability over machines that follow the more conventional design evolution path (See Figure 3). DARPA will coordinate the research and development effort that will produce teraops systems.


The program is structured to focus on technological challenges in the earliest stages of the product development cycle. Critical underlying technologies are developed in prototype form along with associated design tools. This allows empirical evaluation of alternative solutions as the prototype systems mature. Evaluation is performed throughout the development cycle, with experimental results being fed back into the design process to refine successive generations of systems. There is risk inherent in creating new technologies, and each project will be managed according to its proximity to commercial introduction. Larger projects which are close to yielding commercial products are performed on a cost shared basis with industry.

HPCS is composed of the four subcomponents shown below.

Research for Future Generations of Computing Systems. This activity produces the underlying component, packaging, and scaling concepts. These projects ensure that the required advanced technologies will be available for the new systems while providing a foundation for the more powerful systems that need to follow.

Figure 3

Computer System Performance Trends for Grand Challenge Problems – Power and Time to System Development –



System Design Tools. This activity develops computer aided design tools and frameworks for enabling multiple tools to work together to enable design, analysis, simulation, and testing of systems components. The tools will enable rapid prototyping of new system concepts.

Advanced Prototype Systems. This activity consists of focused development of experimental systems that are designed and developed on a cost shared basis with industry. The 100 gigaops systems provide a basis for the teraops systems. New models of computation will be introduced and successive generations of computer systems, along with systems software, will encompass broader grand challenge domains. Modular technologies will enable a wide variety of system configurations using common components. Systems with 100 gigaops sustained performance will be developed by the early 1990's. The teraops level will be reached by the mid 1990's.


Evaluation of Early Systems. Experimental systems will be placed at sites with high levels of expertise to provide feedback to systems architects and software designers. Performance evaluation criteria for systems and results of evaluation will be made widely available. Because of scalability, early systems can be acquired at smaller scales to evaluate their potential performance. As noted below, the ASTA component will support, on a cost shared basis, the acquisition of large scale systems for experimental use in grand challenge applications.

Advanced Software Technology and Algorithms (ASTA)

Dramatic improvements in algorithm design and software technology are essential to achieving sustained teraflops computing system performance. Improvements in hardware, especially for scalable parallel systems, must be matched with new and innovative algorithms and software to enable researchers to expand the boundaries of computational capabilities. In Figure 4, this point is illustrated for broad classes of scientific computing problems by showing that comparable improvements in throughput have resulted from advances in computational methods as from improved hardware technology. In this case, hardware improvements yielded a speed-up of 1000 over 20 years, while software and algorithms improvements yielded a speed-up of 3000 over the same period, for an aggregate speed-up of more than a million.

Figure 4

Performance Improvement for Scientific Computing Problems



The ASTA component of the Federal High Performance Computing and Communications Program has three goals:

Enable solution of grand challenge application problems in science and engineering.

Improve system user-friendliness, reliability, and software productivity.

Use experience gained on leading edge applications to help guide future software efforts.

The emphasis is on the development of advanced algorithms and software technology required to address applications problems on the scale of grand challenges. The ASTA component is comprised of four subcomponents, as shown below.

Software Support for Grand Challenges. The HPCC Program is designed to demonstrate new computing technology capabilities by confronting an expanding number of grand challenge application problems. The goal is to reduce the risks to researchers inherent in adopting innovative high performance computing technologies. Grand challenge application problems will be selected based upon their scientific importance, the potential for cost sharing with sources directly concerned with the specific scientific and engineering applications, and the potential for leveraging across sectors.

Software Components and Tools. The multidisciplinary research teams will have common needs in software technology and programming environments, advanced compiler technology, optimization and parallelization tools, interoperability and data management, visualization, debugging and analysis, and performance measurement. Advances in these generic software areas will minimize the need for specialized researchers to master advanced, complex computing science skills. Coordination of the development of advanced software technology and algorithms among the participating agencies will be important to ensure effective and efficient use of resources. A particular focus of this element will be to develop advanced software applications that exploit the NREN using distributed file systems and national software libraries.

Computational Techniques. The focus of the HPCC Program on scalable parallel computing systems dictates that significant advances in computational techniques will be needed. The design and theory of algorithms are as important as hardware or networking improvements in reaching teraflops computational performance. Research in computational techniques will include parallel algorithms, numerical and mathematical analyses, parallel languages, computational models, and program refinement techniques. Higher level languages will be developed to enable computational scientists to work directly at the level of the abstract computational problem being addressed and to more easily explore specific implementation approaches.

High Performance Computing Research Centers. High performance computing research centers and testbed facilities will provide a large number of researchers access via the NREN to both conventional and innovative supercomputing architectures. Many of these centers will introduce innovative architectures and will be the focus of technology transfer activities and form the base for training new generations of computer and computational scientists.

National Research and Education Network (NREN)

The NREN component of the HPCC Program dramatically expands and enhances the U.S. portion of an existing worldwide infrastructure of interconnected computer networks called the Internet. A substantial fraction of the domestic Internet is supported and loosely coordinated by Federal agencies, principally DARPA, DOE, NASA, and NSF.


Collaboration among scientists is an important and integral facet of the U.S. research environment. It can be greatly improved by increasing the level of network connectivity and by introducing new capabilities into the existing infrastructure. The NREN design will not only address broad network connectivity, but will also provide the basis for necessary higher level capabilities and services.

Many educational institutions, government laboratories, and industrial research facilities are currently connected to the Internet. Yet, it still falls short of a widespread, uniform, and high performance national infrastructure. In order to satisfy the HPCC Program goals, the NREN must not merely provide network access to research and educational institutions at all levels and locations, it must also deliver new capabilities. Some of these, such as distance learning, may initially be extensions of current technology. All capabilities will benefit from, and many will be enabled by, a program of research into very high speed technology. This technology is needed to support access to digital libraries, large scale distributed computing resources, as well as to perform computationally intensive applications that require real time visualization of modeling and simulation results, rapid interrogation and retrieval of scientific data from specialized data bases, remote control of experiments and simulations, and teleconferencing.

In addition to serving the needs of the scientific and research communities, the NREN will provide valuable experience necessary for the successful development of a broader, privately-operated national information infrastructure. Such an infrastructure would allow consumers, businesses, and schools and government at all levels to share quality information and entertainment when and where they want it at a reasonable cost.

Applications conducted over a computer network vary in their flow of information from steady, as in a bulk file transfer, to "bursty," as in human-computer interaction via keyboard. Similarly, some applications can be carried out at relatively low communication rates, while others by their nature require high speeds. A number of scientific networking applications are characterized on the graph of speed versus burstiness in Figure 5. Traffic seen in the early days of networks appears near the bottom of the chart. More advanced applications are furthest from the

origin and require more sophisticated protocols and network capabilities. This chart illustrates that a gigabit network is needed not only to carry the aggregation of low speed traffic, but also to accommodate high speed uses.



The vision of the NREN is of an interconnection of the nation's educational infrastructure to its knowledge and information centers. In this system, elementary schools, high schools, two and four year colleges, and universities will be linked with research centers and

laboratories so that all may share access to: libraries, databases, and diverse scientific instruments such as supercomputers, telescopes, and particle accelerators. The NREN enables communication and fosters collaboration among and within these communities. By reducing the traditional impediments of geographical isolation, the NREN improves the quality and raises the level of education nationwide. The NREN contributes to the success of the Basic Research and Human Resources component of the High Performance Computing and Communications Program. By making unique scientific and informational resources accessible beyond their physical locations, it permits widespread participation in the HPCC Program by scientists, university researchers, and students, and it enables the development of large scale distributed computing resources.

Interagency Interim NREN. The NSF will coordinate the Interim NREN activities by upgrading its backbone network, by assisting regional networks to upgrade facilities, capacity, and bandwidth, and by interconnecting the backbone networks of other agencies. A significant effort in its implementation will be the development and deployment of safeguards to enhance security and control over access to the connected computer systems, and that of the network components themselves in order to minimize vulnerability to careless or malicious attack. Coordination among participating agencies and the non Federal networking community will be expanded through the creation of a National Networking Council.

Gigabit Research and Development. DARPA will coordinate the research and development effort that will culminate in initial deployment of gigabits per second capability. Coordination of research efforts on very high speed switches, protocols, and computer interfaces will be necessary.

Basic Research and Human Resources (BRHR)

This component addresses long term national needs for more skilled personnel, enhancement of education and training, and materials and curriculum development in the high performance computing science and engineering areas. The NREN and ASTA components include support for research in the large scale project environment. The BRHR component is designed to encourage investigator initiated, long term research on experimental projects that will maintain the flow of innovative ideas and talented people into high performance computing areas.

Drawing the best and brightest of our Nation's youth into scientific and technological careers is a formidable challenge that will have profound effects on overall U.S. scientific and technological competitiveness. This component of the program will establish industry, university, and government partnerships to improve the training and utilization of personnel and to expand the base of research and development personnel in high performance computing science, technology, and applications.

The BRHR component of the HPCC Program has five goals:

- Improve the flow of human resources into high performance computing.
- Improve the university infrastructure to stay at the leading edge.
- Expand collaboration and resource sharing among the Federal, academic, and industrial sectors.
- Facilitate multidisciplinary research on high performance computing and communications.
- Create a critical mass of users by building a base with common systems, tools, and interfaces.

The BRHR component is organized across the four sub-components shown below.

Basic Research. These activities support increased participation of individual investigators in the conduct of innovative multidisciplinary research in computer science, computer engineering, and computational science and engineering related to high performance computing. The strategy is to increase the number of multi-disciplinary awards across all disciplines where computational methods are critical to achieving advances or scientific breakthroughs. Program activities include: research on scientific algorithm development for highly parallel computers;

generic computational algorithm development; scaling techniques; prediction techniques for concurrent systems; resource management strategies for highly parallel distributed systems; fault tolerant strategies for parallel and distributed systems; and research on heterogeneous software environments.

Research Participation and Training. These activities address the human resources pipeline in the computer and computational sciences, at post-doctoral (training and re-training), graduate, and undergraduate levels. Program activities include: workshops and seminars; post-doctoral fellowships in computational science and engineering; career training in medical informatics through grants to young investigators; institutional training and post-doctoral programs; knowledge transfer exchange programs at national laboratories, centers, universities, and industry; and software dissemination through national databases and libraries.

Infrastructure. These activities will improve university and government facilities for computer science, computer engineering, and computational science and engineering research related to high performance computing. Program activities include: improvement of equipment in computer science, computer engineering, and computational science and engineering academic departments, centers and institutions; development of scientific databases; and distribution of integrated system building kits and toolsets.

Education, Training and Curriculum. These activities will expand and initiate activities to improve undergraduate and pre-college education and training opportunities in high performance computing and computational science and engineering for both students and educators. The introduction of associated curriculum and training materials at all levels is an integral part of this effort. Program activities include: bringing people to national centers for training, technology transfer, and educational experiences; using professional scientists and engineers to provide curriculum development materials and instruction for high school students in the context of high school supercomputer programs, supercomputer user workshops, summer institutes, and career development informatics for health sciences.

3. PROGRAM DEVELOPMENT AND AGENCY BUDGETS

Program Planning

Leadership for the HPCC Program is provided by the Office of Science and Technology Policy (OSTP), through the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) Committee on Physical, Mathematical, and Engineering Sciences (PMES). The membership of the PMES includes senior executives of many Federal agencies.

The planning process for the HPCC Program was coordinated by a PMES working group through information exchange, the common development of interagency initiatives, and the review of individual agency HPCC proposals and budgets.

Evaluation Criteria

Each participating agency HPCC contribution was reviewed against formal evaluation criteria during the planning and budget process. A review of participating agencies was performed using these evaluation criteria to develop the FY 1992 agency requests for the Program. The evaluation criteria are shown in Figure 6.

Agency Budgets

Over the last three years of the planning process of this initiative, the participating agencies have mutually adjusted their activities within the base to achieve greater efficiency in addressing the goals of the HPCC Program. For FY 1992, \$638 million is being proposed, a \$149 million or 30 percent increase over the FY 1991 enacted level. The budget is shown in Figure 7.

The funds proposed from Federal sources are not intended to carry out the entire HPCC Program. Portions of this program will be cost shared by organizations from the participating sectors. The funding estimates are based on analyses of the practical experience of prior computing and computer networking programs. These estimates were then reviewed as outlined above. Cost sharing will occur with various U.S. industrial and university partners to a large extent in the HPCC and the NREN components. Cost sharing will occur in the ASTA component with agency programs and other computational applications programs, for example, in specific grand challenge areas, via multidisciplinary collaborations. This component also includes deployment of high performance systems to HPCC research centers. The close coupling of

support in the Program will result in significant leverage, accelerate technology transfer, stimulate U.S. industry and markets, and enable the solution of computationally intensive applications. In addition, although HPCC is not intended to include classified programs, the technology produced will have an important impact in these national security areas. Figure 7 illustrates the relative levels of investment in the four program components.

Figure 6

Evaluation Criteria for the Federal High Performance Computing and Communications Program

Relevance/Contribution. The research must significantly contribute to the overall goals and strategy of the Federal High Performance Computing and Communications (HPCC) Program, including computing, software, networking, and basic research, to enable solution of the grand challenges.

Technical/Scientific Merit. The proposed agency program must be technically/scientifically sound and of high quality, and must be the product of a documented technical/scientific planning and review process.

Readiness. A clear agency planning process must be evident, and the organization must have demonstrated capability to carry out the program.

Timeliness. The proposed work must be technically/scientifically timely for one or more of the HPCC program components.

Linkages. The responsible organization must have established policies, programs, and activities promoting effective technical and scientific connections among government, industry, and academic sectors.

Costs. The identified resources must be adequate, represent an appropriate share of the total available HPCC resources (e.g., a balance among program components), promote prospects for joint funding, and address long-term resource implications.


Enhancements to Existing Program Research. Existing agency HPCC programs will receive adequate support before new initiatives are funded.

Agency Approval. The proposed program or activity must have policy-level approval by the submitting agency.

Figure 7

**High Performance Computing and Communications
Budgets by Agency and Program Component
(Dollars in Millions)**

Agency	Base FY 1991	HPCC FY 1992	FY 1992 HPCC Component			
			HPCS	ASTA	NREN	BRHR
DARPA	183.0	232.2	103.3	38.5	32.9	57.5
DOE	65.0	93.0	15.0	58.0	12.0	8.0
NASA	54.0	72.4	14.2	49.8	7.4	1.0
NSF	169.0	213.0	24.0	103.0	32.7	53.3
DOC/NIST	2.1	2.9	0.3	0.6	2.0	0.0
DOC/NOAA	1.4	2.5	0.0	1.8	0.7	0.0
EPA	1.4	5.2	0.0	4.5	0.0	0.7
NIH/NLM	13.5	17.1	0.0	8.9	4.2	4.0
Total	489.4	638.3	156.8	265.1	91.9	124.5



HPCC — High Performance Computing and Communications

Components of HPCC:

- HPCS** — High Performance Computing Systems
- ASTA** — Advanced Software Technology and Algorithms
- NREN** — National Research and Education Network
- BRHR** — Basic Research and Human Resources

Agency Program Descriptions

The agency responsibilities under the Federal HPCC Program are outlined in Figure 8. Several agencies have been assigned a coordinating responsibility in specific technical areas:

- DOE and NASA will coordinate activities in HPC system evaluation, testbed development, and applications software capabilities.
- NASA will coordinate the accumulation of and access to the HPC software base. This will be facilitated by a wide area file system technology that is currently being deployed for early experimental use by DARPA which will be extended to include the NREN as it matures and is deployed by NSF.
- DARPA will coordinate activities in the development of scalable parallel HPC systems, including their basic units of replication, system modules, and the necessary associated systems software. DARPA will also coordinate activities in gigabit network technology research.
- NSF will coordinate activities for the broad deployment of the National Research and Education Network working with all agencies with mission specific requirements.
- NSF will coordinate activities in basic research and human resource development, while each agency retains its role as required to accomplish their own missions.
- NIST will coordinate activities in HPC system instrumentation, evaluation, and in standards issues.

Each agency participates in all of the identified activities to ensure that the resulting capabilities are a good match to user needs. DOC/NOAA, HHS/NIH/NLM, and EPA bring distinctive applications areas of broad interest and network user bases.

Figure 8

HPCC Program: Agency Responsibilities: HPCS, ASTA		
Activity Agency	High Performance Computing Systems	Advanced Software Technology & Algorithms
DARPA	· Technology development and coordination for Teraops systems	· Technology development for parallel algorithms and systems software
DOE	· Technology development · Systems evaluation	· Energy applications centers · Energy computation research · Software tools
NASA	· Aeronautics and space application testbeds · Systems development for space flight	· Software coordination · Research in: · Aerospace computations · Information management
NSF	· Basic architecture research	· Basic research in: · Software tools, databases · Grand challenges · Computer access
DOC/ NIST		
DOC/ NOAA		· Ocean and atmospheric computation research · Software tools
EPA		· Research in environmental computations, databases, and application testbeds
NIH/ NLM		· Medical application testbeds · Software tools · Medical computation research

HPCC Program: Agency Responsibilities: NREN, BRHR		
National Research and Education Network	Basic Research and Human Resources	Activity / Agency
· Technology development and coordination for gigabits networks	· University programs	DARPA
· Energy deployment mission facilities · Gigabits applications research	· Basic research and education programs	DOE
· Aerospace deployment mission facilities	· Research initiation and university block grants	NASA
· Facilities coordination and deployment · Gigabits research	· Programs in: · Basic research · Education · Training / curricula	NSF
· High speed network research and standards		DOC/ NIST
· Ocean and atmospheric mission facilities		DOC/ NOAA
· States environmental mission assimilation	· Technology transfer to States · University programs	EPA
· Medical mission facilities	· University programs	NIH/ NLM

Defense Advanced Research Projects Agency (DARPA)

DARPA, as the DOD lead agency for advanced technology research, will focus on developing the high performance computing and networking technologies needed for the Defense and overall HPCC Programs. DARPA programs have produced both the computing and networking foundation for the HPCC Program, including the first generation of scalable parallel computing systems and large scale computer networks, and the associated system software and supporting technologies. DARPA has worked with industry to pioneer the application of these new technologies within Defense and on a cooperative basis with other Federal agencies.

The DARPA HPCC Program builds upon the DARPA Strategic Computing Program. As the HPCC Program builds up, the Strategic Computing Program integrates its results with Defense specific needs such as embedded systems, accelerators of specific problem domains, and grand challenges problems related to defense. DARPA will continue this mode of executing the HPCC Program, cooperating with various defense organizations, and working closely with AFOSR, ARO, ONR, Defense Service Laboratories, NSA, other Defense organizations, and other Federal agencies as appropriate.

High Performance Computing Systems will be produced in the four main subareas: Research for Future Generations of Computing Systems, System Design Tools, Advanced Prototype Systems, and Evaluation of Early Systems. Systems capable of sustaining 100 gigaops for large problems will be available for deployment by late 1993 and the teraops systems will be available by 1996.

In *Advanced Software Technology and Algorithms*, DARPA projects will produce scalable libraries for Defense problem domains and programming and analysis tools for scalable parallel and distributed heterogeneous systems in a workstation/server configuration that will be open to the integration of embedded systems and accelerators.

For the *National Research and Education Network* component, high performance networking technologies will be produced to satisfy the gigabit technology needs of the NREN and to provide a dual use technology base for Defense. This networking technology includes development of new protocols and switch and transmission technologies, and it will be capable of supporting a wide range of advanced network services.

The *Basic Research and Human Resources* component will focus on fundamental scientific issues in these three technology development areas in cooperation with other basic research programs and provide for smaller individual investigator projects to complement the larger projects. In addition, the relevant and related basic research will also be integrated into the larger projects as it matures.

Department of Energy (DOE)

The Department of Energy will participate in all components of the High Performance Computing and Communications Program. In the area of *High Performance Computing Systems*, the DOE will be an early customer of small versions of systems with advanced architectures and will evaluate these systems on energy related applications. The DOE will consider cooperative development of advanced systems between its national laboratories and vendors, especially integration of very high speed computer and networking hardware with software systems. The DOE will support research and development on algorithms and systems software for parallel computing systems.

The *Advanced Software Technology and Algorithms* effort will include research and development of: parallel algorithms for grand challenge applications, software and tools for early prototypes of 100 gigaflops and teraflops systems, prototype computational science programming environments that meet standards and are transportable, and support for high performance computing research centers to facilitate the transition from research on parallel machines into the applications and the programming environments. The DOE will fund several grand challenge collaborations, initiate a software component and tools program with strong industrial participation, and initiate an applications driven computational research program. The DOE will evaluate proposals and make research awards related to grand challenges in global climate change, molecular biology, human genome research, materials and chemical sciences, combustion research, waste remediation, fusion energy, and other areas within its mission.

The DOE will participate in the cooperative interagency *National Research and Education Network*. The Energy Sciences Net (ESNet) will be incorporated into the NREN and will provide quality network access to the energy research facilities. ESNet will maintain compatibility and will be upgraded in concert with NREN. Gigabit network support technology will be developed for DOE applications distributed across multiple energy research centers at the national laboratories and universities.

The DOE's *Basic Research and Human Resources* activities will include: stimulating research in computational science, expanding training programs at the national laboratories for high school teachers and college students in computing techniques, initiation of a high school supercomputer access program, and provision of fellowships in computational science with internship at national laboratories.

National Aeronautics and Space Administration (NASA)

NASA's HPCC Program participates in all four components of the Federal Program, through a vertically integrated program focused on NASA's grand challenges in: Computational Aerosciences, Earth and Space Sciences, and Remote Exploration and Experimentation.

The goal of NASA's program is to accelerate the development and application of high performance computing technologies to meet NASA's science and engineering requirements. In cooperation with the other Federal agencies, NASA's program will deploy teraflops computer capabilities essential for computational design of integrated aerospace vehicle systems and for predicting long term global change, and will enable the development of massively parallel techniques for spaceborne applications.

NASA's program is focused on bringing together interdisciplinary teams of computer scientists and computational physicists to develop these technologies within three vertically integrated projects that are unique to NASA's missions. These technologies include applications algorithms and programs, systems software, peripherals, networking, and the actual high performance computing hardware. NASA will develop a suite of software tools to enhance productivity. These include: load balancing tools, run time optimizers, monitors, parallelization tools, as well as data management and visualization tools.

NASA's role includes coordinating the *Advanced Software Technology and Algorithm* component for the Federal program; acquiring experimental hardware for testbeds in computational aerosciences, earth and space systems sciences, and remote exploration and experimentation; and supporting the development of the National Research and Education Network. To encourage vigorous research into the underlying theory and concepts of high performance computing, NASA will foster interactions among academia, industry, and national laboratories and will strengthen the basic research in high performance computing at the NASA centers and research institutes and in universities.

NASA's considerable expertise in experimental parallel computer testbeds and small, scalable testbed systems will be used to demonstrate high performance computing technologies as a step toward full-scale computational capabilities. A key to successful exploitation of massively parallel computing power will be the blending of application-driven and architecture-driven computer systems to most effectively meet NASA's needs.

National Science Foundation (NSF)

The NSF HPCC Program impacts the activities of all science and engineering disciplines by providing computing and networking infrastructure support and by developing enabling technologies for advanced computing and communications platforms and paradigms.

In the area of *High Performance Computing Systems*, research will be initiated on new architectures and systems optimized for specific research applications. New tools for systems level automated design and component packaging will be supported to permit the design of application-specific devices and systems.

In *Advanced Software Technology and Algorithms*, research will be initiated on scientific database technology and implementation of prototype networked databases and associated software. Advanced applications tools for research computing environments will be supported to enable nationwide access to the full complement of parallel machines for research on grand challenge problems. Areas of research focus will include numeric and symbolic computing, algorithm development, optimization of applications software for new parallel computers, scientific visualization, automated programming tools, and new methods of scientific and technical information exchanges.

NSF coordination of the *National Research and Education Network* activities will accelerate the harmonizing of multiple agency networks and protocols into a single NREN. The number of nodes will be increased to expand distributed information resources, and to increase redundancy, capacity, control, and security. Mid-level nets will be assisted to upgrade facilities and service very high bandwidth requirements. Support will be provided for research on new protocols for gigabit networks, switch and transmission technology, routing, congestion, and flow control. The exploration of pricing mechanisms for network services and network applications and structured transition to commercial service will be initiated.

In *Basic Research and Human Resources*, support will include multidisciplinary research and university infrastructure for computer science, computer engineering, and computational science and engineering. To increase the human resource pool in computing hardware and software systems areas, support will be expanded for graduate and post-doctoral positions. Curriculum improvement, teacher training, and support for centers to provide education and training in the use of experimental and parallel supercomputers are integral parts of this component.

National Institute for Standards and Technology (DOC/NIST)

The NIST research program is directed toward developing performance monitoring tools and promoting "open systems" software. NIST has proposed to augment its current HPC research program by promoting the commercialization of protocol and security mechanisms for medium speed networks.

The goal of NIST's activities in high performance computing is to develop hardware performance monitoring tools, promote "open system" software, and support a classification system for indexing and distributing scientific software so that industry and the research community can effectively exploit the power of future generations of high performance computers. In support of *The National Research and Education Network*, NIST will develop and speed commercialization of network protocols and security mechanisms that can achieve the desired gigabit speeds on future versions of the network. NIST will participate in the HPCC Program by:

- Developing hardware monitoring methods leading to load characterization and performance measurement techniques for ultra-high-speed systems which will be made available on a publicly accessible database.

- Conducting research on new protocols and related security primitives necessary to sustain gigabit network speeds and standards to provide interoperability, common user interfaces to systems, and enhanced security for the network.

- Establishing a network testbed at NIST instrumented to collect performance data and test new network protocols, management routines, and security mechanisms for gigabit networks.

National Oceanic and Atmospheric Administration (DOC/NOAA)

The National Oceanic and Atmospheric Administration operational and research programs are directed toward weather prediction, ocean sciences, the Climate and Global Change Program, and the Coastal Oceans Program, together with data management activities for all agency programs. The HPCC Program will allow extensive development of new forecast models, studies in computational fluid dynamics, and the incorporation of evolving computer architectures and networks into the systems that carry out agency missions.

The NOAA High Performance Computing Program is focused in two components:

Advanced Software Technology and Algorithms. This component provides support for: grand challenges in atmospheric and oceanic sciences; development of advanced numerical models to simulate the general circulation of the oceans and atmosphere in support of NOAA missions and the activities of collaborating research groups; development of new computational methods for solving atmospheric, oceanic, and related problems on new computer architectures; data management R&D; support for basic research in strategies, techniques, and tools required for the management and analysis of large-scale distributed scientific databases and distributed data handling, including quality control; and development of algorithms for massively parallel processors, together with their standardized software component libraries and tools for the solution of oceanic and atmospheric analysis and forecasting problems.

NOAA will acquire, install, and operate advanced computational facilities for the evaluation of near operational prototype computers having massively parallel architectures and will develop algorithms appropriate for these architectures.

The National Research and Education Network component provides networking in support for NOAA's climate and global change research community and a wide range of agency missions in oceanography, weather prediction, and environmental sciences research. NOAA will evaluate advanced network protocols and hardware technologies related to its missions.

Environmental Protection Agency (EPA)

The EPA research program is directed toward the advancement and dissemination of computational techniques and software tools which form the core of ecosystem, atmospheric chemistry and dynamics models. The models extend the computational capability of environmental assessment tools to handle multi-pollutant interactions and optimization of control strategies.

The *Advanced Software Technology and Algorithms* component will develop new approaches for coupling numerical equation solvers to the high performance features of emerging computer architectures including specific techniques for solution of partial differential equations on massively parallel architectures. These solvers will be designed to operate on a flexible, generic grid system decoupled from specific applications to provide a testbed for evaluation and optimization. The resulting solution framework can serve as the computational engine for a variety of interdisciplinary applications.

Advancement of scientific visualization techniques can provide exceptional benefits in increasing the pace at which environmental problems are solved. EPA's program includes basic research to improve user interfaces and computational algorithms related to image rendering and manipulation of geometric images. The main focus is on simultaneous representations of multiple data sets and real-time visualization in a heterogeneous distributed processing environment. Related research on storage and access techniques for massive environmental data bases across multiple architectures will be conducted.

The *Basic Research and Human Resources* component includes the establishment of technology transfer centers to propagate the use of HPCC technology to State and local environmental groups, and Federal managers for optimization of pollution control and prevention strategies. A main goal of a Technology Transfer Center is to provide non-sophisticated users with training and guidance in the application of a variety high performance computing tools to solve important environmental challenges. The program also supports cross-discipline career training through universities and other institutions to ensure a continuing base of technical professionals knowledgeable in the use of high performance computing technology for environmental problem solving.

National Institutes of Health / National Library of Medicine (HHS/NIH/NLM)

The HHS/NIH program includes molecular biology computing, creation, and transmission of digital electronic images, the linking of academic health centers via computer networks, the creation of advanced methods to retrieve information from life sciences databases, and training in biomedical computer sciences. The HPCC program will complement the Human Genome Project by providing new methods for computer based analysis of normal and disease genes.

The *Advanced Software Technology and Algorithms* component will develop advanced software technology and algorithms in two areas of importance in biomedical research and education: biotechnology computing and digital images. In the area of biotechnology computing, the program will support development of molecular sequence comparison algorithms, new database methods, and algorithms to predict biological structure and function from genetic code. The biomedical images area will support new methods for representing, linking, and rendering images of biological structure

The *National Research and Education Network* component has two sub-components: connections among academic medical centers and their growing array of computerized information sources; and development of intelligent gateways that link conceptually related databanks in the life sciences. The academic medical centers of the country are confronted with a growing array of disconnected computer-based information sources, ranging from patient records, x-rays and laboratory systems, to basic research tools such as protein and DNA sequence databanks. Development of advanced software systems capable of representing and linking these dissimilar data types, and communicating them among centers for research and health care, is the goal of this part of the program. The focus of the Gateways HPCC program is the building of systems that are capable of translating a user's request into multiple computer-based vocabularies, selecting appropriate databases from a wide range of widely-available resources, and retrieving information in a manner that does not require users to understand the structure and syntax of the systems being queried.

The *Basic Research and Human Resources* component addresses the need to train biomedical researchers and health care providers in the use of advanced computing and network communications to aid in their work. The successful predoctoral and postdoctoral grants program for career training in medical informatics administered by the National Institutes of Health will be expanded.

4. GRAND CHALLENGE AND SUPPORTING TECHNOLOGY CASE STUDIES


This section incorporates examples of high performance computing and computer communications technologies and several illustrative grand challenge applications in computational science and engineering, that are presented on the cover of this report. The examples were chosen to illustrate the diversity and significance of application areas that have been addressed to date. The list of grand challenges is too long to allow an example from each area, thus lack of representation of certain areas does not imply lack of importance. A more complete description of the grand challenges can be found in the Federal High Performance Computing Program report issued in September 1989.

Forecasting Severe Weather Events

Making accurate predictions about the behavior of the atmosphere is of critical importance to reducing losses due to storms, hurricanes, floods, and other weather related phenomena. Current operational small scale weather prediction models have been constrained by two broad deficiencies: inadequate size of the data sets necessary to define the detailed atmospheric structures and insufficient computer resources to support high resolution. During the next five years, major progress will be made on the first deficiency by using a combination of new ground-based and remote-observing systems. This leaves inadequate computer power as the primary stumbling block preventing high resolution models from playing operational roles in our national weather prediction efforts.

The high-resolution, operational weather prediction models of the future will represent a new generation of numerical formulations. The primary differences will be in the treatment of vertical motions and small scale processes which in the past have been considered to be "sub-grid scale." This means that physically significant events, such as convection, must be treated as parameterized processes and fine scale details in, for example, thunderstorm evolution that can affect the surroundings cannot be directly addressed by the model. Likewise, fine scale observations of local importance, such as moisture gradients around wet areas, which can provide local forcing, currently cannot be incorporated adequately.

To meet this scientific and national need, new technologies need to be developed and incorporated into existing facilities so that a less than 5 km resolution model can be operational before the end of the century. This would allow updated forecasts on a 6-hourly or shorter basis. Such models have already shown significant advances in the accuracy of predicting a wide variety of weather events, from severe thunderstorms to lake effect snowfalls in research applications. Each reduction of the model resolution by one-half requires an increase of computer power of almost an order of magnitude, as well as comparable increases in supporting memory, mass storage and networks. For example, a 5 km model could require up to a 20 teraflops computer system to meet operational schedules using this model.



Researchers have used supercomputers to model thunderstorms numerically. Wind, temperature, pressure, and other variables are calculated every few seconds at several hundred thousand locations in the area of the developing storm. Mathematical equations are then used to simulate the storm's evolution. In this graphic, particles which are released near the ground at regular intervals are colored orange when rising and blue when sinking. Yellow signifies paths of individual particles.

Cancer Genes

DNA is the blueprint of life, the molecular thread in the nucleus of each of our cells which guides the assembly of molecules and complete living organisms. When the DNA code is altered by mutation, serious diseases can result, such as cancer; this phenomenon was known to scientists studying animal tumors in the 1970's. They isolated cancer-causing genes, called "oncogenes" from animal tumors, and later found that similar genes existed in normal human DNA. This was a profound mystery. Why would we be carrying the seeds of our own destruction in our genetic blueprint?

In 1984, two separate research groups used a computerized searching algorithm to compare a newly discovered oncogene to all known genes. To their astonishment, the cancer causing gene matched a normal gene involved in growth and development. Suddenly, it became clear that cancer might be caused by a normal growth gene being switched on at the wrong time. This fundamental and unexpected insight was an early example of a field that is now known as Computational Biology, the science of using computers to store and analyze data from complex molecules in living cells.

The databases used in 1984 for these comparisons contained information about several thousand molecular units; now they contain over 30 million. Moreover, the current multi-agency genome research programs of NIH, NSF, and DOE will acquire data on tens of billions of molecular units, ranging from simple organisms to human beings. The best computer algorithms for determining the similarity of molecules require time proportional to the length of the molecules being compared; if the methods used to analyze oncogenes in 1984 were applied to the three billion base pairs of the human genome, they would require hundreds of years of computer time on today's fastest supercomputers. New computer designs and software methods will be essential to cope with the explosive growth of molecular data. Functioning as intellectual amplifiers to detect similarities and differences in molecules whose size and complexity are too vast for the unaided human mind, the computer systems to be developed by the HPCC program will be a critical tool for the life sciences in the 1990's, and the health care systems of the 21st Century.




Two views are shown for the binding of a chemical carcinogen to a short stretch of DNA. The stick model on the left shows the overall structure of the DNA–carcinogen complex. The space filling model on the right shows the actual intimate binding of the carcinogen, hiding deep within the DNA. This chemical contact leads to mutations in the DNA code, and ultimately to a tumor.

Predicting New Superconductors

In 1988, the world was excited by the discovery that a particular yttrium–barium–copper–oxide compound superconducts at a temperature of 93° Kelvin, still very cold, but much warmer than any previous superconductor. This discovery sparked a worldwide effort to expand research to discover new superconducting materials. The economic benefit of a high–temperature superconductor is beyond calculation, portending the development of, for example, much more efficient power transmission and lightweight, powerful magnets to revolutionize the transportation business. Advanced computing is a central part of the arsenal of research tools which will be necessary to reach that payoff.

Despite these early successes, many questions remain before it is understood how some materials superconduct when very similar compounds do not. This understanding will be critical to predicting new superconductors, which might work at even higher temperatures, be less expensive, carry more current, or be more amenable to manufacturing processes. Increasing progress in all these areas is required before the impact of these new materials will be felt.

The solution of physical models requires intensive calculations to understand the material structure. High performance computing can shorten the discovery process by allowing the development of accurate simulations to point experimenters in the most promising directions. For example, most of the groups looking for new superconductors are trying various copper–oxide combinations. Researchers are using high performance computers to explore the possibility of various combinations of elements that may lead to new superconducting materials.



Researchers have used visualization techniques to investigate the structure of materials which are thought to be involved in the superconducting mechanism. In this graphic, barium cations (green) and yttrium cations (silver) are shown with their associated oxygen defects. The copper atoms (blue) and oxygen atoms (red) are found in two types of CuO_4 environments, one of which is depicted in yellow, the other in light blue.

Air Pollution

The ability of the atmosphere to absorb and to cleanse itself of pollutant contaminants was taken for granted until the 1950's when "killer fogs" in London, England and Donora, Pennsylvania caused the deaths of hundreds of people. Since then, technological advances in controlling source emissions have reversed the trend of steadily increasing pollution, even in the face of continued industrial growth. However, reduction strategies of individual pollutant types do not always produce the desired results. In fact, these simple solutions can even make air quality worse due to complex chemical interactions of the remaining airborne contaminants. Pollutants may travel long distances from industrial centers to sensitive areas where they are deposited in transformed products such as ozone and acid rain. These complexities of pollutant transport and transformation are costly and difficult to study experimentally, therefore, numerical models of the atmosphere have been developed to assess the effects of man-made emissions on air quality.

The new Clean Air Act mandates the use of numerical models to demonstrate the effectiveness of proposed regulatory control strategies. The potential cost to society to implement these proposed controls is estimated to reach tens of billions of dollars. Current models have been useful in evaluating alternative control strategies, but do not yet have the capability to produce optimum solutions. Present computing limitations on existing supercomputers force simplifications in the scientific descriptions of chemical and physical processes, and slow examination of alternatives. Control strategies for each pollutant are often determined independently with little evaluation of multiple pollutant interactions. Remedial solutions determined for a particular scale of space and time are not readily extendable to other scales of pollutant dispersal.

High performance computing will enable multi-scale numerical explorations with cross-pollutant interactions to be performed in a timely manner so as to be useful to legislated control and prevention requirements. Advanced computer designs and software methods will also enable cost optimization of pollutant control tradeoffs. Improved visualization techniques will enhance the interpretation and evaluation of massive amounts of environmental measurement and computer simulation data. High performance computing models will lead to a better understanding of the actions needed to minimize pollutant damage to materials and environmental damage to crops while making our air safer to breathe for future generations.



Using a supercomputer, atmospheric researchers have simulated the transport, chemical transformation, and surface deposition of sulfur compounds responsible for acid rain. Visualization specialists have depicted the movement of sulfur compounds from major sources in the Ohio river valley to sensitive lakes in the Adirondacks. The yellow cloud represents high concentrations of sulfur compounds. Several sensitive aquatic regions are outlined in green and a typical wind flow pattern is presented in red.

Aerospace Vehicle Design

Being able to predict the aerodynamic characteristics of in-flight vehicles is important for designers. Reproducing such flight regimes in wind tunnels is time consuming, costly, and in some cases impossible. Computational aerodynamics simulations are less costly and much faster than complex wind tunnel tests and are able to simulate many inaccessible flight regimes. These capabilities will be particularly important in the design of supersonic and hypersonic aircraft to serve international markets.

Computational aerosciences directly contributes to maintaining U.S. preeminence in aerospace science and engineering disciplines. The computational technology developed in such computational aerodynamics problems will directly transfer to the U.S. aerospace industry and aircraft manufacturers. Other potential beneficiaries are in diverse fields where fluid dynamics is an important design aspect such as automobile manufacture, ship design, and medicine (e.g., heart/cardiovascular flow simulation).

Massively parallel computing systems and advanced parallel software technology and algorithms will enable the development and validation of multidisciplinary, coupled models. These models will allow the numerical simulation and design optimization of complete aerospace vehicle systems throughout the flight envelope.

Computation

Wind Tunnel



Shown here is a comparison between supercomputer simulation data and actual wind tunnel model data for the pressure distribution along the integrated space shuttle, solid rocket boosters and external tank flying at Mach 1.55. Note the excellent agreement between the two.


Energy Conservation and Turbulent Combustion

For the foreseeable future, 90% of the energy needs of the United States will be met by the combustion of fossil fuels. Two of the largest uses are for electrical power generation and in automobiles. Computational models offer a means of improving the design and efficiency of internal combustion engines.

Automobile engines are most efficient when run at high temperatures, but increased temperatures also lead to increased nitrogen oxide emissions. The burning of alternative fuels such as methanol is complicated by the emission of carcinogens. The effects of the emissions are influenced by local climatic conditions, making it necessary to consider the total system of fuel, engine, and atmosphere when seeking better designs. Our environment is too delicate and cannot be used as a testbed, so the atmospheric effects must be simulated.

A full three dimensional engine design code has been developed and implemented on a supercomputer. The code is designed to handle the most complex engines, such as the stratified charge and the two-stroke engines. The code represents an approximation to reality, because not all of the physical phenomena are well understood. Moreover, even if they were, the limited capabilities of existing computers would not allow this detailed information to be included.

For example, over 400 chemical processes of hydrocarbon and nitrogen chemistry are known to occur in internal combustion engines, yet only ten or less of the most significant reactions are used in simulations, in order to allow the calculation to run in a few hours on today's large supercomputers. Since the 400 hydrocarbon-nitrogen reactions are known, the real problem could be addressed with better algorithms running on a machine 10,000 times more powerful, a teraflops machine. This computational technology is needed by many private industrial engine firms as well as universities and government laboratories.




Researchers have used high performance computers to accomplish the numerical simulation of the properties of a conical fuel jet. The colors of the particles within the jet indicate the droplet size. The smallest particles are blue, intermediate size particles are shown in yellow, while the particles with the largest diameters are depicted in red. The smaller light blue particles around the jet represent the swirling air surrounding the spray.

Microsystems Design and Packaging

Since the invention of the integrated circuit in 1958, the number of transistors fabricated on a microchip has doubled every two years, providing a medium to incorporate ever-increasing complex electronic design into chips, components, and packages. By analogy, the complexity of detail incorporated in a single integrated circuit 1 cm square is equivalent to representing the map features, at a city block level, of the entire Eastern United States. A similar analogy for a 5 cm square module densely packed with a collection of chips would be the equivalent of a map of North and South America. Determining the interconnecting paths, selecting the right modules, testing the interfaces, and choosing the mix of technologies are part of the design process to build the scalable components for workstations to supercomputers.

In the computing world, scalable architectures based on the 1–2 million transistor custom structures of today will evolve in the mid 1990's to system approaches exploiting 10-million-transistor chips, standard component parts, and special interface electronics, all combined in optimally designed modules adhering to standard interfaces. System clock speeds will continue to improve, and the diversity of microchip technologies combined in a single module will allow unprecedented flexibility for designers.

Managing this complexity explosion would be overwhelming without the use of computationally based approaches that enable teams of designers to systematically reduce the time to develop such systems. Today, computational tools enable complex microcircuits to be developed on first pass, at the same time that packaging technologies such as multichip modules can be demonstrated. High performance computing applied to the technology design process will enable the exploration of design alternatives and rapid exploitation of new technologies.




The figure shows both current and future components of the innovation design cycle. In the lower left of the photograph is a multiproject wafer using an industry/academia/government supported prototyping capability. The multiproject wafer shown is 4 inches in diameter, shares resources with 82 projects, and represents approximately 200 million transistors distributed over the wafer. Proven experimental prototypes are used as components of larger modules, interconnected with advanced technology. In the upper right hand corner is an example of an advanced interconnect module of 36 diverse microchips, packed together 1.5 inches on a side. Advanced computationally intensive design tools are essential to realize these high performance components.

Earth's Biosphere

The Earth's biosphere is a complex physical system. There are a multitude of phenomena which can change the state of the biosphere on a local, regional, and global scale. In order to predict the directions and consequences of changes in the state and condition of the biosphere, detailed scientific models are required, which in turn are constructed from massive amounts of experimental and computational data.

Experimental information is derived from satellite and ground based sensors. By the late 1990's, the sensors will have the resolution required to provide data in support of much more accurate and informed decision making on issues such as pollution and global warming. Because the sensors will generate terabytes of data each day, which will be combined with local data sets on the ecosystem, major improvements in the capability for collecting, analyzing, distributing, and archiving data are necessary.

The effort of constructing valid scientific models which describe the dynamics and underlying processes of the biosphere will involve interdisciplinary teams of experts from the geophysical, life, physical, computer, and computational sciences. They will work together to construct computational models which will validate our empirical understanding of the biosphere, and help predict how worldwide activities affect the global ecosystem. Computer and computational scientists will develop the advanced software technology and algorithms for handling massive amounts of data and working with high resolution, coupled models of the Earth's atmospheric-biospheric-oceanic interactions. Efforts in software development and experimentation will predict how local current conditions may impact future global conditions allowing the linking of ecosystem models to global climate models. The result of this collaborative effort will be a much deeper understanding of our environment and the impact of human activities upon it.



A global view of the Earth's ocean chlorophyll and land vegetation is depicted here. The image was derived from accumulated satellite data. Generating this image required over 2 terabytes (10^{12} bytes) of raw data. Data such as these are the baseline "snapshot" of our current biosphere and are vital to understanding short and long term component interactions of the Earth's biosphere for such purposes as crop productivity improvement, fishing and weather prediction.

High Speed Networks

The development of more powerful computers feeds both the demand for, as well as the growth of, more powerful data communications capability. As computing technology progresses, greater demands are placed on network performance as researchers conceive of new tasks and modalities of use that require even higher performance.

Current developments in large scale scientific computing are leading to truly distributed computing, allowing a given job to be executed on several different machines communicating partial results among themselves, sharing in different facets of calculations, and jointly assembling a final result for output. Such intermachine communication is inherently capable of taking place at speeds that stress local area network technology and are a hundred or more times faster than are possible on today's long haul networks. As an example, some of the nation's leading radiology departments have committed to digital transfer of radiology images on broadband local networks operating at speeds of 100 to 1000 million bits per second.

In order for the NREN to support these and other demanding applications, not yet contemplated, a substantial directed research and development effort is needed in the areas of protocols (the formal structure of inter-computer communications), high speed computer interfaces for computers, and network equipment, such as switches. Multi-gigabit networks represent a change in kind, not just in degree, from today's networks. For example, consider that in a coast-to-coast communication at three gigabits per second there are at any instant "in flight" nearly nine megabytes of data, which is more than the memory of most personal computers and workstations.

Some gigabit research has already begun, and several experimental facilities have been established in a productive collaboration of academic, industrial, and governmental organizations, but HPCC support will be needed to carry the research program to the stage that commercial providers can use the technology to install and operate a multi-gigabit NREN.

Education Using the NREN

Computers increasingly fill important niches in all phases of the learning process, providing flexible instruments for interactive instruction and student based learning experiences. Their use in education both provides the skills needed to function in our increasingly technology-intensive world, and aids teaching and learning of all science and engineering topics. The development of the National Research and Education Network will accelerate and transfer the technology of computer communications to the needs of educators and students. The result will be to empower them to share resources and ideas on a national scale.

In a recent project, classes in many locations chose a day to measure the length of the sun's shadow from a vertical measuring stick on the school grounds at 12 noon. Each class consulted maps and geography books to find its school's latitude, and sent the results to a shared database located in the U.S. and Canada. All schools then received the database of measurements from around the world, and each class used the complete database to calculate the curvature of the earth, and from that, the earth's diameter. A normally dry recitation of facts became an engaging problem solving exercise because the students themselves derived the answer from their shared measurements. Along the way they learned geography, geometry, statistics, and how to collect and share data over computer networks.

This project, implemented by the use of the network, is a learning laboratory without walls, similar to the way research scientists take advantage of high speed digital networks to conduct shared research that is "distance independent." The HPCC Program will face the challenge of "scaling-up" today's Internet, making it "user-friendly" and improving its services so that it is readily accessible to all U.S. educational institutions.

Glossary

ASTA

Advanced Software Technology and Algorithms

Bit

Acronym for binary digit

BRHR

Basic Research and Human Resources

Byte

A group of adjacent binary digits operated upon as a unit (usually connotes a group of eight bits)

CAD

Acronym for "computer aided design".

Computer Engineering

The creative application of engineering principles and methods to the design and development of hardware and software systems.

Computer Science

The systematic study of computing systems and computation. The body of knowledge resulting from this discipline contains theories for understanding computing systems and methods; design methodology, algorithms, and tools; methods for the testing of concepts; methods of analysis and verification; and knowledge representation and implementation.

Computational Science and Engineering

The systematic application of computing systems and computational solution techniques to mathematical models formulated to describe and simulate phenomena of scientific and engineering interest.

flops

Acronym for "floating point operations per second". The term "floating point" refers to that format of numbers which is most commonly used for scientific calculation. Flops is used as a measure of a computing system's speed of performing basic arithmetic operations such as adding, subtracting, multiplying, or dividing two numbers.

Giga-

10^9 or billions of ... (e.g.: gigabits)

Grand Challenge

A Grand Challenge is a fundamental problem in science and engineering, with broad economic and scientific impact, whose solution could be advanced by applying high performance computing techniques and resources.

HPCC

High Performance Computing and Communications

HPCS

High Performance Computing Systems

High Performance Computing

High performance computing encompasses advanced computing, communications, and information technologies, including scientific workstations, supercomputer systems, high speed networks, special purpose and experimental systems, the new generation of large scale parallel systems, and applications and systems software with all components well integrated and linked over a high speed national network.

Mega-

10^6 or millions of ... (e.g.: megaflops)

Network

Computer communications technologies that link multiple computers to share information and resources across geographically dispersed locations.

NREN

National Research and Education Network

ops

Acronym for "operations per second". Ops is used as a rating of the speed of computer systems and components. In this report ops is generally taken to mean the usual integer or floating point operations depending on what functional units are included in a particular system configuration.

Parallel Processing

Simultaneous processing by more than one processing unit on a single application.

Peta-

10^{15} or thousands of trillions of ... (e.g.: petabytes)

Supercomputer

At any given time, that class of general-purpose computers that are both faster than their commercial competitors and have sufficient central memory to store the problem sets for which they are designed. Computer memory, throughput, computational rates, and other related computer capabilities contribute to performance. Consequently, a quantitative measure of computer power in large-scale scientific processing does not exist and a precise definition of supercomputers is difficult to formulate.

Tera-

10^{12} or trillions of ... (e.g.: teraops)


ADVANCED MATERIALS AND PROCESSING:

The Federal Program in Materials Science and Technology

*A Report by the FCCSET Committee on
Industry and Technology*

*To Supplement the President's
Fiscal Year 1993 Budget*





Cover:

Materials Research to Application

The cover, front and back, is an artist's rendition of the featured theme of the Advanced Materials and Processing Program - Research to Application. It depicts the breadth of the program through a montage of figures drawn to represent examples of emphasis areas. The back cover contains illustrations of the research and development effort, from the fundamental (crystal structure), to synthesis-processing (phase diagram, injection molder, metals processing ladle), to environmental concerns (fork composed of biodegradable corn starch). Applications of materials technologies are shown on the front cover and include information/communications (transistor cross section, optical fibers), health (MRI image), transportation (automobile, airplane, moon/space exploration) and infrastructure/construction (bridge, roadways). The translation of the R&D to sector beneficiaries is reflected by the optical fiber light pipe that bridges the gap, connecting the front and back covers.

ADVANCED MATERIALS AND PROCESSING:

The Federal Program in Materials Science and Technology

*A Report by the Committee on
Industry and Technology
Federal Coordinating Council for
Science, Engineering and Technology
Office of Science and Technology Policy*

*To Supplement the President's
Fiscal Year 1993 Budget*

**Office of Science and Technology Policy
Federal Coordinating Council for Science,
Engineering and Technology (FCCSET)
Committee on Industry and Technology (CIT)**

Chairman

Robert M. White, Department of Commerce

Vice Chairmen

T S Ary, Department of the Interior

W. Henson Moore, Department of Energy

Members

Curtis Bohlen, Department of State

Joseph Bordogna, National Science Foundation

Erich W. Bretthauer, Environmental Protection Agency

Todd Bucholz, the White House's Economic Policy Council

Michael Deland, Council on Environmental Quality

Gary Denham, Department of Defense

John Diggs, Department of Health and Human Services

Joseph Hezir, Office of Management and Budget

David Kearns, Department of Education

Thomas D. Larson, Department of Transportation

Shellyn McCaffrey, Department of Labor

Hollis McLoughlin, Department of the Treasury

Harry Mussman, United States Department of Agriculture

Gordon Oehler, Central Intelligence Agency

Richard H. Peterson, National Aeronautics and Space Administration

Jack Sommer, Department of Housing and Urban Development

Bruce Wilson, Office of the U.S. Trade Representative

Eugene Wong, Office of Science and Technology Policy

Robert J. Woods, Department of Veterans Affairs

Executive Secretary

Cary Gravatt, Department of Commerce

FCCSET Directorate

Vickie V. Sutton, Assistant Director

Charles H. Dickens, Executive Secretary

Steering Group on Materials

T S Ary, Department of the Interior (Chairman)

Reid Adler, Department of Health and Human Services

Alfred Galli, Environmental Protection Agency

Karl Erb, Office of Science and Technology Policy

Leonard Harris, National Aeronautics and Space Administration

Joseph Hezir, Office of Management and Budget
Harry Mussman, United States Department of Agriculture
David Nelson, Department of Energy
William D. Phillips, Office of Science and Technology Policy
(Past Member)
Lyle H. Schwartz, National Institute of Standards and
Technology, Department of Commerce
Raymond Siewert, Department of Defense (Past Member)
John A. White, National Science Foundation (Past Chairman)
Robert M. White, Department of Commerce

COMAT Working Group on National Planning

Chairman

Lyle H. Schwartz, National Institute of Standards and
Technology, Department of Commerce

Executive Secretary

Samuel J. Schneider, National Institute of Standards and
Technology, Department of Commerce

Members

Andrew Crowson, Department of Defense (Army)
James Eberhardt, Department of Energy
Alfred Galli, Environmental Protection Agency
Garrett Hyde, Department of the Interior, Bureau of Mines
Daniel E. Kugler, United States Department of Agriculture
Stephen Moran, National Aeronautics and Space Administration
Jagdish Narayan, National Science Foundation
Thomas Pasko, Department of Transportation
Bhakta Rath, Department of Defense (Navy)
Vincent Russo, Department of Defense (Air Force)
Iran Thomas, Department of Energy
Samuel Venneri, National Aeronautics and Space Administration
John Watson, National Institutes of Health
Robert Wellek, National Science Foundation
Benjamin Wilcox, Department of Defense (DARPA)

Liaisons/Observers

Robert Civiak, Office of Management and Budget

David Cranmer, Office of Science and Technology Policy

Charles Dickens, FCCSET Directorate

Henry Ehrenreich, Office of Science and Technology Policy

Karl Erb, Office of Science and Technology Policy

David Gold, Office of Management and Budget

Cary Gravatt, Committee on Industry and Technology

David Huber, National Critical Materials Council

John A. White, National Science Foundation

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF SCIENCE AND TECHNOLOGY POLICY
WASHINGTON, D.C. 20506

MEMBERS OF CONGRESS:


I am pleased to forward with this letter "Advanced Materials and Processing: The Federal Program in Materials Science and Technology," prepared by the Committee on Industry and Technology of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), to supplement the President's Fiscal Year 1993 Budget.

This Presidential Initiative recognizes that materials are the basis of a critical enabling technology upon which most other technologies depend for their success. This country has been the world leader in fundamental materials science and technology research, but a number of recent Federal and private sector studies have called attention to the need to strengthen this Nation's commitment to bridging the gap between basic and applied research and development, thus spurring U.S. competitiveness.

The Advanced Materials and Processing Program (AMPP) is a coordinated interagency effort to exploit opportunities in materials research and development to meet significant national goals and to extend U.S. leadership in this area. If aggressively exploited, this effort should lead to materials having properties and performance tailored for specific applications that can be fabricated by cost-effective and environmentally sound processes. Increased investments in materials science and technology should result in major contributions to this Nation's quality of life, national security, industrial productivity, and economic growth.

The AMPP is an integral component of the President's overall approach to challenges in science, technology, and education. Improvements in the processing and performance of materials will impact broad sectors of Federal research activities, including, for example, the Presidential Initiatives for High Performance Computing and Communications and Biotechnology Research.

The coordination and integration of the interagency research and development strategy for this Initiative was led very ably by Dr. Robert White, Chairman of the Committee on Industry and Technology, and his interagency committee members, associates, and staff. They are all to be commended on the excellent work that is manifest in both the Initiative and the report.



D. Allan Bromley
Director

CONTENTS

EXECUTIVE SUMMARY	1
INTRODUCTION TO THE AMPP	5
The Importance of Materials	5
Purpose and Structure of This Initiative	8
The AMPP Planning Framework	8
Figure 1 - AMPP Planning Framework	10
The Origins of the AMPP	9
THE FEDERAL MATERIALS R&D PROGRAM	15
Overview of the Federal Role in Materials R&D	15
Figure 2 - The Total Materials Cycle	16
Defining the Base: Federal Materials R&D Inventory	17
Inventory Highlights	19
THE AMPP BUDGET FOR FY 1993	23
The FY 1993 Budget	23
AMPP Program Enhancements	23
Figure 3 - Evaluation Criteria	24
Table 1 & AMPP R&D Funding By Agency	26
Figure 4	27
Table 2 & AMPP R&D Funding By Research Component	28
Figure 5	29
Table 3 & AMPP R&D Funding By Materials Class	30
Figure 6	31
Summary	33

CONTENTS *continued*

APPENDIX

A: Agency Activities in Fiscal Years 1991 and 1992	35
B: Proposed FY 1993 Program Enhancements	47
C: Breakthrough Opportunities	51
D: Technical Background The Science and Engineering of Materials Classes of Materials	62

EXECUTIVE SUMMARY

Advanced materials and advanced processing of materials are critical to the nation's quality of life, security, and economic strength. Advanced materials are the building blocks of advanced technologies. Everything Americans use is composed of materials. Millions of manufacturing jobs depend on the availability of high quality specialized materials.

Unprecedented opportunities are waiting to be exploited in materials science and technology; however, success in this endeavor requires cooperation throughout the total government, academic, industrial enterprise. The need to strengthen the Federal commitment to materials R&D has been widely recognized. The National Critical Technologies Panel identified materials technology as one of six broad areas of needs, and materials figure prominently in Federal agencies' lists of important technologies.

To address these opportunities and needs, a multi-year, multi-agency program will initiate in FY 1993 to increase the effectiveness of the Federal R&D program in Materials Science and Technology.

THE GOAL OF THIS ADVANCED MATERIALS AND PROCESSING PROGRAM (AMPP) IS TO IMPROVE THE MANUFACTURE AND PERFORMANCE OF MATERIALS TO ENHANCE THE NATION'S QUALITY OF LIFE, SECURITY, INDUSTRIAL PRODUCTIVITY, AND ECONOMIC GROWTH.

The AMPP will focus special attention on the interfaces between universities, government laboratories and industry, and on the process of technology transfer from basic research to application. Continuing input from all parties will be sought and collaboration will be fostered through mechanisms such as consortia and cooperative research and development agreements.

The President's FY 1993 budget includes \$1821.4 million for materials R&D representing an increase of 10 percent over the FY 1992 base. Special emphasis is placed on synthesis and processing, those areas critical to the development of new materials and to the improvement of reliability, cost, and quality of all materials. The budget also increases emphasis on those programs which are aimed at bridging the gap between the innovation of new materials and their application.

INTRODUCTION TO THE AMPP

INTRODUCTION TO THE AMPP

The Importance of Materials

Everything Americans use is composed of materials, from semiconductor chips to flexible concrete skyscrapers, from plastic bags to a ballerina's biomaterial artificial hip, or the composite structures on spacecraft. The impact of materials extends beyond products, in that tens of millions of manufacturing jobs depend on the availability of high-quality, specialized materials that are low in cost.

Throughout human history, new materials have brought about scientific, technological, and cultural evolution. Indeed, periods of history have been named after materials — from the Stone Age of primitive tools, to the Iron Age that spawned machinery, to the present era, sometimes described as the Age of Materials.

For millennia now, humans have molded, fired, and mixed natural commodities to meet personal needs. Today, the greater needs of a complex society can be met with advanced materials designed by science, pieced together atom by atom. This is made possible by detailed knowledge of material composition and structure, coupled with sophisticated techniques of controlling those features. The processing method is critical, because it determines a material's properties as well as the environmental impacts of making it.

Advanced materials are the building blocks of technology. When processed in particular ways, they enable the technological advances that constitute progress. (Examples of Highlights on pages 6 and 7 illustrate this relationship.) Advanced materials and processing methods have become essential to the enhancement of U.S. quality of life, security, industrial productivity, and economic growth. They are tools for addressing urgent problems, such as pollution, declining natural resources, and escalating costs.

Unlimited and unprecedented opportunities are waiting to be exploited in the arena of modern materials science and technology. The defining technologies of the future are likely to be based on materials not yet developed, or even imagined. This is an era of invention as well as discovery.

Some Recent Highlights in Materials Development

In 1977, **silica optical fiber** was an experimental material. Four years later, this light-transmitting material was introduced into commercial service in Chicago, and has since become the carrier of choice for long-distance telecommunications. By 1985, advanced optical fibers had surpassed the signal-carrying capacity of copper wires, which had been in use since the 19th century. Today, optical fibers are replacing the more costly copper in this application. The first trans-Atlantic fiber system was installed in 1988 at a cost per circuit 100 times cheaper than the first trans Atlantic copper cable in 1956, and more powerful systems are yet to come.

The recent discovery of **high-temperature superconductors** will enable numerous technological advances, such as magnetometers that can measure extremely small changes in a magnetic field. An ultra-sensitive magnetometer developed recently can detect magnetic fields nearly 100 million times weaker than that of the Earth. This technology makes use of a SQUID — a superconducting quantum interference device operated with liquid nitrogen — and is sensitive enough for use in geophysical exploration, detection of human brain wave patterns, and non-destructive evaluation.

Composite structural materials, reinforced with ultra-strong fibers such as Kevlar® and graphite, enable designers to reduce the weight and enhance the performance of aircraft, spacecraft, and other systems. Kevlar®, which has a specific strength greater than steel, is a product of basic theoretical studies. Scientists had theorized that as a solution of rod-like molecules becomes more concentrated, crowding would preclude random positioning, and the rods would become oriented in the same direction. The theory later enabled other researchers to use rigid-rod macromolecules to make Kevlar®. This fiber was used to reinforce the bullet-proof helmets worn by troops in Operation Desert Storm.

New knowledge and tools are expanding the range of applications for **semiconductors**, the foundation of the electronics industry. For instance, materials scientists are able to produce atomically thin layers of special semiconductors and then, by introducing impurities that change the electronic structure in specific ways, to engineer solid-state structures that operate according to predetermined specifications. This technology has revolutionized work and leisure environments. Solid-state computer, telecommunications, and information processing technologies are standard in offices, and are integrated into machine tools and processing operations in factories. They are also becoming common in living rooms. Digital recording and playback equipment made from such processed semiconductors are the basis for the high fidelity sound reproductions available today.

In recent years, researchers have fabricated materials with unprecedented magnetic strength — at least twice the strength of traditional magnets. This development has enabled designers to reduce the weight and improve the efficiency of electric motors, by replacing electromagnets with **advanced magnets**. For instance, in some domestic pickup trucks, advanced neodymium-iron-boron magnets have reduced the weight of cranking motors by half. One vehicle may have as many as 70 electric motors; the cranking motor is the largest. The weight reduction is substantial enough to improve fuel economy. The world market for permanent magnets is expected to roughly triple by the year 2000 to \$6.5 billion, of which more than a third will be attributable to neodymium-iron-boron magnets.

The first **artificial heart** valve was implanted in 1952. In the early 1960s, experiments were conducted identifying carbon as a relatively non-thrombogenic material. **Pyrolytic carbon** was simultaneously being developed in the nuclear fuel industry as a coating for nuclear fuel particles. These two unrelated technologies were brought together in tests that showed pyrolytic carbon to have outstanding hemocompatibility. The first artificial heart valve containing this new material was made in 1969. Over 2 million pyrolytic carbon components have now been used in over 1.2 million heart valve replacements resulting in over six million patient-years of safe use.

The need to strengthen the Federal commitment to materials R&D has been widely recognized. In March 1991, the National Critical Technologies Panel identified materials technology as one of six broad areas of needs. Five of the 22 "critical technologies" cited are in the materials category. Materials also are featured in the Department of Defense Critical Technologies Plan and the Department of Commerce list of emerging technologies.

Purpose and Structure of This Initiative

To address the needs and opportunities in materials, the Federal government has initiated a multi-year, multi-agency effort to increase the effectiveness of the Federal R&D program in Materials Science and Technology. This document presents for the first time the framework for this coordinated program entitled the *Advanced Materials and Processing Program (AMPP)*. This interagency program has been developed and is coordinated by the Federal Coordinating Council for Science, Engineering and Technology.

The AMPP incorporates three conceptual tiers: 1) current materials R&D activities (the program base), summarized in the Inventory section and described in more detail in Appendix A; 2) proposed enhancements for FY 1993 (new and expanded programs), summarized in the Budget section and listed individually in Appendix B; and 3) opportunities for technical breakthroughs, which are organized in 10 thematic categories displayed in Appendix C. These glimpses of the future are intended to provide guidance and inspiration for current efforts and future planning.

The AMPP Planning Framework

The AMPP represents a significant effort to exploit tomorrow's opportunities today, both to meet national goals and to extend U.S. leadership in materials science and technology.

Goal. The goal of the AMPP is to improve the manufacture and performance of materials to enhance the nation's quality of life, security, industrial productivity, and economic growth.

The AMPP will optimize federal materials R&D through an interagency strategy, encompassing a research plan and associated funding levels that take advantage of opportunities for significant technical breakthroughs. The program represents the first coordinated Federal approach to materials science and technology. The AMPP is based on a crosscutting analysis of materials R&D at 10 agencies, within a framework of strategic and scientific priorities. (See Figure 1).

Objectives. The AMPP has four strategic objectives:

- Establish and maintain the U.S. scientific and technological leadership position in advanced materials and processing;
- Bridge the gap between innovation and application of advanced materials technologies;
- Support agencies' mission objectives to meet national needs with improvements in advanced materials and processing; and
- Encourage university and private sector R&D activities in materials technologies, their applications, and their implementation.

The AMPP is a research and development program. Success in this technological arena is ultimately measured both by new discoveries and by application of the R&D by industry and government. To achieve these objectives, the AMPP provides a framework within which the interfaces between innovation and application can be strengthened and the participants in the science and technology development enterprise in government, academia, and industry can interact more strongly.

The Origins of the AMPP

The AMPP will add to the momentum of an evolutionary shift in U.S. materials R&D. Once oriented toward modifying natural commodities to make useful technologies, R&D now increasingly concentrates on tailoring materials to achieve specific properties.

Figure 1

U.S. Advanced Materials and Processing Program (AMPP)

GOAL

Improve Manufacture and Performance of National Security and Industrial Productivity

STRATEGIC OBJECTIVES


- A. Establish and Maintain the U.S. Scientific and Technological Leadership Position in Advanced Materials and Processing
- B. Bridge the Gap between Innovation and Application of Advanced Materials Technologies

IMPLEMENTING PRIORITIES

- A. Support the Strategic Objectives Through a Vigorous R&D Effort on New and Improved Materials Technologies
- B. Plan Federal Materials R&D Programs to Incorporate the Needs of Strategic, Industrial, and Social Sectors of U.S. Society
- C. Promote the Application of Advanced Materials and Processing through Cooperative Efforts among Universities, Industry and Government in Generic, Pre-competitive Technology Development

PROGRAM COMPONENTS

Research Programs May Consist of One or More Program Components Linked to Achieve a Specific Goal



Materials to Enhance U.S. Quality of Life,
and Economic Growth

C. Support Agencies' Mission Objectives to Meet National Needs
with Improvements in Advanced Materials and Processing

D. Encourage University and Private Sector R&D Activities on
Materials Technologies, Their Applications and Their
Implementation

D. Support the U.S. Human Resource Base to Meet the Needs of
Industry, Government, and Universities

E. Maintain a Healthy Infrastructure (e.g., Experimental and
Computational Facilities) to Support a Vibrant National R&D
Program in Materials Science and Technology

F. Focus Federal R&D Objectives on Those Materials and Processes
Which are Most Important to Achieving the Strategic Objectives

C.
MATERIALS
CHARACTERIZATION

1. Performance
2. Properties
3. Structure/
Composition

D.
EDUCATION &
HUMAN RESOURCES

1. Technical Training
2. Continuing Education
3. Undergraduate
Education
4. Graduate and Post
Graduate Education

E.
NATIONAL
USER FACILITIES

Photon, neutron and
magnet sources

Increasing Emphasis for Budget Enhancement

The AMPP draws heavily on a series of comprehensive and broad-based studies, particularly a 1989 report by the National Research Council (NRC), "Materials Science and Engineering for the 1990s: Maintaining Competitiveness in the Age of Materials." The impact of this MS&E report has been extensive.

The MS&E report documented the importance of advanced materials to U.S. quality of life, security, industrial productivity, and economic growth. In addition, it brought together more than 100 contributing scientists and engineers from industry, academia, and the government, fostering the cohesion of the materials science community and providing direction. The study was unusual for the NRC in its in-depth review of private sector needs, with eight industries surveyed.

Subsequently, at the request of the NRC and OSTP, a series of four meetings was held across the country. These involved more than 400 participants from industry, academia, and government. The results were summarized in a 1991 publication, "A National Agenda in Materials Science and Engineering: Implementing the MS&E Report."

These two documents¹ provided the key ingredients for preparation of the AMPP, offering many valuable suggestions for research programs and strategies. The key findings influencing the development of the program were:

- The need to enhance materials R&D activities, especially in synthesis and processing;
- The need to bridge the gap between basic understanding of materials and their technological applications;
- The need to involve industry, academia, and the government in mission-oriented planning and the extension of R&D activities.

¹ The MS&E report was published by the National Academy Press, Washington, D.C. The follow-up report was produced by the Materials Research Society, Washington, D.C. Also useful in organizing the AMPP was an extensive literature of detailed studies organized by the NRC's Solid State Sciences Committee and National Materials Advisory Board.

THE FEDERAL MATERIALS R&D PROGRAM

THE FEDERAL MATERIALS R&D PROGRAM

Overview of the Federal Role in Materials R&D


Federal materials R&D programs originated in the early 1900s. Over the years, the essential role of materials in supporting the missions of various agencies stimulated the development of a strong program, although its many components are largely independent. Altogether, the Federal effort is broad, focusing on virtually every aspect of the materials cycle (Figure 2), from extraction through recycling.

For many years, federally funded materials R&D focused on developing materials for defense and nuclear applications. In the 1960s, the range of applications expanded to include the space program and environmental and alternative-energy systems. In recent years, a combination of factors has prompted recognition of the need for a stronger overall Federal commitment to materials R&D. A major concern is growing foreign competition, which already has had a significant impact on manufacturing and national competitiveness. In short, increased national attention to issues associated with economic growth forced re-examination of the Federal research program.

Federal materials R&D supports and complements the activity of industry and state and local governments. Industrial R&D, while much larger overall than the Federal program, focuses heavily on product development. Industry depends on Federal programs to carry out fundamental research, to support the education of scientists and engineers in materials R&D, and, increasingly, to cooperate in the development of generic, pre-competitive technologies. Expanded dialogue with industry to define appropriate areas for cooperation is a central feature of the AMPP.

Materials R&D activities supported by state and local governments generally are based at universities. Although such programs are intended primarily to stimulate local economic development, increasing industrial involvement has created an evolving cooperative public/private network that is a strong base for carrying out the AMPP.

Figure 2
The Total Materials Cycle



The Federal materials program is carried out in diverse locations. Some R&D is conducted in government-owned laboratories operated either by government agencies or by contract with the private sector. Many of these laboratories have strong capabilities in materials R&D and are poised for increased technology transfer activities. The rest of the Federal program is carried out through research grants or contracts by industry, universities, and private research centers.

Defining The Base: The Materials R&D Inventory

The AMPP consists of four scientific and technical R&D program components, which have been identified as critical to sustained progress in materials science and technology. In order of priority, they are:

- synthesis and processing, encompassing the creation of new materials and processes, applied R&D to transfer laboratory achievements to pilot plants, and process integration with design and manufacturing requirements;
- theory, modeling, and simulation, capitalizing on U.S. leadership in computational techniques for quantitative understanding of complex materials and processes;
- materials characterization, focusing on the interrelationships among structure, composition, properties, and performance; and
- education and human resources.

Funding for National User Facilities which are critical to materials R&D, such as sources of photons, neutrons and high magnetic fields, is also included in the AMPP. Within the second component, the AMPP complements activities supported by the presidential initiative in High Performance Computing and Communications. The fourth component complements the activities supported by the presidential initiative in Math and Science Education.

AMPP technical programs are described in terms of a particular class of materials, a category of applications, or both. Broad classes of materials include: biomaterials and biomolecular materials; electronic, optical/

photonic, superconducting, and magnetic materials; and metals, ceramics, polymers, and composites. Brief descriptions of each class may be found in Appendix D, along with a basic explanation of how various materials are synthesized and processed.

Categories of applications are described in terms of opportunities for technical breakthroughs. These are examples that may be used to guide ongoing programs and planning for future AMPP enhancements. Appendix C describes opportunities in agriculture, defense, energy, environment, raw materials extraction and production, health, information and communications, infrastructure and construction, transportation, and materials for the future.

The Federal government supports a broad based program of materials R&D conducted by many agencies, and impacting many areas of application. The term materials R&D has a range of connotations that encompass scientific and engineering research on substances in any form, at any stage of preparation, fabrication, or manufacture. Materials R&D may be conducted as a discrete or multi-component activity clearly identified as materials research, or as a subset of a program with an unrelated title.

Program types vary widely. Materials R&D can be part of: a technology development effort, such as NASA's High Speed Research Program; a device or process development activity such as DOD's ceramic armor work, HHS's development of the artificial heart, the USDA's program for degradable starch polymers or DOE's ceramic engine program; a data project such as DOC's phase-diagram compilations; or individual projects conducted either to advance fundamental understanding or for their intrinsic intellectual interest, such as those sponsored by the NSF.

These variations mean that special surveys are required to obtain a comprehensive overview of the AMPP-related effort. The materials R&D survey covered fiscal years 1991 and 1992. The FY 1992 effort was designated the base for the AMPP and for the FY 1993 enhancement. Each agency was responsible for the identification, fiscal accountability, and tracking of its AMPP R&D, and each employed its own classification scheme. Comparable surveys had been conducted previously for the fiscal years 1976, 1980, and 1982, covering approximately the same agencies and scope of activities.

Inventory Highlights

Of the 21 federal agencies surveyed, 10 reported funding for materials R&D for fiscal years 1991 and 1992. They were:

Department of Commerce (DOC)
Department of Defense (DOD)
Department of Energy (DOE)
Department of Interior (DOI)
Department of Transportation (DOT)
Environmental Protection Agency (EPA)
Department of Health and Human Services (HHS)
National Aeronautics and Space Administration (NASA)
National Science Foundation (NSF)
Department of Agriculture (USDA)

Budgets and research activities varied according to agency missions. Brief descriptions of these agency activities are included in Appendix A. Highlights of the inventory are as follows:


- Total aggregated Federal materials R&D (in current dollars) was \$1663.4 million in FY 1991 and \$1658.6 million in FY 1992. (These totals include funds allocated for construction and operation of the National User Facilities, which are high-intensity photon, neutron, and high magnetic field sources.)
- Analysis of the scope and coverage of the funding surveys indicates that Federal expenditures for AMPP-oriented R&D have remained essentially unchanged since FY 1976, in constant dollars. The proportion of materials R&D relative to total Federal R&D for all fields of science and technology is 2.3% in FY 1992, compared to 3.4% in FY 1976.
- Federal funding for materials R&D is concentrated in a few major agencies. In FY 1992, 79% of all funding was provided by three agencies, DOE (36%), DOD (27%), and NSF (16%).
- Funding levels vary widely for the materials classes covered. In FY 1992, metals received the most funding (13%), followed by composites (11%), electronic materials (10%) and biomaterials (10%).

The Advanced Materials and Processing Program includes materials R&D activities that may benefit from increased interagency and public/private cooperation. Excluded at agency discretion are classified programs for reasons of national security and, some unique mission driven activities, such as the superconducting supercollider and the national aerospace plane. The sorting process becomes more difficult as project size decreases; fine distinctions are left to individual agencies.

Some educational activities are included, although overall Federal enhancements in this area are coordinated through a Presidential Initiative in Mathematics and Science Education. Current activities include an NSF project to develop undergraduate courses and curricula to enhance materials education for science and engineering majors, a NASA initiative to improve the research infrastructure and capabilities of Historically Black Colleges and Universities, and increased DOD support of fellowships for graduate students working in areas important to national security.

The AMPP also includes the National User Facilities, which are critical to materials R&D. These facilities traditionally are budgeted and managed through the materials programs at DOE, NSF, and DOC. However, because these facilities also are used by scientists in many other fields, this category will be monitored and reported through the AMPP.

THE AMPP BUDGET FOR FY 1993



THE AMPP BUDGET FOR FY 1993

The FY 1993 Budget

The FY 1993 budget for the AMPP is \$1821.4 million, an increase of \$162.8 million (10%) over the FY 1992 base. This total includes \$1530.9 million for direct programs, an increase of \$129.5 million (9%), and \$290.5 for construction and operation of the National User Facilities. This budget represents the first year of coordinated interagency planning for the AMPP and is the result of a detailed assessment of needs and opportunities by the 10 participating agencies.

An interagency committee evaluated enhancements to the AMPP proposed by the agencies for technical content and adherence to the goal, strategic objectives and implementing priorities of the AMPP framework (see Evaluation Criteria, Figure 3). Over 85 proposals were considered, of which 38 were included in the President's budget. Enhancements for FY 1993 emphasize synthesis and processing with a 64.8 million increase over the 1992 base.

Tables 1-3 show the breakdown of the FY 1993 budget by agency, AMPP research component, and materials class. Data for materials research for fiscal years 1991 and 1992 are included to illustrate both the increase in overall funding and shifts in priorities. The FY 1991 figures represent actual expenditures (budget obligations), the FY 1992 figures are budget authorizations as set by Congressional appropriations, and the FY 1993 figures reflect the President's budget request.

Figures 4-6 display the identical data in a graphic format.

AMPP Program Enhancements

Specific program enhancements, identified by AMPP technical component and responsible agency, are listed in Appendix B. Among the accomplishments anticipated from these enhancements are:

Synthesis and Processing

- development of the fundamental understanding required for process control of advanced ceramic structures via hot-isostatic-pressing and microwave sintering (DOE);

Figure 3

Evaluation Criteria for the Advanced Materials and Processing Program

Relevance/Contribution. The research and development must significantly contribute to the overall goals and strategy of the Advanced Materials and Processing program (AMPP).

Technical/Scientific Merit. The proposed agency program must be technically/scientifically sound and of high quality, and must be the product of a documented technical/scientific planning and review process.

Environmental/Social Merit. The envisioned products and processes should be designed to take into account national environmental and social goals.

Readiness. A clear agency planning process must be evident, and the organization must have demonstrated capability to carry out the program.

Timeliness. The proposed work must be technically/scientifically timely for one or more of the AMPP components.

Linkages. The responsible organization must have established policies, programs, and activities promoting effective technical and scientific connections among government, industry, and academic sectors.

Costs. The identified resources must be adequate, represent an appropriate share of the total available AMPP resources (e.g., a balance among program components), promote prospects for joint funding, and address long-term resource implications.

Enhancements to Existing Program Research. Existing agency AMPP programs will receive adequate support before new initiatives are funded.

Agency Approval. The proposed program or activity must have policy-level approval by the submitting agency.

- extension of university-based programs emphasizing synthesis and processing (NSF);
- construction of starch/microbial biopolymer composites (USDA);
- expansion of applications for “intelligent” processing of materials, through the formation of multidisciplinary and multifunctional teams of university researchers, systems developers, vendors equipment manufacturers, and end users, with an initial focus on advanced control of injection molding of polymers (DOC); and
- enhancement of basic understanding to allow computation of phase diagrams for covalently and ionically bonded ceramics, to replace the costly empirical strategy (DOE).

Theory, Modeling, and Simulation

- expansion of individual investigator efforts in theory, modeling, and simulation at universities and initiation of large, interdisciplinary groups addressing major problems in mathematics and computational modeling relevant to advanced materials and processing (NSF); and
- use of fundamental theoretical tools to design structural polymers with desirable properties from the atomic scale (DOE).

Characterization

- accelerate the development of atomic-scale characterization techniques and shared facilities (NSF);
- accelerate research on the immunologic response of total joint replacements and bone material interfaces (HHS)


Table 1
AMPP R&D Funding by Agency
 (Dollars in Millions)

Agency	FY91 \$M	FY92 \$M	FY93 \$M	FY93 GROWTH OVER FY92 \$M
DOC	44.5 (33.4)	46.2 (34.2)	48.0 (38.4)	1.8 (4.2)
DOD*	505.9	448.6	431.9	-16.7
DOE*	593.9 (361.8)	602.7 (380.3)	677.9 (423.8)	75.2 (43.5)
DOI	25.0	24.7	23.8	-0.9
DOT	10.0	8.8	15.5	6.7
EPA	3.2	3.5	4.5	1.0
HHS	66.6	76.6	81.7	5.1
NASA	116.3	124.5	153.5	29.0
NSF	246.6 (216.0)	265.9 (243.1)	318.5 (291.7)	52.6 (48.6)
USDA	51.4	57.1	66.1	9.0
TOTAL	1663.4 (1389.6)	1658.6 (1401.4)	1821.4 (1530.9)	162.8 (129.5)

Parentheses indicate total funding excluding major National User Facilities Construction and Operations Cost

**Excludes classified research*

Figure 4
AMPP R&D Funding by Agency




*Includes National User Facilities Construction and Operation Cost

+Excludes classified research

Table 2
AMPP R&D Funding by Research Component
 (Dollars in Millions)

Component	FY91		FY92		FY93		FY93 GROWTH OVER FY92
	\$M	Pct	\$M	Pct	\$M	Pct	\$M
Synthesis and Processing	670.1	40.3	682.9	41.2	747.7	41.0	64.8
Theory, Modeling, Simulation	219.1	13.2	223.6	13.5	253.3	13.9	29.7
Materials Characterization	481.2	28.9	473.8	28.6	503.0	27.6	29.2
Education/Human Resources	19.1	1.2	21.1	1.3	26.9	1.5	5.8
National User Facilities/Other	273.8	16.4	257.2	15.5	290.5	15.9	33.3
TOTAL COMPONENT FUNDING	1663.4	100.0	1658.6	100.0	1821.4	100.0	162.8

Figure 5
AMPP R&D Funding by Research Component



*Includes Construction and Operation Costs

Table 3


AMPP R&D Funding by Materials Class

(Dollars in Millions)

Materials Class	FY91 \$M	FY92 \$M	FY93 \$M	FY93 GROWTH OVER FY92 \$M
Bio/BioMolecular Matls	143.6	166.0	187.4	21.4
Ceramics	137.1	132.4	150.5	18.1
Composites	185.0	182.5	206.8	24.3
Electronic Materials	172.2	162.2	176.7	14.5
Magnetic Materials	21.5	22.6	27.7	5.1
Metals	217.2	211.2	228.9	17.7
Optical/Photonic Matls	143.4	133.0	138.8	5.8
Polymers	83.2	83.0	93.2	10.2
Superconducting Matls	136.8	151.7	142.9	-8.8
Other or Non Materials Specific*	423.4	414.0	468.5	54.5
TOTAL	1663.4	1658.6	1821.4	162.8

**Includes National User Facilities Construction and Operations Cost*

Figure 6
AMPP R&D Funding by Materials Class



**Includes National User Facilities Construction and Operations Cost*

Education and Human Resources

- development of broad-based integrated approaches to undergraduate course and curricula development at intermediate and advanced undergraduate college levels, to improve the depth and breadth of materials education for majors in physics, chemistry, materials science, biology, mathematics and the engineering disciplines (NSF).

AMPP enhancements will emphasize areas of application-oriented research that build on a strong base. Many may contribute toward progress in the Breakthrough Opportunity categories (described in Appendix C). Examples of application areas which will benefit from advanced materials R&D programs include:

Energy

- exploration, through theory and experiment, of multicomponent systems of semiconductors designed to optimize photovoltaic conversion of solar energy (DOE); and
- complete manuals for life cycle analysis and pollution prevention (EPA);
- development of design, synthesis, and processing technologies for large-area semiconductor alloys with variable band gaps to optimize solar energy conversion (DOE) and
- development of advanced functional fluids and greases from oilseed bearing plants (USDA).

Transportation

- development and validation, in cooperation with U.S. industry, of the technical feasibility of high-temperature, lightweight composites for critical components of future high-speed civil transport aircraft (NASA);
- development of processing parameter measurement techniques to assure quality control in thermal spray coating (DOC);

- expansion of the program in ceramic matrix composites for very-high-temperature loading-bearing applications, to emphasize the improvement of powder, fiber, and microchemistry (DOE);
- development of programs (with industry) to acquire the technical knowledge needed to achieve automated processing of advanced structural ceramics (DOE);
- development of a program, with the transportation industry, leading toward improved reliability and reduced processing costs of lightweight materials for ground vehicles (DOE); and
- development of advanced bridge structures through modular construction using concrete composites (DOT).

Information/Communications

- development of measurement technology to assure the quality control of chemical reactions during fabrication of electronic printed wiring boards (DOC);
- expansion of synthesis and processing research on nanostructured materials for electronic, optical, and magnetic applications (DOE).

Summary

In summary, the Advanced Materials and Processing Program initiates a multi-year, multi-agency effort to increase the effectiveness of the Federal R&D program in materials science and technology. The FY 1993 budget for the AMPP is 1821.4 million, an increase of 162.8 million (10%) over the FY 1992 base. This budget will increase the emphasis on synthesis and theory, modeling, and simulation. The budget supports a comprehensive range of materials R&D and gives added emphasis to visionary programs that bridge the gap between the innovation of new materials and their application in several breakthrough opportunity areas.

Success of the AMPP depends on the integration of Federal materials R&D with the efforts of academia and industry. This collaboration will be fostered by the traditional mechanisms of research contracts and grants, as well as new tools introduced in the late 1980s, notably consortia and cooperative research and development agreements. The effectiveness of these collaborations and detailed planning for subsequent years will require expanded interaction with the private sector where many of the resulting technologies will be used.

Exploitation of the programs defined for FY 1993 and extension to the full range of breakthrough opportunities will depend, in part, on further AMPP enhancements and on those of other Presidential initiatives, most notably the High-performance Computing and Communications (HPCC) and Biotechnology Initiatives. Likewise, the HPCC Grand Challenges and the future biotechnology opportunities will benefit from the AMPP Initiative outlined in this report.

APPENDIX A

MATERIALS R&D ACTIVITIES BY AGENCY

Department of Commerce (DOC)

Materials R&D is conducted primarily by the Materials Science and Engineering Laboratory at the National Institute of Standards and Technology (NIST), although four other NIST laboratories are involved. The core program addresses: scientific and measurement issues crucial to industry's success in exploiting the advantages of advanced materials; the development of process technologies to transform new advanced materials into viable products; and manufacturing processes that yield high-quality products at reasonable cost.

NIST conducts R&D in a number of areas featured in the AMPP, including: "intelligent" processing of advanced materials, a computer-based concept for automatic process control that NIST played a leading role in developing; functional materials, such as those used for interconnections in electronic devices; and materials for electronic packaging.

Current program directions include:

- expansion of the NIST reactor, a 20-megawatt reactor with 25 experimental facilities, to include the Cold Neutron Research Facility, with instruments for study of materials structure, materials dynamics, chemical analysis, and neutron physics; and

- new programs in: studies of the structure and properties of oxide superconductors; research on advanced structural ceramics; processing and performance of polymer matrix composites; and functional materials, including magnetic, photonic, and diamond-like materials.

Department of Defense (DOD)

The DOD Materials Science and Technology Program focuses on a range of structural, thermal protection, non-structural, and electronic materials. The program is conducted by the Army, Navy, Air Force, and Defense Advanced Research Projects Agency. Many of the activities are relevant to the AMPP.

The entire DOD program is devoted to projects that are applicable to national security activities. Materials and topics addressed in this effort include: advanced alloys; ceramics; polymers; coatings; corrosion; non-destructive inspection and evaluation; fracture analysis and testing; welding and joining; superconductivity; electronic and electro-optical materials; structural analysis, demonstration, and testing; and performance and durability.

The program also includes discrete projects directed at fulfilling specific national security needs. These efforts focus, for example, on: composite materials for aircraft and missile airframes and gas turbine engines; protection of personnel, sensors, and structures against laser threats; and development of superconducting and electro-optical components.

Where security allows, the results of all DOD science and engineering research are made public. The list of commercial components and systems that originated with the military is almost endless, ranging from modern jet engines to long-lasting packaged food.

Department of Energy (DOE)

Most of the non-classified materials R&D is conducted by the Office of Energy Research, the Office of Nuclear Energy, and the Office of Conservation and Renewable Energy. The program encompasses a broad range of studies, from fundamental research to large-scale development of energy technologies.

The R&D programs are carried out primarily at DOE's National Laboratories, at universities, and, to a lesser extent, by industry. Programs evolve continuously in light of needs and opportunities; the discovery of high-temperature superconductors, for example, led to redirection of about 10% of the budget in the Office of Energy Research, Division of Materials Sciences.

Projects related to the AMPP, in addition to basic research, include the development of:

- catalysts to improve conversion efficiency and selectivity of plentiful domestic coal and gas, so these fuels can be substituted for imported and supply-limited oil;

- corrosion-resistant materials to extend the service life of energy system components; high-temperature materials to improve operating efficiency of energy systems; and electronic materials for solar energy conversion and for sensors;

- practical low-temperature superconductors for magnets to be used in energy storage and in accelerators;

- high-strength magnetic materials to improve the efficiency of electric motors;

- processes for waste minimization, recycling, and resource recovery, as well as materials for waste containment; and

- radiation-resistant alloys for use in nuclear reactors.

Department of Interior (DOI)

The Bureau of Mines (BOM) conducts materials R&D aimed at supplying and conserving both engineered and commodity materials that are important to the nation's defense and economy. The program addresses issues throughout material life cycles, seeking technological solutions to health, safety, and environmental problems. Activities support the primary technical component of the AMPP through development of improved synthesis and processing technologies.

The research carried out by the BOM Division of Minerals and Materials Science has two broad components. Minerals research includes studies of generic technologies leading to innovative methods of minerals recovery and metal extraction. Special effort is made to extract specific raw materials essential to economic growth. Other projects are directed toward recycling and reusing scrap and waste materials.

Materials research focuses on: extending the service life of commodity materials; synthesis and processing of new materials as substitutes for hazardous or environmentally damaging materials; and reducing dependence on insecure supplies of strategic and critical materials.

Department of Transportation (DOT)

The DOT's primary mission is to assure that the nation's transportation system is safe, economical, and efficient, while also considering the impact of the system on the environment and on social and economic conditions. Materials R&D is sponsored by several DOT agencies. Activities related to the AMPP include the development of innovative techniques for facility construction, vehicle design and propulsion.

The Federal Highway Administration administers, with the states, the Federal Aid Highway construction program. Pavements and structures R&D focuses on the use of waste materials for highway applications, enhancement of existing materials for construction, and new understanding of the properties and performance of these materials.

The Federal Railroad Administration's National Maglev Initiative includes materials R&D related to: applications of superconductors and non-magnetic polymer composites for guideways; electro-rheological fluids for oscillation damping; and photonic materials for guideway integrity and intrusion.

Materials R&D also plays a role in the Federal Aviation Administration's Airport Technology R&D program, and in the Coast Guard's efforts to reduce corrosion, improve communications, develop better sensors and inspection tools, reduce energy consumption, and provide better protection. The Research and Special Programs Administration conducts materials R&D related to pipelines, primarily to reduce corrosion.

Environmental Protection Agency (EPA)

Materials R&D at EPA is conducted by two sections of the Office of Research and Development: the Air and Energy Engineering Research Laboratory and the Risk Reduction Engineering Laboratory.

Activities ranges from fundamental research to pilot plant studies. The base program supports the overall goal of the AMPP through R&D on pollution prevention, with relevant projects focusing on waste minimization, product or waste recycling and reuse, life cycle extension, and product substitution. The current EPA program emphasizes:

- development of methods to measure chemical emission rates from indoor pollution sources, and development of low-emitting building materials such as solvents, adhesives and finishes;

- synthesis of advanced sorbents to enhance capture of sulfur dioxide and nitrogen oxide emissions, and to remove chlorinated organics from the emissions stream;

- identification and evaluation of substitutes for ozone-depleting compounds, and study of alternative refrigeration cycles that use such compounds.

Department of Health and Human Services (HHS)

Most of the materials R&D is carried out by the National Institutes of Health (NIH), which conduct and support both basic and applied research in biomaterials. The program supports the primary technical component of the AMPP through a major emphasis on synthesis and processing. Complementary R&D activities are conducted by the Food and Drug Administration, which assesses the long-term effectiveness and safety of new biomaterials and regulates their use.

The NIH supports both extramural and intramural biomaterials research; 70% is supported through investigator-initiated grant applications, and the remainder by grant and contract programs directed at special needs and opportunities. Intramural research is conducted within a few laboratories dedicated to other primary research objectives, such as biomedical engineering or molecular hematology.

Most research supported by the NIH focuses on synthesis and processing, basic research, and applied R&D. Programs range from the study of the interface between blood and a synthetic surface to development of an artificial heart system. Limited research focuses on theory, modeling, and simulation, and on materials characterization. The NIH plans to encourage small business innovation research, including biomaterials process integration.

National Aeronautics and Space Administration (NASA)

NASA's materials R&D supports aeronautics and space applications ranging from airframe and engine materials for future high-speed civil transport aircraft to lightweight thermal protection systems for spacecraft and re-entry vehicles. The R&D program supports the primary AMPP technical components through such efforts as the Enabling Propulsion Materials program which is developing high temperature composites for future high speed civil transport engines.

Aeronautics research emphasizes advanced materials that will improve the performance, durability, economy, and environmental compatibility of civil and military aircraft. In addition to basic research, the strategic plan includes focused R&D programs aimed at subsonic, high-performance, and high-speed aircraft to hasten the applications of advanced materials. As a result of evolving materials and structures requirements, NASA increasingly has emphasized advanced metals and polymer matrix composites for airframe structures, and high-temperature polymeric, intermetallic, and ceramic matrix composites for subsonic and supersonic turbine engines.

Space research focuses on synthesis and development of space- durable materials for launch systems, space platforms, large science instruments, spacecraft, and re-entry and high-speed flight vehicles. Areas of special interest include development of carbon-carbon composites, ceramics, and ceramic matrix composites for lightweight thermal protection systems, and advanced metallic alloys, metal matrix composites, and polymer matrix composites for lightweight cryogenic tanks. Special emphasis is placed on the study of space environmental effects on materials.

National Science Foundation (NSF)

The National Science Foundation supports basic and applied research and science education in the sciences, mathematics, and engineering. As part of the AMPP, NSF supports research to promote the synthesis of advanced materials; enhance fundamental knowledge of the physics and chemistry of materials; further understanding of the links between synthesis or processing and the structure, properties, and performance of advanced materials at molecular, microscopic, and macroscopic levels; and promote development of novel processing and manufacturing methods that result in functional materials with novel or superior properties.

NSF also encourages creative linkages with industry that facilitate knowledge and technology transfer. The agency's emphasis on academic research helps to prepare the next generation of materials scientists and engineers.

NSF supports research on novel electronic, photonic, structural, magnetic, and superconducting materials that provide the foundation for potential economic growth in telecommunications, information processing, and computing systems. NSF also focuses on the emerging field of biomaterials and biomolecular materials and on interdisciplinary programs in biotechnology and solid-state and nanoscale structures.

NSF recently launched a five-year plan to enhance research in synthesis and processing of new materials. An important segment of this program focuses on research at the interfaces of materials research, engineering, chemistry, and biology. Activity is increasing in several areas, including mathematical and computer modeling and simulation, and net-shape processing.

Department of Agriculture (USDA)

Research and development by USDA focuses on the use of renewable non-food agricultural materials and their derivatives in industrial and manufacturing processes and products. Three USDA agencies pursue activities that fall within the AMPP framework, in several materials classes.

The Agricultural Research Service conducts research into the production of biobased materials. Biodegradable plastic films, for instance, can be made from blends of starch and biopolymers synthesized by microorganisms; and soybean oil is replacing petroleum-based products in newsprint ink.

The Cooperative State Research Service focuses on value-added processing and product and market development for renewable agricultural materials. The non-wood fiber kenaf has proven to be useable in commercial newsprint, for example, and guayale, a shrub native to the Southwest, may provide the basis for a domestic rubber industry. Oils and oil derivatives from crops like crambe and lesquerella can be used to make advanced nylons, plastics and functional fluids.

The Forest Service develops and disseminates information and technology to harvest, produce, and use wood products in ways that are efficient, safe, and environmentally sound. The research has two key objectives: to improve the design and performance of wood-based products, and to reduce the adverse environmental impacts of paper manufacturing. Areas of study include the recycling of paper and wood products to make high-value structural composites.

National User Facilities

These facilities, supported by the DOE, NSF, and DOC/NIST, are made available to thousands of industrial, academic, and government researchers each year. There are several types:

Synchrotron light sources provide photons and x-rays for structural and spectroscopic studies of metals, alloys, ceramics, polymers, and various biological materials. Studies range from fundamental to proprietary industrial programs including, for example, the development of x-ray lithography technology for advanced semiconductor integrated circuit fabrication. Six facilities are available, offering the brightest beams in the world, and two are under construction.

Neutron sources can give precise information on the positions and motions of individual atoms. They are useful for studying nuclear physics, radiation effects, and magnetic properties. Five facilities are available, offering the most intense continuous beams in the world and the nation's highest peak neutron flux. A cold neutron facility is under construction, which when completed, will provide research capabilities now available only in Europe and Japan. Another such facility is in the design phase.

High magnetic field laboratories are the premier sites for magnet design and development. Research using high magnetic fields contributes significantly to the development and understanding of materials such as semiconductors, superconductors, and magnetic materials. An existing facility offers the world's highest steady-state magnetic fields. Another facility, a collaborative effort involving the state of Florida, academia, and a national laboratory is under construction.

APPENDIX B

FY 1993 DIRECTORY

NEW AND EXPANDED AGENCY

PROJECTS BY PROGRAM

COMPONENT

Synthesis and Processing

- DOC Intelligent Processing of Materials
- DOC Materials for Electronic Interconnects
- DOC Functionally Gradient Materials
- DOE Synthesis and Processing of Materials for Photovoltaic Applications
- DOE Synthesis and Processing of Materials for Corrosive Environments
- DOE Advanced Techniques for Near Net Shape Processing of Bulk Ceramics
- DOE Synthesis and Processing of Catalytic Materials
- DOE Synthesis and Processing of Magnetic Materials
- DOE Advanced Techniques for Processing - Welding and Joining
- DOE Synthesis and Processing of Nanometer Scale Artificial Structured Materials
- DOE Synthesis and Processing of Conducting and Semiconducting Polymers
- DOE Synthesis and Processing - Materials for Radiation Environments
- DOE Synthesis and Processing of Structural Polymers
- DOE Synthesis and Processing of Polymeric and Organic Superconducting Materials
- DOE Basic Research on Tribological Studies
- DOE High Temperature Load Bearing Structural Ceramic Matrix Composites
- DOE Novel and Improved Materials for Photovoltaic Devices: Applied Research

DOE	Lightweight Transportation Materials
DOE	Ceramic Manufacturing
DOT	Pavements and Structures Technology
EPA	Environmentally Benign Materials and Processes
HHS	Programs of Excellence in Orthopaedic Biomaterials
HHS	Alternatives to Dental Amalgams for Posterior Restorations
HHS	Ocular Biomaterials Development
HHS	Genetically Enhanced Cardiovascular Implants
HHS	NCRR Advanced Biomaterials Laboratory
HHS	Biomaterials for Neural Prostheses
HHS	Development of New Artificial Kidney Membranes and Access Devices
HHS	Vaccine Microencapsulation Facility
NASA	High Speed Research: Enabling Propulsion Materials
NSF	Synthesis and Processing
USDA	Agricultural Materials From Vegetable Oils
USDA	National Paper and Wood Recycling for New Materials
USDA	Bioplastics

Theory, Modeling and Simulation

DOC	Intelligent Processing of Materials
DOE	Synthesis and Processing of Magnetic Materials
DOE	Synthesis and Processing: Modeling and Computation
DOE	Synthesis and Processing of Polymeric and Organic Superconducting Materials

DOE	Basic Research on Tribological Studies
DOE	High Temperature Load Bearing Structural Ceramic Matrix Composites
DOE	Novel and Improved Materials for Photovoltaic Devices: Applied Research
DOE	Lightweight Transportation Materials
DOE	Ceramic Manufacturing
DOT	Aircraft Safety Technology
EPA	Environmentally Benign Materials and Processes
HHS	Programs of Excellence in Orthopaedic Biomaterials
HHS	Alternatives to Dental Amalgams for Posterior Restorations
HHS	Ocular Biomaterials Development
HHS	Genetically Enhanced Cardiovascular Implants
HHS	NCRR Advanced Biomaterials Laboratory
HHS	Development of New Artificial Kidney Membranes and Access Devices
NSF	Theory, Modeling and Simulation

Materials Characterization

DOE	Synthesis and Processing of Magnetic Materials
DOE	Advanced Techniques for Processing - Welding and Joining
DOE	Synthesis and Processing of Structural Polymers
DOE	Synthesis and processing of Polymeric and Organic Superconducting Materials
DOE	Basic Research on Tribological Studies
DOE	High Temperature Load Bearing Structural Ceramic Matrix Composites

DOE	Novel and Improved Materials for Photovoltaic Devices
DOE	Lightweight Transportation Materials
DOE	Ceramic Manufacturing
DOT	Pavements and Structures Technology
DOT	Aircraft Safety Technology
EPA	Environmentally Benign Materials and Processes
HHS	Programs of Excellence in Orthopaedic Biomaterials
HHS	Alternatives to Dental Amalgams for Posterior Restorations
HHS	Ocular Biomaterials Development
HHS	Genetically Enhanced Cardiovascular Implants
HHS	NCRRL Advanced Biomaterials Laboratory
HHS	Biomaterials for Neural Prostheses
HHS	Development of New Artificial Kidney Membranes and Access Devices
NASA	High Speed Research: Enabling Propulsion Materials
NSF	Materials Characterization
USDA	Agricultural Materials from Vegetable Oils
USDA	National Paper and Wood Recycling for New Materials
USDA	Bioplastics

Education and Human Resources

EPA	Environmentally Benign Materials and Processes
HHS	Programs of Excellence in Orthopaedic Biomaterials
NSF	Education and Human Resources

APPENDIX C

BREAKTHROUGH OPPORTUNITIES

This section suggests a vision of the future in which technologies can be improved dramatically through materials R&D, with significant benefits for the United States. These visions represent an attempt to organize the technological possibilities, for the purpose of guiding and inspiring current R&D efforts and future planning. Each page focuses on one economic sector or other thematic category of activities, identifying technologies for which materials R&D can contribute to resolving problems and improving applications cost-effectively. These are complex issues that must be addressed through collaborative efforts, and as such are ideal subjects for the AMPP.

Opportunities beckon in virtually every direction; major investments would be required to pursue all the applications listed, but success in even a few offers large payoffs. In many of the R&D areas, success could have an impact on multiple segments of U.S. society. The issues are discussed, however, in the context of the sector or category that would be the principal beneficiary of advances.

This compilation is based on extensive background information contained in a number of studies, including the MS&E report and the follow-up report on its implementation. Also useful was an extensive literature of detailed studies organized by the National Research Council's Solid State Sciences Committee and National Materials Advisory Board. All of these reports were based on broad input from the government, academia, and industry.

Agriculture

The United States produces 60 billion pounds of plastics a year. Though valuable, these synthetic polymers pose serious concerns. First, they are made primarily from a non-renewable resource, petroleum. Second, they do not degrade readily. This is unfortunate, because plastics comprise 7% of municipal solid waste by weight and 18% by volume. Plastics recycling is catching on, but the impact is not yet significant.

An opportunity to address these problems lies in development of advanced natural polymers, which would be made from renewable resources and also would biodegrade. Agriculture could provide the raw materials to make the advanced polymers and also would be among the sectors to benefit from these innovative materials.

Advanced natural polymers would be made from plant material, such as cellulose, starch, and proteins. (This could create new markets for agricultural commodities and for new individual crops.) The technical challenge is to develop polymers that achieve environmental goals while retaining low cost and other desirable properties. In addition, new technologies would be required to ensure that the polymers would degrade when desired, i.e. *after* use, perhaps in response to deliberate changes in light, temperature, or other conditions.

Biodegradable polymers would be useful in farming, as mulches, for example, or as sustained-release matrices for pest-control agents. Business and industry could use natural polymers in fast-food containers, loose packing materials for shipping, and other limited-use consumer products.

Defense

U.S. national security strategy must adapt to evolving global conditions, including the diffusion of advanced weapon technologies, and the potential for small, rapidly developing conflicts. Major R&D advances are needed to ensure that advanced high quality weapon systems can be designed and manufactured field ready in a timely and affordable manner.

New materials are a primary enabling technology for the DOD science and technology strategy, which relies on evolutionary weapons improvements as well as highly leveraged breakthroughs. Of the 21 technologies described in the DOD Critical Technologies Plan, 15 are materials-related.

Synthesis and processing of improved, lower-cost materials and components would enable advances in electronics and electro-optical technology, for example, and in composite structures. A major opportunity lies in development and application of theories, mathematical models, and supercomputer simulations that would reduce the need for expensive and time-consuming laboratory work.

Material yields and performance could be enhanced — and costs reduced — through development of “intelligent” processing technologies. The challenge is to develop non-intrusive sensors that respond to optical, chemical, temperature, and pressure changes. Such sensors also could be incorporated into “smart” materials and structures to improve weapon systems performance. Composites could be designed to change shape, for instance, in response to external stimuli. This could trigger automatic airfoil changes in aircraft or dampen vibrations of vehicles.

Energy

Inexpensive and readily available energy encourages economic prosperity. New materials and processing technologies could reduce energy costs significantly and conserve resources in the generation, transmission, storage, and use of energy.

Solar energy, for instance, could become economically competitive through the development of new or improved materials for more efficient photovoltaic technology. This opportunity could be realized through improved techniques for growing crystals, and by achievement of the precise material compositions that optimize electronic behavior. In addition, processing technologies might have to be developed to ensure that large-scale manufacturing of photovoltaic devices did not generate hazardous waste.

The development of superconducting electric power lines and devices could save \$5 billion annually in energy now lost due to resistance. This could be enabled through the development of flexible superconductors that operate at high temperatures and carry large electric currents, as well as methods for fabricating the materials and devices.

Additional billions of dollars could be saved by improving the efficiency of electric motors. Two-thirds of the nation's electric power is used in some sort of electric motor, from which 5-10% of the energy is typically lost. New materials with much greater magnetic strength could make this possible.

The safety and reliability of nuclear reactors also could be enhanced through materials research, specifically the study of radiation-induced swelling and embrittlement of metallic alloys and ceramics.

Environment

In the past, technological progress often resulted in serious harm to the environment, during the procurement or use of raw materials, manufacturing, or the disposal of waste or obsolete products. It is possible, however, to develop materials and processes that will satisfy both economic and environmental goals.

Materials science and technology offers the opportunity for the design of a materials cycle of indefinite re-use, such that all materials, processes, and products have life cycles that lead into other life cycles. This vision could be realized through the replacement of some commercial materials, the design of new materials that could be recycled endlessly, and the development of processes with few harmful waste products — or none at all.

Opportunities lie, for example, in the development of:

- aerogels that can be substituted for chlorofluorocarbons (CFCs) in foam plastics, now built into commercial and household refrigerators as thermal insulators;

- advanced sorbents that enhance removal of sulfur dioxide and nitrogen oxide emissions from the ambient environment and combustion furnaces;

- advanced processes to recycle or reuse waste ash from power plants and municipal waste incinerators; and

- plastic materials that biodegrade more rapidly in landfills.

Extraction/Production of Raw Materials

The economic advantages afforded by advanced materials are limited by the supply of raw materials and the environmental impact of the entire materials cycle, from extraction to disposal. Research on advanced processes could overcome both barriers.

In particular, new processes could revolutionize extraction and production of minerals, metals, and other natural materials, reducing costs and enabling the production of competitively priced commodities. This is a case where advanced processing of traditional materials could have a greater impact than development of new materials. Advanced processes that may offer advantages include:

- technologies to lower the cost of lightweight titanium for bulk applications, such as armor and ground vehicles;

- ultra-high-field magnetic separation of minerals from waste rock, a process that could eliminate the hazardous chemicals used in flotation extraction and also could reduce water consumption;

- biological treatment of surfaces to make metal extraction more selective;

- supercritical fluids as a means of chemically destroying toxic organic materials in waste streams; and

- unconventional energy sources (microwave and shock wave) and novel pretreatments (biological attack) as means of improving efficiency and reducing costs of crushing and grinding.

Other opportunities lie in the development of intelligent process control systems for complex operations, and computer models of physical and chemical separation.

Health

More than 11 million Americans have biomaterial implants, including artificial joints and heart valves, intra-ocular lens implants, pacemakers, infusion pumps, and dental and silicone implants.

Implants can be useful for five years or more, but most have a shorter effective life cycle because of wear, degradation, and incompatibility with human body tissues. Thus, most recipients, over time, lose some function and face implant replacement with the attendant discomfort and high costs.

Advanced materials and processes offer enormous potential for improving health care and reducing costs through the custom design of biomaterial implants that last a lifetime. For example:

Hybrid organs, made of a biomaterial lattice impregnated with genetically engineered cells, could provide more cost-effective delivery of drugs to treat chronic diseases such as diabetes.

Advanced artificial kidney membranes could improve the efficacy of dialysis, thus reducing risks and costs.

Artificial blood, made of oxygen-carrying biomaterials or hemoglobin encased in gas-permeable polymers, would eliminate the risk of contracting infections from blood transfusions.

Processes to fabricate and implant neural prostheses could help persons disabled by brain injuries.

Information/Communications

The burgeoning \$400 billion world electronics market is powered by integrated circuitry based on silicon. Advances in semiconductor materials, design, and processing technologies would enhance the viability of the U.S. electronics industry and open new market opportunities.

Silicon technology could be advanced through improvement of lithographic and processing techniques, such that more information could be contained on a chip. Such advances would improve the function/cost ratios of new circuit technologies and would lead to higher performance by means of smaller feature sizes, higher speeds, reduced power requirements, and higher production levels.

For the future, opportunities lie in developing new electronic and optoelectronic materials with properties superior to silicon, coupled with new processing methods. Compound semiconductors such as gallium arsenide have superior electronic and photonic properties; the challenge is to develop methods of producing high-quality integrated circuits in large amounts at a low cost. Other materials with novel or superior properties also need to be identified and developed. One attractive concept entails the monolithic integration of materials with complementary properties.

Enhanced performance of devices and circuits hinges on the development of more sophisticated methods of manipulating and controlling the structure of semiconductor materials. Other opportunities for enhancing performance lie in the development of advanced packaging and interconnect-device materials.

Infrastructure/Construction

Over the next 20 years, hundreds of billions of dollars will be spent to improve the nation's infrastructure — buildings, highways, bridges, airport runways, and other civil works. U.S. quality of life, industrial competitiveness, and productivity throughout the economy will depend on this far-reaching effort.

Advanced materials — new steel alloys and high-performance concrete — would extend the service life of virtually all structures, and thereby reduce costs and conserve energy and resources. The pay-off could be enormous. If the performance and durability of roads and bridges were improved by just 1 percent, for example, the direct savings could be as much as \$30 billion.

The construction industry now has the opportunity to exploit advanced steels and concrete for transportation and other projects. This would entail the development of relevant performance criteria and new specifications to assure compliance with — or revisions in — construction codes, standards, and practices.

Other opportunities for strengthening the infrastructure lie in improvement of fabrication processes, welding materials, design and analysis procedures, methods for predicting durability, and quality assurance systems. Computerized knowledge systems also would be useful, to distribute all of this information to architectural firms, materials suppliers, fabrication shops, and construction companies.

Transportation

One of the remaining positive contributors to the U.S. trade balance is overseas sales of commercial aircraft. On the negative side, the two largest components of the trade deficit are imported oil and automobiles. Advanced metal alloys, polymers, ceramics, and composites could help maintain U.S. leadership in the increasingly competitive world aircraft market, and at the same time reduce imports of oil and automobiles.

Specifically, materials advances would reduce the weight and increase the operating efficiency of future commercial aircraft. For example, advanced metals and composites, innovative, structural concepts, and improved fabrication processes would offer opportunities to make stronger and more cost-effective primary wing and fuselage structures.

Materials advances also would improve prospects for a new generation of supersonic commercial aircraft, with projected sales exceeding \$200 billion. Advanced ceramic and intermetallic composites would enable the development of affordable engines operating with reduced emissions and at acceptable noise levels, and advanced aluminum alloys and composites would help reduce airframe weight.

Other advances would enable the development of more fuel-efficient cars, thus reducing the need for imported oil and potentially increasing domestic car sales and exports. Opportunities lie in the development of processing and fabrication technologies for lightweight composites, high-strength metals, and tough ceramics, that will enable mass production and recycling at acceptable costs. This would permit manufacturers to produce lighter, more reliable, and more fuel efficient vehicles without sacrificing vehicle size and associated passenger comfort and safety.

Materials for the Future

This is an era of visionary research. Scientists now have the knowledge and tools to synthesize, understand, and exploit highly complex materials such as semiconductors, diamond and diamond-like materials, ceramics, novel structural materials, new alloys, polymers, and biomolecular materials. The potential technological breakthroughs can only be imagined, but the possibilities are illustrated by the recent synthesis of new superconducting materials, and the subsequent discovery of high-temperature superconducting phenomena previously thought to be impossible.

A research program aimed at fostering such innovation on a broader scale would pursue new synthetic approaches to new materials. This, combined with fundamental understanding of the novel phenomena exhibited by these materials and structures, would lead to advanced materials engineering breakthroughs and new technologies. Promising material systems to be developed include:

Nanostructured materials, a new class of metals, ceramics, semiconductors, and composites composed of structural units less than 100 nanometers (billionths of a meter) in size. The potential economic impact has been established by numerous reports of these materials exhibiting greatly enhanced electronic, magnetic, or mechanical properties.

“Intelligent” materials, which by virtue of their sensing capability will enable diverse technologies — including those used in processing and manufacturing — to respond dynamically to environmental stimuli. Optical, electronic, magnetic, or other types of materials may be used, depending on the desired application.

APPENDIX D

TECHNICAL BACKGROUND

The Science and Engineering of Materials

Materials science and technology encompasses a wide range of activities, from invention to application, from understanding to prediction, from the mundane to the esoteric. The field draws from a number of disciplines including physics, chemistry, biology, and engineering, and it builds upon its predecessors of metallurgy, ceramics, and polymer science.

All materials are, of course, composed of atoms. It is the specific combinations of atoms and their geometric arrangements that generate the myriad variations in material properties. Thus, control of structure and chemical composition, from the atomic to the macroscopic scale, is the key to designing a useful material. Such a material has well-defined and reliable properties and can be included safely in the list for product manufacture.

Properties determine how a material can be used. In the case of one traditional material, the brilliance and transparency of diamond give rise to its use as gemstones and optical coatings. Diamond is also the hardest material known and it resists corrosion as well, qualities that permit quite different applications in cutting tools or corrosion-resistant coatings. Depending on the application, the processing technique varies. To make a diamond gemstone, carbon is subjected to very high pressures and temperatures. To make a coating, scientists deposit thin layers of chemicals.

Properties depend on how a material is prepared. Advanced synthesis and processing techniques are crucial to the precise control of material structure and composition. Such techniques include **molecular beam epitaxy** (in which beams of atoms or molecules are sprayed into ultra-thin layers) and **rapid solidification** (cooling from liquid to solid by millions of degrees per second). Advances in processing of materials also can offer significant cost, quality, and environmental benefits.

Diverse types of studies have advanced understanding of how processing relates to material structure, how structure relates to properties, and how all three factors influence material performance in service. Analytical tools include **high-performance computers**, which are becoming powerful enough to predict material structures given only the types of atoms

present, and to simulate complex materials process technologies. Other tools include **optical, electron, X-ray, and scanning tunneling microscopes**, which provide images of structural levels, from the atomic through the macroscopic; **X-ray and neutron diffraction**, which map the spatial arrangement of atoms or molecules in a crystal; and a host of specialized tools to determine material properties at conditions approximating the low temperatures of outer space and the high pressures at the depths of the oceans.

Classes of Materials²

Biomaterials and biomolecular materials - include diverse materials that are compatible with human tissues and/or mimic biologic phenomena, as well as materials made from products of biological origin. Traditional materials in this category include dental fillings or crowns. Advanced materials are made from metals, ceramics, and polymers; applications range from artificial hearts to ultra-tough ceramic tank armor modeled on the molecular structure of abalone shells, and biodegradable plastics for use in packaging and fast-food containers.

Ceramics - include materials made from nonmetallic inorganic minerals. Ceramics are noted for their light weight, hardness, corrosion resistance, and tolerance to high temperatures. Spark plug insulators are a traditional example. Advanced ceramics are used for thermal coatings, and in structures of high-temperature engines.

Composites - include hybrids of two or more materials, usually reinforced ceramics, metals, or organic-matrix materials. These materials combine the most useful properties of each component. Fiberglass is a traditional composite, composed of glass fibers in an epoxy matrix. Advanced composites have both structural and non-structural applications and often are used in air and land vehicles.

² This list encompasses both traditional and advanced materials. For purposes of this report, "advanced" means the most recent evolutionary developments within a materials class or an applications category.

Electronic materials - include active materials that transmit signals by way of electrons, such as semiconductors. These materials are processed to form elements in circuitry. Current electronic technology is based on silicon. A newer material is gallium arsenide. This class also includes metals, ceramics, and polymers used in electronic wiring, interconnections, and packaging.

Magnetic materials - include metallic, organic, or ceramic-based materials in which electron “spins” are arranged in certain structures, such that the materials either attract or repel each other, depending on their orientation. When magnets are moved relative to a coil of wire, an electrical field is generated. Iron metal is a traditional magnet. The much stronger iron-rare earth alloy magnets are among the advanced varieties.

Metals - include ductile, strong structural materials and electrical conductors. Traditional metals include commodity alloys of elements such as iron and aluminum. Advanced metals tailored for specialty applications are predominantly lightweight aluminum alloys and high-temperature/high-strength intermetallics.

Optical/photonic materials - include materials that transmit light, materials used for light sources, such as lasers, and materials that are used to switch and modulate light. Glass in numerous forms, from window panes to optical fibers, continues to be an important material. Compound semiconductors, such as gallium arsenide, and complex structures made from such materials play an important role in modern communications systems.

Polymers - include a variety of large molecules consisting of long chains of repeated units. Polymers are noted for unique combinations of properties and have a range of applications, from plastic containers to liquid crystal displays. Plastic wrap is a traditional example; polyimides are advanced, high-temperature polymers used for electronic packaging and aircraft skins.

Superconducting materials - include materials that carry electrical currents with no resistance. Some metals and alloys exhibit this characteristic only at temperatures approaching absolute zero. Advanced varieties, including oxides, organics, and some intermetallics, superconduct at higher temperatures.

To obtain a copy of this document - send request to:
Committee on Industry and Technology/COMAT
c/o National Institute of Standards and Technology
Room B309, Materials Building
Gaithersburg MD 20899
[Phone: (301) 975-5655]




Grand Challenges: High Performance Computing and Communications

The FY 1992 U.S. Research and Development Program

A Report by the Committee on
Physical, Mathematical, and Engineering Sciences
Federal Coordinating Council for Science,
Engineering, and Technology
Office of Science and Technology Policy

To Supplement the President's Fiscal Year 1992 Budget



On the Cover:


1. Numerically modelled thunderstorm.
 2. Computational model of chemical carcinogen binding with DNA molecule.
 3. Visualization of structure of superconducting material.
 4. Simulation of acid rain pollutants over Ohio River basin.
 5. Aerodynamic characteristics of space vehicle: computer simulated versus wind tunnel data.
 6. Photo of wafer and multi-chip prototypes.
 7. Numerical simulation of fuel jet.
 8. Computer image of earth's biosphere components generated from satellite data.
- The images used in this report were produced by ongoing scientific projects in areas of the planned HPCC program. They were selected to illustrate the breadth of subject matter of the HPCC program, and are elaborated upon in Chapter 4. The U.S. map suggests how the National Research and Education Network supports geographically distributed collaborative research activities.

Grand Challenges: High Performance Computing and Communications

The FY 1992 U.S. Research and Development Program

A Report by the Committee on
Physical, Mathematical, and Engineering Sciences
Federal Coordinating Council for Science,
Engineering, and Technology
Office of Science and Technology Policy

To Supplement the President's Fiscal Year 1992 Budget



On the Cover:

1. Numerically modelled thunderstorm.
 2. Computational model of chemical carcinogen binding with DNA molecule.
 3. Visualization of structure of superconducting material.
 4. Simulation of acid rain pollutants over Ohio River basin.
 5. Aerodynamic characteristics of space vehicle: computer simulated versus wind tunnel data.
 6. Photo of wafer and multi-chip prototypes.
 7. Numerical simulation of fuel jet.
 8. Computer image of earth's biosphere components generated from satellite data.
- The images used in this report were produced by ongoing scientific projects in areas of the planned HPCC program. They were selected to illustrate the breadth of subject matter of the HPCC program, and are elaborated upon in Chapter 4. The U.S. map suggests how the National Research and Education Network supports geographically distributed collaborative research activities.

Office of Science and Technology Policy
Federal Coordinating Council for Science,
Engineering, and Technology
Committee on Physical, Mathematical, and
Engineering Sciences

Acting Chairman

Charles Herzfeld, Department of Defense

Members

Barry Williamson, Department of the Interior
Charles E. Hess, Department of Agriculture
Robert White, Department of Commerce
William Raub, Department of Health and Human Services
James Decker, Department of Energy
John B. Childers, Department of Education
Norine Noonan, Office of Management and Budget
Eugene Wong, Office of Science and Technology Policy
Arnold Aldrich, National Aeronautics and Space Administration
Erich Bretthauer, Environmental Protection Agency
Fredrick M. Bernthal, National Science Foundation

Executive Secretary

Jane Stutsman, National Science Foundation

FCCSET Directorate

Maryanne C. Bach, Executive Director
Charles H. Dickens, Senior Staff Associate
Jean Grace, Executive Assistant

High Performance Computing and Communications Working Group

David Nelson, Department of Energy, Chairman
Barry Boehm, Defense Advanced Research Projects Agency, Co-Chairman
Charles N. Brownstein, National Science Foundation, Co-Chairman
Lee Holcomb, National Aeronautics and Space Administration
Lawrence E. Brandt, National Science Foundation
John S. Cavallini, Department of Energy
Melvyn Ciment, National Science Foundation
Jack D. Fellows, Office of Management and Budget
Stephen M. Griffin, National Science Foundation
Paul E. Hunter, National Aeronautics and Space Administration
Gary M. Johnson, Department of Energy
Thomas A. Kitchens, Department of Energy
Norman H. Kreisman, Department of Energy
Albert T. Landberg, Jr., National Institute of Standards and Technology
Fred Scoresby Long, National Oceanic and Atmospheric Administration
Daniel R. Masys, National Library of Medicine
Joan H. Novak, Environmental Protection Agency
William L. Scherlis, Defense Advanced Research Projects Agency
Paul H. Smith, National Aeronautics and Space Administration
K. Speierman, National Security Agency
Stephen L. Squires, Defense Advanced Research Projects Agency
Stephen Wolff, National Science Foundation

Office of Science and Technology Policy
Federal Coordinating Council for Science,
Engineering, and Technology

Committee on Physical, Mathematical, and
Engineering Sciences

Acting Chairman

Charles Herzfeld, Department of Defense

Members

Barry Williamson, Department of the Interior
Charles E. Hess, Department of Agriculture
Robert White, Department of Commerce
William Raub, Department of Health and Human Services
James Decker, Department of Energy
John B. Childers, Department of Education
Norine Noonan, Office of Management and Budget
Eugene Wong, Office of Science and Technology Policy
Arnold Aldrich, National Aeronautics and Space Administration
Erich Bretthauer, Environmental Protection Agency
Fredrick M. Bernthal, National Science Foundation

Executive Secretary

Jane Stutsman, National Science Foundation

FCCSET Directorate

Maryanne C. Bach, Executive Director
Charles H. Dickens, Senior Staff Associate
Jean Grace, Executive Assistant

High Performance Computing and Communications Working Group

David Nelson, Department of Energy, Chairman
Barry Boehm, Defense Advanced Research Projects Agency, Co-Chairman
Charles N. Brownstein, National Science Foundation, Co-Chairman
Lee Holcomb, National Aeronautics and Space Administration
Lawrence E. Brandt, National Science Foundation
John S. Cavallini, Department of Energy
Melvyn Ciment, National Science Foundation
Jack D. Fellows, Office of Management and Budget
Stephen M. Griffin, National Science Foundation
Paul E. Hunter, National Aeronautics and Space Administration
Gary M. Johnson, Department of Energy
Thomas A. Kitchens, Department of Energy
Norman H. Kreisman, Department of Energy
Albert T. Landberg, Jr., National Institute of Standards and Technology
Fred Scoresby Long, National Oceanic and Atmospheric Administration
Daniel R. Masys, National Library of Medicine
Joan H. Novak, Environmental Protection Agency
William L. Scherlis, Defense Advanced Research Projects Agency
Paul H. Smith, National Aeronautics and Space Administration
K. Speierman, National Security Agency
Stephen L. Squires, Defense Advanced Research Projects Agency
Stephen Wolff, National Science Foundation

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF SCIENCE AND TECHNOLOGY POLICY
WASHINGTON, D. C. 20506

MEMBERS OF CONGRESS:

I am pleased to forward with this letter "Grand Challenges: High Performance Computing and Communications, The FY 1992 U. S. Research and Development Program," a report by the Committee on Physical, Mathematical, and Engineering Sciences of the Federal Coordinating Council for Science, Engineering, and Technology, a supplement to the President's Fiscal Year 1992 Budget.

The report presents an ambitious and well-coordinated research and development program designed to sustain and extend U.S. leadership in all advanced areas of computing and networking. The program not only provides a far-sighted vision for the underlying technologies but also gives recognition to the importance of both human resources and those applications that serve major national needs. This is a program of national investment that will bring both economic and social dividends.

The program is strategically related to other key components of the President's overall approach to challenges in science, technology, and education. It provides for the use of improved computational and communications technologies to contribute to more effective solutions of grand challenge problems.

The goal of the Federal High Performance Computing and Communications (HPCC) Program is to accelerate significantly the commercial availability and utilization of the next generation of high performance computers and networks. Recent advances offer the potential for a thousand-fold improvement in useful computing capability and a hundred-fold improvement in available computer communications capability by 1996. These advances will come through improvements in hardware and software. This increased capability will greatly expand the availability of these resources for research and education. It is my personal view, moreover, that the successful implementation of this program will lay the foundation for changes in education at all levels.

Several years of effort on the part of senior government, industry, and academic scientists and managers are reflected in this program. Acting Chairman Charles Herzfeld and his interagency committee members, associates, and staff are to be commended on the excellent work that is manifest in both the program and the report.


D. Allan Bromley
Director

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF SCIENCE AND TECHNOLOGY POLICY
WASHINGTON, D.C. 20506

MEMBERS OF CONGRESS:

I am pleased to forward with this letter "Grand Challenges: High Performance Computing and Communications, The FY 1992 U. S. Research and Development Program," a report by the Committee on Physical, Mathematical, and Engineering Sciences of the Federal Coordinating Council for Science, Engineering, and Technology, a supplement to the President's Fiscal Year 1992 Budget.

The report presents an ambitious and well-coordinated research and development program designed to sustain and extend U.S. leadership in all advanced areas of computing and networking. The program not only provides a far-sighted vision for the underlying technologies but also gives recognition to the importance of both human resources and those applications that serve major national needs. This is a program of national investment that will bring both economic and social dividends.

The program is strategically related to other key components of the President's overall approach to challenges in science, technology, and education. It provides for the use of improved computational and communications technologies to contribute to more effective solutions of grand challenge problems.

The goal of the Federal High Performance Computing and Communications (HPCC) Program is to accelerate significantly the commercial availability and utilization of the next generation of high performance computers and networks. Recent advances offer the potential for a thousand-fold improvement in useful computing capability and a hundred-fold improvement in available computer communications capability by 1996. These advances will come through improvements in hardware and software. This increased capability will greatly expand the availability of these resources for research and education. It is my personal view, moreover, that the successful implementation of the program will lay the foundation for changes in education at all levels.

Several years of effort on the part of senior government, industry, and academic scientists and managers are reflected in this program. Acting Chairman Charles Herzfeld and his interagency committee members, associates, and staff are to be commended on the excellent work that is manifest in both the program and the report.


D. Allan Bromley
Director

Table of Contents

	Page
Executive Summary	2
1. Program Goals and Overview	5
Introduction	
Needs and Benefits	
Program Description	
Goals	
Strategy	
Program Execution Strategy	
2. HPCC Program Components	12
High Performance Computing Systems	
Advanced Software Technology and Algorithms	
National Research and Education Network	
Basic Research and Human Resources	
3. Program Development and Agency Budgets	22
Program Planning	
Evaluation Criteria	
Agency Budgets	
Agency Program Descriptions	
Defense Advanced Research Projects Agency	
Department of Energy	
National Aeronautics and Space Administration	
National Science Foundation	
National Institute for Standards and Technology	
National Oceanic and Atmospheric Administration	
Environmental Protection Agency	
National Library of Medicine	
4. Grand Challenge and Supporting Technology Case Studies	37
Forecasting Severe Weather Events	
Cancer Genes	
Predicting New Superconductors	
Air Pollution	
Aerospace Vehicle Design	
Energy Conservation and Turbulent Combustion	
Microsystems Design and Packaging	
Earth's Biosphere	
High Speed Networks	
Education Using the NREN	
Glossary	56

To obtain a copy of this document — send request to:

Committee on Physical, Mathematical, and Engineering Sciences
c/o National Science Foundation
Computer and Information Science and Engineering
1800 G Street, N.W.
Washington, D.C. 20550

Table of Contents

	Page
Executive Summary	2
1. Program Goals and Overview	5
Introduction	
Needs and Benefits	
Program Description	
Goals	
Strategy	
Program Execution Strategy	
2. HPCC Program Components	12
High Performance Computing Systems	
Advanced Software Technology and Algorithms	
National Research and Education Network	
Basic Research and Human Resources	
3. Program Development and Agency Budgets	22
Program Planning	
Evaluation Criteria	
Agency Budgets	
Agency Program Descriptions	
Defense Advanced Research Projects Agency	
Department of Energy	
National Aeronautics and Space Administration	
National Science Foundation	
National Institute for Standards and Technology	
National Oceanic and Atmospheric Administration	
Environmental Protection Agency	
National Library of Medicine	
4. Grand Challenge and Supporting Technology Case Studies	37
Forecasting Severe Weather Events	
Cancer Genes	
Predicting New Superconductors	
Air Pollution	
Aerospace Vehicle Design	
Energy Conservation and Turbulent Combustion	
Microsystems Design and Packaging	
Earth's Biosphere	
High Speed Networks	
Education Using the NREN	
Glossary	56

To obtain a copy of this document — send request to:

Committee on Physical, Mathematical, and Engineering Sciences
c/o National Science Foundation
Computer and Information Science and Engineering
1800 G Street, N.W.
Washington, D.C. 20550

List of Figures

Figure 1. The High Performance Computing and Communications Program	Page 3
Figure 2. Performance Requirements for Grand Challenge Problems	7
Figure 3. Computer System Performance Trends for Grand Challenge Problems	13
Figure 4. Performance Improvement	15
Figure 5. NREN Applications	18
Figure 6. Evaluation Criteria for the HPCC Program	23
Figure 7. HPCC Budgets by Agency and Program Component	24
Figure 8. Agency Responsibilities	26

Grand Challenges: High Performance Computing and Communications

The FY 1992 U.S. Research and Development Program

Executive Summary

List of Figures

Figure 1. The High Performance Computing and Communications Program	Page 3
Figure 2. Performance Requirements for Grand Challenge Problems	7
Figure 3. Computer System Performance Trends for Grand Challenge Problems	13
Figure 4. Performance Improvement	15
Figure 5. NREN Applications	18
Figure 6. Evaluation Criteria for the HPCC Program	23
Figure 7. HPCC Budgets by Agency and Program Component	24
Figure 8. Agency Responsibilities	26

Grand Challenges: High Performance Computing and Communications

The FY 1992 U.S. Research and Development Program

Executive Summary

EXECUTIVE SUMMARY

- High performance computing and computer communications networks are becoming increasingly important to scientific advancement, economic competition, and national security. The technology is reaching the point of having a transforming effect on our society, industries, and educational institutions. The goal of the Federal High Performance Computing and Communications (HPCC) Program is to accelerate significantly the commercial availability and utilization of the next generation of high performance computers and networks in a manner consistent with the Strategic and Integrating Priorities shown in Figure 1.
- The HPCC Program is the result of several years of effort on the part of senior government, industry, and academic scientists and managers to design a research agenda to extend U.S. leadership in high performance computing and networking technologies.
- For FY 1992 the HPCC Program proposes to invest \$638 million in the four complementary and coordinated components shown in Figure 1. This investment represents a \$149 million, or 30%, increase over the FY 1991 enacted level.
- The HPCC Program is driven by the recognition that unprecedented computational power and capability is needed to investigate and understand a wide range of scientific and engineering "grand challenge" problems. These are fundamental problems whose solution is critical to national needs. Progress toward solution of these problems is essential to fulfilling many of the missions of the participating agencies. Examples of grand challenges addressed include: prediction of weather, climate, and global change; determination of molecular, atomic, and nuclear structure; understanding turbulence, pollution dispersion, and combustion systems; mapping the human genome and understanding the structure of biological macromolecules; improving research and education communications; understanding the nature of new materials; and problems applicable to national security needs.
- The HPCC Program nurtures the educational process at all levels by improving academic research and teaching capabilities. Advanced computing and computer communications technologies will accelerate the research process in all disciplines and enable educators to integrate new knowledge and methodologies directly into course curricula. Students at all levels will be drawn into learning and participating in a wide variety of research experiences in all components of this program.
- The FY 1992 Program and this document were developed by the HPCC Working Group under the direction of the Committee on Physical, Mathematical, and Engineering Sciences of the Federal Coordinating Council for Science, Engineering, and Technology.

Figure 1
The High Performance Computing
and Communications Program

Goals: Strategic Priorities Extend U.S. technological leadership in high performance computing and computer communications. Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment. Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.
Strategy: Integrating Priorities Support solutions to important scientific and technical challenges through a vigorous R&D effort. Reduce the uncertainties to industry for R&D and use of this technology through increased cooperation between government, industry, and universities and by the continued use of government and government-funded facilities as a prototype user for early commercial HPCC products. Support the underlying research, network, and computational infrastructures on which U.S. high performance computing technology is based. Support the U.S. human resource base to meet the needs of industry, universities, and government.
Program Components High Performance Computing Systems Research for Future Generations of Computing Systems System Design Tools Advanced Prototype Systems Evaluation of Early Systems Advanced Software Technology and Algorithms Software Support for Grand Challenges Software Components and Tools Computational Techniques High Performance Computing Research Centers National Research and Education Network Interagency Interim NREN Gigabits Research and Development Basic Research and Human Resources Basic Research Research Participation and Training Infrastructure Education, Training, and Curriculum

EXECUTIVE SUMMARY

- High performance computing and computer communications networks are becoming increasingly important to scientific advancement, economic competition, and national security. The technology is reaching the point of having a transforming effect on our society, industries, and educational institutions. The goal of the Federal High Performance Computing and Communications (HPCC) Program is to accelerate significantly the commercial availability and utilization of the next generation of high performance computers and networks in a manner consistent with the Strategic and Integrating Priorities shown in Figure 1.
- The HPCC Program is the result of several years of effort on the part of senior government, industry, and academic scientists and managers to design a research agenda to extend U.S. leadership in high performance computing and networking technologies.
- For FY 1992 the HPCC Program proposes to invest \$638 million in the four complementary and coordinated components shown in Figure 1. This investment represents a \$149 million, or 30%, increase over the FY 1991 enacted level.
- The HPCC Program is driven by the recognition that unprecedented computational power and capability is needed to investigate and understand a wide range of scientific and engineering "grand challenge" problems. These are fundamental problems whose solution is critical to national needs. Progress toward solution of these problems is essential to fulfilling many of the missions of the participating agencies. Examples of grand challenges addressed include: prediction of weather, climate, and global change; determination of molecular, atomic, and nuclear structure; understanding turbulence, pollution dispersion, and combustion systems; mapping the human genome and understanding the structure of biological macromolecules; improving research and education communications; understanding the nature of new materials; and problems applicable to national security needs.
- The HPCC Program nurtures the educational process at all levels by improving academic research and teaching capabilities. Advanced computing and computer communications technologies will accelerate the research process in all disciplines and enable educators to integrate new knowledge and methodologies directly into course curricula. Students at all levels will be drawn into learning and participating in a wide variety of research experiences in all components of this program.
- The FY 1992 Program and this document were developed by the HPCC Working Group under the direction of the Committee on Physical, Mathematical, and Engineering Sciences of the Federal Coordinating Council for Science, Engineering, and Technology.

Figure 1 The High Performance Computing and Communications Program

Goals: Strategic Priorities Extend U.S. technological leadership in high performance computing and computer communications. Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment. Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.	Strategy: Integrating Priorities Support solutions to important scientific and technical challenges through a vigorous R&D effort. Reduce the uncertainties to industry for R&D and use of this technology through increased cooperation between government, industry, and universities and by the continued use of government and government-funded facilities as a prototype user for early commercial HPCC products. Support the underlying research, network, and computational infrastructures on which U.S. high performance computing technology is based. Support the U.S. human resource base to meet the needs of industry, universities, and government.	Program Components High Performance Computing Systems Research for Future Generations of Computing Systems System Design Tools Advanced Prototype Systems Evaluation of Early Systems Advanced Software Technology and Algorithms Software Support for Grand Challenges Software Components and Tools Computational Techniques High Performance Computing Research Centers National Research and Education Network Interagency Interim NREN Gigabits Research and Development Basic Research and Human Resources Basic Research Research Participation and Training Infrastructure Education, Training, and Curriculum
--	---	--

1. PROGRAM GOALS AND OVERVIEW

Introduction

High performance computing (HPC) is emerging as a powerful technology for industrial design and manufacturing, scientific research, communications, and information management. A robust U.S. high performance computing and computer communications capability contributes to leadership in critical technology and national security areas. Improved computational and communications technologies contribute to more effective approaches to problem solving, new products and services, and enhanced national competitiveness across broad sectors of the economy.

Recent advances offer the potential for a thousand-fold improvement in useful computing capability and a hundred-fold improvement in available computer communications capability by 1996. Based on several years of planning, under the auspices of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), Federal agencies and the technical community have developed the Federal High Performance Computing and Communications (HPCC) Program to realize this potential and to meet the challenges of advancing computing and associated communications technology and practices. Agencies have realigned and enhanced their high performance computing research and development programs, coordinated their activities with other agencies, and shared common resources to develop the program presented in this document.

Needs and Benefits

High performance computing has become a vital enabling force in the conduct of science and engineering research over the past three decades. Computational science and engineering has joined, and in some areas displaced, the traditional methods of theory and experiment. For example, in the design of commercial aircraft, many engineering issues are resolved through computer simulation rather than through costly wind tunnel experiments. This trend has been powered by computing hardware and software, computational methodologies and algorithms, availability and access to high performance computing systems and infrastructure, and the growth of a trained pool of scientists and engineers. This process has been nurtured by Federal investment in advanced research, agency supercomputer centers, and national networks through DARPA, DOE, NASA, NSA, and NSF. These facilities have contributed to national mission areas such as energy, space, health, defense, environment, weather, and basic science and

1. PROGRAM GOALS AND OVERVIEW

Introduction

High performance computing (HPC) is emerging as a powerful technology for industrial design and manufacturing, scientific research, communications, and information management. A robust U.S. high performance computing and computer communications capability contributes to leadership in critical technology and national security areas. Improved computational and communications technologies contribute to more effective approaches to problem solving, new products and services, and enhanced national competitiveness across broad sectors of the economy.

Recent advances offer the potential for a thousand-fold improvement in useful computing capability and a hundred-fold improvement in available computer communications capability by 1996. Based on several years of planning, under the auspices of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), Federal agencies and the technical community have developed the Federal High Performance Computing and Communications (HPCC) Program to realize this potential and to meet the challenges of advancing computing and associated communications technology and practices. Agencies have realigned and enhanced their high performance computing research and development programs, coordinated their activities with other agencies, and shared common resources to develop the program presented in this document.

Needs and Benefits

High performance computing has become a vital enabling force in the conduct of science and engineering research over the past three decades. Computational science and engineering has joined, and in some areas displaced, the traditional methods of theory and experiment. For example, in the design of commercial aircraft, many engineering issues are resolved through computer simulation rather than through costly wind tunnel experiments. This trend has been powered by computing hardware and software, computational methodologies and algorithms, availability and access to high performance computing systems and infrastructure, and the growth of a trained pool of scientists and engineers. This process has been nurtured by Federal investment in advanced research, agency supercomputer centers, and national networks through DARPA, DOE, NASA, NSA, and NSF. These facilities have contributed to national mission areas such as energy, space, health, defense, environment, weather, and basic science and


technology that could not be effectively addressed without the use of such advanced facilities.

High performance computing technology is knowledge and innovation intensive. Its development and use engages the entire scientific and engineering community. Building upon fundamental research of the early 1980's, a new computing technology of scalable parallel processing computers emerged. By the mid-1990's, this innovative approach to high performance computing systems promises to achieve sustained performance improvements of a thousand-fold compared to current systems.

In a growing number of science and technology fields, progress and productivity in modern research are increasingly dependent on the close interaction of people located in distant places, sharing and accessing computational resources across networks. Although the U.S. is the world leader in most of the critical aspects of computing technology, this lead is being challenged.

The Federal HPCC Program is a strategic Federal investment in the frontiers of computing and computer communications technologies and is formulated to satisfy national needs from a variety of perspectives including: technology, science applications, human resources, and technology transition. Needs are derived from the agency missions and based on the underlying science, engineering, and technology base required to carry out these missions. Many of these mission needs are related to solving very intensive large scale computing problems. These fundamental problems often cut across various agencies and missions and are called grand challenge problems (Figure 2).

The industrial and academic sectors provide major sources of innovation, cost effective development, and support of information technologies and their application to grand challenge problems. As these technologies are developed, the results support the Federal agency missions and become available nationally. The program provides for development of these revolutionary technologies within a framework of a partnership among government, industry, and academe and allows for rapid transition of laboratory results into new products that will then be applied within the program.



technology that could not be effectively addressed without the use of such advanced facilities.


High performance computing technology is knowledge and innovation intensive. Its development and use engages the entire scientific and engineering community. Building upon fundamental research of the early 1980's, a new computing technology of scalable parallel processing computers emerged. By the mid-1990's, this innovative approach to high performance computing systems promises to achieve sustained performance improvements of a thousand-fold compared to current systems.

In a growing number of science and technology fields, progress and productivity in modern research are increasingly dependent on the close interaction of people located in distant places, sharing and accessing computational resources across networks. Although the U.S. is the world leader in most of the critical aspects of computing technology, this lead is being challenged.

The Federal HPCC Program is a strategic Federal investment in the frontiers of computing and computer communications technologies and is formulated to satisfy national needs from a variety of perspectives including: technology, science applications, human resources, and technology transition. Needs are derived from the agency missions and based on the underlying science, engineering, and technology base required to carry out these missions. Many of these mission needs are related to solving very intensive large scale computing problems. These fundamental problems often cut across various agencies and missions and are called grand challenge problems (Figure 2).

The industrial and academic sectors provide major sources of innovation, cost effective development, and support of information technologies and their application to grand challenge problems. As these technologies are developed, the results support the Federal agency missions and become available nationally. The program provides for development of these revolutionary technologies within a framework of a partnership among government, industry, and academe and allows for rapid transition of laboratory results into new products that will then be applied within the program.

Figure 2
Performance Requirements
for Grand Challenge Problems



Program Description

The Program consists of four integrated components representing the key areas of high performance computing and communications:

High Performance Computing Systems (HPCS)—the development of the underlying technology required for scalable parallel computing systems capable of sustaining trillions of operations per second on large problems.

Advanced Software Technology and Algorithms (ASTA)—the development of generic software technology and algorithms for grand challenge research applications to realize the performance potential of high performance computing systems in a networked environment.

National Research and Education Network (NREN)—the development of a national high speed network to provide distributed computing capability to research and educational institutions and to further advanced research on very high speed networks and applications.

Basic Research and Human Resources (BRHR)—support for individual investigator and multidisciplinary long term research drawn from diverse disciplines, including computer science, computer engineering, and computational science and engineering; initiation of activities to significantly increase the pool of trained personnel; and support for efforts leading to accelerated technology transition.

Advances in high performance computing enable advances in almost every other science and engineering discipline. There is a complex web of research interdependencies among the four components, and each area contributes to progress in other areas. Because of these dependencies, achieving and maintaining balance between the research components is a primary goal and the most important priority in the current context and environment. The HPCC Program is designed to provide balanced support both for technology areas including components, systems, software, and algorithms, and for applications, infrastructure, and human resources to achieve rapid overall research progress and productivity.

The component activities are planned to produce a succession of intermediate benefits on the way to meeting the long range programmatic goals. The HPCC Program builds on Federal programs already in place, providing additional resources in selected areas. Computational science and engineering grand challenges as illustrated in Figure 2 are the focal points for these efforts.

Goals

The goals of the High Performance Computing and Communications Program are to:

Extend U.S. technological leadership in high performance computing and computer communications.

Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment.

Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.

These goals will be realized by achieving: computational performance of one trillion operations per second (10^{12} ops, or teraops) on a wide range of important applications; development of associated system software, tools, and improved algorithms for a wide range of problems; a national research network capable of one billion bits per second (10^9 bits, or gigabits); sufficient production of Ph.D.'s and other trained professionals per year in computational science and engineering to enable effective use and application of these new technologies.

Strategy

The goals will be met through coordinated government, industry, and university collaboration to:

Support solutions to important scientific and technical challenges through a vigorous R&D effort.

Reduce the uncertainties to industry for R&D and use of this technology through increased cooperation between government, industry, and universities and by the continued use of government and government-funded facilities as a prototype user for early commercial HPCC products.

Support the underlying research, network, and computational infrastructures on which U.S. high performance computing technology is based.

Support the U.S. human resource base to meet the needs of industry, universities, and government.

At the program component level the strategy will: exploit and extend scalable parallelism and engage in intensive software development in HPCS; use common requirements of the grand challenges to foster HPC

Program Description

The Program consists of four integrated components representing the key areas of high performance computing and communications:

High Performance Computing Systems (HPCS)—the development of the underlying technology required for scalable parallel computing systems capable of sustaining trillions of operations per second on large problems.

Advanced Software Technology and Algorithms (ASTA)—the development of generic software technology and algorithms for grand challenge research applications to realize the performance potential of high performance computing systems in a networked environment.

National Research and Education Network (NREN)—the development of a national high speed network to provide distributed computing capability to research and educational institutions and to further advanced research on very high speed networks and applications.

Basic Research and Human Resources (BRHR)—support for individual investigator and multidisciplinary long term research drawn from diverse disciplines, including computer science, computer engineering, and computational science and engineering; initiation of activities to significantly increase the pool of trained personnel; and support for efforts leading to accelerated technology transition.

Advances in high performance computing enable advances in almost every other science and engineering discipline. There is a complex web of research interdependencies among the four components, and each area contributes to progress in other areas. Because of these dependencies, achieving and maintaining balance between the research components is a primary goal and the most important priority in the current context and environment. The HPCC Program is designed to provide balanced support both for technology areas including components, systems, software, and algorithms, and for applications, infrastructure, and human resources to achieve rapid overall research progress and productivity.

The component activities are planned to produce a succession of intermediate benefits on the way to meeting the long range programmatic goals. The HPCC Program builds on Federal programs already in place, providing additional resources in selected areas. Computational science and engineering grand challenges as illustrated in Figure 2 are the focal points for these efforts.

Goals

The goals of the High Performance Computing and Communications Program are to:

Extend U.S. technological leadership in high performance computing and computer communications.

Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment.

Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.

These goals will be realized by achieving: computational performance of one trillion operations per second (10^{12} ops, or teraops) on a wide range of important applications; development of associated system software, tools, and improved algorithms for a wide range of problems; a national research network capable of one billion bits per second (10^9 bits, or gigabits); sufficient production of Ph.D.'s and other trained professionals per year in computational science and engineering to enable effective use and application of these new technologies.

Strategy

The goals will be met through coordinated government, industry, and university collaboration to:

Support solutions to important scientific and technical challenges through a vigorous R&D effort.

Reduce the uncertainties to industry for R&D and use of this technology through increased cooperation between government, industry, and universities and by the continued use of government and government-funded facilities as a prototype user for early commercial HPCC products.

Support the underlying research, network, and computational infrastructures on which U.S. high performance computing technology is based.

Support the U.S. human resource base to meet the needs of industry, universities, and government.

At the program component level the strategy will: exploit and extend scalable parallelism and engage in intensive software development in HPCS; use common requirements of the grand challenges to foster HPC

software progress in ASTA to strengthen HPC software development and coordination; evolve from the current Internet network to the NREN using a series of testbed systems; and strengthen academic activities in computer science and computational science and engineering as part of BRHR.

Program Execution Strategy

The strategy is based on the strengths of partnerships among the Federal agencies and other organizations. Major portions of the program will be cost-shared and leveraged by the participation of industry and universities. This general approach operates well today and provides strong evidence that the HPCC Program can be successful in the future. The specific elements of the approach are to:

Create a balanced, critical-mass program. The program must achieve sufficient scope and balance among the components. A technology program that created extremely fast processors without comparable memory, input-output, and mass storage systems would not succeed. Neither would a program that created powerful computers without adequate software, network access, and capable people. Similarly, a program that created only high performance networks would not satisfy the increased performance requirements needed for grand challenges. The HPCC Program must operate at a sufficient scale and coverage of technology areas that the new technologies can be effectively applied to grand challenge problems with acceptable levels of risk. The HPCC Program achieves balance by the extensive participation of experienced users, applications developers, and researchers in the HPCC disciplines throughout the design, development, and implementation process.

Build on agency strengths. The strategy builds on agency strengths by giving appropriate agencies the responsibility to coordinate activities for areas of demonstrated capability. It also ensures that the strengths of the other agencies are included by integrating their participation in various task areas. DOE, NASA, NSF, DOC (NOAA, NIST), and DOD (DARPA, NSA) have decades of experience in applying the world's most powerful computers, and thus provide a valuable perspective on high performance computing requirements and applications. NSF has the demonstrated technical and operational expertise needed for deploying high performance national networks in the research community, and is uniquely positioned to support basic research in computer science, computational science, and other scientific areas that can benefit from high performance computing in

interdisciplinary programs. DARPA, having pioneered technologies for high performance computing systems and microelectronic components, has a strong existing research program in teraops computing and gigabit networking, and offers rapid technology transition into commercial products in support of DOD science, technology, and applications base. EPA, NIH, and NOAA provide complementary HPCC requirements perspectives and application bases. NIST has extensive experience in HPC systems and networks instrumentation and evaluation, and provides a means for standards development.

Accelerate Technology Transfer. The transfer of technology from research to development and to application can be a very slow process due to a number of barriers to the use of new technology: high initial cost, inadequate and user-unfriendly software and systems, and lack of standards. The HPCC Program relies upon substantial industry participation to overcome these barriers and yield the benefits derived from moving new technologies to industry. The strategy accelerates technology transition by using a participative development process for each of the task areas, and by stimulating the growth of shared knowledge and capabilities. An example of this is the creation of network-accessible repositories of scalable HPC software and associated user groups to provide usage feedback and improvement of the HPC software base.

Overcome Barriers. To overcome high costs of creating successive generations of high performance computers, the program will emphasize scalable computer designs. Scalable computing and networking technologies enable exploratory use with small, lower cost prototype systems needed to eventually support the acquisition of larger systems. The development of user-friendly software and systems as part of the investment in the HPC software base is a major, integral part of the program. The central role of mission agencies, and the broad academic scientific research community involvement, ensures that the hardware, software, and networking technology developed will be responsive to user needs. A strategy of cooperative government, industry, and university activities is used to manage and coordinate the coupling of these sectors to achieve maximum synergy.

software progress in ASTA to strengthen HPC software development and coordination; evolve from the current Internet network to the NREN using a series of testbed systems; and strengthen academic activities in computer science and computational science and engineering as part of BRHR.

Program Execution Strategy

The strategy is based on the strengths of partnerships among the Federal agencies and other organizations. Major portions of the program will be cost-shared and leveraged by the participation of industry and universities. This general approach operates well today and provides strong evidence that the HPCC Program can be successful in the future. The specific elements of the approach are to:

Create a balanced, critical-mass program. The program must achieve sufficient scope and balance among the components. A technology program that created extremely fast processors without comparable memory, input-output, and mass storage systems would not succeed. Neither would a program that created powerful computers without adequate software, network access, and capable people. Similarly, a program that created only high performance networks would not satisfy the increased performance requirements needed for grand challenges. The HPCC Program must operate at a sufficient scale and coverage of technology areas that the new technologies can be effectively applied to grand challenge problems with acceptable levels of risk. The HPCC Program achieves balance by the extensive participation of experienced users, applications developers, and researchers in the HPCC disciplines throughout the design, development, and implementation process.

Build on agency strengths. The strategy builds on agency strengths by giving appropriate agencies the responsibility to coordinate activities for areas of demonstrated capability. It also ensures that the strengths of the other agencies are included by integrating their participation in various task areas. DOE, NASA, NSF, DOC (NOAA, NIST), and DOD (DARPA, NSA) have decades of experience in applying the world's most powerful computers, and thus provide a valuable perspective on high performance computing requirements and applications. NSF has the demonstrated technical and operational expertise needed for deploying high performance national networks in the research community, and is uniquely positioned to support basic research in computer science, computational science, and other scientific areas that can benefit from high performance computing in

interdisciplinary programs. DARPA, having pioneered technologies for high performance computing systems and microelectronic components, has a strong existing research program in teraops computing and gigabit networking, and offers rapid technology transition into commercial products in support of DOD science, technology, and applications base. EPA, NIH, and NOAA provide complementary HPCC requirements perspectives and application bases. NIST has extensive experience in HPC systems and networks instrumentation and evaluation, and provides a means for standards development.

Accelerate Technology Transfer. The transfer of technology from research to development and to application can be a very slow process due to a number of barriers to the use of new technology: high initial cost, inadequate and user-unfriendly software and systems, and lack of standards. The HPCC Program relies upon substantial industry participation to overcome these barriers and yield the benefits derived from moving new technologies to industry. The strategy accelerates technology transition by using a participative development process for each of the task areas, and by stimulating the growth of shared knowledge and capabilities. An example of this is the creation of network-accessible repositories of scalable HPC software and associated user groups to provide usage feedback and improvement of the HPC software base.

Overcome Barriers. To overcome high costs of creating successive generations of high performance computers, the program will emphasize scalable computer designs. Scalable computing and networking technologies enable exploratory use with small, lower cost prototype systems needed to eventually support the acquisition of larger systems. The development of user-friendly software and systems as part of the investment in the HPC software base is a major, integral part of the program. The central role of mission agencies, and the broad academic scientific research community involvement, ensures that the hardware, software, and networking technology developed will be responsive to user needs. A strategy of cooperative government, industry, and university activities is used to manage and coordinate the coupling of these sectors to achieve maximum synergy.

2. HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM COMPONENTS

The HPCC Program is composed of four integrated and coordinated components that are designed to enhance scientific productivity and support long term agency needs. The emerging scientific computing environment is that of advanced workstations with high resolution color displays connected to a high speed computer communications network and high performance computing resources. The regional and national networks provide a means to gain access to additional high performance computing systems and research resources. Realizing the full potential of these teraops computing and gigabits networking systems will require advanced software technology and algorithms and people educated and trained to use these tools and resources in this dynamic field.

High Performance Computing Systems (HPCS)


The HPCS component produces scalable parallel computing systems through the development of prototype systems. The program is designed to attack computational science problems by developing innovative systems that will provide a one-hundred to one-thousand-fold increase of sustained computing capability over machines that follow the more conventional design evolution path (See Figure 3). DARPA will coordinate the research and development effort that will produce teraops systems.

The program is structured to focus on technological challenges in the earliest stages of the product development cycle. Critical underlying technologies are developed in prototype form along with associated design tools. This allows empirical evaluation of alternative solutions as the prototype systems mature. Evaluation is performed throughout the development cycle, with experimental results being fed back into the design process to refine successive generations of systems. There is risk inherent in creating new technologies, and each project will be managed according to its proximity to commercial introduction. Larger projects which are close to yielding commercial products are performed on a cost shared basis with industry.

HPCS is composed of the four subcomponents shown below.

Research for Future Generations of Computing Systems. This activity produces the underlying component, packaging, and scaling concepts. These projects ensure that the required advanced technologies will be available for the new systems while providing a foundation for the more powerful systems that need to follow.

Figure 3 Computer System Performance Trends for Grand Challenge Problems
– Power and Time to System Development –



2. HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM COMPONENTS

The HPCC Program is composed of four integrated and coordinated components that are designed to enhance scientific productivity and support long term agency needs. The emerging scientific computing environment is that of advanced workstations with high resolution color displays connected to a high speed computer communications network and high performance computing resources. The regional and national networks provide a means to gain access to additional high performance computing systems and research resources. Realizing the full potential of these teraops computing and gigabits networking systems will require advanced software technology and algorithms and people educated and trained to use these tools and resources in this dynamic field.

High Performance Computing Systems (HPCS)


The HPCS component produces scalable parallel computing systems through the development of prototype systems. The program is designed to attack computational science problems by developing innovative systems that will provide a one-hundred to one-thousand-fold increase of sustained computing capability over machines that follow the more conventional design evolution path (See Figure 3). DARPA will coordinate the research and development effort that will produce teraops systems.

The program is structured to focus on technological challenges in the earliest stages of the product development cycle. Critical underlying technologies are developed in prototype form along with associated design tools. This allows empirical evaluation of alternative solutions as the prototype systems mature. Evaluation is performed throughout the development cycle, with experimental results being fed back into the design process to refine successive generations of systems. There is risk inherent in creating new technologies, and each project will be managed according to its proximity to commercial introduction. Larger projects which are close to yielding commercial products are performed on a cost shared basis with industry.

HPCS is composed of the four subcomponents shown below.

Research for Future Generations of Computing Systems. This activity produces the underlying component, packaging, and scaling concepts. These projects ensure that the required advanced technologies will be available for the new systems while providing a foundation for the more powerful systems that need to follow.

Figure 3 Computer System Performance Trends for Grand Challenge Problems – Power and Time to System Development –



System Design Tools. This activity develops computer aided design tools and frameworks for enabling multiple tools to work together to enable design, analysis, simulation, and testing of systems components. The tools will enable rapid prototyping of new system concepts.


Advanced Prototype Systems. This activity consists of focused development of experimental systems that are designed and developed on a cost shared basis with industry. The 100 gigapops systems provide a basis for the teraops systems. New models of computation will be introduced and successive generations of computer systems, along with systems software, will encompass broader grand challenge domains. Modular technologies will enable a wide variety of system configurations using common components. Systems with 100 gigapops sustained performance will be developed by the early 1990's. The teraops level will be reached by the mid 1990's.

Evaluation of Early Systems. Experimental systems will be placed at sites with high levels of expertise to provide feedback to systems architects and software designers. Performance evaluation criteria for systems and results of evaluation will be made widely available. Because of scalability, early systems can be acquired at smaller scales to evaluate their potential performance. As noted below, the ASTA component will support, on a cost shared basis, the acquisition of large scale systems for experimental use in grand challenge applications.

Advanced Software Technology and Algorithms (ASTA)

Dramatic improvements in algorithm design and software technology are essential to achieving sustained teraflops computing system performance. Improvements in hardware, especially for scalable parallel systems, must be matched with new and innovative algorithms and software to enable researchers to expand the boundaries of computational capabilities. In Figure 4, this point is illustrated for broad classes of scientific computing problems by showing that comparable improvements in throughput have resulted from advances in computational methods as from improved hardware technology. In this case, hardware improvements yielded a speed-up of 1000 over 20 years, while software and algorithms improvements yielded a speed-up of 3000 over the same period, for an aggregate speed-up of more than a million.

Figure 4
Performance Improvement
for Scientific Computing Problems



The ASTA component of the Federal High Performance Computing and Communications Program has three goals:

- Enable solution of grand challenge application problems in science and engineering.
- Improve system user-friendliness, reliability, and software productivity.
- Use experience gained on leading edge applications to help guide future software efforts.

System Design Tools. This activity develops computer aided design tools and frameworks for enabling multiple tools to work together to enable design, analysis, simulation, and testing of systems components. The tools will enable rapid prototyping of new system concepts.


Advanced Prototype Systems. This activity consists of focused development of experimental systems that are designed and developed on a cost shared basis with industry. The 100 gigaflops systems provide a basis for the teraflops systems. New models of computation will be introduced and successive generations of computer systems, along with systems software, will encompass broader grand challenge domains. Modular technologies will enable a wide variety of system configurations using common components. Systems with 100 gigaflops sustained performance will be developed by the early 1990's. The teraflops level will be reached by the mid 1990's.

Evaluation of Early Systems. Experimental systems will be placed at sites with high levels of expertise to provide feedback to systems architects and software designers. Performance evaluation criteria for systems and results of evaluation will be made widely available. Because of scalability, early systems can be acquired at smaller scales to evaluate their potential performance. As noted below, the ASTA component will support, on a cost shared basis, the acquisition of large scale systems for experimental use in grand challenge applications.

Advanced Software Technology and Algorithms (ASTA)

Dramatic improvements in algorithm design and software technology are essential to achieving sustained teraflops computing system performance. Improvements in hardware, especially for scalable parallel systems, must be matched with new and innovative algorithms and software to enable researchers to expand the boundaries of computational capabilities. In Figure 4, this point is illustrated for broad classes of scientific computing problems by showing that comparable improvements in throughput have resulted from advances in computational methods as from improved hardware technology. In this case, hardware improvements yielded a speed-up of 1000 over 20 years, while software and algorithms improvements yielded a speed-up of 3000 over the same period, for an aggregate speed-up of more than a million.

Figure 4
Performance Improvement
for Scientific Computing Problems



The ASTA component of the Federal High Performance Computing and Communications Program has three goals:

- Enable solution of grand challenge application problems in science and engineering.
- Improve system user-friendliness, reliability, and software productivity.
- Use experience gained on leading edge applications to help guide future software efforts.

The emphasis is on the development of advanced algorithms and software technology required to address applications problems on the scale of grand challenges. The ASTA component is comprised of four subcomponents, as shown below.

Software Support for Grand Challenges. The HPCC Program is designed to demonstrate new computing technology capabilities by confronting an expanding number of grand challenge applications problems. The goal is to reduce the risks to researchers inherent in adopting innovative high performance computing technologies. Grand challenge application problems will be selected based upon their scientific importance, the potential for cost sharing with sources directly concerned with the specific scientific and engineering applications, and the potential for leveraging across sectors.

Software Components and Tools. The multidisciplinary research teams will have common needs in software technology and programming environments, advanced compiler technology, optimization and parallelization tools, interoperability and data management, visualization, debugging and analysis, and performance measurement. Advances in these generic software areas will minimize the need for specialized researchers to master advanced, complex computing science skills. Coordination of the development of advanced software technology and algorithms among the participating agencies will be important to ensure effective and efficient use of resources. A particular focus of this element will be to develop advanced software applications that exploit the NREN using distributed file systems and national software libraries.

Computational Techniques. The focus of the HPCC Program on scalable parallel computing systems dictates that significant advances in computational techniques will be needed. The design and theory of algorithms are as important as hardware or networking improvements in reaching teraflops computational performance. Research in computational techniques will include parallel algorithms, numerical and mathematical analyses, parallel languages, computational models, and program refinement techniques. Higher level languages will be developed to enable computational scientists to work directly at the level of the abstract computational problem being addressed and to more easily explore specific implementation approaches.

High Performance Computing Research Centers. High performance computing research centers and testbed facilities will provide a large number of researchers access via the NREN to both conventional and innovative supercomputing architectures. Many of these centers will introduce innovative architectures and will be the focus of technology transfer activities and form the base for training new generations of computer and computational scientists.

National Research and Education Network (NREN)

The NREN component of the HPCC Program dramatically expands and enhances the U.S. portion of an existing worldwide infrastructure of interconnected computer networks called the Internet. A substantial fraction of the domestic Internet is supported and loosely coordinated by Federal agencies, principally DARPA, DOE, NASA, and NSF.

Collaboration among scientists is an important and integral facet of the U.S. research environment. It can be greatly improved by increasing the level of network connectivity and by introducing new capabilities into the existing infrastructure. The NREN design will not only address broad network connectivity, but will also provide the basis for necessary higher level capabilities and services.

Many educational institutions, government laboratories, and industrial research facilities are currently connected to the Internet. Yet, it still falls short of a widespread, uniform, and high performance national infrastructure. In order to satisfy the HPCC Program goals, the NREN must not merely provide network access to research and educational institutions at all levels and locations, it must also deliver new capabilities. Some of these, such as distance learning, may initially be extensions of current technology. All capabilities will benefit from, and many will be enabled by, a program of research into very high speed technology. This technology is needed to support access to digital libraries, large scale distributed computing resources, as well as to perform computationally intensive applications that require real time visualization of modeling and simulation results, rapid interrogation and retrieval of scientific data from specialized data bases, remote control of experiments and simulations, and teleconferencing.

In addition to serving the needs of the scientific and research communities, the NREN will provide valuable experience necessary for the successful development of a broader, privately-operated national information infrastructure. Such an infrastructure would allow consumers, businesses, and schools and government at all levels to share quality information and entertainment when and where they want it at a reasonable cost.

Applications conducted over a computer network vary in their flow of information from steady, as in a bulk file transfer, to "bursty," as in human-computer interaction via keyboard. Similarly, some applications can be carried out at relatively low communication rates, while others by their nature require high speeds. A number of scientific networking applications are characterized on the graph of speed versus burstiness in Figure 5. Traffic seen in the early days of networks appears near the bottom of the chart. More advanced applications are furthest from the

The emphasis is on the development of advanced algorithms and software technology required to address applications problems on the scale of grand challenges. The ASTA component is comprised of four subcomponents, as shown below.

Software Support for Grand Challenges. The HPCC Program is designed to demonstrate new computing technology capabilities by confronting an expanding number of grand challenge application problems. The goal is to reduce the risks to researchers inherent in adopting innovative high performance computing technologies. Grand challenge application problems will be selected based upon their scientific importance, the potential for cost sharing with sources directly concerned with the specific scientific and engineering applications, and the potential for leveraging across sectors.

Software Components and Tools. The multidisciplinary research teams will have common needs in software technology and programming environments, advanced compiler technology, optimization and parallelization tools, interoperability and data management, visualization, debugging and analysis, and performance measurement. Advances in these generic software areas will minimize the need for specialized researchers to master advanced, complex computing science skills. Coordination of the development of advanced software technology and algorithms among the participating agencies will be important to ensure effective and efficient use of resources. A particular focus of this element will be to develop advanced software applications that exploit the NREN using distributed file systems and national software libraries.

Computational Techniques. The focus of the HPCC Program on scalable parallel computing systems dictates that significant advances in computational techniques will be needed. The design and theory of algorithms are as important as hardware or networking improvements in reaching teraflops computational performance. Research in computational techniques will include parallel algorithms, numerical and mathematical analyses, parallel languages, computational models, and program refinement techniques. Higher level languages will be developed to enable computational scientists to work directly at the level of the abstract computational problem being addressed and to more easily explore specific implementation approaches.

High Performance Computing Research Centers. High performance computing research centers and testbed facilities will provide a large number of researchers access via the NREN to both conventional and innovative supercomputing architectures. Many of these centers will introduce innovative architectures and will be the focus of technology transfer activities and form the base for training new generations of computer and computational scientists.

National Research and Education Network (NREN)

The NREN component of the HPCC Program dramatically expands and enhances the U.S. portion of an existing worldwide infrastructure of interconnected computer networks called the Internet. A substantial fraction of the domestic Internet is supported and loosely coordinated by Federal agencies, principally DARPA, DOE, NASA, and NSF.


Collaboration among scientists is an important and integral facet of the U.S. research environment. It can be greatly improved by increasing the level of network connectivity and by introducing new capabilities into the existing infrastructure. The NREN design will not only address broad network connectivity, but will also provide the basis for necessary higher level capabilities and services.

Many educational institutions, government laboratories, and industrial research facilities are currently connected to the Internet. Yet, it still falls short of a widespread, uniform, and high performance national infrastructure. In order to satisfy the HPCC Program goals, the NREN must not merely provide network access to research and educational institutions at all levels and locations, it must also deliver new capabilities. Some of these, such as distance learning, may initially be extensions of current technology. All capabilities will benefit from, and many will be enabled by, a program of research into very high speed technology. This technology is needed to support access to digital libraries, large scale distributed computing resources, as well as to perform computationally intensive applications that require real time visualization of modeling and simulation results, rapid interrogation and retrieval of scientific data from specialized data bases, remote control of experiments and simulations, and teleconferencing.

In addition to serving the needs of the scientific and research communities, the NREN will provide valuable experience necessary for the successful development of a broader, privately-operated national information infrastructure. Such an infrastructure would allow consumers, businesses, and schools and government at all levels to share quality information and entertainment when and where they want it at a reasonable cost.

Applications conducted over a computer network vary in their flow of information from steady, as in a bulk file transfer, to "bursty," as in human-computer interaction via keyboard. Similarly, some applications can be carried out at relatively low communication rates, while others by their nature require high speeds. A number of scientific networking applications are characterized on the graph of speed versus burstiness in Figure 5. Traffic seen in the early days of networks appears near the bottom of the chart. More advanced applications are furthest from the

origin and require more sophisticated protocols and network capabilities. This chart illustrates that a gigabit network is needed not only to carry the aggregation of low speed traffic, but also to accommodate high speed uses.




The vision of the NREN is of an interconnection of the nation's educational infrastructure to its knowledge and information centers. In this system, elementary schools, high schools, two and four year colleges, and universities will be linked with research centers and

laboratories so that all may share access to: libraries, databases, and diverse scientific instruments such as supercomputers, telescopes, and particle accelerators. The NREN enables communication and fosters collaboration among and within these communities. By reducing the traditional impediments of geographical isolation, the NREN improves the quality and raises the level of education nationwide. The NREN contributes to the success of the Basic Research and Human Resources component of the High Performance Computing and Communications Program. By making unique scientific and informational resources accessible beyond their physical locations, it permits widespread participation in the HPCC Program by scientists, university researchers, and students, and it enables the development of large scale distributed computing resources.

Interagency Interim NREN. The NSF will coordinate the Interim NREN activities by upgrading its backbone network, by assisting regional networks to upgrade facilities, capacity, and bandwidth, and by interconnecting the backbone networks of other agencies. A significant effort in its implementation will be the development and deployment of safeguards to enhance security and control over access to the connected computer systems, and that of the network components themselves in order to minimize vulnerability to careless or malicious attack. Coordination among participating agencies and the non Federal networking community will be expanded through the creation of a National Networking Council.

Gigabit Research and Development. DARPA will coordinate the research and development effort that will culminate in initial deployment of gigabits per second capability. Coordination of research efforts on very high speed switches, protocols, and computer interfaces will be necessary.

origin and require more sophisticated protocols and network capabilities. This chart illustrates that a gigabit network is needed not only to carry the aggregation of low speed traffic, but also to accommodate high speed uses.



The vision of the NREN is of an interconnection of the nation's educational infrastructure to its knowledge and information centers. In this system, elementary schools, high schools, two and four year colleges, and universities will be linked with research centers and

laboratories so that all may share access to: libraries, databases, and diverse scientific instruments such as supercomputers, telescopes, and particle accelerators. The NREN enables communication and fosters collaboration among and within these communities. By reducing the traditional impediments of geographical isolation, the NREN improves the quality and raises the level of education nationwide. The NREN contributes to the success of the Basic Research and Human Resources component of the High Performance Computing and Communications Program. By making unique scientific and informational resources accessible beyond their physical locations, it permits widespread participation in the HPCC Program by scientists, university researchers, and students, and it enables the development of large scale distributed computing resources.

Interagency Interim NREN. The NSF will coordinate the Interim NREN activities by upgrading its backbone network, by assisting regional networks to upgrade facilities, capacity, and bandwidth, and by interconnecting the backbone networks of other agencies. A significant effort in its implementation will be the development and deployment of safeguards to enhance security and control over access to the connected computer systems, and that of the network components themselves in order to minimize vulnerability to careless or malicious attack. Coordination among participating agencies and the non Federal networking community will be expanded through the creation of a National Networking Council.

Gigabit Research and Development. DARPA will coordinate the research and development effort that will culminate in initial deployment of gigabits per second capability. Coordination of research efforts on very high speed switches, protocols, and computer interfaces will be necessary.

Basic Research and Human Resources (BRHR)

This component addresses long term national needs for more skilled personnel, enhancement of education and training, and materials and curriculum development in the high performance computing science and engineering areas. The NREN and ASTA components include support for research in the large scale project environment. The BRHR component is designed to encourage investigator initiated, long term research on experimental projects that will maintain the flow of innovative ideas and talented people into high performance computing areas.

Drawing the best and brightest of our Nation's youth into scientific and technological careers is a formidable challenge that will have profound effects on overall U.S. scientific and technological competitiveness.

This component of the program will establish industry, university, and government partnerships to improve the training and utilization of personnel and to expand the base of research and development personnel in high performance computing science, technology, and applications.

The BRHR component of the HPCC Program has five goals:

- Improve the flow of human resources into high performance computing.
- Improve the university infrastructure to stay at the leading edge.
- Expand collaboration and resource sharing among the Federal, academic, and industrial sectors.
- Facilitate multidisciplinary research on high performance computing and communications.
- Create a critical mass of users by building a base with common systems, tools, and interfaces.

The BRHR component is organized across the four sub-components shown below.

Basic Research. These activities support increased participation of individual investigators in the conduct of innovative multidisciplinary research in computer science, computer engineering, and computational science and engineering related to high performance computing. The strategy is to increase the number of multi-disciplinary awards across all disciplines where computational methods are critical to achieving advances or scientific breakthroughs. Program activities include: research on scientific algorithm development for highly parallel computers;

generic computational algorithm development; scaling techniques; prediction techniques for concurrent systems; resource management strategies for highly parallel distributed systems; fault tolerant strategies for parallel and distributed systems; and research on heterogeneous software environments.

Research Participation and Training. These activities address the human resources pipeline in the computer and computational sciences, at post-doctoral (training and re-training), graduate, and undergraduate levels. Program activities include: workshops and seminars; post-doctoral fellowships in computational science and engineering; career training in medical informatics through grants to young investigators; institutional training and post-doctoral programs; knowledge transfer exchange programs at national laboratories, centers, universities, and industry; and software dissemination through national databases and libraries.

Infrastructure. These activities will improve university and government facilities for computer science, computer engineering, and computational science and engineering research related to high performance computing. Program activities include: improvement of equipment in computer science, computer engineering, and computational science and engineering academic departments, centers and institutions; development of scientific databases; and distribution of integrated system building kits and toolsets.

Education, Training and Curriculum. These activities will expand and initiate activities to improve undergraduate and pre-college education and training opportunities in high performance computing and computational science and engineering for both students and educators. The introduction of associated curriculum and training materials at all levels is an integral part of this effort. Program activities include: bringing people to national centers for training, technology transfer, and educational experiences; using professional scientists and engineers to provide curriculum development materials and instruction for high school students in the context of high school supercomputer programs, supercomputer user workshops, summer institutes, and career development informatics for health sciences.

Basic Research and Human Resources (BRHR)

This component addresses long term national needs for more skilled personnel, enhancement of education and training, and materials and curriculum development in the high performance computing science and engineering areas. The NREN and ASTA components include support for research in the large scale project environment. The BRHR component is designed to encourage investigator initiated, long term research on experimental projects that will maintain the flow of innovative ideas and talented people into high performance computing areas.

Drawing the best and brightest of our Nation's youth into scientific and technological careers is a formidable challenge that will have profound effects on overall U.S. scientific and technological competitiveness. This component of the program will establish industry, university, and government partnerships to improve the training and utilization of personnel and to expand the base of research and development personnel in high performance computing science, technology, and applications.

The BRHR component of the HPCC Program has five goals:

- Improve the flow of human resources into high performance computing.
- Improve the university infrastructure to stay at the leading edge.
- Expand collaboration and resource sharing among the Federal, academic, and industrial sectors.
- Facilitate multidisciplinary research on high performance computing and communications.
- Create a critical mass of users by building a base with common systems, tools, and interfaces.

The BRHR component is organized across the four sub-components shown below.

Basic Research. These activities support increased participation of individual investigators in the conduct of innovative multidisciplinary research in computer science, computer engineering, and computational science and engineering related to high performance computing. The strategy is to increase the number of multi-disciplinary awards across all disciplines where computational methods are critical to achieving advances or scientific breakthroughs. Program activities include: research on scientific algorithm development for highly parallel computers;

generic computational algorithm development; scaling techniques; prediction techniques for concurrent systems; resource management strategies for highly parallel distributed systems; fault tolerant strategies for parallel and distributed systems; and research on heterogeneous software environments.

Research Participation and Training. These activities address the human resources pipeline in the computer and computational sciences, at post-doctoral (training and re-training), graduate, and undergraduate levels. Program activities include: workshops and seminars; post-doctoral fellowships in computational science and engineering; career training in medical informatics through grants to young investigators; institutional training and post-doctoral programs; knowledge transfer exchange programs at national laboratories, centers, universities, and industry; and software dissemination through national databases and libraries.

Infrastructure. These activities will improve university and government facilities for computer science, computer engineering, and computational science and engineering research related to high performance computing. Program activities include: improvement of equipment in computer science, computer engineering, and computational science and engineering academic departments, centers and institutions; development of scientific databases; and distribution of integrated system building kits and toolsets.

Education, Training and Curriculum. These activities will expand and initiate activities to improve undergraduate and pre-college education and training opportunities in high performance computing and computational science and engineering for both students and educators. The introduction of associated curriculum and training materials at all levels is an integral part of this effort. Program activities include: bringing people to national centers for training, technology transfer, and educational experiences; using professional scientists and engineers to provide curriculum development materials and instruction for high school students in the context of high school supercomputer programs, supercomputer user workshops, summer institutes, and career development informatics for health sciences.

3. PROGRAM DEVELOPMENT AND AGENCY BUDGETS

Program Planning

Leadership for the HPCC Program is provided by the Office of Science and Technology Policy (OSTP), through the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) Committee on Physical, Mathematical, and Engineering Sciences (PMES). The membership of the PMES includes senior executives of many Federal agencies.

The planning process for the HPCC Program was coordinated by a PMES working group through information exchange, the common development of interagency initiatives, and the review of individual agency HPCC proposals and budgets.

Evaluation Criteria

Each participating agency HPCC contribution was reviewed against formal evaluation criteria during the planning and budget process. A review of participating agencies was performed using these evaluation criteria to develop the FY 1992 agency requests for the Program. The evaluation criteria are shown in Figure 6.

Agency Budgets

Over the last three years of the planning process of this initiative, the participating agencies have mutually adjusted their activities within the base to achieve greater efficiency in addressing the goals of the HPCC Program. For FY 1992, \$638 million is being proposed, a \$149 million or 30 percent increase over the FY 1991 enacted level. The budget is shown in Figure 7.

The funds proposed from Federal sources are not intended to carry out the entire HPCC Program. Portions of this program will be cost shared by organizations from the participating sectors. The funding estimates are based on analyses of the practical experience of prior computing and computer networking programs. These estimates were then reviewed as outlined above. Cost sharing will occur with various U.S. industrial and university partners to a large extent in the HPCCS and the NREN components. Cost sharing will occur in the ASTA component with agency programs and other computational applications programs, for example, in specific grand challenge areas, via multidisciplinary collaborations. This component also includes deployment of high performance systems to HPCC research centers. The close coupling of

support in the Program will result in significant leverage, accelerate technology transfer, stimulate U.S. industry and markets, and enable the solution of computationally intensive applications. In addition, although HPCC is not intended to include classified programs, the technology produced will have an important impact in these national security areas. Figure 7 illustrates the relative levels of investment in the four program components.

Figure 6
Evaluation Criteria
for the
Federal High Performance Computing
and Communications Program

Relevance/Contribution. The research must significantly contribute to the overall goals and strategy of the Federal High Performance Computing and Communications (HPCC) Program, including computing, software, networking, and basic research, to enable solution of the grand challenges.

Technical/Scientific Merit. The proposed agency program must be technically/scientifically sound and of high quality, and must be the product of a documented technical/scientific planning and review process.

Readiness. A clear agency planning process must be evident, and the organization must have demonstrated capability to carry out the program.

Timeliness. The proposed work must be technically/scientifically timely for one or more of the HPCC program components.

Linkages. The responsible organization must have established policies, programs, and activities promoting effective technical and scientific connections among government, industry, and academic sectors.

Costs. The identified resources must be adequate, represent an appropriate share of the total available HPCC resources (e.g., a balance among program components), promote prospects for joint funding, and address long-term resource implications.

Enhancements to Existing Program Research. Existing agency HPCC programs will receive adequate support before new initiatives are funded.

Agency Approval. The proposed program or activity must have policy-level approval by the submitting agency.

3. PROGRAM DEVELOPMENT AND AGENCY BUDGETS

Program Planning

Leadership for the HPCC Program is provided by the Office of Science and Technology Policy (OSTP), through the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) Committee on Physical, Mathematical, and Engineering Sciences (PMES). The membership of the PMES includes senior executives of many Federal agencies.

The planning process for the HPCC Program was coordinated by a PMES working group through information exchange, the common development of interagency initiatives, and the review of individual agency HPCC proposals and budgets.

Evaluation Criteria

Each participating agency HPCC contribution was reviewed against formal evaluation criteria during the planning and budget process. A review of participating agencies was performed using these evaluation criteria to develop the FY 1992 agency requests for the Program. The evaluation criteria are shown in Figure 6.

Agency Budgets

Over the last three years of the planning process of this initiative, the participating agencies have mutually adjusted their activities within the base to achieve greater efficiency in addressing the goals of the HPCC Program. For FY 1992, \$638 million is being proposed, a \$149 million or 30 percent increase over the FY 1991 enacted level. The budget is shown in Figure 7.

The funds proposed from Federal sources are not intended to carry out the entire HPCC Program. Portions of this program will be cost shared by organizations from the participating sectors. The funding estimates are based on analyses of the practical experience of prior computing and computer networking programs. These estimates were then reviewed as outlined above. Cost sharing will occur with various U.S. industrial and university partners to a large extent in the HPCS and the NREN components. Cost sharing will occur in the ASTA component with agency programs and other computational applications programs, for example, in specific grand challenge areas, via multidisciplinary collaborations. This component also includes deployment of high performance systems to HPCC research centers. The close coupling of

support in the Program will result in significant leverage, accelerate technology transfer, stimulate U.S. industry and markets, and enable the solution of computationally intensive applications. In addition, although HPCC is not intended to include classified programs, the technology produced will have an important impact in these national security areas. Figure 7 illustrates the relative levels of investment in the four program components.

Figure 6
Evaluation Criteria
for the
Federal High Performance Computing
and Communications Program

Relevance/Contribution. The research must significantly contribute to the overall goals and strategy of the Federal High Performance Computing and Communications (HPCC) Program, including computing, software, networking, and basic research, to enable solution of the grand challenges.

Technical/Scientific Merit. The proposed agency program must be technically/scientifically sound and of high quality, and must be the product of a documented technical/scientific planning and review process.

Readiness. A clear agency planning process must be evident, and the organization must have demonstrated capability to carry out the program.

Timeliness. The proposed work must be technically/scientifically timely for one or more of the HPCC program components.

Linkages. The responsible organization must have established policies, programs, and activities promoting effective technical and scientific connections among government, industry, and academic sectors.

Costs. The identified resources must be adequate; represent an appropriate share of the total available HPCC resources (e.g., a balance among program components); promote prospects for joint funding, and address long-term resource implications.

Enhancements to Existing Program Research. Existing agency HPCC programs will receive adequate support before new initiatives are funded.


Agency Approval. The proposed program or activity must have policy-level approval by the submitting agency.

Figure 7

High Performance Computing and Communications Budgets by Agency and Program Component

(Dollars in Millions)

Agency	Base FY 1991	HPCC FY 1992	FY 1992 HPCC Component			
			HPCC	ASTA	NREN	BRHR
DARPA	183.0	232.2	103.3	38.5	32.9	57.5
DOE	65.0	93.0	15.0	58.0	12.0	8.0
NASA	54.0	72.4	14.2	49.8	7.4	1.0
NSF	169.0	213.0	24.0	103.0	32.7	53.3
DOC/NIST	2.1	2.9	0.3	0.6	2.0	0.0
DOC/NOAA	1.4	2.5	0.0	1.8	0.7	0.0
EPA	1.4	5.2	0.0	4.5	0.0	0.7
NIH/NLM	13.5	17.1	0.0	8.9	4.2	4.0
Total	489.4	638.3	156.8	265.1	91.9	124.5



HPCC — High Performance Computing and Communications

Components of HPCC:

HPCC — High Performance Computing Systems
 ASTA — Advanced Software Technology and Algorithms
 NREN — National Research and Education Network
 BRHR — Basic Research and Human Resources

Agency Program Descriptions

The agency responsibilities under the Federal HPCC Program are outlined in Figure 8. Several agencies have been assigned a coordinating responsibility in specific technical areas:


- DOE and NASA will coordinate activities in HPC system evaluation, tested development, and applications software capabilities.
- NASA will coordinate the accumulation of and access to the HPC software base. This will be facilitated by a wide area file system technology that is currently being deployed for early experimental use by DARPA which will be extended to include the NREN as it matures and is deployed by NSF.
- DARPA will coordinate activities in the development of scalable parallel HPC systems, including their basic units of replication, system modules, and the necessary associated systems software. DARPA will also coordinate activities in gigabit network technology research.
- NSF will coordinate activities for the broad deployment of the National Research and Education Network working with all agencies with mission specific requirements.
- NSF will coordinate activities in basic research and human resource development, while each agency retains its role as required to accomplish their own missions.
- NIST will coordinate activities in HPC system instrumentation, evaluation, and in standards issues.

Each agency participates in all of the identified activities to ensure that the resulting capabilities are a good match to user needs. DOC/NOAA, HHS/NIH/NLM, and EPA bring distinctive applications areas of broad interest and network user bases.

Figure 7

High Performance Computing and Communications Budgets by Agency and Program Component (Dollars in Millions)

Agency	Base FY 1991	HPCC FY 1992	FY 1992 HPCC Component			
			HPCS	ASTA	NREN	BRHR
DARPA	183.0	232.2	103.3	38.5	32.9	57.5
DOE	65.0	93.0	15.0	58.0	12.0	8.0
NASA	54.0	72.4	14.2	49.8	7.4	1.0
NSF	169.0	213.0	24.0	103.0	32.7	53.3
DOC/NIST	2.1	2.9	0.3	0.6	2.0	0.0
DOC/NOAA	1.4	2.5	0.0	1.8	0.7	0.0
EPA	1.4	5.2	0.0	4.5	0.0	0.7
NIH/NLM	13.5	17.1	0.0	8.9	4.2	4.0
Total	489.4	638.3	156.8	265.1	91.9	124.5



HPCC — High Performance Computing and Communications

Components of HPCC:

- HPCC — High Performance Computing Systems
- ASTA — Advanced Software Technology and Algorithms
- NREN — National Research and Education Network
- BRHR — Basic Research and Human Resources

Agency Program Descriptions

The agency responsibilities under the Federal HPCC Program are outlined in Figure 8. Several agencies have been assigned a coordinating responsibility in specific technical areas:

- DOE and NASA will coordinate activities in HPC system evaluation, testbed development, and applications software capabilities.
- NASA will coordinate the accumulation of and access to the HPC software base. This will be facilitated by a wide area file system technology that is currently being deployed for early experimental use by DARPA which will be extended to include the NREN as it matures and is deployed by NSF.
- DARPA will coordinate activities in the development of scalable parallel HPC systems, including their basic units of replication, system modules, and the necessary associated systems software. DARPA will also coordinate activities in gigabit network technology research.
- NSF will coordinate activities for the broad deployment of the National Research and Education Network working with all agencies with mission specific requirements.
- NSF will coordinate activities in basic research and human resource development, while each agency retains its role as required to accomplish their own missions.
- NIST will coordinate activities in HPC system instrumentation, evaluation, and in standards issues.

Each agency participates in all of the identified activities to ensure that the resulting capabilities are a good match to user needs. DOC/NOAA, HHS/NIH/NLM, and EPA bring distinctive applications areas of broad interest and network user bases.

Figure 8

HPCC Program: Agency Responsibilities: HPCS, ASTA		
Activity Agency	High Performance Computing Systems	Advanced Software Technology & Algorithms
DARPA	Technology development and coordination for Teraops systems	Technology development for parallel algorithms and systems software
DOE	- Technology development - Systems evaluation	- Energy applications centers - Energy computation research - Software tools
NASA	- Aeronautics and space application testbeds - Systems development for space flight	- Software coordination - Research in: - Aerospace computations - Information management
NSF	Basic architecture research	- Basic research in: - Software tools, databases - Grand challenges - Computer access
DOC/ NIST		
DOC/ NOAA		- Ocean and atmospheric computation research - Software tools
EPA		Research in environmental computations, databases, and application testbeds
NIH/ NLM		- Medical application testbeds - Software tools - Medical computation research

HPCC Program: Agency Responsibilities: NREN, BRHR		
National Research and Education Network	Basic Research and Human Resources	Activity Agency
Technology development and coordination for gigabits networks	University programs	DARPA
- Energy deployment mission facilities - Gigabits applications research	Basic research and education programs	DOE
Aerospace deployment mission facilities	Research initiation and university block grants	NASA
- Facilities coordination and deployment - Gigabits research	- Programs in: - Basic research - Education - Training / curricula	NSF
High speed network research and standards		DOC/ NIST
Ocean and atmospheric mission facilities		DOC/ NOAA
States environmental mission assimilation	- Technology transfer to States - University programs	EPA
Medical mission facilities	University programs	NIH/ NLM

Figure 8

HPCC Program: Agency Responsibilities: HPCS, ASTA		
Activity Agency	High Performance Computing Systems	Advanced Software Technology & Algorithms
DARPA	Technology development and coordination for Teraops systems	Technology development for parallel algorithms and systems software
DOE	- Technology development - Systems evaluation	- Energy applications centers - Energy computation research - Software tools
NASA	- Aeronautics and space application testbeds - Systems development for space flight	- Software coordination - Research in: - Aerospace computations - Information management
NSF	Basic architecture research	- Basic research in: - Software tools, databases - Grand challenges - Computer access
DOC/ NIST		
DOC/ NOAA		- Ocean and atmospheric computation research - Software tools
EPA		Research in environmental computations, databases, and application testbeds
NIH/ NLM		- Medical application testbeds - Software tools - Medical computation research

HPCC Program: Agency Responsibilities: NREN, BRHR		
National Research and Education Network	Basic Research and Human Resources	Activity Agency
Technology development and coordination for gigabits networks	University programs	DARPA
- Energy deployment mission facilities - Gigabits applications research	Basic research and education programs	DOE
Aerospace deployment mission facilities	Research initiation and university block grants	NASA
- Facilities coordination and deployment - Gigabits research	- Programs in: - Basic research - Education - Training / curricula	NSF
High speed network research and standards		DOC/ NIST
Ocean and atmospheric mission facilities		DOC/ NOAA
States environmental mission assimilation	- Technology transfer to States - University programs	EPA
Medical mission facilities	University programs	NIH/ NLM

Defense Advanced Research Projects Agency (DARPA)

DARPA, as the DOD lead agency for advanced technology research, will focus on developing the high performance computing and networking technologies needed for the Defense and overall HPCC Programs. DARPA programs have produced both the computing and networking foundation for the HPCC Program, including the first generation of scalable parallel computing systems and large scale computer networks, and the associated system software and supporting technologies. DARPA has worked with industry to pioneer the application of these new technologies within Defense and on a cooperative basis with other Federal agencies.

The DARPA HPCC Program builds upon the DARPA Strategic Computing Program. As the HPCC Program builds up, the Strategic Computing Program integrates its results with Defense specific needs such as embedded systems, accelerators of specific problem domains, and grand challenges problems related to defense. DARPA will continue this mode of executing the HPCC Program, cooperating with various defense organizations, and working closely with AFOSR, ARO, ONR, Defense Service Laboratories, NSA, other Defense organizations, and other Federal agencies as appropriate.

High Performance Computing Systems will be produced in the four main subareas: Research for Future Generations of Computing Systems, System Design Tools, Advanced Prototype Systems, and Evaluation of Early Systems. Systems capable of sustaining 100 gigaflops for large problems will be available for deployment by late 1993 and the teraflops systems will be available by 1996.

In *Advanced Software Technology and Algorithms*, DARPA projects will produce scalable libraries for Defense problem domains and programming and analysis tools for scalable parallel and distributed heterogeneous systems in a workstation/server configuration that will be open to the integration of embedded systems and accelerators.

For the *National Research and Education Network* component, high performance networking technologies will be produced to satisfy the gigabit technology needs of the NREN and to provide a dual use technology base for Defense. This networking technology includes development of new protocols and switch and transmission technologies, and it will be capable of supporting a wide range of advanced network services.

The *Basic Research and Human Resources* component will focus on fundamental scientific issues in these three technology development areas in cooperation with other basic research programs and provide for smaller individual investigator projects to complement the larger projects. In addition, the relevant and related basic research will also be integrated into the larger projects as it matures.

Department of Energy (DOE)

The Department of Energy will participate in all components of the High Performance Computing and Communications Program. In the area of *High Performance Computing Systems*, the DOE will be an early customer of small versions of systems with advanced architectures and will evaluate these systems on energy related applications. The DOE will consider cooperative development of advanced systems between its national laboratories and vendors, especially integration of very high speed computer and networking hardware with software systems. The DOE will support research and development on algorithms and systems software for parallel computing systems.

The *Advanced Software Technology and Algorithms* effort will include research and development of: parallel algorithms for grand challenge applications, software and tools for early prototypes of 100 gigaflops and teraflops systems, prototype computational science programming environments that meet standards and are transportable, and support for high performance computing research centers to facilitate the transition from research on parallel machines into the applications and the programming environments. The DOE will fund several grand challenge collaborations, initiate a software component and tools program with strong industrial participation, and initiate an applications driven computational research program. The DOE will evaluate proposals and make research awards related to grand challenges in global climate change, molecular biology, human genome research, materials and chemical sciences, combustion research, waste remediation, fusion energy, and other areas within its mission.

The DOE will participate in the cooperative interagency *National Research and Education Network*. The Energy Sciences Net (ESNet) will be incorporated into the NREN and will provide quality network access to the energy research facilities. ESNet will maintain compatibility and will be upgraded in concert with NREN. Gigabit network support technology will be developed for DOE applications distributed across multiple energy research centers at the national laboratories and universities.

The DOE's *Basic Research and Human Resources* activities will include: stimulating research in computational science, expanding training programs at the national laboratories for high school teachers and college students in computing techniques, initiation of a high school supercomputer access program, and provision of fellowships in computational science with internship at national laboratories.

Defense Advanced Research Projects Agency (DARPA)

DARPA, as the DOD lead agency for advanced technology research, will focus on developing the high performance computing and networking technologies needed for the Defense and overall HPCC Programs. DARPA programs have produced both the computing and networking foundation for the HPCC Program, including the first generation of scalable parallel computing systems and large scale computer networks, and the associated system software and supporting technologies. DARPA has worked with industry to pioneer the application of these new technologies within Defense and on a cooperative basis with other Federal agencies.

The DARPA HPCC Program builds upon the DARPA Strategic Computing Program. As the HPCC Program builds up, the Strategic Computing Program integrates its results with Defense specific needs, such as embedded systems, accelerators of specific problem domains, and grand challenges problems related to defense. DARPA will continue this mode of executing the HPCC Program, cooperating with various defense organizations, and working closely with AFOSR, ARO, ONR, Defense Service Laboratories, NSA, other Defense organizations, and other Federal agencies as appropriate.

High Performance Computing Systems will be produced in the four main subareas: Research for Future Generations of Computing Systems, System Design Tools, Advanced Prototype Systems, and Evaluation of Early Systems. Systems capable of sustaining 100 gigaflops for large problems will be available for deployment by late 1993 and the teraflops systems will be available by 1996.

In *Advanced Software Technology and Algorithms*, DARPA projects will produce scalable libraries for Defense problem domains and programming and analysis tools for scalable parallel and distributed heterogeneous systems in a workstation/server configuration that will be open to the integration of embedded systems and accelerators.

For the *National Research and Education Network* component, high performance networking technologies will be produced to satisfy the gigabit technology needs of the NREN and to provide a dual use technology base for Defense. This networking technology includes development of new protocols and switch and transmission technologies, and it will be capable of supporting a wide range of advanced network services.

The *Basic Research and Human Resources* component will focus on fundamental scientific issues in these three technology development areas in cooperation with other basic research programs and provide for smaller individual investigator projects to complement the larger projects. In addition, the relevant and related basic research will also be integrated into the larger projects as it matures.

Department of Energy (DOE)

The Department of Energy will participate in all components of the High Performance Computing and Communications Program. In the area of *High Performance Computing Systems*, the DOE will be an early customer of small versions of systems with advanced architectures and will evaluate these systems on energy related applications. The DOE will consider cooperative development of advanced systems between its national laboratories and vendors, especially integration of very high speed computer and networking hardware with software systems. The DOE will support research and development on algorithms and systems software for parallel computing systems.

The *Advanced Software Technology and Algorithms* effort will include research and development of: parallel algorithms for grand challenge applications, software and tools for early prototypes of 100 gigaflops and teraflops systems, prototype computational science programming environments that meet standards and are transportable, and support for high performance computing research centers to facilitate the transition from research on parallel machines into the applications and the programming environments. The DOE will fund several grand challenge collaborations, initiate a software component and tools program with strong industrial participation, and initiate an applications driven computational research program. The DOE will evaluate proposals and make research awards related to grand challenges in global climate change, molecular biology, human genome research, materials and chemical sciences, combustion research, waste remediation, fusion energy, and other areas within its mission.

The DOE will participate in the cooperative interagency *National Research and Education Network*. The Energy Sciences Net (ESNet) will be incorporated into the NREN and will provide quality network access to the energy research facilities. ESNet will maintain compatibility and will be upgraded in concert with NREN. Gigabit network support technology will be developed for DOE applications distributed across multiple energy research centers at the national laboratories and universities.

The DOE's *Basic Research and Human Resources* activities will include: stimulating research in computational science, expanding training programs at the national laboratories for high school teachers and college students in computing techniques, initiation of a high school supercomputer access program, and provision of fellowships in computational science with internship at national laboratories.

National Aeronautics and Space Administration (NASA)

NASA's HPCC Program participates in all four components of the Federal Program, through a vertically integrated program focused on NASA's grand challenges in: Computational Aerosciences, Earth and Space Sciences, and Remote Exploration and Experimentation.

The goal of NASA's program is to accelerate the development and application of high performance computing technologies to meet NASA's science and engineering requirements. In cooperation with the other Federal agencies, NASA's program will deploy teraflops computer capabilities essential for computational design of integrated aerospace vehicle systems and for predicting long term global change, and will enable the development of massively parallel techniques for spaceborne applications.

NASA's program is focused on bringing together interdisciplinary teams of computer scientists and computational physicists to develop these technologies within three vertically integrated projects that are unique to NASA's missions. These technologies include applications algorithms and programs, systems software, peripherals, networking, and the actual high performance computing hardware. NASA will develop a suite of software tools to enhance productivity. These include: load balancing tools, run time optimizers, monitors, parallelization tools, as well as data management and visualization tools.

NASA's role includes coordinating the *Advanced Software Technology and Algorithm* component for the Federal program; acquiring experimental hardware for testbeds in computational aerosciences, earth and space systems sciences, and remote exploration and experimentation; and supporting the development of the National Research and Education Network. To encourage vigorous research into the underlying theory and concepts of high performance computing, NASA will foster interactions among academia, industry, and national laboratories and will strengthen the basic research in high performance computing at the NASA centers and research institutes and in universities.

NASA's considerable expertise in experimental parallel computer testbeds and small, scalable testbed systems will be used to demonstrate high performance computing technologies as a step toward full-scale computational capabilities. A key to successful exploitation of massively parallel computing power will be the blending of application-driven and architecture-driven computer systems to most effectively meet NASA's needs.

National Science Foundation (NSF)

The NSF HPCC Program impacts the activities of all science and engineering disciplines by providing computing and networking infrastructure support and by developing enabling technologies for advanced computing and communications platforms and paradigms.

In the area of *High Performance Computing Systems*, research will be initiated on new architectures and systems optimized for specific research applications. New tools for systems level automated design and component packaging will be supported to permit the design of application-specific devices and systems.

In *Advanced Software Technology and Algorithms*, research will be initiated on scientific database technology and implementation of prototype networked databases and associated software. Advanced applications tools for research computing environments will be supported to enable nationwide access to the full complement of parallel machines for research on grand challenge problems. Areas of research focus will include numeric and symbolic computing, algorithm development, optimization of applications software for new parallel computers, scientific visualization, automated programming tools, and new methods of scientific and technical information exchanges.

NSF coordination of the *National Research and Education Network* activities will accelerate the harmonizing of multiple agency networks and protocols into a single NREN. The number of nodes will be increased to expand distributed information resources, and to increase redundancy, capacity, control, and security. Mid-level nets will be assisted to upgrade facilities and service very high bandwidth requirements. Support will be provided for research on new protocols for gigabit networks, switch and transmission technology, routing, congestion, and flow control. The exploration of pricing mechanisms for network services and network applications and structured transition to commercial service will be initiated.

In *Basic Research and Human Resources*, support will include multidisciplinary research and university infrastructure for computer science, computer engineering, and computational science and engineering. To increase the human resource pool in computing hardware and software systems areas, support will be expanded for graduate and post-doctoral positions. Curriculum improvement, teacher training, and support for centers to provide education and training in the use of experimental and parallel supercomputers are integral parts of this component.

National Aeronautics and Space Administration (NASA)

NASA's HPCC Program participates in all four components of the Federal Program, through a vertically integrated program focused on NASA's grand challenges in: Computational Aerosciences, Earth and Space Sciences, and Remote Exploration and Experimentation.

The goal of NASA's program is to accelerate the development and application of high performance computing technologies to meet NASA's science and engineering requirements. In cooperation with the other Federal agencies, NASA's program will deploy teraflops computer capabilities essential for computational design of integrated aerospace vehicle systems and for predicting long term global change, and will enable the development of massively parallel techniques for spaceborne applications.

NASA's program is focused on bringing together interdisciplinary teams of computer scientists and computational physicists to develop these technologies within three vertically integrated projects that are unique to NASA's missions. These technologies include applications algorithms and programs, systems software, peripherals, networking, and the actual high performance computing hardware. NASA will develop a suite of software tools to enhance productivity. These include: load balancing tools, run time optimizers, monitors, parallelization tools, as well as data management and visualization tools.

NASA's role includes coordinating the *Advanced Software Technology and Algorithm* component for the Federal program; acquiring experimental hardware for testbeds in computational aerosciences, earth and space systems sciences, and remote exploration and experimentation; and supporting the development of the National Research and Education Network. To encourage vigorous research into the underlying theory and concepts of high performance computing, NASA will foster interactions among academia, industry, and national laboratories and will strengthen the basic research in high performance computing at the NASA centers and research institutes and in universities.

NASA's considerable expertise in experimental parallel computer testbeds and small, scalable testbed systems will be used to demonstrate high performance computing technologies as a step toward full-scale computational capabilities. A key to successful exploitation of massively parallel computing power will be the blending of application-driven and architecture-driven computer systems to most effectively meet NASA's needs.

National Science Foundation (NSF)

The NSF HPCC Program impacts the activities of all science and engineering disciplines by providing computing and networking infrastructure support and by developing enabling technologies for advanced computing and communications platforms and paradigms.

In the area of *High Performance Computing Systems*, research will be initiated on new architectures and systems optimized for specific research applications. New tools for systems level automated design and component packaging will be supported to permit the design of application-specific devices and systems.

In *Advanced Software Technology and Algorithms*, research will be initiated on scientific database technology and implementation of prototype networked databases and associated software. Advanced applications tools for research computing environments will be supported to enable nationwide access to the full complement of parallel machines for research on grand challenge problems. Areas of research focus will include numeric and symbolic computing, algorithm development, optimization of applications software for new parallel computers, scientific visualization, automated programming tools, and new methods of scientific and technical information exchanges.

NSF coordination of the *National Research and Education Network* activities will accelerate the harmonizing of multiple agency networks and protocols into a single NREN. The number of nodes will be increased to expand distributed information resources, and to increase redundancy, capacity, control, and security. Mid-level nets will be assisted to upgrade facilities and service very high bandwidth requirements. Support will be provided for research on new protocols for gigabit networks, switch and transmission technology, routing, congestion, and flow control. The exploration of pricing mechanisms for network services and network applications and structured transition to commercial service will be initiated.

In *Basic Research and Human Resources*, support will include multidisciplinary research and university infrastructure for computer science, computer engineering, and computational science and engineering. To increase the human resource pool in computing hardware and software systems areas, support will be expanded for graduate and post-doctoral positions. Curriculum improvement, teacher training, and support for centers to provide education and training in the use of experimental and parallel supercomputers are integral parts of this component.

National Institute for Standards and Technology (DOC/NIST)

The NIST research program is directed toward developing performance monitoring tools and promoting "open systems" software. NIST has proposed to augment its current HPC research program by promoting the commercialization of protocol and security mechanisms for medium speed networks.

The goal of NIST's activities in high performance computing is to develop hardware performance monitoring tools, promote "open system" software, and support a classification system for indexing and distributing scientific software so that industry and the research community can effectively exploit the power of future generations of high performance computers. In support of *The National Research and Education Network*, NIST will develop and speed commercialization of network protocols and security mechanisms that can achieve the desired gigabit speeds on future versions of the network. NIST will participate in the HPCC Program by:

- Developing hardware monitoring methods leading to load characterization and performance measurement techniques for ultra-high-speed systems which will be made available on a publicly accessible database.

- Conducting research on new protocols and related security primitives necessary to sustain gigabit network speeds and standards to provide interoperability, common user interfaces to systems, and enhanced security for the network.

- Establishing a network testbed at NIST instrumented to collect performance data and test new network protocols, management routines, and security mechanisms for gigabit networks.

National Oceanic and Atmospheric Administration (DOC/NOAA)

The National Oceanic and Atmospheric Administration operational and research programs are directed toward weather prediction, ocean sciences, the Climate and Global Change Program, and the Coastal Oceans Program, together with data management activities for all agency programs. The HPCC Program will allow extensive development of new forecast models, studies in computational fluid dynamics, and the incorporation of evolving computer architectures and networks into the systems that carry out agency missions.

The NOAA High Performance Computing Program is focused in two components:

Advanced Software Technology and Algorithms. This component provides support for: grand challenges in atmospheric and oceanic sciences; development of advanced numerical models to simulate the general circulation of the oceans and atmosphere in support of NOAA missions and the activities of collaborating research groups; development of new computational methods for solving atmospheric, oceanic, and related problems on new computer architectures; data management R&D; support for basic research in strategies, techniques, and tools required for the management and analysis of large-scale distributed scientific databases and distributed data handling, including quality control; and development of algorithms for massively parallel processors, together with their standardized software component libraries and tools for the solution of oceanic and atmospheric analysis and forecasting problems.

NOAA will acquire, install, and operate advanced computational facilities for the evaluation of near operational prototype computers having massively parallel architectures and will develop algorithms appropriate for these architectures.

The National Research and Education Network component provides networking in support for NOAA's climate and global change research community and a wide range of agency missions in oceanography, weather prediction, and environmental sciences research. NOAA will evaluate advanced network protocols and hardware technologies related to its missions.

National Institute for Standards and Technology (DOC/NIST)

The NIST research program is directed toward developing performance monitoring tools and promoting "open systems" software. NIST has proposed to augment its current HPC research program by promoting the commercialization of protocol and security mechanisms for medium speed networks.

The goal of NIST's activities in high performance computing is to develop hardware performance monitoring tools, promote "open system" software, and support a classification system for indexing and distributing scientific software so that industry and the research community can effectively exploit the power of future generations of high performance computers. In support of *The National Research and Education Network*, NIST will develop and speed commercialization of network protocols and security mechanisms that can achieve the desired gigabit speeds on future versions of the network. NIST will participate in the HPCC Program by:

- Developing hardware monitoring methods leading to load characterization and performance measurement techniques for ultra-high-speed systems which will be made available on a publicly accessible database.

- Conducting research on new protocols and related security primitives necessary to sustain gigabit network speeds and standards to provide interoperability, common user interfaces to systems, and enhanced security for the network.

- Establishing a network testbed at NIST instrumented to collect performance data and test new network protocols, management routines, and security mechanisms for gigabit networks.

National Oceanic and Atmospheric Administration (DOC/NOAA)

The National Oceanic and Atmospheric Administration operational and research programs are directed toward weather prediction, ocean sciences, the Climate and Global Change Program, and the Coastal Oceans Program, together with data management activities for all agency programs. The HPCC Program will allow extensive development of new forecast models, studies in computational fluid dynamics, and the incorporation of evolving computer architectures and networks into the systems that carry out agency missions.

The NOAA High Performance Computing Program is focused in two components:

Advanced Software Technology and Algorithms. This component provides support for: grand challenges in atmospheric and oceanic sciences; development of advanced numerical models to simulate the general circulation of the oceans and atmosphere in support of NOAA missions and the activities of collaborating research groups; development of new computational methods for solving atmospheric, oceanic, and related problems on new computer architectures; data management R&D; support for basic research in strategies, techniques, and tools required for the management and analysis of large-scale distributed scientific databases and distributed data handling, including quality control; and development of algorithms for massively parallel processors, together with their standardized software component libraries and tools for the solution of oceanic and atmospheric analysis and forecasting problems.

NOAA will acquire, install, and operate advanced computational facilities for the evaluation of near operational prototype computers having massively parallel architectures and will develop algorithms appropriate for these architectures.

The National Research and Education Network component provides networking in support for NOAA's climate and global change research community and a wide range of agency missions in oceanography, weather prediction, and environmental sciences research. NOAA will evaluate advanced network protocols and hardware technologies related to its missions.

Environmental Protection Agency (EPA)

The EPA research program is directed toward the advancement and dissemination of computational techniques and software tools which form the core of ecosystem, atmospheric chemistry and dynamics models. The models extend the computational capability of environmental assessment tools to handle multi-pollutant interactions and optimization of control strategies.

The *Advanced Software Technology and Algorithms* component will develop new approaches for coupling numerical equation solvers to the high performance features of emerging computer architectures including specific techniques for solution of partial differential equations on massively parallel architectures. These solvers will be designed to operate on a flexible, generic grid system decoupled from specific applications to provide a testbed for evaluation and optimization. The resulting solution framework can serve as the computational engine for a variety of interdisciplinary applications.

Advancement of scientific visualization techniques can provide exceptional benefits in increasing the pace at which environmental problems are solved. EPA's program includes basic research to improve user interfaces and computational algorithms related to image rendering and manipulation of geometric images. The main focus is on simultaneous representations of multiple data sets and real-time visualization in a heterogeneous distributed processing environment. Related research on storage and access techniques for massive environmental data bases across multiple architectures will be conducted.

The *Basic Research and Human Resources* component includes the establishment of technology transfer centers to propagate the use of HPCC technology to State and local environmental groups, and Federal managers for optimization of pollution control and prevention strategies. A main goal of a Technology Transfer Center is to provide non-sophisticated users with training and guidance in the application of a variety of high performance computing tools to solve important environmental challenges. The program also supports cross-discipline career training through universities and other institutions to ensure a continuing base of technical professionals knowledgeable in the use of high performance computing technology for environmental problem solving.

National Institutes of Health / National Library of Medicine (HHS/NIH/NLM)

The HHS/NIH program includes molecular biology computing, creation, and transmission of digital electronic images, the linking of academic health centers via computer networks, the creation of advanced methods to retrieve information from life sciences databases, and training in biomedical computer sciences. The HPCC program will complement the Human Genome Project by providing new methods for computer based analysis of normal and disease genes.

The *Advanced Software Technology and Algorithms* component will develop advanced software technology and algorithms in two areas of importance in biomedical research and education: biotechnology computing and digital images. In the area of biotechnology computing, the program will support development of molecular sequence comparison algorithms, new database methods, and algorithms to predict biological structure and function from genetic code. The biomedical images area will support new methods for representing, linking, and rendering images of biological structure.

The *National Research and Education Network* component has two sub-components: connections among academic medical centers and their growing array of computerized information sources; and development of intelligent gateways that link conceptually related databanks in the life sciences. The academic medical centers of the country are confronted with a growing array of disconnected computer-based information sources, ranging from patient records, x-rays and laboratory systems, to basic research tools such as protein and DNA sequence databanks. Development of advanced software systems capable of representing and linking these dissimilar data types, and communicating them among centers for research and health care, is the goal of this part of the program. The focus of the Gateways HPCC program is the building of systems that are capable of translating a user's request into multiple computer-based vocabularies, selecting appropriate databases from a wide range of widely-available resources, and retrieving information in a manner that does not require users to understand the structure and syntax of the systems being queried.

The *Basic Research and Human Resources* component addresses the need to train biomedical researchers and health care providers in the use of advanced computing and network communications to aid in their work. The successful predoctoral and postdoctoral grants program for career training in medical informatics administered by the National Institutes of Health will be expanded.

Environmental Protection Agency (EPA)

The EPA research program is directed toward the advancement and dissemination of computational techniques and software tools which form the core of ecosystem, atmospheric chemistry and dynamics models. The models extend the computational capability of environmental assessment tools to handle multi-pollutant interactions and optimization of control strategies.

The *Advanced Software Technology and Algorithms* component will develop new approaches for coupling numerical equation solvers to the high performance features of emerging computer architectures including specific techniques for solution of partial differential equations on massively parallel architectures. These solvers will be designed to operate on a flexible, generic grid system decoupled from specific applications to provide a testbed for evaluation and optimization. The resulting solution framework can serve as the computational engine for a variety of interdisciplinary applications.

Advancement of scientific visualization techniques can provide exceptional benefits in increasing the pace at which environmental problems are solved. EPA's program includes basic research to improve user interfaces and computational algorithms related to image rendering and manipulation of geometric images. The main focus is on simultaneous representations of multiple data sets and real-time visualization in a heterogeneous distributed processing environment. Related research on storage and access techniques for massive environmental data bases across multiple architectures will be conducted.

The *Basic Research and Human Resources* component includes the establishment of technology transfer centers to propagate the use of HPCC technology to State and local environmental groups, and Federal managers for optimization of pollution control and prevention strategies. A main goal of a Technology Transfer Center is to provide non-sophisticated users with training and guidance in the application of a variety of high performance computing tools to solve important environmental challenges. The program also supports cross-discipline career training through universities and other institutions to ensure a continuing base of technical professionals knowledgeable in the use of high performance computing technology for environmental problem solving.

National Institutes of Health / National Library of Medicine (HHS/NIH/NLM)

The HHS/NIH program includes molecular biology computing, creation, and transmission of digital electronic images, the linking of academic health centers via computer networks, the creation of advanced methods to retrieve information from life sciences databases, and training in biomedical computer sciences. The HPCC program will complement the Human Genome Project by providing new methods for computer based analysis of normal and disease genes.

The *Advanced Software Technology and Algorithms* component will develop advanced software technology and algorithms in two areas of importance in biomedical research and education: biotechnology computing and digital images. In the area of biotechnology computing, the program will support development of molecular sequence comparison algorithms, new database methods, and algorithms to predict biological structure and function from genetic code. The biomedical images area will support new methods for representing, linking, and rendering images of biological structure.

The *National Research and Education Network* component has two sub-components: connections among academic medical centers and their growing array of computerized information sources; and development of intelligent gateways that link conceptually related databanks in the life sciences. The academic medical centers of the country are confronted with a growing array of disconnected computer-based information sources, ranging from patient records, x-rays and laboratory systems, to basic research tools such as protein and DNA sequence databanks.

Development of advanced software systems capable of representing and linking these dissimilar data types, and communicating them among centers for research and health care, is the goal of this part of the program. The focus of the Gateways HPCC program is the building of systems that are capable of translating a user's request into multiple computer-based vocabularies, selecting appropriate databases from a wide range of widely-available resources, and retrieving information in a manner that does not require users to understand the structure and syntax of the systems being queried.

The *Basic Research and Human Resources* component addresses the need to train biomedical researchers and health care providers in the use of advanced computing and network communications to aid in their work. The successful predoctoral and postdoctoral grants program for career training in medical informatics administered by the National Institutes of Health will be expanded.

4. GRAND CHALLENGE AND SUPPORTING TECHNOLOGY CASE STUDIES

This section incorporates examples of high performance computing and computer communications technologies and several illustrative grand challenge applications in computational science and engineering, that are presented on the cover of this report. The examples were chosen to illustrate the diversity and significance of application areas that have been addressed to date. The list of grand challenges is too long to allow an example from each area, thus lack of representation of certain areas does not imply lack of importance. A more complete description of the grand challenges can be found in the Federal High Performance Computing Program report issued in September 1989.

4. GRAND CHALLENGE AND SUPPORTING TECHNOLOGY CASE STUDIES


This section incorporates examples of high performance computing and computer communications technologies and several illustrative grand challenge applications in computational science and engineering, that are presented on the cover of this report. The examples were chosen to illustrate the diversity and significance of application areas that have been addressed to date. The list of grand challenges is too long to allow an example from each area, thus lack of representation of certain areas does not imply lack of importance. A more complete description of the grand challenges can be found in the Federal High Performance Computing Program report issued in September 1989.

Forecasting Severe Weather Events

Making accurate predictions about the behavior of the atmosphere is of critical importance to reducing losses due to storms, hurricanes, floods, and other weather related phenomena. Current operational small scale weather prediction models have been constrained by two broad deficiencies: inadequate size of the data sets necessary to define the detailed atmospheric structures and insufficient computer resources to support high resolution. During the next five years, major progress will be made on the first deficiency by using a combination of new ground-based and remote-observing systems. This leaves inadequate computer power as the primary stumbling block preventing high resolution models from playing operational roles in our national weather prediction efforts.

The high-resolution, operational weather prediction models of the future will represent a new generation of numerical formulations. The primary differences will be in the treatment of vertical motions and small scale processes which in the past have been considered to be "sub-grid scale." This means that physically significant events, such as convection, must be treated as parameterized processes and fine scale details in, for example, thunderstorm evolution that can affect the surroundings cannot be directly addressed by the model. Likewise, fine scale observations of local importance, such as moisture gradients around wet areas, which can provide local forcing, currently cannot be incorporated adequately.

To meet this scientific and national need, new technologies need to be developed and incorporated into existing facilities so that a less than 5 km resolution model can be operational before the end of the century. This would allow updated forecasts on a 6-hourly or shorter basis. Such models have already shown significant advances in the accuracy of predicting a wide variety of weather events, from severe thunderstorms to lake effect snowfalls in research applications. Each reduction of the model resolution by one-half requires an increase of computer power of almost an order of magnitude, as well as comparable increases in supporting memory, mass storage and networks. For example, a 5 km model could require up to a 20 teraflops computer system to meet operational schedules using this model.




Researchers have used supercomputers to model thunderstorms numerically. Wind, temperature, pressure, and other variables are calculated every few seconds at several hundred thousand locations in the area of the developing storm. Mathematical equations are then used to simulate the storm's evolution. In this graphic, particles which are released near the ground at regular intervals are colored orange when rising and blue when sinking. Yellow signifies paths of individual particles.

Forecasting Severe Weather Events

Making accurate predictions about the behavior of the atmosphere is of critical importance to reducing losses due to storms, hurricanes, floods, and other weather related phenomena. Current operational small scale weather prediction models have been constrained by two broad deficiencies: inadequate size of the data sets necessary to define the detailed atmospheric structures and insufficient computer resources to support high resolution. During the next five years, major progress will be made on the first deficiency by using a combination of new ground-based and remote-observing systems. This leaves inadequate computer power as the primary stumbling block preventing high resolution models from playing operational roles in our national weather prediction efforts.

The high-resolution, operational weather prediction models of the future will represent a new generation of numerical formulations. The primary differences will be in the treatment of vertical motions and small scale processes which in the past have been considered to be "sub-grid scale." This means that physically significant events, such as convection, must be treated as parameterized processes and fine scale details in, for example, thunderstorm evolution that can affect the surroundings cannot be directly addressed by the model. Likewise, fine scale observations of local importance, such as moisture gradients around wet areas, which can provide local forcing, currently cannot be incorporated adequately.

To meet this scientific and national need, new technologies need to be developed and incorporated into existing facilities so that a less than 5 km resolution model can be operational before the end of the century. This would allow updated forecasts on a 6-hourly or shorter basis. Such models have already shown significant advances in the accuracy of predicting a wide variety of weather events, from severe thunderstorms to lake effect snowfalls in research applications. Each reduction of the model resolution by one-half requires an increase of computer power of almost an order of magnitude, as well as comparable increases in supporting memory, mass storage and networks. For example, a 5 km model could require up to a 20 teraflops computer system to meet operational schedules using this model.




Researchers have used supercomputers to model thunderstorms numerically. Wind, temperature, pressure, and other variables are calculated every few seconds at several hundred thousand locations in the area of the developing storm. Mathematical equations are then used to simulate the storm's evolution. In this graphic, particles which are released near the ground at regular intervals are colored orange when rising and blue when sinking. Yellow signifies paths of individual particles.

Cancer Genes

DNA is the blueprint of life, the molecular thread in the nucleus of each of our cells which guides the assembly of molecules and complete living organisms. When the DNA code is altered by mutation, serious diseases can result, such as cancer; this phenomenon was known to scientists studying animal tumors in the 1970's. They isolated cancer-causing genes, called "oncogenes" from animal tumors, and later found that similar genes existed in normal human DNA. This was a profound mystery. Why would we be carrying the seeds of our own destruction in our genetic blueprint?

In 1984, two separate research groups used a computerized searching algorithm to compare a newly discovered oncogene to all known genes. To their astonishment, the cancer causing gene matched a normal gene involved in growth and development. Suddenly, it became clear that cancer might be caused by a normal growth gene being switched on at the wrong time. This fundamental and unexpected insight was an early example of a field that is now known as Computational Biology, the science of using computers to store and analyze data from complex molecules in living cells.

The databases used in 1984 for these comparisons contained information about several thousand molecular units; now they contain over 30 million. Moreover, the current multi-agency genome research programs of NIH, NSF, and DOE will acquire data on tens of billions of molecular units, ranging from simple organisms to human beings. The best computer algorithms for determining the similarity of molecules require time proportional to the length of the molecules being compared; if the methods used to analyze oncogenes in 1984 were applied to the three billion base pairs of the human genome, they would require hundreds of years of computer time on today's fastest supercomputers. New computer designs and software methods will be essential to cope with the explosive growth of molecular data. Functioning as intellectual amplifiers to detect similarities and differences in molecules whose size and complexity are too vast for the unaided human mind, the computer systems to be developed by the HPCC program will be a critical tool for the life sciences in the 1990's, and the health care systems of the 21st Century.




Two views are shown for the binding of a chemical carcinogen to a short stretch of DNA. The stick model on the left shows the overall structure of the DNA-carcinogen complex. The space filling model on the right shows the actual intimate binding of the carcinogen, hiding deep within the DNA. This chemical contact leads to mutations in the DNA code, and ultimately to a tumor.

Cancer Genes

DNA is the blueprint of life, the molecular thread in the nucleus of each of our cells which guides the assembly of molecules and complete living organisms. When the DNA code is altered by mutation, serious diseases can result, such as cancer; this phenomenon was known to scientists studying animal tumors in the 1970's. They isolated cancer-causing genes, called "oncogenes" from animal tumors, and later found that similar genes existed in normal human DNA. This was a profound mystery. Why would we be carrying the seeds of our own destruction in our genetic blueprint?

In 1984, two separate research groups used a computerized searching algorithm to compare a newly discovered oncogene to all known genes. To their astonishment, the cancer causing gene matched a normal gene involved in growth and development. Suddenly, it became clear that cancer might be caused by a normal growth gene being switched on at the wrong time. This fundamental and unexpected insight was an early example of a field that is now known as Computational Biology, the science of using computers to store and analyze data from complex molecules in living cells.

The databases used in 1984 for these comparisons contained information about several thousand molecular units; now they contain over 30 million. Moreover, the current multi-agency genome research programs of NIH, NSF, and DOE will acquire data on tens of billions of molecular units, ranging from simple organisms to human beings. The best computer algorithms for determining the similarity of molecules require time proportional to the length of the molecules being compared; if the methods used to analyze oncogenes in 1984 were applied to the three billion base pairs of the human genome, they would require hundreds of years of computer time on today's fastest supercomputers. New computer designs and software methods will be essential to cope with the explosive growth of molecular data. Functioning as intellectual amplifiers to detect similarities and differences in molecules whose size and complexity are too vast for the unaided human mind, the computer systems to be developed by the HPCC program will be a critical tool for the life sciences in the 1990's, and the health care systems of the 21st Century.




Two views are shown for the binding of a chemical carcinogen to a short stretch of DNA. The stick model on the left shows the overall structure of the DNA-carcinogen complex. The space filling model on the right shows the actual intimate binding of the carcinogen, hiding deep within the DNA. This chemical contact leads to mutations in the DNA code, and ultimately to a tumor.

Predicting New Superconductors

In 1988, the world was excited by the discovery that a particular yttrium-barium-copper-oxide compound superconducts at a temperature of 93° Kelvin, still very cold, but much warmer than any previous superconductor. This discovery sparked a worldwide effort to expand research to discover new superconducting materials. The economic benefit of a high-temperature superconductor is beyond calculation, portending the development of, for example, much more efficient power transmission and lightweight, powerful magnets to revolutionize the transportation business. Advanced computing is a central part of the arsenal of research tools which will be necessary to reach that payoff.

Despite these early successes, many questions remain before it is understood how some materials superconduct when very similar compounds do not. This understanding will be critical to predicting new superconductors, which might work at even higher temperatures, be less expensive, carry more current, or be more amenable to manufacturing processes. Increasing progress in all these areas is required before the impact of these new materials will be felt.

The solution of physical models requires intensive calculations to understand the material structure. High performance computing can shorten the discovery process by allowing the development of accurate simulations to point experimenters in the most promising directions. For example, most of the groups looking for new superconductors are trying various copper-oxide combinations. Researchers are using high performance computers to explore the possibility of various combinations of elements that may lead to new superconducting materials.




Researchers have used visualization techniques to investigate the structure of materials which are thought to be involved in the superconducting mechanism. In this graphic, barium cations (green) and yttrium cations (silver) are shown with their associated oxygen defects. The copper atoms (blue) and oxygen atoms (red) are found in two types of CuO_4 environments, one of which is depicted in yellow, the other in light blue.

Predicting New Superconductors

In 1988, the world was excited by the discovery that a particular yttrium-barium-copper-oxide compound superconducts at a temperature of 93° Kelvin, still very cold, but much warmer than any previous superconductor. This discovery sparked a worldwide effort to expand research to discover new superconducting materials. The economic benefit of a high-temperature superconductor is beyond calculation, portending the development of, for example, much more efficient power transmission and lightweight, powerful magnets to revolutionize the transportation business. Advanced computing is a central part of the arsenal of research tools which will be necessary to reach that payoff.

Despite these early successes, many questions remain before it is understood how some materials superconduct when very similar compounds do not. This understanding will be critical to predicting new superconductors, which might work at even higher temperatures, be less expensive, carry more current, or be more amenable to manufacturing processes. Increasing progress in all these areas is required before the impact of these new materials will be felt.

The solution of physical models requires intensive calculations to understand the material structure. High performance computing can shorten the discovery process by allowing the development of accurate simulations to point experimenters in the most promising directions. For example, most of the groups looking for new superconductors are trying various copper-oxide combinations. Researchers are using high performance computers to explore the possibility of various combinations of elements that may lead to new superconducting materials.




Researchers have used visualization techniques to investigate the structure of materials which are thought to be involved in the superconducting mechanism. In this graphic, barium cations (green) and yttrium cations (silver) are shown with their associated oxygen defects. The copper atoms (blue) and oxygen atoms (red) are found in two types of CuO_4 environments, one of which is depicted in yellow, the other in light blue.

Air Pollution

The ability of the atmosphere to absorb and to cleanse itself of pollutant contaminants was taken for granted until the 1950's when "killer fogs" in London, England and Donora, Pennsylvania caused the deaths of hundreds of people. Since then, technological advances in controlling source emissions have reversed the trend of steadily increasing pollution, even in the face of continued industrial growth. However, reduction strategies of individual pollutant types do not always produce the desired results. In fact, these simple solutions can even make air quality worse due to complex chemical interactions of the remaining airborne contaminants. Pollutants may travel long distances from industrial centers to sensitive areas where they are deposited in transformed products such as ozone and acid rain. These complexities of pollutant transport and transformation are costly and difficult to study experimentally, therefore, numerical models of the atmosphere have been developed to assess the effects of man-made emissions on air quality.

The new Clean Air Act mandates the use of numerical models to demonstrate the effectiveness of proposed regulatory control strategies. The potential cost to society to implement these proposed controls is estimated to reach tens of billions of dollars. Current models have been useful in evaluating alternative control strategies, but do not yet have the capability to produce optimum solutions. Present computing limitations on existing supercomputers force simplifications in the scientific descriptions of chemical and physical processes, and slow examination of alternatives. Control strategies for each pollutant are often determined independently with little evaluation of multiple pollutant interactions. Remedial solutions determined for a particular scale of space and time are not readily extendable to other scales of pollutant dispersal.

High performance computing will enable multi-scale numerical explorations with cross-pollutant interactions to be performed in a timely manner so as to be useful to legislated control and prevention requirements. Advanced computer designs and software methods will also enable cost optimization of pollutant control tradeoffs. Improved visualization techniques will enhance the interpretation and evaluation of massive amounts of environmental measurement and computer simulation data. High performance computing models will lead to a better understanding of the actions needed to minimize pollutant damage to materials and environmental damage to crops while making our air safer to breathe for future generations.




Using a supercomputer, atmospheric researchers have simulated the transport, chemical transformation, and surface deposition of sulfur compounds responsible for acid rain. Visualization specialists have depicted the movement of sulfur compounds from major sources in the Ohio river valley to sensitive lakes in the Adirondacks. The yellow cloud represents high concentrations of sulfur compounds. Several sensitive aquatic regions are outlined in green and a typical wind flow pattern is presented in red.

Air Pollution

The ability of the atmosphere to absorb and to cleanse itself of pollutant contaminants was taken for granted until the 1950's when "killer fogs" in London, England and Donora, Pennsylvania caused the deaths of hundreds of people. Since then, technological advances in controlling source emissions have reversed the trend of steadily increasing pollution, even in the face of continued industrial growth. However, reduction strategies of individual pollutant types do not always produce the desired results. In fact, these simple solutions can even make air quality worse due to complex chemical interactions of the remaining airborne contaminants. Pollutants may travel long distances from industrial centers to sensitive areas where they are deposited in transformed products such as ozone and acid rain. These complexities of pollutant transport and transformation are costly and difficult to study experimentally, therefore, numerical models of the atmosphere have been developed to assess the effects of man-made emissions on air quality.

The new Clean Air Act mandates the use of numerical models to demonstrate the effectiveness of proposed regulatory control strategies. The potential cost to society to implement these proposed controls is estimated to reach tens of billions of dollars. Current models have been useful in evaluating alternative control strategies, but do not yet have the capability to produce optimum solutions. Present computing limitations on existing supercomputers force simplifications in the scientific descriptions of chemical and physical processes, and slow examination of alternatives. Control strategies for each pollutant are often determined independently with little evaluation of multiple pollutant interactions. Remedial solutions determined for a particular scale of space and time are not readily extendable to other scales of pollutant dispersal.

High performance computing will enable multi-scale numerical explorations with cross-pollutant interactions to be performed in a timely manner so as to be useful to legislated control and prevention requirements. Advanced computer designs and software methods will also enable cost optimization of pollutant control tradeoffs. Improved visualization techniques will enhance the interpretation and evaluation of massive amounts of environmental measurement and computer simulation data. High performance computing models will lead to a better understanding of the actions needed to minimize pollutant damage to materials and environmental damage to crops while making our air safer to breathe for future generations.




Using a supercomputer, atmospheric researchers have simulated the transport, chemical transformation, and surface deposition of sulfur compounds responsible for acid rain. Visualization specialists have depicted the movement of sulfur compounds from major sources in the Ohio river valley to sensitive lakes in the Adirondacks. The yellow cloud represents high concentrations of sulfur compounds. Several sensitive aquatic regions are outlined in green and a typical wind flow pattern is presented in red.

Aerospace Vehicle Design

Being able to predict the aerodynamic characteristics of in-flight vehicles is important for designers. Reproducing such flight regimes in wind tunnels is time consuming, costly, and in some cases impossible. Computational aerodynamics simulations are less costly and much faster than complex wind tunnel tests and are able to simulate many inaccessible flight regimes. These capabilities will be particularly important in the design of supersonic and hypersonic aircraft to serve international markets.

Computational aerosciences directly contributes to maintaining U.S. preeminence in aerospace science and engineering disciplines. The computational technology developed in such computational aerodynamics problems will directly transfer to the U.S. aerospace industry and aircraft manufacturers. Other potential beneficiaries are in diverse fields where fluid dynamics is an important design aspect such as automobile manufacture, ship design, and medicine (e.g., heart/cardiovascular flow simulation).

Massively parallel computing systems and advanced parallel software technology and algorithms will enable the development and validation of multidisciplinary, coupled models. These models will allow the numerical simulation and design optimization of complete aerospace vehicle systems throughout the flight envelope.




Shown here is a comparison between supercomputer simulation data and actual wind tunnel model data for the pressure distribution along the integrated space shuttle, solid rocket boosters and external tank flying at Mach 1.55. Note the excellent agreement between the two.

Aerospace Vehicle Design

Being able to predict the aerodynamic characteristics of in-flight vehicles is important for designers. Reproducing such flight regimes in wind tunnels is time consuming, costly, and in some cases impossible. Computational aerodynamics simulations are less costly and much faster than complex wind tunnel tests and are able to simulate many inaccessible flight regimes. These capabilities will be particularly important in the design of supersonic and hypersonic aircraft to serve international markets.

Computational aerosciences directly contributes to maintaining U.S. preeminence in aerospace science and engineering disciplines. The computational technology developed in such computational aerodynamics problems will directly transfer to the U.S. aerospace industry and aircraft manufacturers. Other potential beneficiaries are in diverse fields where fluid dynamics is an important design aspect such as automobile manufacture, ship design, and medicine (e.g., heart/cardiovascular flow simulation).

Massively parallel computing systems and advanced parallel software technology and algorithms will enable the development and validation of multidisciplinary, coupled models. These models will allow the numerical simulation and design optimization of complete aerospace vehicle systems throughout the flight envelope.



Shown here is a comparison between supercomputer simulation data and actual wind tunnel model data for the pressure distribution along the integrated space shuttle, solid rocket boosters and external tank flying at Mach 1.55. Note the excellent agreement between the two.


Energy Conservation and Turbulent Combustion

For the foreseeable future, 90% of the energy needs of the United States will be met by the combustion of fossil fuels. Two of the largest uses are for electrical power generation and in automobiles. Computational models offer a means of improving the design and efficiency of internal combustion engines.

Automobile engines are most efficient when run at high temperatures, but increased temperatures also lead to increased nitrogen oxide emissions. The burning of alternative fuels such as methanol is complicated by the emission of carcinogens. The effects of the emissions are influenced by local climatic conditions, making it necessary to consider the total system of fuel, engine, and atmosphere when seeking better designs. Our environment is too delicate and cannot be used as a testbed, so the atmospheric effects must be simulated.

A full three dimensional engine design code has been developed and implemented on a supercomputer. The code is designed to handle the most complex engines, such as the stratified charge and the two-stroke engines. The code represents an approximation to reality, because not all of the physical phenomena are well understood. Moreover, even if they were, the limited capabilities of existing computers would not allow this detailed information to be included.

For example, over 400 chemical processes of hydrocarbon and nitrogen chemistry are known to occur in internal combustion engines, yet only ten or less of the most significant reactions are used in simulations, in order to allow the calculation to run in a few hours on today's large supercomputers. Since the 400 hydrocarbon-nitrogen reactions are known, the real problem could be addressed with better algorithms running on a machine 10,000 times more powerful, a teraflops machine. This computational technology is needed by many private industrial engine firms as well as universities and government laboratories.



Researchers have used high performance computers to accomplish the numerical simulation of the properties of a conical fuel jet. The colors of the particles within the jet indicate the droplet size. The smallest particles are blue, intermediate size particles are shown in yellow, while the particles with the largest diameters are depicted in red. The smaller light blue particles around the jet represent the swirling air surrounding the spray.


Energy Conservation and Turbulent Combustion

For the foreseeable future, 90% of the energy needs of the United States will be met by the combustion of fossil fuels. Two of the largest uses are for electrical power generation and in automobiles. Computational models offer a means of improving the design and efficiency of internal combustion engines.

Automobile engines are most efficient when run at high temperatures, but increased temperatures also lead to increased nitrogen oxide emissions. The burning of alternative fuels such as methanol is complicated by the emission of carcinogens. The effects of the emissions are influenced by local climatic conditions, making it necessary to consider the total system of fuel, engine, and atmosphere when seeking better designs. Our environment is too delicate and cannot be used as a testbed, so the atmospheric effects must be simulated.

A full three dimensional engine design code has been developed and implemented on a supercomputer. The code is designed to handle the most complex engines, such as the stratified charge and the two-stroke engines. The code represents an approximation to reality, because not all of the physical phenomena are well understood. Moreover, even if they were, the limited capabilities of existing computers would not allow this detailed information to be included.

For example, over 400 chemical processes of hydrocarbon and nitrogen chemistry are known to occur in internal combustion engines, yet only ten or less of the most significant reactions are used in simulations, in order to allow the calculation to run in a few hours on today's large supercomputers. Since the 400 hydrocarbon-nitrogen reactions are known, the real problem could be addressed with better algorithms running on a machine 10,000 times more powerful, a teraflops machine. This computational technology is needed by many private industrial engine firms as well as universities and government laboratories.




Researchers have used high performance computers to accomplish the numerical simulation of the properties of a conical fuel jet. The colors of the particles within the jet indicate the droplet size. The smallest particles are blue, intermediate size particles are shown in yellow, while the particles with the largest diameters are depicted in red. The smaller light blue particles around the jet represent the swirling air surrounding the spray.

Microsystems Design and Packaging

Since the invention of the integrated circuit in 1958, the number of transistors fabricated on a microchip has doubled every two years, providing a medium to incorporate ever-increasing complex electronic design into chips, components, and packages. By analogy, the complexity of detail incorporated in a single integrated circuit 1 cm square is equivalent to representing the map features, at a city block level, of the entire Eastern United States. A similar analogy for a 5 cm square module densely packed with a collection of chips would be the equivalent of a map of North and South America. Determining the interconnecting paths, selecting the right modules, testing the interfaces, and choosing the mix of technologies are part of the design process to build the scalable components for workstations to supercomputers.

In the computing world, scalable architectures based on the 1–2 million transistor custom structures of today will evolve in the mid 1990's to system approaches exploiting 10–million-transistor chips, standard component parts, and special interface electronics, all combined in optimally designed modules adhering to standard interfaces. System clock speeds will continue to improve, and the diversity of microchip technologies combined in a single module will allow unprecedented flexibility for designers.

Managing this complexity explosion would be overwhelming without the use of computationally based approaches that enable teams of designers to systematically reduce the time to develop such systems. Today, computational tools enable complex microcircuits to be developed on first pass, at the same time that packaging technologies such as multichip modules can be demonstrated. High performance computing applied to the technology design process will enable the exploration of design alternatives and rapid exploitation of new technologies.




The figure shows both current and future components of the innovation design cycle. In the lower left of the photograph is a multiproject wafer using an industry/academia/government supported prototyping capability. The multiproject wafer shown is 4 inches in diameter, shares resources with 82 projects, and represents approximately 200 million transistors distributed over the wafer. Proven experimental prototypes are used as components of larger modules, interconnected with advanced technology. In the upper right hand corner is an example of an advanced interconnect module of 36 diverse microchips, packed together 1.5 inches on a side. Advanced computationally intensive design tools are essential to realize these high performance components.

Microsystems Design and Packaging

Since the invention of the integrated circuit in 1958, the number of transistors fabricated on a microchip has doubled every two years, providing a medium to incorporate ever-increasing complex electronic design into chips, components, and packages. By analogy, the complexity of detail incorporated in a single integrated circuit 1 cm square is equivalent to representing the map features, at a city block level, of the entire Eastern United States. A similar analogy for a 5 cm square module densely packed with a collection of chips would be the equivalent of a map of North and South America. Determining the interconnecting paths, selecting the right modules, testing the interfaces, and choosing the mix of technologies are part of the design process to build the scalable components for workstations to supercomputers.

In the computing world, scalable architectures based on the 1-2 million transistor custom structures of today will evolve in the mid 1990's to system approaches exploiting 10-million-transistor chips, standard component parts, and special interface electronics, all combined in optimally designed modules adhering to standard interfaces. System clock speeds will continue to improve, and the diversity of microchip technologies combined in a single module will allow unprecedented flexibility for designers.

Managing this complexity explosion would be overwhelming without the use of computationally based approaches that enable teams of designers to systematically reduce the time to develop such systems. Today, computational tools enable complex microcircuits to be developed on first pass, at the same time that packaging technologies such as multichip modules can be demonstrated. High performance computing applied to the technology design process will enable the exploration of design alternatives and rapid exploitation of new technologies.



The figure shows both current and future components of the innovation design cycle. In the lower left of the photograph is a multiproject wafer using an industry/academia/government supported prototyping capability. The multiproject wafer shown is 4 inches in diameter, shares resources with 82 projects, and represents approximately 200 million transistors distributed over the wafer. Proven experimental prototypes are used as components of larger modules, interconnected with advanced technology. In the upper right hand corner is an example of an advanced interconnect module of 36 diverse microchips, packed together 1.5 inches on a side. Advanced computationally intensive design tools are essential to realize these high performance components.

Earth's Biosphere

The Earth's biosphere is a complex physical system. There are a multitude of phenomena which can change the state of the biosphere on a local, regional, and global scale. In order to predict the directions and consequences of changes in the state and condition of the biosphere, detailed scientific models are required, which in turn are constructed from massive amounts of experimental and computational data.

Experimental information is derived from satellite and ground based sensors. By the late 1990's, the sensors will have the resolution required to provide data in support of much more accurate and informed decision making on issues such as pollution and global warming. Because the sensors will generate terabytes of data each day, which will be combined with local data sets on the ecosystem, major improvements in the capability for collecting, analyzing, distributing, and archiving data are necessary.

The effort of constructing valid scientific models which describe the dynamics and underlying processes of the biosphere will involve interdisciplinary teams of experts from the geophysical, life, physical, computer, and computational sciences. They will work together to construct computational models which will validate our empirical understanding of the biosphere, and help predict how worldwide activities affect the global ecosystem. Computer and computational scientists will develop the advanced software technology and algorithms for handling massive amounts of data and working with high resolution, coupled models of the Earth's atmospheric-biospheric-oceanic interactions. Efforts in software development and experimentation will predict how local current conditions may impact future global conditions allowing the linking of ecosystem models to global climate models. The result of this collaborative effort will be a much deeper understanding of our environment and the impact of human activities upon it.



A global view of the Earth's ocean chlorophyll and land vegetation is depicted here. The image was derived from accumulated satellite data. Generating this image required over 2 terabytes (10^{12} bytes) of raw data. Data such as these are the baseline "snapshot" of our current biosphere and are vital to understanding short and long term component interactions of the Earth's biosphere for such purposes as crop productivity improvement, fishing and weather prediction.

Earth's Biosphere

The Earth's biosphere is a complex physical system. There are a multitude of phenomena which can change the state of the biosphere on a local, regional, and global scale. In order to predict the directions and consequences of changes in the state and condition of the biosphere, detailed scientific models are required, which in turn are constructed from massive amounts of experimental and computational data.

Experimental information is derived from satellite and ground based sensors. By the late 1990's, the sensors will have the resolution required to provide data in support of much more accurate and informed decision making on issues such as pollution and global warming. Because the sensors will generate terabytes of data each day, which will be combined with local data sets on the ecosystem, major improvements in the capability for collecting, analyzing, distributing, and archiving data are necessary.

The effort of constructing valid scientific models which describe the dynamics and underlying processes of the biosphere will involve interdisciplinary teams of experts from the geophysical, life, physical, computer, and computational sciences. They will work together to construct computational models which will validate our empirical understanding of the biosphere, and help predict how worldwide activities affect the global ecosystem. Computer and computational scientists will develop the advanced software technology and algorithms for handling massive amounts of data and working with high resolution, coupled models of the Earth's atmospheric-biospheric-oceanic interactions. Efforts in software development and experimentation will predict how local current conditions may impact future global conditions allowing the linking of ecosystem models to global climate models. The result of this collaborative effort will be a much deeper understanding of our environment and the impact of human activities upon it.



A global view of the Earth's ocean chlorophyll and land vegetation is depicted here. The image was derived from accumulated satellite data. Generating this image required over 2 terabytes (10^{12} bytes) of raw data. Data such as these are the baseline "snapshot" of our current biosphere and are vital to understanding short and long term component interactions of the Earth's biosphere for such purposes as crop productivity improvement, fishing and weather prediction.

High Speed Networks

The development of more powerful computers feeds both the demand for, as well as the growth of, more powerful data communications capability. As computing technology progresses, greater demands are placed on network performance as researchers conceive of new tasks and modalities of use that require even higher performance.

Current developments in large scale scientific computing are leading to truly distributed computing, allowing a given job to be executed on several different machines communicating partial results among themselves, sharing in different facets of calculations, and jointly assembling a final result for output. Such intermachine communication is inherently capable of taking place at speeds that stress local area network technology and are a hundred or more times faster than are possible on today's long haul networks. As an example, some of the nation's leading radiology departments have committed to digital transfer of radiology images on broadband local networks operating at speeds of 100 to 1000 million bits per second.

In order for the NREN to support these and other demanding applications, not yet contemplated, a substantial directed research and development effort is needed in the areas of protocols (the formal structure of inter-computer communications), high speed computer interfaces for computers, and network equipment, such as switches. Multi-gigabit networks represent a change in kind, not just in degree, from today's networks. For example, consider that in a coast-to-coast communication at three gigabits per second there are at any instant "in flight" nearly nine megabytes of data, which is more than the memory of most personal computers and workstations.

Some gigabit research has already begun, and several experimental facilities have been established in a productive collaboration of academic, industrial, and governmental organizations, but HPCC support will be needed to carry the research program to the stage that commercial providers can use the technology to install and operate a multi-gigabit NREN.

Education Using the NREN

Computers increasingly fill important niches in all phases of the learning process, providing flexible instruments for interactive instruction and student based learning experiences. Their use in education both provides the skills needed to function in our increasingly technology-intensive world, and aids teaching and learning of all science and engineering topics. The development of the National Research and Education Network will accelerate and transfer the technology of computer communications to the needs of educators and students. The result will be to empower them to share resources and ideas on a national scale.

In a recent project, classes in many locations chose a day to measure the length of the sun's shadow from a vertical measuring stick on the school grounds at 12 noon. Each class consulted maps and geography books to find its school's latitude, and sent the results to a shared database located in the U.S. and Canada. All schools then received the database of measurements from around the world, and each class used the complete database to calculate the curvature of the earth, and from that, the earth's diameter. A normally dry recitation of facts became an engaging problem solving exercise because the students themselves derived the answer from their shared measurements. Along the way they learned geography, geometry, statistics, and how to collect and share data over computer networks.

This project, implemented by the use of the network, is a learning laboratory without walls, similar to the way research scientists take advantage of high speed digital networks to conduct shared research that is "distance independent." The HPCC Program will face the challenge of "scaling-up" today's Internet, making it "user-friendly" and improving its services so that it is readily accessible to all U.S. educational institutions.

High Speed Networks

The development of more powerful computers feeds both the demand for, as well as the growth of, more powerful data communications capability. As computing technology progresses, greater demands are placed on network performance as researchers conceive of new tasks and modalities of use that require even higher performance.

Current developments in large scale scientific computing are leading to truly distributed computing, allowing a given job to be executed on several different machines communicating partial results among themselves, sharing in different facets of calculations, and jointly assembling a final result for output. Such intermachine communication is inherently capable of taking place at speeds that stress local area network technology and are a hundred or more times faster than are possible on today's long haul networks. As an example, some of the nation's leading radiology departments have committed to digital transfer of radiology images on broadband local networks operating at speeds of 100 to 1000 million bits per second.

In order for the NREN to support these and other demanding applications, not yet contemplated, a substantial directed research and development effort is needed in the areas of protocols (the formal structure of inter-computer communications), high speed computer interfaces for computers, and network equipment, such as switches. Multi-gigabit networks represent a change in kind, not just in degree, from today's networks. For example, consider that in a coast-to-coast communication at three gigabits per second there are at any instant "in flight" nearly nine megabytes of data, which is more than the memory of most personal computers and workstations.

Some gigabit research has already begun, and several experimental facilities have been established in a productive collaboration of academic, industrial, and governmental organizations, but HPCC support will be needed to carry the research program to the stage that commercial providers can use the technology to install and operate a multi-gigabit NREN.

Education Using the NREN

Computers increasingly fill important niches in all phases of the learning process, providing flexible instruments for interactive instruction and student based learning experiences. Their use in education both provides the skills needed to function in our increasingly technology-intensive world, and aids teaching and learning of all science and engineering topics. The development of the National Research and Education Network will accelerate and transfer the technology of computer communications to the needs of educators and students. The result will be to empower them to share resources and ideas on a national scale.

In a recent project, classes in many locations chose a day to measure the length of the sun's shadow from a vertical measuring stick on the school grounds at 12 noon. Each class consulted maps and geography books to find its school's latitude, and sent the results to a shared database located in the U.S. and Canada. All schools then received the database of measurements from around the world, and each class used the complete database to calculate the curvature of the earth, and from that, the earth's diameter. A normally dry recitation of facts became an engaging problem solving exercise because the students themselves derived the answer from their shared measurements. Along the way they learned geography, geometry, statistics, and how to collect and share data over computer networks.

This project, implemented by the use of the network, is a learning laboratory without walls, similar to the way research scientists take advantage of high speed digital networks to conduct shared research that is "distance independent." The HPCC Program will face the challenge of "scaling-up" today's Internet, making it "user-friendly" and improving its services so that it is readily accessible to all U.S. educational institutions.

Glossary

ASTA

Advanced Software Technology and Algorithms

Bit

Acronym for binary digit

BRHR

Basic Research and Human Resources

Byte

A group of adjacent binary digits operated upon as a unit (usually connotes a group of eight bits)

CAD

Acronym for "computer aided design".

Computer Engineering

The creative application of engineering principles and methods to the design and development of hardware and software systems.

Computer Science

The systematic study of computing systems and computation. The body of knowledge resulting from this discipline contains theories for understanding computing systems and methods; design methodology, algorithms, and tools; methods for the testing of concepts; methods of analysis and verification; and knowledge representation and implementation.

Computational Science and Engineering

The systematic application of computing systems and computational solution techniques to mathematical models formulated to describe and simulate phenomena of scientific and engineering interest.

flops

Acronym for "floating point operations per second". The term "floating point" refers to that format of numbers which is most commonly used for scientific calculation. Flops is used as a measure of a computing system's speed of performing basic arithmetic operations such as adding, subtracting, multiplying, or dividing two numbers.

Giga-

10^9 or billions of ... (e.g.: gigabits)

Grand Challenge

A Grand Challenge is a fundamental problem in science and engineering, with broad economic and scientific impact, whose solution could be advanced by applying high performance computing techniques and resources.

HPCC

High Performance Computing and Communications

HPCS

High Performance Computing Systems

High Performance Computing

High performance computing encompasses advanced computing, communications, and information technologies, including scientific workstations, supercomputer systems, high speed networks, special purpose and experimental systems, the new generation of large scale parallel systems, and applications and systems software with all components well integrated and linked over a high speed national network.

Mega-

10^6 or millions of ... (e.g.: megaflops)

Network

Computer communications technologies that link multiple computers to share information and resources across geographically dispersed locations.

NREN

National Research and Education Network

ops

Acronym for "operations per second". Ops is used as a rating of the speed of computer systems and components. In this report ops is generally taken to mean the usual integer or floating point operations depending on what functional units are included in a particular system configuration.

Parallel Processing

Simultaneous processing by more than one processing unit on a single application.

Peta-

10^{15} or thousands of trillions of ... (e.g.: petabytes)

Supercomputer

At any given time, that class of general-purpose computers that are both faster than their commercial competitors and have sufficient central memory to store the problem sets for which they are designed. Computer memory, throughput, computational rates, and other related computer capabilities contribute to performance. Consequently, a quantitative measure of computer power in large-scale scientific processing does not exist and a precise definition of supercomputers is difficult to formulate.

Tera-

10^{12} or trillions of ... (e.g.: teraops)

Glossary

ASTA

Advanced Software Technology and Algorithms

Bit

Acronym for binary digit

BRHR

Basic Research and Human Resources

Byte

A group of adjacent binary digits operated upon as a unit (usually connotes a group of eight bits)

CAD

Acronym for "computer aided design".

Computer Engineering

The creative application of engineering principles and methods to the design and development of hardware and software systems.

Computer Science

The systematic study of computing systems and computation. The body of knowledge resulting from this discipline contains theories for understanding computing systems and methods; design methodology, algorithms, and tools; methods for the testing of concepts; methods of analysis and verification; and knowledge representation and implementation.

Computational Science and Engineering

The systematic application of computing systems and computational solution techniques to mathematical models formulated to describe and simulate phenomena of scientific and engineering interest.

flops

Acronym for "floating point operations per second". The term "floating point" refers to that format of numbers which is most commonly used for scientific calculation. Flops is used as a measure of a computing system's speed of performing basic arithmetic operations such as adding, subtracting, multiplying, or dividing two numbers.

Giga-

10^9 or billions of ... (e.g.: gigabits)

Grand Challenge

A Grand Challenge is a fundamental problem in science and engineering, with broad economic and scientific impact, whose solution could be advanced by applying high performance computing techniques and resources.

HPCC

High Performance Computing and Communications

HPCS

High Performance Computing Systems

High Performance Computing

High performance computing encompasses advanced computing, communications, and information technologies, including scientific workstations, supercomputer systems, high speed networks, special purpose and experimental systems, the new generation of large scale parallel systems, and applications and systems software with all components well integrated and linked over a high speed national network.

Mega-

10^6 or millions of ... (e.g.: megaflops)

Network

Computer communications technologies that link multiple computers to share information and resources across geographically dispersed locations.

NREN

National Research and Education Network

ops

Acronym for "operations per second". Ops is used as a rating of the speed of computer systems and components. In this report ops is generally taken to mean the usual integer or floating point operations depending on what functional units are included in a particular system configuration.

Parallel Processing

Simultaneous processing by more than one processing unit on a single application.

Peta-

10^{15} or thousands of trillions of ... (e.g.: petabytes)

Supercomputer

At any given time, that class of general-purpose computers that are both faster than their commercial competitors and have sufficient central memory to store the problem sets for which they are designed. Computer memory, throughput, computational rates, and other related computer capabilities contribute to performance. Consequently, a quantitative measure of computer power in large-scale scientific processing does not exist and a precise definition of supercomputers is difficult to formulate.


Tera-

10^{12} or trillions of ... (e.g.: teraops)

Grand Challenges 1993:

High Performance Computing and Communications

The FY 1993 U.S. Research and Development Program



A Report by the Committee on
Physical, Mathematical, and Engineering Sciences
Federal Coordinating Council for Science, Engineering,
and Technology

To Supplement the President's Fiscal Year 1993 Budget

1	2	3
4		5
6	7	8

On the Cover:

1. Fuel Combustion
2. High Speed Civil Transports
3. Rational Drug Design
4. Air Pollution
5. Venus Imaging
6. Magnetic Recording Technology
7. Ocean Modeling
8. Design of Protein Structures

The images used in this report were produced by ongoing scientific projects in areas of the HPCC Program. They were selected to illustrate the breadth of subject matter of the HPCC Program, and are elaborated upon in Chapter 4. The center figure is a visualization of the emerging National Research and Education Network, showing the interconnection of the network backbone with selected regional and campus-area networks.

Grand Challenges 1993:

High Performance Computing and Communications

The FY 1993 U.S. Research and Development Program

A Report by the Committee on
Physical, Mathematical, and Engineering Sciences

Federal Coordinating Council for Science,
Engineering, and Technology

Office of Science and Technology Policy

To Supplement the President's Fiscal Year 1993 Budget

Office of Science and Technology Policy
Federal Coordinating Council for Science,
Engineering, and Technology
Committee on Physical, Mathematical, and
Engineering Sciences

Chairman

Walter Massey, National Science Foundation

Members

William Happer, Department of Energy, Vice Chairman

Victor H. Reis, Department of Defense, Vice Chairman

Scott Sewell, Department of the Interior

Harry C. Mussman, Department of Agriculture

Robert White, Department of Commerce

Donald A. B. Lindberg, National Institutes of Health

John B. Childers, Department of Education

Norine Noonan, Office of Management and Budget

Karl Erb, Office of Science and Technology Policy

Richard H. Petersen, National Aeronautics and Space Administration

Erich Bretthauer, Environmental Protection Agency

Vernon Pankonin, National Science Foundation, Executive Secretary

FCCSET Directorate

Vickie V. Sutton, Assistant Director

Charles H. Dickens, Executive Secretary

**High Performance Computing and Communications and
Information Technology Subcommittee**

David Nelson, Department of Energy, Chairman

Stephen L. Squires, Defense Advanced Research Projects Agency, Co-Chairman

Lee Holcomb, National Aeronautics and Space Administration, Co-Chairman

A. Nico Habermann, National Science Foundation, Co-Chairman

Norman H. Kreisman, Department of Energy, Executive Secretary

James Burrows, National Institute of Standards and Technology

Steven Isakowitz, Office of Management and Budget

Daniel R. Masys, National Institutes of Health

Fred Scoresby Long, National Oceanic and Atmospheric Administration

Joan H. Novak, Environmental Protection Agency

Kirk Winters, Department of Education

Eugene Wong, Office of Science and Technology Policy

Editorial Group for Grand Challenges 1993

Melvyn Ciment, National Science Foundation, Co-Editor

William L. Scherlis, Defense Advanced Research Projects Agency, Co-Editor

Stephen M. Griffin, National Science Foundation, Executive Secretary

Charles N. Brownstein, National Science Foundation

John S. Cavallini, Department of Energy

Robin L. Dennis, Environmental Protection Agency

Daniel A. Hitchcock, Department of Energy

Paul E. Hunter, National Aeronautics and Space Administration

R.J. (Jerry) Linn, National Institute of Standards and Technology

Fred Scoresby Long, National Oceanic and Atmospheric Administration

David Rodbard, National Institutes of Health

Elliot R. Siegel, National Institutes of Health

Paul H. Smith, National Aeronautics and Space Administration

Judith Vaitukaitis, National Institutes of Health

Anthony Villasenor, National Aeronautics and Space Administration

Stephen Wolff, National Science Foundation

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF SCIENCE AND TECHNOLOGY POLICY
WASHINGTON, D. C. 20506

MEMBERS OF CONGRESS:

I am pleased to forward with this letter "Grand Challenges 1993: High Performance Computing and Communications, The FY 1993 U.S. Research and Development Program" prepared by the Committee on Physical, Mathematical, and Engineering Sciences of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), to supplement the President's Fiscal Year 1993 Budget. This report, the second in a series begun in 1991, describes the Presidential Initiative in High Performance Computing and Communications.

This Initiative is a well-coordinated, interagency research and development effort designed to sustain and extend U.S. leadership in all advanced areas of computing and networking. The Initiative, which represents over four years of planning, not only provides a far-sighted vision for investment in technology but also recognizes the importance of human resources and applications that serve major national needs. This sound national investment will bring both economic and social dividends, including advances in education, productivity, basic science, and technological innovation.

High performance computing and communications are pervasive tools for science and technology. Therefore, planning for this Initiative has included close cooperation with other Presidential Initiatives developed through the interagency FCCSET process, including Global Change Research, Advanced Materials and Processing, Biotechnology Research, and Mathematics and Science Education.

During the past year, substantial technical advances have been made in several areas of advanced computing, all of which strengthen the rationale for moving forward with this Initiative as rapidly as possible. The strong support for the President's Initiative by the Congress, academia, and leaders of the U.S. computer industry is greatly appreciated.

The coordination and integration of the interagency research and development strategy for this Initiative was led very ably by Dr. Walter E. Massey, Chairman of the Committee on Physical, Mathematical, and Engineering Sciences, and his interagency committee members, associates, and staff. They are all to be commended on the excellent work that is manifest in both the Initiative and the report.


D. Allan Bromley
Director

Table of Contents

	Page
Executive Summary	1
1. Program Goals and Overview	5
Introduction	
Needs and Benefits	
Program Description	
Goals	
Strategy	
Program Management Approach	
2. HPCC Program Components	13
High Performance Computing Systems	
Advanced Software Technology and Algorithms	
National Research and Education Network	
Basic Research and Human Resources	
3. Program Development and Agency Budgets	25
Program Planning	
Evaluation Criteria	
Agency Budgets	
Agency Program Descriptions	
Defense Advanced Research Projects Agency	
National Science Foundation	
Department of Energy	
National Aeronautics and Space Administration	
National Institutes of Health	
National Oceanic and Atmospheric Administration	
Environmental Protection Agency	
National Institute of Standards and Technology	
4. Grand Challenge and Supporting Technology Case Studies	41
Magnetic Recording Technology	
Rational Drug Design	
High Speed Civil Transports	
Catalysis	
Fuel Combustion	
Ocean Modeling	
Ozone Depletion	
Digital Anatomy	
Air Pollution	
Design of Protein Structures	
Venus Imaging	
Technology Links Research to Education	
Glossary	66

List of Figures

	Page
Figure 1. The High Performance Computing and Communications Program	3
Figure 2. Availability of Applications Software for High Performance Computing Systems	17
Figure 3. Network Applications	19
Figure 4. Network Connectivity	21
Figure 5. Evaluation Criteria for HPCC Program	26
Figure 6. HPCC Budgets by Agency and Program Component	28
Figure 7. HPCC Agency Responsibilities	30-31

Grand Challenges 1993:

High Performance Computing and Communications

The FY 1993 U.S. Research and Development Program

EXECUTIVE SUMMARY

High performance computing and computer communications networks are increasingly important to scientific advancement, economic strength, and national security. The goal of the Federal High Performance Computing and Communications (HPCC) Program is to accelerate significantly the availability and utilization of the next generation of high performance computers and networks in a manner consistent with the Strategic and Integrating Priorities shown in Figure 1.

The HPCC Program is the result of several years of effort on the part of senior government, industry, and academic scientists and managers to design a research agenda to extend U.S. leadership in high performance computing and networking technologies. The program is planned, funded, and managed with close cooperation among Federal agencies and laboratories, private industry, and academe to ensure that the fruits of this research program are brought into the educational and commercial marketplaces as rapidly as possible.

There have been steady accomplishments over the past year. Several participating agencies have begun to fund HPCC research groups, centers, and consortia on various grand challenge application problems. Major new scalable high performance systems have been announced and delivered. New software applications have been developed or ported to the emerging high performance systems. Traffic on the operational network has doubled in the past year, as has the number of interconnected local and regional networks. In the past year, a large number of researchers, scholars, students, educators, scientists, and engineers have been trained to use these emerging technologies.

Numerous organizations have recently undertaken studies that have provided valuable feedback on the structure and content of the program. Some of the organizations that supported studies include the National Academy of Sciences; EDUCOM and the Computing Research Association, each representing numerous

universities; the Computer Systems Policy Project, representing leading U.S. computer companies; the Office of Technology Assessment; professional societies including Association for Computing Machinery, Institute of Electrical and Electronic Engineers, American Mathematical Society, and others; and the Federal Networking Advisory Committee.

For FY 1993 the HPCC Program proposes to invest \$803 million in the four complementary and coordinated components shown in Figure 1. This investment represents a \$148 million, or 23%, increase over the FY 1992 enacted level.

The HPCC Program is driven by the recognition that unprecedented computational power and its creative use are needed to investigate and understand a wide range of scientific and engineering "grand challenge" problems. These are fundamental problems whose solution requires significant increases in computational capability and is critical to national needs. Progress toward solution of these problems is essential to fulfilling many of the missions of the participating agencies. Examples of grand challenges addressed include: prediction of weather, climate, and global change; determination of molecular, atomic, and nuclear structure; understanding turbulence, pollution dispersion, and combustion systems; understanding the structure of biological macromolecules; improving research and education communications; understanding the nature of new materials; and solving problems applicable to national security needs.

Other Presidential Initiatives including Global Change Research, Advanced Materials and Processing, Biotechnology, and Mathematics and Science Education depend on the capabilities that this initiative will produce. In addition, the generic technologies developed will make possible advances in many other areas of direct benefit to millions of Americans.

The HPCC Program nurtures the educational process at all levels by improving academic research and teaching capabilities. Advanced computing and computer communications technologies will accelerate the research process in all disciplines and enable educators to integrate new knowledge and methodologies directly into course curricula. Students at all levels will be drawn into learning and participating in a wide variety of research experiences in all components of this program.

The FY 1993 Program and this document were developed by the High Performance Computing and Communications and Information Technology Subcommittee under the direction of the Committee on Physical, Mathematical, and Engineering Sciences (PMES) of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET).

FIGURE 1

The High Performance Computing and Communications Program

Goals: Strategic Priorities

- Extend U.S. technological leadership in high performance computing and computer communications.
- Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment.
- Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.

Strategy: Integrating Priorities

- Support solutions to important scientific and technical challenges through a vigorous R&D effort.
- Reduce the uncertainties to industry for R&D and use of this technology through increased cooperation among government, industry, and universities and by the continued use of government and government-funded facilities as a prototype user for early commercial HPCC products.
- Support the underlying research, network, and computational infrastructures on which U.S. high performance computing technology is based.
- Support the U.S. human resource base to meet the needs of industry, universities, and government.

Program Components

High Performance Computing Systems (HPCS)

Research for Future Generations of Computing Systems
System Design Tools
Advanced Prototype Systems
Evaluation of Early Systems

Advanced Software Technology and Algorithms (ASTA)

Software Support for Grand Challenges
Software Components and Tools
Computational Techniques
High Performance Computing Research Centers

National Research and Education Network (NREN)

Interagency Interim NREN
Gigabits Research and Development

Basic Research and Human Resources (BRHR)

Basic Research
Research Participation and Training
Infrastructure
Education, Training, and Curriculum

1. PROGRAM GOALS AND OVERVIEW

Introduction

High performance computing continues to gain recognition as a powerful technology for industrial design and manufacturing, scientific research, communications, and information management. A robust U.S. high performance computing and computer communications capability contributes to leadership in critical technology and national security areas. Improved computational and communications technologies contribute to more effective approaches to problem solving, new products and services, and enhanced national competitiveness across broad sectors of the economy. The High Performance Computing and Communications (HPCC) Program was initiated last year in the President's Fiscal Year 1992 Budget. It will accelerate the development of a thousand-fold improvement in useful computing capability and a hundred-fold improvement in available computer communications capability by 1996, and it will enhance the range of scientific and engineering disciplines that can effectively exploit this computational capability.

Based on several years of planning under the auspices of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), Federal agencies in collaboration with scientists and managers from U.S. industry, universities, and laboratories have developed the HPCC Program to meet the challenges of advancing computing and associated communications technology and practices. Agencies have realigned and enhanced their high performance computing research and development programs, coordinated their activities with other agencies, and shared common resources to develop the program described in this document.

Planning for the HPCC Program has proceeded under the leadership of the President's Office of Science and Technology Policy. This program received added impetus by the passage of a multi-year authorization in the High Performance Computing Act of 1991 (P.L. 102-194), signed by President Bush on December 9, 1991. On that occasion, the President stated, "The development of high performance computing and communications technology offers the potential to transform radically the way in which all Americans will work, learn and communicate in the future. It holds the promise of changing society as much as the other great inventions of the 20th century, including the telephone, air travel and radio and TV."

There have been steady accomplishments over the past year. Several participating agencies have begun to fund application groups, centers, and consortia on various grand challenge problems. Major new scalable high performance systems have been announced and delivered. New software applications have been developed or ported to the emerging high performance systems. Traffic on the operational network has doubled in the past year, as has the number of interconnected local and regional networks. In addition, large numbers of researchers, scholars, students, educators, scientists, and engineers have been trained to use these emerging technologies.

Needs and Benefits

High performance computing has become a vital enabling force in the conduct of science and engineering research over the past three decades. Computational science and engineering has joined, and in some areas displaced, the traditional methods of theory and experiment. This trend has been powered by computing hardware and software, computational methodologies and algorithms, availability and access to high performance computing systems and infrastructure, and the growth of a trained pool of scientists and engineers. This process has been nurtured by Federal investment in advanced research, agency supercomputer centers, and national networks through DARPA, DOD, DOE, ED, EPA, NASA, HHS/NIH, DOC/NIST, DOC/NOAA, NSA, and NSF. These facilities have contributed to national mission areas such as energy, space, health, defense, environment, weather, and basic science and technology that could not be effectively addressed without the use of such advanced facilities.

High performance computing technology is knowledge and innovation intensive. Its development and use engages the entire scientific and engineering community. Building upon fundamental research of the early 1980's, a new computing technology of advanced parallel processing computers has emerged. This innovative approach to high performance computing systems promises to achieve sustained performance improvements of a thousand-fold compared to the systems at the start of the HPCC Program.

In a growing number of science and technology fields, progress and productivity in modern research are increasingly dependent on the close interaction of people located in distant places, sharing and accessing computational resources across networks. Although the U.S. is the world leader in most of the critical aspects of computing and computer communications technology, this lead is being challenged.

The HPCC Program is a strategic Federal investment in the frontiers of computing and computer communications technologies and is formulated to satisfy national needs from a variety of perspectives including: technology, science applications, human resources, and technology transition. Needs are derived from the agency missions and based on the underlying science, engineering, and technology base required to carry out these missions. Many of these mission needs are related to solving very intensive large scale computing problems. These fundamental problems often cut across agencies and missions and are called "grand challenge" problems.

The industrial and academic sectors provide major sources of innovation, cost-effective development, and support of information technologies and their application to grand challenge problems. As these technologies are developed, the results support the Federal agency missions and become available nationally. The program provides for development of these revolutionary technologies within a framework of a partnership among government, industry, and academe, which promotes the rapid transition of laboratory results into new products that can then be applied within the program.

During the past year substantial advice and recommendations have come from industrial, professional, and scientific organizations including the National Academy of Sciences; EDUCOM and the Computing Research Association, each representing numerous universities; the Computer Systems Policy Project, representing leading U.S. computer companies; the Office of Technology Assessment; professional societies including Association for Computing Machinery, Institute of Electrical and Electronic Engineers, American Mathematical Society, and others; and the Federal Networking Advisory Committee. This has helped focus planning and management activities and strengthen coordination with these other sectors.

Program Description

The Program consists of four integrated components representing the key areas of high performance computing and communications:

High Performance Computing Systems (HPCS) - the development of the underlying technology required for scalable high performance computing systems capable of sustaining trillions of operations per second on large problems. Research in very high performance systems is focusing both on increasing the absolute level of performance attainable and on reducing the cost and size of these very high performance systems in order to make them accessible to a broader range of applications.

Advanced Software Technology and Algorithms (ASTA) - the development of generic software technology and algorithms and the deployment of the most innovative systems for grand challenge research applications in a networked environment.

National Research and Education Network (NREN) - the development of a national high speed network to provide distributed computing capability to research and educational institutions and to further advanced research on very high speed networks and applications.

Basic Research and Human Resources (BRHR) - support for individual investigator and multidisciplinary long term research drawn from diverse disciplines, including computer science, computer engineering, and computational science and engineering; initiation of activities to significantly increase the pool of trained personnel; and support for efforts leading to accelerated technology transition.

Advances in high performance computing enable advances in almost every other science and engineering discipline. There is a complex web of research interdependencies among the four components, and each area contributes to progress in other areas. Because of these dependencies, achieving and maintaining balance between the research components is a primary goal and the most important priority in the current context and environment. The HPCC Program is designed to provide balanced support both for technology areas (including components, systems, software, and algorithms) and for applications, infrastructure, and human resources to achieve rapid overall research progress and productivity.

The component activities are planned to produce a succession of intermediate benefits on the way to meeting the long range programmatic goals. The HPCC Program builds on Federal programs already in place, providing additional resources in selected areas. Computational science and engineering grand challenges that emerge from all science and engineering disciplines are the focal points for these efforts.

Goals

The goals of the High Performance Computing and Communications Program are to:

- Extend U.S. technological leadership in high performance computing and computer communications.
- Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment.
- Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.

These goals will be realized by achieving: computational performance of one trillion operations per second (10^{12} ops, or teraops) on a wide range of important applications; development of associated system software, tools, and improved algorithms for a wide range of problems; a national research network capable of one billion bits per second (10^9 bits, or gigabits); sufficient production of Ph.D.s and other trained professionals per year in computational science and engineering to permit effective use and application of these new technologies.

Strategy

The goals will be met through coordinated government, industry, and university collaboration to:

- Support solutions to important scientific and technical challenges through a vigorous R&D effort.
- Reduce the uncertainties to industry for R&D and use of this technology through increased cooperation among government, industry, and universities and by the continued use of government and government-funded facilities as a prototype user for early commercial HPCC products.
- Support the underlying research, network, and computational infrastructures on which U.S. high performance computing technology is based.
- Support the U.S. human resource base to meet the needs of industry, universities, and government.

At the program component level the strategy will: exploit and extend scalable parallelism and engage in intensive system software development in HPCS; use common requirements of the grand challenges to foster software progress in ASTA; evolve from the current Internet and other Federally managed networks supporting research and education to the NREN using an evolutionary series of systems; and strengthen academic and human resources development activities in computer science and computational science and engineering as part of BRHR.

Program Management Approach

The High Performance Computing and Communications Program is implemented as a partnership among Federal agencies and other organizations. Major portions of the program are cost-shared and leveraged by the participation of industry and universities.

Leadership for the Program is provided by the Office of Science and Technology Policy, through the Federal Coordinating Council on Science, Engineering, and Technology (FCCSET) Committee on Physical, Mathematical, and Engineering Sciences (PMES). The membership of PMES includes senior executives of many Federal agencies. Planning for the HPCC Program is coordinated by the PMES High Performance Computing, Communications, and Information Technology (HPCCIT) Subcommittee. The HPCCIT, currently led by DOE, meets regularly to coordinate agency HPCC programs through information exchanges, the common development of interagency programs, and the review of individual agency plans and budgets.

Under the HPCCIT there are currently several task groups that coordinate activities in specific areas, including Networking, Research, Education, and Applications. From time to time, individual agencies are assigned responsibility to lead the coordination of the HPCCIT and the task groups. The Applications group, currently led by NASA, coordinates activities related to grand challenge applications, software tools needed for applications development, and software development at high performance computing centers. The Networking group, currently led by NSF, coordinates network integration activities and works closely with the Federal Networking Council (FNC), which consists of representatives from interested federal agencies. The FNC is responsible for coordinating the efforts of government HPCC participants and other NREN constituents, in addition to providing a liaison to others interested in the Federal program. The Research group, currently led by DARPA, addresses a broad range of issues including basic research progress, technology trends, and alternative approaches to deal with technology limits in information technology. The Education group, currently led by NIH, coordinates education and training aspects of HPCC and provides liaison with other education-related activities under the FCCSET.

In order to improve collaboration with U.S. industry, universities, state and local governments, and other non-Federal communities, a Federal Networking Advisory Committee has been created to support the FNC, and pursuant to P.L. 102-194, a

High Performance Computing Advisory Committee will be established to support the overall HPCC Program.

The participating agencies have altered their organizational or management structure in order to facilitate participation in the HPCC Program. Each agency has established a focal point for matters related to the HPCC Program. (See inside back cover for agency contact addresses.) DARPA has established a Computing Systems Technology Office which serves as the DOD focal point for HPCC. In addition, DARPA has created a High Performance Computing Joint Program Office to coordinate advanced technology development within the DOD and cooperatively with other agencies. NSF has established an HPCC Coordination Committee with budget, planning, and oversight responsibilities. The DOE has expanded its Scientific Computing Staff to respond to this initiative and has assigned to it overall responsibility for the DOE HPCC program component. NASA has responded by creating an HPCC Program Office within the Office of Aeronautics and Space Technology to coordinate, plan, and implement HPCC technology development.

Most of the participating agencies have published solicitations for research in HPCC areas. U.S. industry, academia, and other developers and users of HPCC technology are involved in program planning and execution through advisory committees, commissioned reviews, and self-generated commentary; as well as through direct participation as performers of R&D and as suppliers of technology.

The HPCC Program has benefitted from the interest, advice, and specific recommendations that have come from various industrial, professional, and scientific organizations. Some of these organizations include the National Academy of Sciences; EDUCOM and the Computer Research Association, each representing numerous universities; the Computer Systems Policy Project (CSPP), representing leading U.S. computer companies; the Office of Technology Assessment; professional societies including Association for Computing Machinery, Institute of Electrical and Electronics Engineers, American Mathematical Society, and others. As an example of the dynamism and flexibility of the HPCC program, many of these recommendations are already being incorporated and responded to in the current HPCC Program plans. For example, the CSPP has conducted an intensive study of the HPCC Program structure and published their views in December 1991. The CSPP has noted that the HPCC program is a significant, critical, and necessary undertaking by the Federal government, and has provided recommendations that have helped to focus discussion on four important issues below:

Expanding the vision of the HPCC Program to include research on generic, enabling technologies to support a wider range of applications such as better health care, lifelong learning, improved services for senior citizens, enhanced industrial design, and broad access to public and private databases. The HPCC program emphasizes investment in many of the supporting advanced technologies such as open scalable systems, portable software, mass data storage, and advanced network protocols. In addition, the HPCC program is making additional efforts to inform industry and academia to propose and enter collaborations on specific grand challenge applications.

Establishing a technology and policy foundation for an information and communications infrastructure for the future. To address this issue, the Federal Networking Council, an interagency coordination group that supports the development of the National Research and Education Network, has formed a policy committee to address critical network issues such as security and privacy, intellectual property rights, and network access, interoperability, and technology transfer.

Improving management and governance of the HPCC program and increasing opportunities for industry participation. The HPCC program will continue to strengthen its current coordinated management structure as the program matures and more activities shift from planning to execution phase. This should provide for validation of policies and plans, enable detailed monitoring of progress towards goals and objectives, better satisfy Congressional reporting requirements, and serve as a single point-of-contact for industry, academia, and other government agencies. Over the coming months, the Administration will be exploring alternative management approaches that might bring additional coordination and accountability to the program. To improve technology transfer, the Department of Energy is supporting the development of a model Cooperative Research and Development Agreement (CRADA) that will ease significantly the exchange of computing technology between private industry and government laboratories.

Formulating the HPCC budget priorities to achieve a balanced program. The budget proposes the largest dollar increase for ASTA, the program component which addresses grand challenge applications and software research. The program will continue to support basic research, generic technologies, and broad applications, but will work in the coming years to sustain an appropriate funding balance between software research and computer equipment and facilities.

The HPCC Program is executed according to management principles that have proven effective in other focused research programs in government and industry. Some of these principles include:

Sustain balance and critical mass. The program must achieve sufficient scope and balance among the components. The HPCC Program must operate at a sufficient scale and coverage of technology areas that the new technologies can be effectively applied to grand challenge problems with acceptable levels of risk.

Build on agency strengths. The strategy builds on agency strengths by giving appropriate agencies the responsibility to coordinate activities in areas of demonstrated capability. Joint programs are developed where agencies have complementary strengths in particular technology areas.

Accelerate technology transfer. The transfer of technology from research to development and to application is a principal challenge addressed by the HPCC Program. Barriers to acceptance of new technologies can include high initial cost, inadequate and user-unfriendly software, and lack of standards. The HPCC Program addresses barriers such as these by relying on substantial industry participation, by sharing development among agencies, and by involving users early in the process. Risk is minimized by adopting a

mixed strategy, in which multiple approaches are pursued to technology issues such as high performance architecture, microelectronics, and system software.

Seek broadly applicable technology. The program emphasizes general-purpose solutions with wide applicability, rather than promoting the development of special-purpose systems tailored to specific problem domains. Wide diffusion of new technologies into the national technology base will result in cost-effective support for the broadest range of applications on a diverse base of computing resources, facilitating the application of high performance technologies to major national problems such as health and education.

Overcome barriers. The program emphasizes scalable, interoperable technologies in order to overcome costs in creating successive generations of high performance computers and software, thereby facilitating evolution of applications to increasingly powerful and flexible computing, networking, and software technologies.

In the coming year, the HPCC Program will continue to strengthen its management approach for improving planning, assessment, budgeting, and reporting.

2. HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM COMPONENTS

The HPCC Program is composed of four integrated and coordinated components that are designed to enhance scientific and engineering productivity and support long term agency needs. The emerging computing environment is that of advanced workstations with high resolution displays connected to a high speed computer communications network and high performance computing resources. The regional and national networks provide access to additional high performance computing systems and research resources. Realizing the full potential of teraops computing and gigabits networking systems will require advanced software technology and algorithms and people educated and trained to use these tools and resources in this dynamic field.

High Performance Computing Systems (HPCS)

The HPCS component produces scalable parallel computing systems through the development of prototype systems. The program is designed to attack computational science problems by developing innovative systems that will provide a one-hundred- to one-thousand-fold increase of sustained computing capability over machines that were in conventional use at the start of the program.

The program is structured to focus on technological challenges in the earliest stages of the product development cycle. Critical underlying technologies are developed in prototype form along with associated design tools. This allows empirical evaluation of alternative solutions as the prototype systems mature. Evaluation is performed throughout the development cycle, with experimental results being fed back into the design process to refine successive generations of systems. There is risk inherent in creating new technologies, and each project will be managed according to its proximity to commercial introduction. For instance, larger projects which are close to yielding commercial products are performed on a cost-shared basis with industry.

HPCS is composed of the four subcomponents described below:

Research for Future Generations of Computing Systems. This activity focuses on the underlying architecture, components, packaging, systems software, and scaling concepts to achieve affordable high performance systems. These projects ensure that the required advanced technologies will be available for the new systems while providing a foundation for the more powerful systems that need to follow. This activity also produces system software to support heterogeneous configurations of workstations and high performance servers.

System Design Tools. This activity develops computer aided design tools and frameworks for enabling multiple tools to work together to enable design, analysis, simulation, and testing of systems components. The tools will enable rapid prototyping of new system concepts.

Advanced Prototype Systems. Systems capable of scaling to 100 gigaops performance have already been produced. The teraops level will be reached by the mid-1990's. Research in very high performance systems is focusing both on increasing the absolute level of performance attainable and on reducing the cost and size of these very high performance systems in order to make them accessible to a broader range of applications.

Evaluation of Early Systems. Experimental systems will be placed at sites with high levels of expertise to provide feedback to systems architects and software designers. Performance evaluation criteria for systems and results of evaluation will be made widely available. Because of scalability, early systems can be acquired in smaller configurations to evaluate their potential performance. As noted below, the ASTA component will support, on a cost shared basis, the acquisition of large scale systems for experimental use in grand challenge applications.

Accomplishments:

- High performance computing systems capable of scaling to 100 gigaops have been announced. Systems in this performance class will eventually be available in multiple architectures, many of which represent hybridizations of computational models used in earlier systems.
- New models of computation that may yield higher performance for many classes of problems, including more advanced hybrid models, are being explored through joint university-industry projects.
- Scalable mass storage systems have been designed, prototyped, and demonstrated in collaboration with industry. The scalable mass storage approach is now being used in most new high performance computing systems. Very large capacity (10^{15} bits) archival storage systems have also been developed.
- Operating system technology has been developed that is capable of efficiently supporting parallel and heterogeneous distributed applications involving configurations of high performance systems and workstations.

Advanced Software Technology and Algorithms (ASTA)

The ASTA component of the HPCC Program produces generic software technology, new algorithms, and early prototype applications software to support a broad range of high performance applications of heterogeneous and networked high performance computing systems. Significant improvements in algorithm design and software technology are essential to achieving sustained high levels of computing system performance. In addition, improvements in algorithms and software will open up major new areas of application of high performance computing technology in science, engineering, health care, education, national security, and other areas. Parallel and distributed systems provide new opportunities and challenges for software and algorithms. Advances in generic software are needed to minimize the time researchers in other scientific disciplines spend to master advanced, complex computing science skills.

A typical user environment can include various kinds of advanced workstations with high resolution color displays connected to a high performance network with diverse local high performance computing resources. Access to additional high performance computing systems is available through regional and national networks. Systems software will support these heterogeneous hardware and network configurations, with software support for library, user interaction, and other functions common across applications.

Support for Grand Challenges. Grand challenge problems currently addressed by the HPCC Program are of such a magnitude and complexity that they will require continuous advances in computational power and improvements in computational models for the next decade or more. Grand challenge problems will be selected for their national importance, for their potential to challenge computational technologists to advance their solution, for the extent of cost sharing with sources directly concerned with the specific applications, and for the leverage they can exert on other related technology developments. Grand challenge problems that will take advantage of the capabilities developed by the HPCC Program are contained in many other Federal initiatives, including Global Change Research, Advanced Materials and Processing, Biotechnology, and Mathematics and Science Education. Support will be provided to improve algorithms for hybrid computational models, multidisciplinary applications, and heterogeneous and distributed problems.

Software Components and Tools. Users of high performance systems have common needs in software technology and programming environments, including software tools, systems software, and support for reuse and for software libraries. Software tools include, for example, higher level languages, advanced compiler technology, optimization and parallelization tools, interoperability support, data management, visualization, debugging and analysis, instrumentation and performance measurement, and other tools. The availability of supported applications software is generally less for emerging new architectures than it is for established ones. Increasing support for software components and tools, with an emphasis on generic multi-architecture capability, will help accelerate software product development. (See Figure 2.)

Because of this very wide range of software tools required to develop and manage high performance applications, tool development efforts in HPCC will emphasize the integration of multiple tools developed by different organizations for use on multiple hardware platforms. Integration technology increases opportunity for commercial participation in tool development. Systems software efforts will focus on high performance support for heterogeneous configurations of systems in order to maintain flexibility and growth potential in user environments.

Finally, software library efforts will focus on development of common architectures and interfaces to increase the potential for reusability despite multiple underlying models of computation, the diversity of programming languages in use, and the varying degree of assurance provided with software components. Library efforts will lead to a software interchange capability that can support exchange of software components among researchers, enabling collaborative development of systems through appropriate sharing of software assets.

Computational Techniques. Scalable and portable software libraries are being developed to enable software to move from one generation to the next and across different computational platforms. A variety of experimental applications are being developed to evaluate the performance and generality of the new computing technologies. Research will include support for new parallel algorithms, numerical and mathematical analyses, parallel languages, and conceptual models of scientific applications. To support visualization of data and processes, standard systems level tools will be developed. Data interpretation and communications will be vigorously addressed to enable real-time visualization in heterogeneous workstation/server configurations.


High Performance Computing Research Centers. The HPCC Program supports the deployment of innovative high performance computing architectures to high performance computing research centers for use by computational scientists working on grand challenge and disciplinary applications. High performance computing research centers generally acquire early high performance computing systems which are integrated into prototype heterogeneous computing configurations and connected to the developing NREN. This allows for a wide spectrum of experiments for scalability studies, as well as an opportunity to provide grand challenge researchers access to the largest possible advanced systems. As part of their efforts, the high performance computing research centers also support researchers in areas such as algorithms development, software environment and tool development, mass storage, and scientific visualization.

Accomplishments:

- Collaborations between agencies and industry have been initiated on grand challenge problems in quantum chemistry, environmental remediation, ground-water simulation, materials science problems, air pollution models, design of catalysts for chemical processes, and other areas.
- Consortia, centers, and groups are being established involving industrial, academic, and government participants to address applications software development, generic systems software, and effective sharing of leading-edge computational resources.
- Progress in algorithms and applications with direct applications to improved health care has been significant, ranging from the design of pharmaceuticals to enhancing understanding of fundamental biological processes. Already, order of magnitude speedups have been achieved for codes such as three-dimensional reconstruction of two-dimensional electron micrographs of such viruses as the herpes simplex virus—the cause of cold sores or fever blisters. Other health-care related applications include prosthesis design.
- A parallel high performance computing system has been used successfully to develop a method for enhancing the data from the Hubble Space Telescope. This method, called the Maximum Entropy Method, increases resolution by a factor of four over the degraded data.

FIGURE 2

Availability of Applications Software for High Performance Computing Systems



National Research and Education Network (NREN)


The NREN is both a goal of the HPCC Program and a key enabling technology for success in the other components. As used in this report, the NREN is the future realization of an interconnected gigabit computer network system supporting HPCC. The NREN is intended to revolutionize the ability of U.S. researchers and educators to carry out collaborative research and education activities, regardless of the physical location of the participants or the computational resources to be used. (See Figure 3.) As its name implies, NREN is a network for research and education, not general purpose communication. Nonetheless, its use as a testbed for new communications technologies is vital. A fundamental goal of the HPCC is to develop and transfer advanced computing and communications technologies to the private sector of the U.S. as rapidly as possible, and to enhance the nation's research and education enterprise. The development and deployment of advanced applications, such as image visualization and distributed computing, will be applied to problems such as medical diagnosis, aerodynamics, advanced materials, and global change, and will provide the impetus necessary for transferring the supporting technologies and capabilities throughout the U.S. science, technology, and education infrastructure. These capabilities and technologies will be developed through the cooperative effort of U.S. industry, the Federal government, and the educational community.

In order to ensure a coordinated Federal program in this area, a Federal Networking Council (FNC) has been established consisting of agency representatives. The FNC has the primary responsibility for coordinating the efforts of government HPCC participants and other NREN constituents, in addition to providing a liaison with others involved or interested in the program. To further collaboration with non-Federal communities such as state and local government, U.S. industry, and private education, the FNC has established the Federal Network Advisory Committee, consisting of distinguished members from those communities.

The NREN component of the HPCC Program is comprised of two interrelated and complementary subprograms, the Interagency Interim NREN subcomponent, and the Gigabit Research and Development subcomponent.

Interagency Interim NREN. The Interagency Interim NREN is an evolving operating network system. Near term (1992-1996) research and development activities will provide for the smooth evolution of this networking infrastructure into the future gigabit NREN. Interagency Interim NREN activities will achieve this goal by expanding the connectivity and enhancing the capabilities of the Federally funded portion of today's research and education networks, and by deploying advanced technologies and services as they mature. The Interagency Interim NREN, which is based on DARPA's internet technology, builds on NSF's NSFNET, DOE's Energy Sciences Network, NASA's Science Internet and other networks supporting research and education. (See Figure 4.) During 1992, the Interagency Interim NREN will accelerate the introduction of commercial 45 megabit transmission technologies and services into operational use, including the Switched Multi-Megabit Data Services (SMDS). These advances will support

FIGURE 3
Network Applications



NREN supported activities will promote enhancements in networking infrastructure both for the connectivity to many geographically distributed science facilities, information resources, and research teams and also for the technologies and tools to perform the distributed computing, data access, and multi-media processes as illustrated above.

both HPCC and non-HPCC related research and education, extend our understanding of the uses of moderately high capability networks and promote the evolution of Interagency Interim NREN performance toward gigabit speeds.

Additional research and development activities include the following:

- Assist upgrades of regional and community-of-interest networks, where appropriate, especially where these upgrades enhance end-to-end reliability.
- Produce improved user-level tools to enable scientists and educators to take advantage of the network's capabilities.
- Refine the understanding of the requirements for high capability networks: bandwidth, latency, predictability, and stability.
- Improve the technologies necessary for policy control, resource allocation, accounting, security, and routing coordination. Integrate and interconnect commercial services and other relevant networks where appropriate.
- Provide for a network information service that acts as a primary source of information on access to and use of the network.
- Enhance the current interconnected multi-agency architecture to provide for the interoperability of Federal and non-Federal networks, to the extent appropriate, in a way that allows for the autonomy of each network component.


Gigabit Research and Development. The Gigabit Research and Development subcomponent is aimed at providing the technology base needed to achieve, at a minimum, gigabit speeds and advanced capabilities in the NREN. It will include efforts to develop technology to implement, operate, and effectively use a very high capability communications infrastructure.

Specific research and development activities include the following:

- Develop technology for advanced high throughput, low latency switches. Develop advanced computer-network interface technology to enable end users to access the capabilities of the network.
- Develop methods for control of congestion and network flow for high capability networks where the transit time across the network has a significant impact on the problem.
- Support gigabit testbed networks to provide early evaluation of research concepts and advanced data communications equipment in the context of applications. These testbeds include collaborations of government laboratories, university researchers, and industrial partners.
- Incorporate the advanced communications technology now under development in coordination with industry, in testbeds and gigabit research efforts. Transfer these advanced technologies to the Interagency Interim NREN as they mature and become stable.

FIGURE 4

Network Connectivity



The image represents the interconnected “backbone” networks of NSF, NASA, and DOE, together with selected client regional and campus-area networks. Nodes of the backbones are represented as connected spheres on a plane above the outline of the United States; the client networks are represented as dendritic lines from the backbone nodes to the geographic locations where the client networks attach.

This program subcomponent will include industrial partners, thereby assisting them to develop prototypes for a future high capability commercial communications infrastructure and reduce the risks for U.S. communications vendors in several important ways: defining and developing the market for the high capability network services; implementing, operating, and managing leading edge networks; promoting the prompt and timely adoption of standards for advanced communication protocols; and providing commercial communications vendors with successful examples of high capability networks which can form a basis for economic analyses of their future offerings.

Accomplishments:

- Packet traffic on the major Interagency Interim NREN backbone has doubled in the past year, as has the number of attached networks.
- Interagency collaboration has been initiated on gigabit networking research, and several testbeds have been established, each focusing on specific applications for high capacity networks.
- Multi-protocol routers are now in preliminary use to support interoperability and integration of diverse networks.

Basic Research and Human Resources (BRHR)

This component addresses long term national needs related both to the development of high performance computing and communications technology and to its effective application. The BRHR component includes activities designed to increase the base of skilled personnel, to enhance the basic research foundation for high performance computing and communications, and to more effectively apply high performance computing and communications technologies to education and training.

Basic Research. These activities support increased participation of investigators in innovative basic and multidisciplinary research in computer science, computer engineering, mathematics, and computational science and engineering. The strategy is to increase the number of basic research and multidisciplinary awards across all disciplines for which computational methods are critical to achieving advances or scientific breakthroughs. Example program activities in basic research include: research on scientific algorithm development; research on scalable numeric and symbolic parallel algorithms for common computational problems; performance prediction techniques for concurrent systems; fault tolerant strategies for parallel and distributed systems; and development of optimization techniques that support heterogeneous software configurations involving multiple programming languages. Basic research may relate to any of the other three HPCC components (HPCS, ASTA, NREN).

Research Participation and Training. These activities address the human resource needs in the computer and computational sciences at the postdoctoral, graduate, and undergraduate levels. In addition, senior investigators and practitioners are trained in the use of advanced computing and network communica-

tions to aid in their work. Program activities include: workshops and seminars; postdoctoral fellowships in computational science and engineering; institutional training and postdoctoral programs; internships; exchange programs for knowledge transfer through national laboratories, centers, universities, and industry; and software dissemination through centers, national databases and libraries.

Infrastructure. These activities will improve university facilities for computer science, computer engineering, and computational science and engineering research related to high performance computing and its underlying basic research disciplines. Program activities include improvement of equipment in computer science, computer engineering, and computational science and engineering academic departments, centers, groups, and institutions; development and expansion of scientific databases; and distribution of integrated system-building kits and toolsets.

Education, Training and Curriculum. These activities will expand and initiate activities to improve undergraduate and precollege education and training opportunities in high performance computing and computational science and engineering, for both students and educators. The introduction of associated curriculum and training materials at all levels is an integral part of this effort. Program activities include training scientists at national centers; supercomputer workshops for established investigators to gain "hands on" expertise and tools for their research problems; summer internship programs for promising undergraduate and high school students; and integrating the experience of professional scientists and engineers in the development of curriculum and instruction materials for high school students and teachers.

Accomplishments:

- Initiated Computational Science Graduate Fellowships with 22 recipients in the first year to schools in 14 states.
- Significantly increased the number of awards for computational science and engineering postdoctoral research associateships.
- Provided multiagency support for the supercomputing applications contest programs for teams of high school students, such as SuperQuest.
- Augmented University Centers of Excellence and Research Institutes to pursue HPCC basic research. A Historically Black Colleges and Universities program has been established.
- Increased the Graduate Student Researchers Program to support fellowships in high performance computing and computation at major centers.

3. PROGRAM DEVELOPMENT AND AGENCY BUDGETS

Program Planning

Leadership for the HPCC Program is provided by the Office of Science and Technology Policy (OSTP), through the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) Committee on Physical, Mathematical, and Engineering Sciences (PMES). The membership of PMES includes the senior executives of many Federal agencies. The High Performance Computing, Communications, and Information Technology Subcommittee (HPCCIT) of PMES, the successor to the Executive Committee of the High Performance Computing and Communications Working Group, was chartered in March 1991 by PMES and FCCSET. The basic duties of the Subcommittee are to establish task groups, prepare budgets, and coordinate program planning and execution. The membership of the HPCCIT Subcommittee includes a representative from each member agency. One new member agency, Department of Education, was added for a total of nine. Currently there are four task groups that address Networking, Research, Education, and Applications. In addition, because of the close coordination required in the implementation of the NREN, a Federal Networking Council consisting of agency representatives has been established.

Evaluation Criteria

Each participating agency's HPCC contribution was reviewed against formal evaluation criteria during the planning and budget process. A review of participating agencies was performed using these evaluation criteria to develop the FY 1993 agency requests for the Program. The evaluation criteria are shown in Figure 5.

FIGURE 5

Evaluation Criteria for the Federal High Performance Computing and Communications Program

Relevance/Contribution. The research must significantly contribute to the overall goals and strategy of the Federal High Performance Computing and Communications (HPCC) Program, including computing, software, networking, and basic research, to enable solution of the grand challenges.

Technical/Scientific Merit. The proposed agency program must be technically/scientifically sound and of high quality, and must be the product of a documented technical/scientific planning and review process.

Readiness. A clear agency planning process must be evident, and the organization must have demonstrated capability to carry out the program.

Timeliness. The proposed work must be technically/scientifically timely for one or more of the HPCC program components.

Linkages. The responsible organization must have established policies, programs, and activities promoting effective technical and scientific connections among government, industry, and academic sectors.

Costs. The identified resources must be adequate, represent an appropriate share of the total available HPCC resources (e.g., a balance among program components), promote prospects for joint funding, and address long-term resource implications.

Enhancements to Existing Program Research. Existing agency HPCC programs will receive adequate support before new initiatives are funded.

Agency Approval. The proposed program or activity must have policy-level approval by the submitting agency.

Agency Budgets

For FY 1993, \$803 million is being proposed, a \$148 million or 23% increase over the FY 1992 enacted level. The budget is shown in Figure 6.


The funds proposed from Federal sources are not intended to carry out the entire HPCC Program. Portions of this Program will be cost shared by organizations from the participating sectors. Cost sharing will occur with various U.S. industrial and university partners to a large extent in the HPCCS and the NREN components, and to a lesser but increasing extent, in ASTA. Cost sharing will occur in the ASTA component with agency programs and other computational applications programs, for example, in specific grand challenge areas, via multidisciplinary collaborations. This component also includes deployment of high performance systems to HPCC research centers. The close coupling of support in the Program will result in significant leverage, accelerate technology transfer, stimulate U.S. industry and markets, and enable the solution of computationally intensive applications. In addition, although the HPCC Program is not intended to include national security applications, the technology produced will have an important impact in these national security areas.

FIGURE 6

High Performance Computing and Communications Budgets by Agency and Program Component

(Dollars in millions)

Agency	FY 1992	FY 1993	FY 1993 HPCC Component			
			HPCS	ASTA	NREN	BRHR
DARPA	232.2	275.0	119.5	49.7	43.6	62.2
NSF	200.9	261.9	28.6	125.6	45.1	62.6
DOE	92.3	109.1	10.9	69.2	14.0	15.0
NASA	71.2	89.1	14.1	61.4	9.8	3.8
HHS/NIH	41.3	44.9	4.2	22.6	7.2	10.9
DOC/NOAA	9.8	10.8	0.0	10.4	0.4	0.0
EPA	5.0	8.0	0.0	6.1	0.4	1.5
DOC/NIST	2.1	4.1	1.1	1.0	2.0	0.0
Total	654.8	802.9	178.4	346.0	122.5	156.0



HPCC — High Performance Computing and Communications

Components of HPCC:

HPCS — High Performance Computing Systems

ASTA — Advanced Software Technology and Algorithms

NREN — National Research and Education Network

BRHR — Basic Research and Human Resources

Agency Program Descriptions

The agency responsibilities under the HPCC Program are outlined in Figure 7. Several agencies have been assigned a coordinating responsibility in specific technical areas:

- DOE and NASA will coordinate activities in high performance computing system evaluation, testbed development, and applications software capabilities.
- NASA will coordinate the accumulation and access to the high performance computing software base. This will be facilitated by a wide area file system technology currently being deployed for early experimental use by DARPA, which will be extended to include the NREN as it matures and is deployed by NSF.
- DARPA will coordinate activities in the development of high performance computing systems and associated operating systems software. NASA and DARPA will coordinate the development of associated software tools. DARPA will also coordinate activities in gigabit network technology research.
- NSF will coordinate the broad deployment of the Interagency Interim NREN working with universities, industry, and agencies having mission specific requirements. NSF will also coordinate basic research and infrastructure activities to foster balance in research participation and training.
- HHS/NIH will coordinate educational activities.
- DOC/NIST will coordinate standards, network security, and systems instrumentation and methodology for performance measurement.

Each agency participates in all of the identified activities to ensure that the resulting capabilities are a good match to user needs. DOC/NOAA, ED, HHS/NIH, and EPA bring distinctive applications areas of broad interest and network user bases.

FIGURE 7

HPCC Program: Agency Responsibilities: HPCS, ASTA

Activity Agency	High Performance Computing Systems	Advanced Software Technology & Algorithms
DARPA	• Technology development and coordination for affordable teraops systems	• Technology development for parallel algorithms and software tools
NSF	• Basic architecture research • Prototyping experimental systems	• Research in: - Software tools, databases - Grand challenges • Computer access
DOE	• Technology development • Systems evaluation	• Energy applications research centers • Energy grand challenge and computation research • Software tools
NASA	• Aeronautics and space application testbeds	• Software coordination • Computational research in: - Aerosciences - Earth and space sciences
HHS/NIH	• System evaluation and performance analysis	• Medical application testbeds for medical computation research
DOC/NOAA		• Ocean and atmospheric computation research • Software tools • Computational techniques
EPA		• Research in environmental computations, databases, and application testbeds
DOC/NIST	• Research in systems instrumentation and performance measurement • Research in interfaces and standards	• Research in: - Software indexing and exchange - Scalable parallel algorithms

FIGURE 7 (continued)

HPCC Program: Agency Responsibilities: NREN, BRHR

National Research and Education Network	Basic Research and Human Resources	Activity Agency
• Technology development and coordination for gigabits networks	• University programs	DARPA
• Facilities coordination and deployment • Gigabits research	• Programs in: - Basic research - Education/training/curricula - Infrastructure	NSF
• Gigabits applications research • Access to energy research facilities and databases	• Basic research and education programs	DOE
• Access to aeronautic and spaceflight research centers	• Research institutes and university block grants	NASA
• Development of intelligent gateways • Access for academic medical centers	• Basic Research • Internships for parallel algorithm development • Training and career development	HHS/NIH
• Ocean and atmospheric mission facilities • Access to environmental databases		DOC/NOAA
• Environmental mission assimilation by the States	• Technology transfer to States • University programs	EPA
• Coordinate performance assesment and standards • Programs in protocols and security		DOC/NIST

Defense Advanced Research Projects Agency (DARPA)

The DARPA HPCC Program is balanced across the four HPCC Program components and is designed to stimulate the aggressive advance of the high performance computing and communications technology base. These technologies include scalable technologies for high performance computing and networking along with system software to support the effective application of the high performance technologies to large scale problems.

DARPA is the lead DOD agency for advanced technology research and has the leadership responsibility for HPCC within DOD. Complementing the DARPA HPCC program is the DARPA Strategic Computing Program, which addresses defense specific needs such as embedded systems, accelerators for specific problems domains, and grand challenge problems related to defense. DARPA executes its portion of the HPCC program in close cooperation with the Office of Naval Research, the Air Force Office of Scientific Research, the Army Research Office, Service Laboratories, the National Security Agency, and other DOD organizations and Federal agencies.

High Performance Computing Systems efforts are organized into four areas: research for future generations of computing systems, system design tools, advanced prototype systems, and evaluation of early systems. Technologies are being developed to support multiple models of computation, high performance interconnection, scalable mass storage, and supporting system software. Systems capable of 100 gigaops for large problems have already been announced and several of these systems are expected to be operational during FY 1993. The research challenge is not just to develop architecture and systems capable of scaling to teraops performance and beyond, but to do this in such a way that very large scale systems will become more affordable over time, thus increasing the range of applications that can effectively exploit the high performance technologies. DARPA has also developed and demonstrated system software approaches that support open heterogeneous workstation/server approaches to high performance computing.

In *Advanced Software Technology and Algorithms*, DARPA is producing scalable libraries for Defense-related problem domains and programming and analysis tools for scalable parallel and distributed heterogeneous systems.

DARPA efforts in the *National Research and Education Network* focus on the development of high performance networking technologies to satisfy gigabit technology needs for NREN and for Defense. Several gigabit networking testbeds have already been established to explore issues related to switch design, protocols, and interoperability.

The *Basic Research and Human Resources* component focuses on the basic scientific issues underlying the other three areas, in cooperation with other basic research programs. The DARPA program, along with NSF, provides the major Federal support to universities in computer science. The program addresses the fundamental scientific challenges underlying HPCC technology, including, for example, very high performance systems, flexible and reliable software, and high performance algorithms.

National Science Foundation (NSF)

The NSF HPCC program supports activities in all four components through individual investigator, group, center, and infrastructure awards. The activities are supported and coordinated through existing program structures to provide close cooperation with ongoing computational science and engineering programs throughout the NSF disciplinary directorates. The NSF HPCC program supports technology transfer by developing HPCC technologies through cooperative projects involving industries and academic research institutions.

The NSF *High Performance Computing Systems* component activities focus on support of basic research on heterogeneous and distributed high performance computing systems, coupled with prototyping and evaluation of promising new systems. NSF will emphasize research on new architectures, novel computing systems optimized for specific research applications, and new systems level tools for automated design and evaluation.

In *Advanced Software Technology and Algorithms*, NSF emphasizes developing the elements of a scientifically productive computational research environment by supporting research on generic software and programming environments, programming and applications tools for disciplinary scientists, computational techniques, and new conceptual models for scientific computing. Support will be provided for multidisciplinary grand challenge applications groups centered at academic institutions, in cooperation with other agencies and industry. NSF high performance computing research centers will: provide researchers access to a variety of high performance computing systems; be the focus for technology transfer and training activities; and provide a rich environment for coordinated approaches to grand challenge problems.

The NSF *National Research and Education Network* component activities include: coordinating the upgrading of the Interagency Interim NREN backbone network; assisting regional networks to upgrade facilities, capacity, and bandwidth; interconnecting the backbone networks of other agencies; and developing safeguards to enhance security of access to the connected computer systems. NSF supports the coordinated research and development effort that will culminate in initial deployment of a gigabit networking capability, including work on very high speed switches. NSF educational activities will expand grants for high school educational activities, upgrade operational services for educational users, and extend information access facilities and tools.

The NSF *Basic Research and Human Resources* activities support innovative basic research in scientific algorithm development for highly parallel computers; prediction techniques for concurrent systems; and heterogeneous computing. Research participation and training activities will address HPCC human resources at the postdoctoral, graduate, and undergraduate levels. Infrastructure activities will provide equipment and develop general use scientific databases, system-building kits, and toolsets. Education, training, and curriculum activities will be expanded to improve undergraduate and precollege education and opportunities will be provided for students and educators to participate in ongoing computational science and engineering research activities, such as the NSF-led SuperQuest.

Department of Energy (DOE)

In order to sustain a balanced program the DOE has taken steps to ensure that it makes contributions to all four components of the HPCC Program.

The DOE *High Performance Computing Systems* component emphasizes the evaluation and characterization of the applicability of advanced computers to large scientific and engineering problems. To perform this evaluation, DOE through its laboratories and other contractors and grantees is integrating the latest computers, mass storage devices, and other tools to further the knowledge base in computational science.

The DOE *Advanced Software Technology and Algorithms* component supports, in conjunction with energy related industrial partners, the application of high performance computers in computational chemistry, porous flow, groundwater transport, and high temperature superconductors, catalysis, quantum chemistry-based pharmaceutical design, and environmental remediation. The DOE recently conducted an evaluation of proposals for research centers which will support collaborations of computational scientists and engineers. These centers will provide early evaluation of advanced computer architectures and a coordinated program focused on supporting grand challenge computational research which is cost shared with other scientific research programs in DOE or with industrial partners to ensure the diffusion of technology. In addition, to ensure that technology transfer occurs effectively between DOE and its industrial collaborators, the DOE laboratories have initiated over 10 HPCC Cooperative Research and Development Agreements in the past year.

The DOE *National Research and Education Network* activity will ensure the availability of advanced network capabilities for access to its researchers and facilities. DOE plans to improve connectivity and network capabilities available to DOE researchers and research facilities nationwide through an upgrade to the Energy Sciences Network and through collaboration with other Federal agencies and regional networks. DOE will also help to define, jointly fund with other agencies, and in some cases carry out the research which will be required to evolve from today's infrastructure to the NREN. This will include participation in gigabit network testbeds to evaluate how they impact the solution of large science and engineering programs, development of tools for building and testing large distributed applications, and research on network congestion and flow control.

The DOE *Basic Research and Human Resources* component is expanding its program in basic research in the broad range of mathematical and computer sciences necessary to underpin all its other scientific research. DOE will expand the graduate fellowship programs and postdoctoral support in computational sciences. The first year of the computational science fellowship activity was highly competitive, with awards at 15 universities in 14 states. DOE will continue to expand its programs for access to high performance computers for students and teachers at high schools as well as its summer internship programs for promising undergraduate students at DOE laboratories.

National Aeronautics and Space Administration (NASA)

NASA's HPCC Program participates in all four components of the Federal HPCC Program through an integrated program focused on NASA's grand challenges in Computational Aerosciences and in Earth and Space Sciences. The goal of the NASA program is to accelerate the development and application of high performance computing technologies to meet NASA science and engineering requirements.

NASA's program brings together teams of computer and computational scientists to develop the necessary technologies within two vertically integrated NASA grand challenge projects that are unique to the NASA mission. These technologies include applications algorithms and programs, systems software, peripherals, networking, and high performance computing hardware.

In cooperation with the other Federal agencies, NASA's *High Performance Computing Systems* component will deploy experimental scalable computer capabilities essential for computational design of integrated aerospace vehicle systems and for predicting long term global changes.

In *Advanced Software Technology and Algorithms*, NASA will acquire hardware for computational science testbeds and develop software tools to enhance productivity, including load balancing tools, run time optimizers, monitors, parallelization tools, as well as data management and visualization tools. NASA has supported planning for an HPCC Software Sharing Experiment to identify issues related to software sharing and reuse. Cooperation and coordination with academia has been enhanced through a series of workshops recently conducted by NASA. These workshops addressed the Federal HPCC program and NASA's needs in software technologies and algorithms, computational aerosciences, and supercomputer testbed.

The NASA *National Research and Education Network* activity will provide high-speed network connections among NASA, industry and academic researchers.

NASA's *Basic Research and Human Resources* component will encourage research into the underlying theory and concepts of high performance computing and will foster research in high performance computing at NASA centers, research institutes, and universities.

Several computational advances have already been achieved using high performance computing testbeds at NASA centers. A unique computational method for enhancing imagery from the Hubble Space Telescope (HST) has been developed. This technique, called the Maximum Entropy Method, increases resolution by a factor of four over that of an unprocessed HST image. A new turbomachinery simulation code has been implemented on an advanced parallel system. This academically-developed code simulates three-dimensional inviscid, time-accurate fluid flow through multistage turbomachinery (i.e., jet engines). A multidisciplinary approach to simulating the flow field surrounding a powered-lift aircraft has been accomplished. The simulation used over 2.8 million grid points and incorporates an engine simulator to model the effects of experimental variables.

National Institutes of Health (HHS/NIH)

The HHS/NIH HPCC Program integrates high performance computing with computationally intensive biomedical research applications, links academic health centers via computer networks, creates advanced methods to retrieve information from life sciences databases, and provides training in biomedical computer sciences. NIH operates the first supercomputer dedicated to biomedical research in the public domain. NIH is evaluating advanced supercomputer architectures and system software for applicability to biomedical problems.

The *High Performance Computing Systems* component focuses on instrumentation, performance analysis, and evaluation of high performance systems, with emphasis on algorithms important to biomedical research.

The *Advanced Software Technology and Algorithms* component focuses on two areas of importance in biomedical research and education: biotechnology computing and digital imaging. In biotechnology computing, the program will support development of molecular sequence comparison algorithms, new database methods, structure determination from x-ray and NMR data and algorithms to predict biological structure and function from genetic code. The HPCC program will complement the Human Genome Project by providing new methods for computer based analysis of normal and disease associated genes. Accomplishments already realized in this area include the development of a similarity searching algorithm capable of rapidly analyzing hundreds of millions of molecular units, and the development of a promising new retrieval system to access DNA and protein sequence databases, along with literature citations in which the sequences were published. The biomedical images area will support new methods for representing, linking, and rendering images of biological structure.

The *National Research and Education Network* component has two subcomponents: connections among academic medical centers and their associated computerized information sources, and development of intelligent gateways that link conceptually related databanks in the life sciences. The academic medical centers of the country are confronted with a growing array of disconnected computer-based information sources, ranging from patient records, x-rays, and laboratory systems, to basic research tools such as protein and DNA sequence databanks. The focus of the Gateways program is the building of systems that are capable of translating a user's request into queries to appropriate databases from a broad range of widely available resources.

The *Basic Research and Human Resources* component addresses the need to train biomedical researchers and health care providers in the use of advanced computing and network communications to aid in their work. The successful predoctoral and postdoctoral grants program for career training in medical informatics will be expanded and a new postdoctoral training program will be started for the cross-training of biomedical investigators in advanced computer techniques and communications systems.

National Oceanic and Atmospheric Administration (DOC/NOAA)

DOC/NOAA formulates and carries out operational and research programs in weather prediction, ocean sciences, the Climate and Global Change Research Program, and the Coastal Oceans Program, together with data management activities for all agency programs. The HPCC Program will allow extensive development of new forecast models, studies in computational fluid dynamics, and the incorporation of evolving computer architectures and networks into the systems that carry out agency missions. These activities are carried on in collaboration with both the HPCC agencies and others.

The NOAA HPCC Program focuses on two program components:

Advanced Software Technology and Algorithms supports grand challenges in atmospheric and oceanic sciences; development of advanced numerical models simulating the general circulation of the oceans and atmosphere in support of NOAA missions and the activities of collaborating research groups; development of new computational methods for solving atmospheric, oceanic, and related problems on new computer architectures; data management R&D; support for basic research in strategies, techniques, and tools required for the management and analysis of large-scale distributed scientific databases and distributed data handling, including quality control; and development of algorithms for massively parallel processors, together with their standardized software component libraries and tools for the solution of oceanic and atmospheric analysis and forecasting problems.

NOAA uses advanced computational facilities for the evaluation of near operational prototype computers having novel architectures, and is developing algorithms appropriate for these new architectures. As the program continues to grow, advanced studies of software systems will be undertaken to determine the most advantageous production deployment of these systems.

The *National Research and Education Network* component provides networking for the climate and global change research community and a wide range of agency missions in oceanography, weather prediction, and environmental sciences research including distribution of high quality data sets and information. NOAA will expand use of advanced network strategies and hardware technologies related to its missions.

Environmental Protection Agency (EPA)

The EPA research program has three main goals: incorporation of advances in computing and communications technology into EPA's environmental assessment applications, demonstration of the advanced techniques on several grand challenge problems, and development of techniques to transfer these high performance tools to key State, Federal, and industrial users with decision making responsibility. These advanced environmental assessment tools will be capable of handling multiple pollutant interactions including the air/water interchange and optimization of control strategies.

The *Advanced Software Technology and Algorithms* component focuses on development and optimization of key environmental algorithms for computers with several to thousands of parallel processors. An advanced modeling framework is being designed to allow rapid integration of these new algorithms and emerging high performance technology to address grand challenge problems of air and water quality. Advances in scientific visualization center on provision of critical means for scientists and policy analysts to better interpret and understand results of environmental assessment tools. EPA has demonstrated the use of an object oriented applications building package to control air pollution model computation and visualization in a distributed, heterogeneous computing environment.

The *National Research and Educational Network* component was added to EPA's FY 1993 program to ensure that a critical mass of environmental problem solving groups are connected to the network to support use of advanced environmental assessment tools and collaboration research studies.

The *Basic Research and Human Resources* component includes the establishment of technology training programs targeted to educate State users, industrial environmental groups, and Federal policy analysts on the use of HPCC technology for solving environmental problems. The approach provides early prototype applications to key user communities to obtain feedback on ease of use and functionality for environmental assessment analysis. The program also supports cross-discipline career training through universities and encourages high school students to explore HPCC technology through programs like SuperQuest.

National Institute of Standards and Technology (DOC/NIST)

DOC/NIST supports three components of the HPCC Program.

In *High Performance Computing Systems*, NIST develops instrumentation and methodology for performance measurement of high performance networks and massively parallel computer systems. During the past year, two VLSI based instrumentation and monitoring methods were developed for shared memory and distributed memory computer architectures. Both chip sets provide a passive, programmable, non-perturbing means to capture events and trace system level activities. Data acquired is time stamped and sent by hardware in real-time to a separate computer via an integrated network interface for subsequent processing and transformation into performance metrics. Both were designed to be incorporated in scalable, massively parallel systems. Data analysis has led to identification of a means for succinct representation of state information of dominant modes of operation within massively parallel computer systems, and has established the importance of the succinct representation of these performance metrics. This information is important in system component load balancing and support of user built analytic tools.

In *Advanced Software Technology and Algorithms*, NIST will focus on: scalable, parallel algorithms in the field of computational geometry, and software classification and exchange via the Guide to Available Mathematical Software (GAMS) system. The GAMS system has been incorporated in the NASA Software Sharing Experiment as a means to classify, select and exchange software for scientific and engineering applications. Work has been initiated to improve the utility and portability of the system.

For the *National Research and Education Network*, NIST will support, coordinate and promote the development of standards within the Federal Government to provide interoperability, common user interfaces to systems, and enhanced security. As advanced requirements flow from HPCC research and development, NIST will promote "Open Systems" standards to assist industry commercialization of the R&D, with the aid of other agencies. NIST also conducts research and development on new high performance communications protocols and will deploy a gigabit testbed to assess protocol performance. NIST will coordinate and contribute to the development of security technology, guidelines and standards which are related to components of the Program. A gigabit network testbed is being deployed to support this activity and includes a high performance parallel interface switch, a massively-parallel, video supercomputer and a vector class supercomputer. Related communications hardware interfaces and high performance communications protocols are in the design stage.

4. GRAND CHALLENGE AND SUPPORTING TECHNOLOGY CASE STUDIES

This section incorporates examples of high performance computing and computer communications technologies and several illustrative grand challenge applications in computational science and engineering that are presented on the cover of this report. The list of grand challenges is too long to allow an example from each possible area, and lack of representation of certain areas does not imply lack of importance. In addition to these examples, there are many important applications of high performance computing in national security areas that are supported by agency mission funding.

The examples below were chosen to illustrate the diversity and significance of application areas that have been addressed to date.

- Magnetic Recording Technology
- Rational Drug Design
- High Speed Civil Transports
- Catalysis
- Fuel Combustion
- Ocean Modeling
- Ozone Depletion
- Digital Anatomy
- Air Pollution
- Design of Protein Structures
- Venus Imaging
- Technology Links Research to Education

Magnetic Recording Technology



The knowledge and information explosion in every aspect of modern life is creating rapidly increasing requirements for efficient storage of information. Currently, the dominant methods for information storage employ magnetic recording media such as the thin magnetic films that are used in computer disk drives and tape systems.

Future advances in magnetic recording technology will underlie new technological advances and scientific breakthroughs, from laptop supercomputers to data collection from earth orbiting satellites downloading terabits of information per day.

One of the goals of magnetic media research is to pack more and more information onto available media. The fundamental problem is to avoid or compensate for magnetic noise in the metallic thin-films used to coat high density disks. The films are crystalline, rather than continuous, and are composed of tiny magnetic grains of cobalt alloys deposited on a substrate of chromium. The grainy, particulate nature of the medium is the fundamental cause of noise.

Researchers are using high performance computers to study the effects of the two primary magnetic interactions: magnetostatic and exchange. First, all magnetic dipoles (grains) in the films interact with each other via long range magnetostatic coupling. It is this long-range interaction that makes the analysis computationally intensive. Second, exchange coupling occurs between neighboring grains when, for example, cobalt diffuses into the intergranular boundaries, causing direct magnetic interaction. Even minuscule exchange coupling, the simulations show, dramatically increases noise.

The smaller that the distinctly separable grains can be made, the higher will be the density at which digital bits can be recorded and retrieved. The key is to devise alloys and fabrication processes to assure nonmagnetic grain boundaries, minimizing exchange-coupling noise. Based on theoretical and computational studies, researchers have recently shown that one billion bits can be packed onto a square inch of disk surface, nearly 30 times the storage density of current media.




The pictures represent magnetization patterns for recorded bit boundaries, or transitions, in thin film recording media. The color, from blue to red, represents the magnetization component in the recording direction. The bottom picture is the magnetization pattern for a typical film medium with intergranular exchange coupling, which is a noisy medium. The top picture is the magnetization pattern for a typical film medium with zero intergranular exchange coupling and the resulting transition noise is much less.

Rational Drug Design


Recent progress in the development of therapeutic agents for two medical problems — cancer and acquired immunodeficiency syndrome (AIDS) — illustrates the value and potential of modern computer-assisted, receptor-based drug design as a means to markedly reduce the time and cost of the drug development cycle. Parallel developments in x-ray crystallography and nuclear magnetic resonance spectroscopy along with high performance computing have provided extraordinary and rapid progress in our understanding of the three-dimensional structure determinations of biological macromolecules. This provides the basis for rational drug design, the logical search for a specific therapeutic agent developed from an improved understanding of molecular structure and function.

The most commonly used alkylating agent to treat some cancers is cyclophosphamide. The active metabolite of that drug, phosphoramidate mustard (PM), induces the anticancer effects. It crosslinks with DNA molecules in cancerous cells to inhibit their growth. The mechanism through which PM determines which cells are cancerous is unknown, but it is widely believed to be due to a unique biochemical property of the cancer cells. More detailed chemical calculations have challenged existing concepts of how PM induces its effects on cancer cells. The complete elucidation of the mechanisms for that anticancer drug and others requires a computing capability many times greater than that which currently exists. This may lead to the development of more effective drugs to treat several different types of cancer.

The development of therapeutics for AIDS is another example of active-site-based drug design. The structure of the human immunodeficiency virus (HIV) includes a protease, an enzyme which catalyzes a mandatory step in the maturation of the virus so that it can infect cells. Using a high performance computer, new potential agents have been identified to block the action of the HIV protease. These compounds have anti-HIV activity, but further modifications of candidate agents need to be made to induce a greater therapeutic potential and lower unwanted side effects. Higher performance systems are needed to fully understand these new candidate agents and to develop and design new drugs.



The most widely used anti-cancer drug, cyclophosphamide (inset) acts by binding (cross-linking) two (guanine) bases of the DNA of a cancerous cell. Two chlorine atoms from the drug (green in the inset) are stripped off in the body and replaced by two nitrogen atoms (shown by arrows) of the DNA. This forms the bonding that prevents DNA replication, thus killing the cancerous cell. When broken down by the body, the side chain shown on the lower left of the inset is released, causing toxic side effects. Though this research required massive amounts of computation, calculations thousands of times more complex are needed to fully understand the activity of this drug.




Atomic structure of a rationally-designed inhibitor bound to the active site of HIV protease ("ribbon" structure). The structure of the complex was determined by x-ray crystallographic analysis using a supercomputer networked to a high speed graphics workstation.

High Speed Civil Transports

The U.S. has initiated a research program to examine the feasibility of a new High Speed Civil Transport (HSCT). This type of aircraft will significantly impact the world's transportation system. Significant technological challenges must be addressed to assure environmental compatibility and to realize the commercial potential. Among the most important of these are: aerodynamic efficiency, advanced materials, and engines with acceptable environmental impact. High performance computing systems are essential to bringing HSCT to reality at the earliest possible date.

Computational fluid dynamics codes running on high performance computing systems will enable the rapid evaluation of alternative HSCT designs for aerodynamic efficiency and stability, as well as predicting the associated sonic boom. In addition, other programs will account for the atmospheric heating of the vehicle and the physical stresses of supersonic flight to provide the necessary information for the development of advanced materials and fabrication of new structures to commercialize HSCTs. High performance computing technologies are crucial to the design and evaluation of advanced propulsion systems that will efficiently and safely power HSCTs while significantly reducing community noise and exhaust emissions harmful to the upper atmosphere.




The figure above represents a numerical simulation model of a proposed low sonic boom commercial High Speed Civil Transport. The program is simulating almost level flight at twice the speed of sound (Mach 2). The intensity and directional characteristics of the sonic boom are depicted by coloring the atmospheric pressure surrounding the aircraft. Highest intensity is shown in red; lowest is in blue.

Catalysis

Catalysts for chemical reactions are currently used in one-third of all manufacturing processes, producing hundreds of billions of dollars worth of products each year. While most catalysts are still designed empirically, high performance computing, along with state-of-the-art algorithms and physical models, is being used to develop a computer design and analysis capability for commercially important catalytic processes and catalysts. The goal is to reduce the time required to design catalysts and to optimize their properties. Important catalysis applications include enzymes, biomimetic catalysts, microporous solids, and metals.

Many biological processes are catalytically controlled by enzymes, which give a high level of selectivity to a desired reaction over competing reactions. While enzymes are fundamental in medical science, environmental restoration, and food processing, the way in which enzymes function in their catalytic role is not completely understood for any system. High performance computing technology will enable larger simulations that more realistically model the substrate, cofactor, and local portions of the enzyme.

Enzymatic reactions are attractive because they exhibit an extremely high degree of selectivity to specific reactants and products. However, they also can be limited to chemical and physical environments different from those in many industrial processes. Alternatives to enzymes that may be more useful industrially are biomimetic catalysts, so called because they mimic biological catalysts (enzymes). These catalysts are synthetic materials that generally are much simpler chemically and can function at higher temperatures than their biological counterparts. Computer modeling of these biomimetic catalysts using relatively simple force field models has been crucial in the development of molecules with the correct size and shape properties. Significant increases in computational power will be necessary to determine the effect of molecular structure on electronic states in these catalysts, and to evaluate the chemical reactions that will occur. Massively parallel quantum models are being developed to analyze these complex systems.



The enzyme dihydrofolate reductase catalyses the reaction of dihydrofolate to tetrahydrofolate, an important one-carbon donor in a variety of biosynthesis reactions including the synthesis of amino acids and nucleic acids. The figure shows the effect of the enzyme in shifting the electron density in dihydrofolate towards that required to make the reaction take place. Red indicates increases of electron density, blue indicates decreases of electron density.

Fuel Combustion

A detailed understanding of fuel burning inside an internal combustion engine requires the comprehension of a number of complex processes. One example is an experimental engine design whose goal is to run leaner and cooler than conventional engines, thereby reducing emissions while increasing the compression ratio and control of ignition for greater efficiency and fuel economy. The burning of the fuel-air mixture in such an engine brings into play a few hundred different chemical reactions between numerous short and long-lived chemical species. The rates at which these reactions occur depend on the temperature and concentrations of the species. For instance, as combustion proceeds, various pollutants may be formed, such as nitrogen oxides and unburned hydrocarbons. These hydrocarbons may form solid particles known as soot. Also, the fuel may autoignite in some region before the flame arrives, giving rise to the small explosion known as engine knock.


Because engine operation is so complex, designing an optimum engine by experimentation (varying all design parameters independently) is prohibitively expensive. A computer model of the engine allows us to effect design changes numerically on a scale of hours at a fraction of the cost.

The engine components depicted in the illustration are a cupped piston, fuel injector, and spark plug. The piston draws in fresh air and expels exhaust gases as the combustion cycle progresses. At an appropriate point in the cycle (upper right in the figure), liquid fuel is injected into the engine cylinder. The fuel jet breaks up into small droplets, which interact with the air. The droplets alternately elongate and flatten; some of them break up into multiple new droplets, others collide and coalesce to form larger drops. The spray is caught up by the air flow and swirls around in the cylinder. It simultaneously evaporates to form a gaseous fuel-air mixture.

Present engine models are deficient in one way or another because of limitations in computing resources. Computer codes can provide information on large scale fluid flows, but are limited in their ability to describe the small scale flow structures in turbulent flow. Furthermore, fuel chemistry models are necessarily crude. Current chemical kinetics calculations approach the detail necessary for an adequate description of the chemistry, but even these models require serious improvements. Such calculations are capable of only the most rudimentary fluid mechanical effects.

The envisioned increase in computing power will make improved models of turbulence and chemistry computationally tractable. These improvements are critical to the effective development of clean, efficient engines.

Fuel injection and combustion in a stratified charge engine cylinder at four crank angles



Shown above is a computer simulation depicting four different times during the combustion cycle of a Direct-Injection Stratified Charge engine. The multicolored surface indicates the fuel's location. When the spray has evaporated sufficiently a spark ignites the mixture, and combustion begins as a flame propagating through the chamber. As the engine proceeds through its cycle (indicated by the increase in crank angle from -23 to +17 degrees), the fuel burns, resulting in high temperatures (see color code).


Ocean Modeling

The ocean is the other turbulent fluid in the climate system. By contrast with the atmosphere, the ocean has much smaller spatial structures and much longer time responses. The ocean currents, which transport almost as much heat as atmospheric flows, are only about 50 kilometers across, and ocean eddies that are the counterparts of synoptic weather systems are one-hundredth the size of those in the atmosphere. These oceanic phenomena evolve over periods from years to centuries to profoundly influence the climate. The small space scales and long time scales of the ocean present a massive challenge to high performance computing.

The ocean is particularly important in maintaining the climate of the Northern Hemisphere through an organized system of currents which has been likened to a conveyor belt. Currents near the surface transport heat from the Tropical Pacific all the way to the North Atlantic, where the heat enters the atmosphere through evaporation. The evaporation makes the surface waters heavy enough to sink and return to the Tropical Pacific, where they rise and are freshened by rains and warmed by the sun for another trip to the North Atlantic. However, the entire system is very sensitive to changes in atmospheric conditions, such as might occur from increasing greenhouse gases.

The first global ocean simulation which is capable of barely resolving the eddies and strong currents, including those that act as a conveyor belt, has only recently been accomplished. An integration of just a decade required 3000 processor hours on a gigaflop supercomputer. If multi-century ocean simulations with adequate spatial resolution are to be performed in support of climate prediction, then teraflop machines will be required.

Because of the turbulent nature of oceanic flows, there should be demands for archiving, visualization, and communication that parallel those of the most ambitious atmospheric and hydrodynamic applications. For example, the first global simulation with barely resolved eddies produced more than 250 gigabytes of processed output. A two hour animation was made to help understand the resulting phenomena, and the animation required the transfer of gigabytes of model output across high-speed networks.




Salinity of the upper Pacific Ocean.

The atmospheric hydrological cycle strongly influences the North Pacific, where an excess of precipitation freshens (and lightens) the water that will transport heat all the way to the North Atlantic. This effect is seen in upper-ocean salinity from a global ocean model with barely resolved eddies. The blue zones in the figure are relatively fresh, while the red ones are more saline. Steady ocean currents set up a sharp frontal barrier along the equator. However, transient flows, such as a vigorous current reversal near New Guinea and waves on either side of the equator, transport salty water toward the North Pacific to restore a balance. This figure is part of an animation that graphically portrays these features, as well as other vigorous flows along the path to the North Atlantic.

Ozone Depletion

The depletion of the Earth's ozone has an important impact on animal and plant life, including higher incidence of skin cancer in humans, because of the increased ultraviolet radiation reaching the surface through the reduced ozone shield. Recent observations of unpredicted ozone losses in the Earth's stratosphere have led to intense research efforts to understand the cause. The "ozone hole" that has been observed in the stratosphere above Antarctica during the southern hemisphere winter is a dramatic example of this ozone depletion process. Recently, reduced ozone levels have been observed in the northern hemisphere stratosphere at high latitudes around the perimeter of the polar vortex. Scientists are now studying the dynamical and chemical mechanisms associated with this newly discovered ozone depletion to determine whether an ozone hole similar to that over Antarctica can also form in the northern hemisphere.

The chemical and dynamical mechanisms that control the ozone depletion process are extremely complex. For example, the chemistry of the problem involves hundreds of different chemical processes. Moreover, the efficiency of this chemical activity depends upon the detailed transport and mixing of air containing very different chemical compositions. Because of the complex interaction between dynamics and chemistry, an essential element in this investigation is the use of comprehensive global, high-resolution computer models to simulate the ozone chemistry and associated dynamics that occurs in the Earth's atmosphere. For example, current models require as much as ten supercomputer processor hours to model one day of interaction. One of the major requirements for continuing advancements in this field will be the incorporation of new technologies to meet the tremendous computational demands of future global climate models.




The figure shows an instantaneous field of nitrous oxide (N_2O) for a model date in February on a constant entropy surface corresponding to an altitude of roughly 20 kilometers. Nitrous oxide is a biologically produced trace gas that is a fundamental determinant of the natural ozone balance. Its long lived character also makes it an excellent tracer of atmospheric motions. The figure shows very low values of N_2O in the polar regions; the N_2O structure essentially defines the winter polar vortex that determines the chemistry typical of the polar ozone hole phenomenon. Note the virtual wall of N_2O gradient that occurs between the polar vortex and the outside region. At times, the wall structure becomes complicated and folds upon itself, forming what appear to be double and triple wall regions.

Digital Anatomy

New computer-based imaging techniques are making it possible to explore the structure and function of the human body with unprecedented accuracy. Clinical imaging methods such as computed tomography, magnetic resonance imaging, and positron emission tomography yield two-dimensional images which may be transformed by computer algorithms into three-dimensional pictures. Reference standard normals, such as the Visible Human project, provide three-dimensional numerical coordinates from which internal as well as external structure can be depicted, rotated, viewed from any angle, and reversibly "dissected."

Automated correlation of images derived from different imaging techniques is a formidable technical problem whose solutions are computation-intensive. These image datasets are also large (the Visible Human project, for example, will generate a four trillion byte image library), and high speed networks are the only feasible method for allowing researchers and health professionals to browse such electronic image collections interactively, in a manner analogous to today's interactive browsing of textual databases.

The computer and network tools applied to the challenge of multimodal imaging promise an era of new insight into the anatomy and physiology of human health and disease in the three dimensions of biological structure and the fourth dimension of time.




The volumetric rendering of the human head is capable of depicting deep as well as surface structures. This three-dimensional view was reconstructed from two-dimensional Magnetic Resonance Imaging data acquired from a living person. Conventional computers require several seconds to several hours to compute a single three-dimensional view; advanced architecture machines developed under HPCC will allow such images to be rotated, "dissected," and viewed from any angle under real-time control.

Air Pollution

Reduction of pollutant and pollutant precursor emissions is costing billions of dollars. The early optimism of the past two decades about the ease of controlling and reducing air pollution, especially ozone, has given way to a realization of the complexity and extreme difficulty of achieving the reductions necessary to attain the health standards in many urban areas. In addition, connections between pollutant media are increasingly being recognized. For example, nitrogen deposition from air pollutants is now thought to be a major contributor to eutrophication of coastal estuaries, adding to the agricultural and urban sources that have been studied and characterized.

Computational models are powerful and necessary tools to study the transport and transformation of pollutants and to provide guidance on the expected effectiveness of emission control strategies. These air quality models must incorporate descriptions of physical, chemical and meteorological processes that encompass scales from regional (several thousands of kilometers) to local (several kilometers) to adequately represent the production of pollutants. Interactions between many physical processes, such as clouds and precipitation, must be incorporated in the computational descriptions, leading to models that are, of necessity, highly sophisticated and which must be run on supercomputers. Yet, the complexity and completeness of the air quality models must be advanced to more accurately simulate the effectiveness of emissions controls and find effective and efficient solutions to air pollution.

High performance computing will enable additional and more complete descriptions of the physical and chemical processes to be simulated in the air quality models. Advanced visualization techniques will allow the scientists to better study and understand the complex, nonlinear interactions of the atmospheric pollutant system, leading to more informed assessments. High performance computing will enable scientists to explore the complex interactions in the models more quickly and provide the type of guidance needed to reduce air pollution to meet the air quality standards.




Depicting the interactions between several processes simulated by air quality models leads researchers to more understanding of the complex atmospheric system. High concentrations of sulfate (shown as a grey cloud) are followed as they are simultaneously transported and scavenged by precipitation (the silver pill boxes). The resultant accumulated wet deposition of sulfur due to several rain storms passing over the eastern United States is shown as a color-coded, two-dimensional concentration map “painted” on the landscape.

Design of Protein Structures

Although the information contained in the amino acid sequence of a protein is sufficient to determine the three-dimensional or “folded” structure of a protein, the rules by which the protein assumes its final configuration — its three-dimensional structure — are not yet understood. Success, measured even by a partial understanding of these rules, is often of immediate importance to biomedical research, and scientists are using high performance computing to learn more about the process.

An understanding of protein folding could help to design new proteins with a specific structure, for a specific need, perhaps to catalyze a biochemical reaction. If the design of a protein *de novo* seems far away, a more reasonable goal of intelligently modifying a protein to yield a desired change in stability, reactivity, or specificity seems more achievable. Simulation of protein activity, tracking the movements as the protein goes from unfolded to a folded state, has proved particularly fruitful. But, since most biochemical reactions proceed in fractions of microseconds, current computationally intensive modelling methods are thousands to millions of times short of simulating interesting kinetic reactions.

Computational biologists are looking to massively parallel computing architectures to overcome this barrier, anticipating order of magnitude improvements in the calculations, and the ability to engineer the complete structural properties of a protein to desired specifications.




Scientists have been forming models of channels, like the one found in the acetylcholine receptor, to see how alpha-helices (a pattern for folding amino-acids residues in proteins in a circular pattern) might pack together to control the flow of ions through the protein. The figure shows a channel with a space filling technique. In the center, an ion surrounded by several water molecules is shown in blue. Construction of the model was based on a four-helix bundle.

Venus Imaging

The Magellan spacecraft has been mapping the surface of Venus since September 15, 1990. Magellan uses a radar instrument to penetrate the thick clouds surrounding the planet to produce 120-meter resolution images of the Venus surface. Venus is the most Earth-like of the planets of the inner solar system, with about the same size, mass and density. However, the "greenhouse effect" due to the clouds surrounding Venus results in a surface temperature of about 900° F and pressure at the surface about 90 times that of Earth. Despite these differences, Magellan data for Venus have revealed a wide range of geologic features from volcanoes to mountain ranges that are providing insight into how geologic processes on Earth and Venus operate. Magellan has provided images and topographic information for over 90 percent of the Venus surface. The total volume of data returned by Magellan exceeds 3×10^{12} bits or 3 terabits (all previous U.S. planetary missions combined have returned about 9×10^{11} bits).

The Magellan radar data go through several processing steps, starting at a Synthetic Aperture Radar (SAR) processor that turns the echoes returned from the planet's surface into images. The data then go through various image processing steps, where Magellan image strips are combined to form mosaics which can then be color enhanced or combined with topographic data to provide scientists with tools for understanding the geology of Venus. The large data volume associated with Magellan has necessitated new and increasingly more complex methods of data processing, handling and storage that will be significant to future data-intensive projects such as NASA's Mission to Planet Earth. Magellan images, such as that shown in the figure, provide insights into the physics of volcanism and its significance to planetary evolution.

Using the large parallel supercomputers for producing the rendered images dramatically increases our capabilities. For example, a single image produced on a conventional high performance workstation takes several minutes to complete. Animations, even short ones, typically take weeks to produce. Massively parallel supercomputers hold the promise of being able to store the huge datasets on-line and render animations in real-time.



Maat Mons, the largest volcano on Venus, is shown in the three-dimensional perspective view of the Venus surface rendered on a 528 node parallel supercomputer. The viewpoint is located 560 kilometers north of the volcano at an elevation of 1.7 kilometers. Vertical exaggeration is about 22 times. Lava flows extend for hundreds of kilometers across the fractured plains shown in the foreground, to the base of Maat Mons. The volcano is about 8 kilometers high and over 300 kilometers across, similar in size to Mauna Loa on Earth. To produce the image, Magellan SAR data were combined with radar altimetry to develop a three-dimensional map of the surface. Rays cast in the computer intersect the surface to create a three-dimensional perspective view. Simulated color and a digital elevation map are used to enhance small-scale structure. The simulated hues are based on color images recorded by Soviet landers.

Technology Links Research to Education

It can take ten years or more for new ideas in science to be incorporated into the curriculum of schools. Given the rapid advances in science, new ways must be found to link researchers, educators and learners.

The technologies that make possible the accelerated pace of scientific advances are also a key to creating the needed linkages with education. High performance computer applications such as simulations and visualizations used to advance scientific knowledge are now being tailored to help teachers and students understand both the methodologies of modern scientific inquiry and specific concepts in these scientific domains. The models and simulations used by researchers are mathematically complex. Through advances in visualization technology, it becomes possible for students who are not yet sophisticated in quantitative reasoning to get a qualitative understanding of a model of complex phenomena. The phenomena under investigation are often too small (e.g., molecules, ions) or too abstract (e.g., fractals, chaos) for students to experience directly in the "real world."

The tools of high performance computers and workstations, computer networks, and software make it possible for researchers to share their current research activities with high school teachers and students. While high performance computers are not physically located at the schools, communications networks enable students and teachers working at their own workstations to access the very large and computationally intensive simulation software and data bases operating at a remote university or research center. The network also makes it feasible for teachers and students in schools to interact with the scientific research teams located at the remote sites.

For example, researchers in computational physics develop and use software models of microscopic molecular dynamics. Students in a variety of different junior high and high schools are now using those same software models to develop their understanding of molecular behavior as well as the underlying probabilistic nature of such phenomena. Connected to the simulations is special front end software designed to guide students' learning, investigations, and experiments. One such simulation concerns the complex web of interactions between individual water molecules. Parameters such as temperature and pressure directly affect the hydrogen bonds in the network.



Students using an advanced workstation linked to a massively parallel supercomputer. Shown is a real-time molecular dynamics simulation of 216 water molecules. The blue lines represent a view of the spanning network of hydrogen bonds in liquid water. The simulation provides a “hands on” experience enabling the student to discover for herself the principles involved in hydrogen bond network formation in liquid water.

GLOSSARY

Algorithm

A procedure designed to solve a problem. Scientific computing programs contain algorithms.

ASTA

Advanced Software Technology and Algorithms

Bit

Acronym for binary digit

BRHR

Basic Research and Human Resources

Byte

A group of adjacent binary digits operated upon as a unit (usually connotes a group of eight bits).

CAD

Acronym for "computer aided design."

Computer Engineering

The creative application of engineering principles and methods to the design and development of hardware and software systems.

Computer Science

The systematic study of computing systems and computation. The body of knowledge resulting from this discipline contains theories for understanding computing systems and methods; design methodology, algorithms, and tools; methods for the testing of concepts; methods of analysis and verification; and knowledge representation and implementation.

Computational Science and Engineering

The systematic application of computing systems and computational solution techniques to mathematical models formulated to describe and simulate phenomena of scientific and engineering interest.

DARPA

Defense Advanced Research Projects Agency

Distributed systems

A distributed system consists of multiple interconnected computers, each potentially containing data and programs, that can together carry out a single computation.

DOC

Department of Commerce

DOD

Department of Defense

DOE

Department of Energy

ED

Department of Education

EPA

Environmental Protection Agency

FCCSET

Federal Coordinating Council for Science, Engineering, and Technology

Flops

Acronym for floating point operations per second. The term "floating point" refers to that format of numbers which is most commonly used for scientific calculation. Flops is used as a measure of a computing system's speed of performing basic arithmetic operations such as adding, subtracting, multiplying, or dividing two numbers.

Giga-

10⁹ or billions of ... (e.g.: gigabits)

Grand challenge

A grand challenge is a fundamental problem in science and engineering, with broad economic and scientific impact, whose solution could be advanced by applying high performance computing techniques and resources.

Heterogeneous system

A heterogeneous system is a distributed system that contains more than one kind of computer.

HHS

Health and Human Services

HPCC

High Performance Computing and Communications

HPCCIT

High Performance Computing, Communications, and Information Technology

HPCS

High Performance Computing Systems

High Performance Computing

High performance computing encompasses advanced computing, communications, and information technologies, including scientific workstations, supercomputer systems, high speed networks, special purpose and experimental systems, the new generation of large scale parallel systems, and applications and systems software with all components well integrated and linked over a high speed network.

Interagency Interim NREN

Interagency Interim National Research and Education Network. The Interagency Interim NREN is an evolving operating network system. Near term (1992-1996) research and development activities will provide for the smooth evolution of this networking infrastructure into the future gigabit NREN.

Interoperability

Interoperability is the effective interconnection of two or more different computer systems, databases, or networks in order to support distributed computing and/or data exchange.

Mega-

10⁶ or millions of ... (e.g.: megaflops)

NASA

National Aeronautics and Space Administration

Network

Computer communications technologies that link multiple computers to share information and resources across geographically dispersed locations.

NIST

National Institute of Standards and Technology

NOAA

National Oceanic and Atmospheric Administration

NREN

National Research and Education Network. The NREN is the realization of an interconnected gigabit computer network system devoted to HPCC.

NSA

National Security Agency

NSF

National Science Foundation

Ops

Acronym for operations per second. Ops is used as a rating of the speed of computer systems and components. In this report ops is generally taken to mean the usual integer or floating point operations depending on what functional units are included in a particular system configuration.

Parallel processing

Simultaneous processing by more than one processing unit on a single application.

Peta-

10^{15} or thousands of trillions of ... (e.g.: petabytes)

PMES

Committee on Physical, Mathematical, and Engineering Sciences

Portable

Portable computer programs can be run with little or no change on many kinds of computer systems.

Prototype

The original demonstration model of what is expected to be a series of systems. Prototypes are used to prove feasibility, but often are not as efficient or well-designed as later production models.

Scalable

Scalable computing systems have the property that the size of the system can be increased with an expected increase in performance proportional to the extent of the increase in size.

Supercomputer

At any given time, that class of general-purpose computers that are both faster than their commercial competitors and have sufficient central memory to store the problem sets for which they are designed. Computer memory, throughput, computational rates, and other related computer capabilities contribute to performance. In addition, performance is often dependent on details of algorithms and on data throughput requirements. Consequently, a quantitative measure of computer power in large-scale scientific processing is difficult to formulate and may change as the technology progresses.

High Performance Computing and Communications Agency Points of Contact

Director
Computing Systems Technology Office
Defense Advanced Research
Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203

HPCC Coordinator
Office of the Assistant Director
Computer & Information Science
& Engineering
National Science Foundation
1800 G Street, N.W.
Washington, DC 20550

HPCC Program Coordinator
Scientific Computing Staff
MS-G217, GTN
Office of Energy Research
Department of Energy
Washington, DC 20585

Director for HPCC
Office of Aeronautics and
Space Technology
National Aeronautics and
Space Administration
600 Independence Avenue, S.W.
Code RC
Washington, DC 20546

Associate Director
Division of Extramural Programs
National Library of Medicine
Bethesda, MD 20892

Office of the Chief Scientist
National Oceanic and
Atmospheric Administration
Herbert C. Hoover Building
Room 5128
14th and Constitution Avenue
Washington, DC 20230

HPCC Program Coordinator
Atmospheric Research and Exposure
Assessment Laboratory
ACMD, MD-80
U.S. Environmental Protection Agency
Research Triangle Park, NC 27711

Computer Systems Laboratory
Office of the Director
HPCC Coordinator
B166 Technology Building
National Institute of Standards
and Technology
Gaithersburg, MD 20899

To obtain a copy of this document — send request to:

Federal Coordinating Council for Science, Engineering, and Technology
Committee on Physical, Mathematical, and Engineering Sciences
c/o National Science Foundation
Computer and Information Science and Engineering Directorate
1800 G Street, N.W.
Washington, D.C. 20550