

DOE/ER/25232-77-Pl.1

HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS:

**ADVANCING THE FRONTIERS OF INFORMATION
TECHNOLOGY**

**A Report by the Committee on Computing, Information, and
Communications**

National Science and Technology Council

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED



MASTER

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

THE WHITE HOUSE

WASHINGTON

MEMBERS OF CONGRESS:

I am pleased to forward with this letter *High Performance Computing and Communications: Advancing the Frontiers of Information Technology*, which was prepared by the Computing, Information, and Communications (CIC) R&D Subcommittee of the National Science and Technology Council's (NSTC) Committee on Computing, Information, and Communications (CCIC). This report, which supplements the President's Fiscal Year 1997 Budget, describes the interagency High Performance Computing and Communications (HPCC) Program.

America is leading the way into the Information Age with its preeminence in the development and deployment of information technologies. Under the HPCC Program, which is often cited as a model for successful Federal multi-agency program coordination, academia, industry, and Government have developed and exploited leading edge computing and communications technologies to provide for more effective Government operations and to boost the economic competitiveness of the U.S.

The HPCC Program will celebrate its fifth anniversary in October 1996 with an impressive array of accomplishments to its credit. These range from underlying technologies in computing and networking, to applications that meet the mission needs of the participating Federal agencies.

The successes of the Program coupled with the changing role of information technology have led to the creation of a broader R&D agenda that reflects long-term efforts in advanced computing, information, and communications technologies built on the strong foundation of the Federal HPCC Program. The new program focus areas continue the sustained, long-term Federal commitment to accelerate the development of advanced technologies for the Information Age. Interacting with the private sector and academia, the Federal departments and agencies who participate in these programs have demonstrated their ability to closely coordinate Government investments in these technologies to provide for the Nation's defense and enhance the well-being of its citizens.

The CIC R&D Subcommittee and its member agencies have developed numerous successful collaborative relationships with academia and industry. Many of these are highlighted in this document. John C. Toole, Chair of the CIC R&D Subcommittee, many people in the HPCC agencies, their associates, and staff are to be commended for their efforts.



John H. Gibbons
Assistant to the President
for
Science and Technology

Table of Contents

I. Executive Summary	1
II. Research Accomplishments	
Introduction	3
Enabling Technologies	
Networking	5
<i>Networks and the Internet</i>	6
<i>Information-Wide-Area-Year (I-WAY)</i>	7
High Performance Systems	8
Software for High Performance Computing	10
Security and Privacy, and Electronic Commerce	11
Visualization, Virtual Reality, and Human-Machine Interfaces	12
Digital Libraries	13
Applications	
Weather Modeling	14
The Environment and Energy Management	15
Manufacturing — Design, Processes, and Products	17
Health Care and Biomedical Imaging	18
Education, Lifelong Learning, and Access to Information	19
Biomedicine, Biomechanics, and Molecular Biology	20
<i>Advances in Membrane Simulations</i>	21
<i>Progress in Predicting Protein Folding</i>	22
Computational Chemistry	23
Materials Sciences	24
Computational Physics	25
III. HPCC Research Centers and Facilities	27
IV. CIC R&D Program Organization	30
V. Future Directions	32
<i>Program Component Areas</i>	32
Planned FY 1997 CIC R&D Activities by Program Component Area	34
Agency HPCC Budgets by Program Component Area	35
VI. CIC R&D Summary	
CIC R&D Goals and Agencies	36
Evaluation Criteria for CIC R&D Programs	37
Committee on Computing, Information, and Communications	38
CIC R&D Subcommittee	39
VII. Glossary	40
VIII. Contacts	48

I. Executive Summary

INTRODUCTION

The Federal High Performance Computing and Communications (HPCC) Program will celebrate its fifth anniversary in October 1996 with an impressive array of accomplishments to its credit. Throughout its existence, the HPCC Program has conducted long-term research and development in advanced computing, communications, and information technologies, and in applying those technologies to the missions of the participating Federal departments and agencies. The use of advanced information technologies across the Federal government and throughout the economy demonstrates the significant impact of the HPCC Program. The HPCC Program continues to receive bipartisan Congressional support, as it has since the passage of the High Performance Computing Act of 1991 (Public Law 102-194).

Under the leadership of the National Science and Technology Council's (NSTC) Committee on Computing, Information, and Communications (CCIC), the HPCC Program is evolving. For the past two years, Federal agencies, academia, and industry have undertaken an intensive planning process to develop an R&D agenda that builds on the firm foundation of HPCC Program accomplishments and extends the goals of the original Program. The new Computing, Information, and Communications (CIC) R&D programs will focus on long-term interagency efforts that address R&D needs critical to the goals of providing more secure and higher quality lives in 21st century America.

ACCOMPLISHMENTS

Over its five-year history, the HPCC Program has focused on developing high performance computing and communications technologies that can be applied to computation-intensive applications—the Grand Challenges of science and engineering. In FY 1994, the Administration added the goal of developing key enabling technologies for the National Information Infrastructure to demonstrate information-intensive applications—the National Challenges. Major highlights for FY 1996:

- High performance computing systems enable practical solutions to complex problems with accuracies not possible five years ago. Weather models use these systems' increased computa-

tional power and expanded memory to employ significantly more detailed physics algorithms. These models allow meteorologists to track severe storms and hurricanes with sufficient accuracy to produce forecasts dependable enough to implement the more precise evacuation plans required to save lives and property.

- HPCC-funded research in very large scale networking techniques has been instrumental in the evolution of the Internet, which continues exponential growth in size, speed, and availability of information. Today's R&D in large scale wireline, optical, wireless, and mobile network technology is required to satisfy the explosive demand for new generations of information technology.
- The combination of hardware capability measured in gigaflop/s (billions of floating-point operations per second), networking technology measured in gigabit/s (billions of bits per second), and new computational science techniques for modeling phenomena has demonstrated that very large scale accurate scientific calculations can be executed across heterogeneous parallel processing systems located thousands of miles apart. The integration of these technologies is enabling new modes of science and engineering collaborations that provide valuable new knowledge benefiting Federal agency missions.
- Federal investments in HPCC software R&D support researchers who pioneered the development of parallel languages and compilers, high performance mathematical, engineering, and scientific libraries, and software tools—technologies that allow scientists to use powerful parallel systems to focus on Federal agency mission applications.
- HPCC support for virtual environments has enabled the development of immersive technologies, where researchers can explore and manipulate multi-dimensional scientific and engineering problems.
- Educational programs fostered by the HPCC Program have brought into classrooms new science and engineering curricula designed to teach computational science. The Program has been responsible for training scientists to devel-



op and exploit new HPCC technologies in their quest to discover new knowledge for the benefit of all.

This document contains a small sample of the significant HPCC Program accomplishments in FY 1996. A more comprehensive view is accessible from the HPCC Program's Web site at <http://www.hpcc.gov>.

NEW DIRECTIONS

The HPCC Program has provided a national focus for Federally funded R&D in high performance computing, communications, and information technologies. It is critically important that HPCC R&D continue. However, as a result of the successes of the Program and the changing role of information technology in Federal agency mission applications, broader collaborative R&D investments in computing, information, and communications are needed.

The NSTC's Committee on Computing, Information, and Communications (CCIC) has organized its collaborative program efforts into five Program Component Areas (PCAs). HPCC Program activities form the core of the new Computing, Information, and Communications (CIC) R&D programs. The PCA structure evolved from the five original components of the HPCC Program (HPCS, NREN, ASTA, IITA, and BRHR, described in the Glossary). The PCAs represent areas of high priority investments by the Federal agencies that participate in the coordinated R&D programs. The PCAs and their goals are:

- **High End Computing and Computation (HECC):** The goal of HECC R&D is to assure U.S. leadership in computing through investments in leading-edge hardware and software innovations, and in algorithms and software for modeling and simulation needed to address Grand Challenge-class applications.
- **Large Scale Networking (LSN):** The goal of LSN R&D is to assure U.S. leadership in communications in high performance network components; technologies that enable wireless, optical, mobile, and wireline communications; large scale network engineering, management, and services; and systems software and program development environments for network-centric computing.

- **High Confidence Systems (HCS):** The goal of HCS R&D is to develop technologies that provide users with high levels of security, protection of privacy and data, reliability, and restorability of information services.

- **Human Centered Systems (HuCS):** The goal of HuCS R&D is to make computing and networking more useful through laboratories, technologies that provide knowledge from distributed repositories, multi-modal interactive systems, and virtual reality environments.

- **Education, Training, and Human Resources, (ETHR):** The goal of ETHR R&D is to support research that enables modern education and training technologies, including technologies that support lifelong and distance learning, and curriculum development.

In addition, the CCIC has established an Applications Council to promote multi-agency leadership in the early application of advanced computing, information, and communications technologies, with special focus on projects that are widely applicable to Federal agency missions. Also during this year, the Federal Networking Council, which acts as a forum for networking collaborations among Federal agencies to meet their research, education, and operational mission goals, is operating under the purview of the CCIC.

The estimated FY1996 HPCC Program budget for the participating Federal agencies was \$1,023 million. For FY 1997 the President requested \$1,017 million.

The world has entered the Information Age and the U.S. is leading the way. Federal R&D in advanced computing, communications, and information technology is directly responsible for this leadership. In order to remain effective leaders, U.S. academia, industry, and government must be able to exploit the most advanced technologies in the most affordable way. Computing, Information, and Communications R&D programs will be key driving forces for advancing these technologies and their application to a more secure, better-quality life in 21st century America.



II. Research Accomplishments

Introduction

As we approach the 21st century, we are clearly leaving behind the Industrial Revolution and embarking on the Information Age. Advances in the sciences that used to take decades now take years or months, as research and implementation move ahead with dizzying speed. As part of the Committee on Computing, Information, and Communications (CCIC) of the National Science and Technology Council (NSTC), the High Performance Computing and Communications (HPCC) Program is a driving force in information technologies, computing, communications, and information infrastructure, and a major component of America's investment in its future, helping to maintain and widen the competitive lead that will keep our citizens productive well into the next century.

Through the HPCC Program, Federal departments and agencies fund computational projects in the fundamental sciences in areas of scientific research where computational capabilities limit innovation. Efforts that focus on computationally intensive problems in science and technology are called "Grand Challenges," while those that focus on informationally intensive problems are called "National Challenges."

The HPCC Program was authorized with bipartisan Congressional support in 1991 through passage of the High Performance Computing Act. Created as a dynamic R&D program to extend U.S. leadership in high performance computing and communications, it has provided the sustained focus needed for developing these technologies and has adapted to the needs and opportunities of a changing world. As we enter the second five years, the Program continues to support outstanding science and foster innovative solutions to complex problems. This year's book highlights some of the Program's most notable achievements, including:

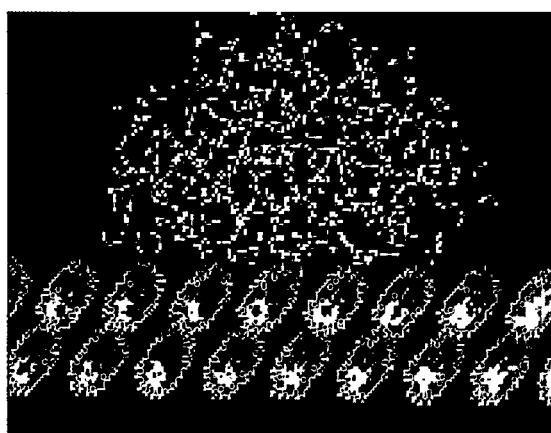
- High performance computing systems have brought scalable parallel computing into a new era. Today's focus is on scalable architectures that deliver high performance, as well as on new component technologies, embedded systems, and wireless communications technology.

- Advanced software technologies R&D continues to make marked progress in microkernel operating systems, programming languages, scalable input/output (I/O), and tools for developing software for parallel computing systems.
- Interdisciplinary R&D has led to new science and engineering in a wide range of disciplines. By modeling weather, the environment, pollution, climate, manufacturing, the human body, galaxies, and the way new proteins are made, high performance computing and communications have brought the Nation new knowledge and new abilities.
- Enabling technologies for information infrastructure have focused on distributed computing, reliability, mobility, security and privacy, and tools for building distributed applications.
- Basic research and education remain a foundation of the Program. Fundamental new ideas are being investigated across the participating organizations.

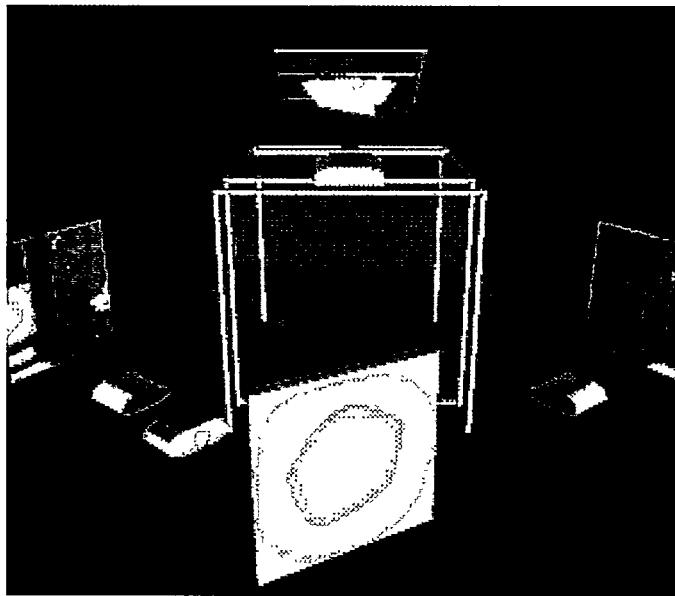
This book describes more fully these and other exciting FY 1996 research areas as well as the FY 1997 research directions of dedicated and creative scientists, engineers, individuals, and organizations throughout the country. The examples summarized here comprise only a small fraction of many impressive accomplishments. Others can be accessed via the World Wide Web, either starting with <http://www.hpcc.gov>, or using the URLs available at the end of each section.



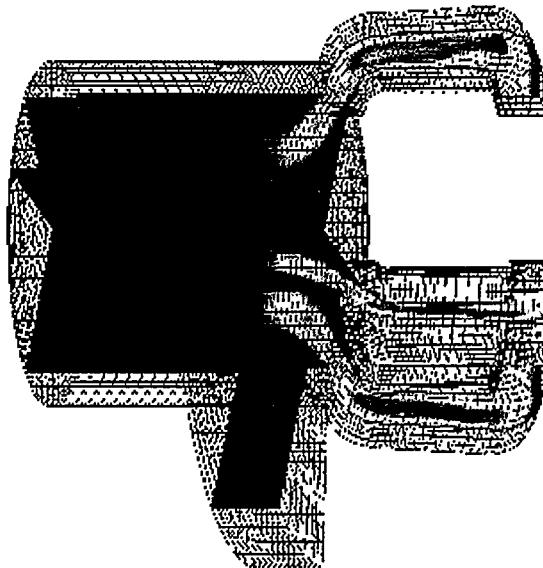
World Wide Web traffic is the largest and fastest growing segment of the Internet. Globally available resources create access patterns for which the geographical, as well as temporal, aspects must be considered.



Molecular forces create tension at the surface of liquid films that makes them contract to a minimum surface area. Two researchers have simulated the formation of water droplets on the surface of such widely used polymers as polyethylene and Teflon. The above image represents the simulated droplet of water formed from a monolayer of 324 water molecules on the surface of crystal teflon.



The CAVE is a surround-screen, surround-sound, projection-based virtual reality (VR) system developed under the HPCC Program. The illusion of immersion is created by projecting 3-D computer graphics into a 10'x10'x9' cube composed of display screens that completely surround the viewer. It is coupled with head and hand tracking systems to produce the correct stereo perspective and to isolate the position and orientation of a 3-D input device. A sound system provides audio feedback. The viewer explores the virtual world by moving around inside the cube and grabbing objects with a three-button, wand-like device.



An image of an incompressible fluid flow in a two-stroke engine, computed using the CMPGRD and OVERTURE libraries to support computations involving complex geometries. The wire frame shows the location of the grids, including transverse ports, cylinder, cylinder head, and exhaust port. The colors indicate the magnitude of one velocity component going left to right in this picture, with blue representing the maximum positive velocity and violet representing the maximum negative velocity.



Networking

DARPA, NSF, NASA, DOE, NSA, NIST

HPCC Program support for large scale networking research has had a significant impact in advancing communications technology and has led to important advances in the mode of communications now being adopted by the commercial sector. In the past few years, the Federally supported communications infrastructure has taken a vital lead in every area of research and education. Advances in networking technology have changed the way education, science, and business are conducted, and everyday activities of Americans are pursued.

Several HPCC networking projects have advanced the availability and effectiveness of distributed applications. The NSFNET Connections Program, for example, supports the acquisition of high performance network connections at research institutions to provide capabilities that are required for state-of-the-art scientific and engineering applications. This Program also supports connections at lower speeds that facilitate the extension of the Internet into K-12, public libraries, and museums.

Two National Laboratories are collaborating with universities and private industry on the ESnet and GigaNet projects to interconnect multiple parallel

supercomputers via a high-speed ATM network. The enhanced network will allow researchers to work concurrently on a single computation-intensive or information-intensive application. This joint agency metacomputing experiment provides the computing power to solve very large scale problems.

The Air Pollution Distance Learning Network is used in training researchers to employ advanced climate models as well as to disseminate information for environmental problems. Researchers have used this network to collaborate through a national video broadcast, using the Multicast Backbone (MBONE) on the Internet, while providing simultaneous distribution to an interested international audience. Use of this Network and the MBONE for such collaborations has provided state and local agencies with evidence of the practical uses of the Internet.

A collaborative project at the Advanced Light Source—a new, high-powered synchrotron light source for producing copious, focused X-rays—is providing investigators from nine geographically distributed institutions remote access to analytical tools for lithography and crystallography for research in material sciences.

Advances in wireless and mobile information systems have spurred a revolution in microsystems technology needed to make portable devices usable as research tools. These techniques range from the development of highly capable, low power, small size “smart radios” to distributed computing systems that can deal with mobility and sporadic connectivity patterns. Researchers have also developed an experimental system to investigate issues of seamless and transparent migration across different kinds of networks, variable bandwidth availability, data consistency problems caused by frequent disconnections, and access to resources in the host environment.



The next step beyond parallel computing is wide area computing. As wide area networks close the bandwidth gap with local area networks, it becomes possible to treat geographically dispersed resources as a tightly linked system. Wide area computing uses high-speed networks to connect supercomputing and data resources to create virtual computers with teraflop-class computing capacity, thus providing the potential for petabytes of information resources, and immersive collaborative environments.

www.hpcc.gov/blue97/net

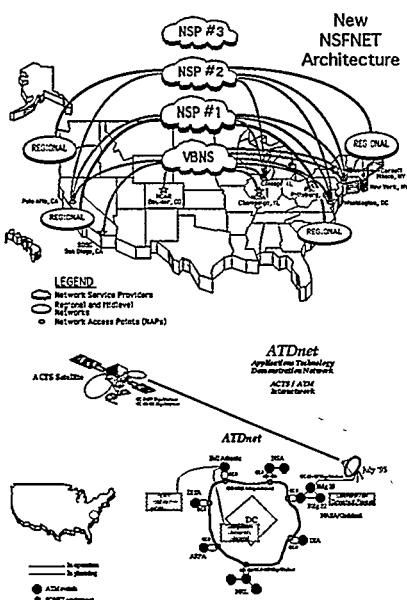


The Old NSFNET Backbone

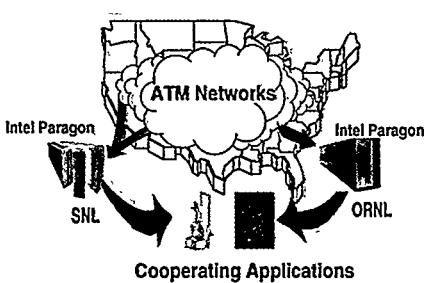


Networks and the Internet

The Internet began with the development of ARPANET in the late 1960s. It was initially an experimental network that connected a few researchers and provided a platform from which to develop packet-switching techniques and technologies that are still in use today. Various Federal agencies established their own networks to support their respective research communities.



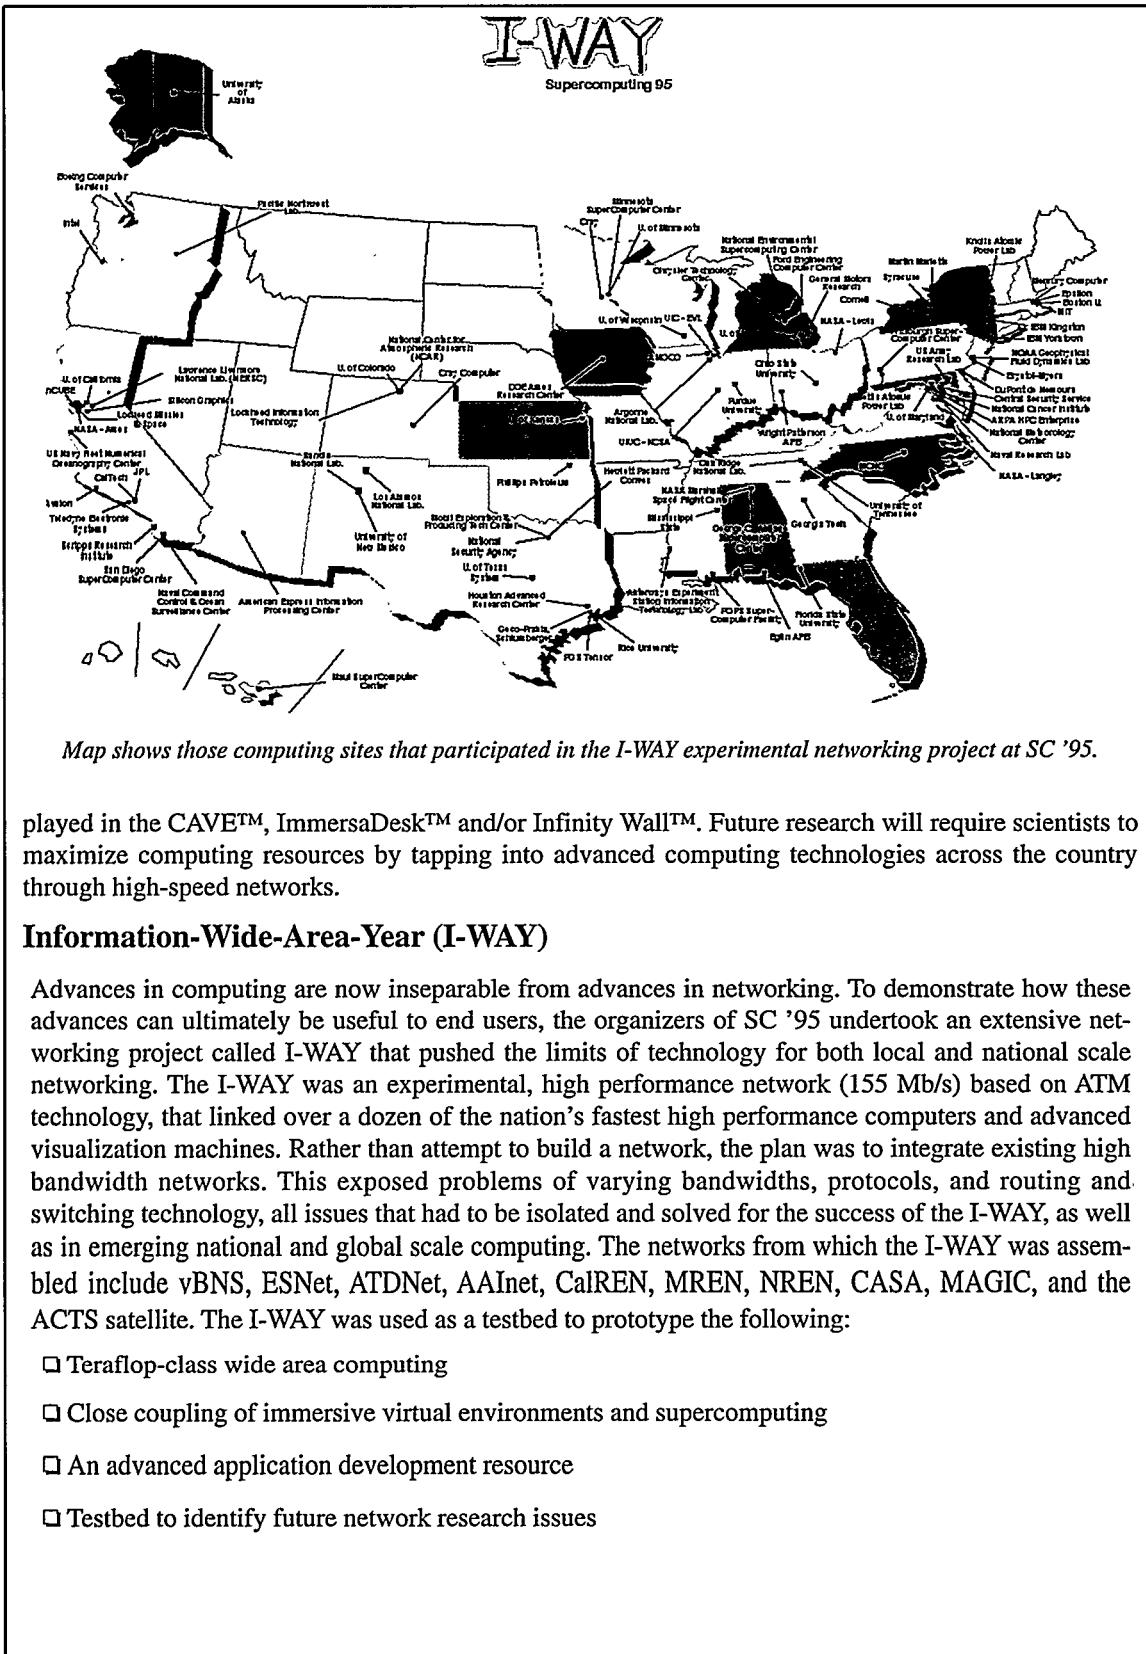
The NSFNET was created in 1985 to provide a network infrastructure for the research and education community within the U.S. and also as a testbed for network-related research. The community-level service was turned over to private sector Internet service providers in April 1995, at the same time the NSF created the vBNS network, a commercial infrastructure that supports the very high performance needs of the research and education communities (initially at 155 Mb/s). DOE has acquired commercial services and capabilities for its ESnet to provide high bandwidth — up to OC-3 (155 Mb/s) ATM connectivity — to the worldwide energy research community. ESnet also provides a collaborative testbed for merging its commercially procured domestic and international ATM technology sources with the Internet. ATDnet, a collaboration among NASA and various Department of Defense agencies, provides the HPCC ESS and NREN Projects with an effective testbed for gigabit network R&D connecting metropolitan area networks. It is implemented using 2.4 Gb/s SONET/ATM networks.



The Internet's phenomenal growth worldwide has been leveraged through industry investment that has been much larger than the small but persistent base of support from various government agencies and the HPCC Program. The unprecedented growth of the Internet worldwide is a result of stimulation provided by the HPCC Program and from educational, public service, private sector, and personal investment. A more detailed history of the development of the Internet was published in the FY 1996 Blue Book.



At SC '95, an annual conference held to showcase high performance computing and communications, several large networking projects were featured — including the GII Testbed and the I-WAY. The GII Testbed was created for SC '95 to accelerate the development of the teams and tools required for the types of distributed computing necessary to meet Grand and National Challenges and beyond. The GII Testbed featured interactive 2- and 3-D scientific visualization and virtual reality demonstrations, where the simulations were computed on remote supercomputers and the data sent via high-speed networks to the conference site, where it was rendered and dis-





High Performance Systems

DARPA, NSF, NASA, DOE, NSA, NIST

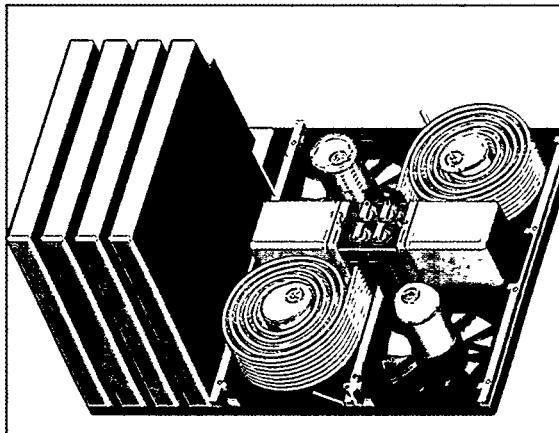
Over the past year, advances in high performance systems include system prototypes that demonstrate an order of magnitude reduction in cost per megaflop; workstation clusters using novel operating system techniques that achieve an order of magnitude reduction in application level latency (where the latency is the time required for an application running on one processor to retrieve data located on another processor); unprecedeted performance levels of graphics rendering (50 million polygons/sec); and dramatic improvements in low latency communications rates for local area distributed computing applications (2 Gb/s). Advances also include formal theoretical methods for the verification of complex chip design, technologies required to design, fabricate, and test computer system microarchitectures, and the infrastructure required to allow heterogenous architectures to function together in an integrated system.

Heterogeneous computing environments combine computational engines with different architectures, file systems, and high speed interconnects. These heterogeneous systems provide support for solving computational problems that require different capabilities of their constituent subsystems. HPCC-supported researchers used a collection of heterogeneous supercomputers on the Internet to simulate and render the collision of two galaxies. This process is too large to implement on any single system; therefore, researchers combined several geographically separated high end systems by high speed networks to simulate appropriate parts of this single model. The successful simulation yielded graphical display capabilities not possible before, and resulted in the first computer-generated scientific IMAX (wide screen) movie, currently showing at the Smithsonian National Air and Space Museum.

Networks of workstations (NOWs)—systems that use a newly developed communication fabric to connect collections of heterogeneous, commercial workstations—demonstrate the potential for supercomputer performance with workstations already in use. Research focuses on communications microarchitectures and high performance memory interfaces that improve bandwidth and latency for dis-

tributed applications. NOWs require research in distributed operating systems, fault detection and recovery, resource discovery and allocation across multiple administrative domains, real-time response guarantees, and multi-level security. This research demonstrated a promising first generation distributed resource management system that uses techniques for resource discovery and allocation, and advanced fault management. Researchers using a high performance interconnect of clusters of up to 50 workstations have demonstrated applications that engender memory latencies of only a few microseconds. Driving applications for NOWs include the fastest known World Wide Web indexer and the world's fastest connected components kernel for solid-state physics.

Researchers using scalable, high performance computing systems for embedded applications focus on techniques to ensure that the highest possible performance scalable computing technology is openly available from a commercial technology base. R&D



The MARQUISE program will repackage the J90 processor into a system with four parallel vector processors and 1024 Mbytes of DRAM, providing a peak performance of 800 Mflops. The prototype will employ four advanced packaging techniques to transform a conventional computer room-sized machine into a pizza box-sized package. The techniques are (1) multi-chip modules for weight and volume reduction; (2) diamond substrate to support high power and current density with robust reliability; (3) innovative phase change spray cooling to remove heat with very little added weight, volume, and power overhead; and (4) high density flex interconnect to support the signal-rich network interconnect.



investments include research into novel architectures, real-time operating systems, standard library interfaces, program development environments, and demonstrations of the early insertion of new technologies into agency mission driven applications. The two level multicomputer architectural concept cleanly separates computation from communication to facilitate the development of heterogeneous systems critical to embedded applications. Researchers demonstrated a multicomputer capable of reducing costs to \$10/Mflop, and this design is being transferred to industry for future insertions into embedded applications.

University researchers developed the PixelFlow system prototype that produced unprecedented performance levels for graphics rendering. The PixelFlow system uses a massively parallel processor-per-pixel approach capable of achieving a record 3-D graphics rendering performance of 50 million polygons per second. A major vendor will acquire the license to this technology and produce a series of new products that could set new standards in scalability, programmability, and performance.

In addition to the well known computing paradigms of shared memory and distributed memory, there is a third paradigm—distributed shared memory (DSM)—in which memory is physically distributed but logically shared. The DSM technique enables the user to view distributed memory resources as a single address space and provides transparent access to computational resources in scalable systems. DSM research efforts investigate both software and hardware approaches to implement a logically shared view. For example, the FLASH project designed a very high performance protocol engine to implement DSM on a Very Large Scale Integration (VLSI) chip. This programmable chip supports experimentation with various shared and distributed memory protocols emerging from the research community.

An exciting new area of research is Biomolecular Computing (BC). Researchers have applied innovative molecular biology methods to a Hamiltonian path-finding problem to determine a feasible path through a complex search space using molecules of DNA in a test tube. BC has the potential to solve problems much faster than traditional silicon based supercomputers and to provide technologies capable of storing information in a tiny fraction of the

spatial volume required for today's storage media. This is only one of several high risk technologies being explored under the HPCC Program over the next few years. Other novel approaches include computing systems based on quantum, molecular, and amorphous techniques. These systems explore entirely new algorithms to address problems in cryptography, very large distributed information systems, and symbolic computing.

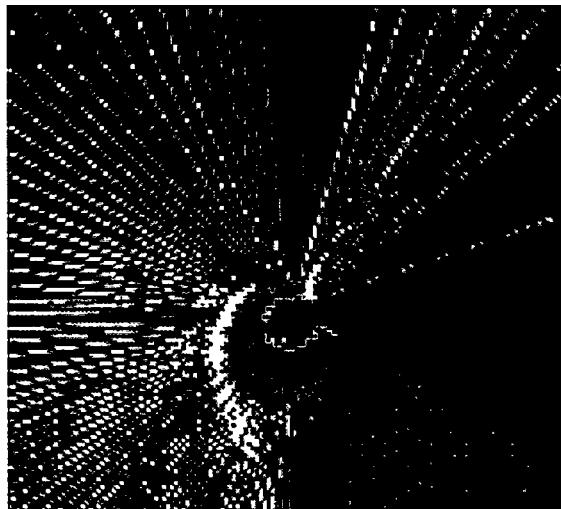


Software for High Performance Computing

DARPA, NSF, NASA, DOE, NIH, NSA, NIST, NOAA, EPA

Widespread efficient use of high performance computing is possible only when sufficient software infrastructure and tools are available to enable users to harness transparently the capabilities of these systems. Issues such as languages, operating systems, algorithms, compilers and libraries, portability and scalability of software, and domain specific tools must all be addressed to ensure that users are able to use these systems effectively.

Languages and libraries for scalable high performance computing have seen the development of pC++, a data parallel language, and ZPL, an array-based programming language for scientific computations. Compiler development has been aided by new techniques for conventional programming languages for parallel computers that deliver substantial performance improvements on small scale shared memory multiprocessors. Critical to the continued and widespread use of scalable parallel processors is the ability to generate efficient parallel programs from sequential code. There are several approaches to this problem. In one, a parallelizer called SUIF uses whole program analysis to achieve that goal. Another employs interprocedural compile time parallelization analysis on programming languages such as FORTRAN, C, and C++ to obtain significant speed-ups. A



The Time Tunnel metaphor in PABLO provides the user with a graphical display of events that occur in a parallel code such as computations, communications, or I/O behavior. Event time lines are mapped in a circle, with each color representing a certain behavior.

third effort has the goal of advancing the art of programming language design and implementation, and the art of systems building. Library development during the past year has seen the first draft specification of the High Performance C++ Library by the HPC Forum, and SCPLib, a portable programming library for complex irregular problems.

Current and future scalable parallel computing can only succeed if there are sufficient tools and infrastructure to enable their efficient utilization. Tools such as Paradyne or PABLO that monitor the behavior of application code; the D-System project that supports development tools and debugging aids for data parallel programs; a tool (PGRT) for instrumentation and visualization of realtime embedded systems; and new file system policies that use the results of input/output (I/O) characterization tools — all contribute to increased acceptance and efficient use of parallel high performance computing.

Operating systems development is critical to the efficient use of scalable heterogeneous systems, and several are under development. An improved microkernel based operating system was adopted for the Scalable Realtime Operating System (SRTOS) and for the Hiper-D program. The exokernel is a new paradigm in the relationship between application programs and operating systems. The Synthetix OS creates highly modular operating systems through incremental specialization, and an OS technique realized through SPIN allows application code to be loaded into the kernel at runtime, thereby decreasing overhead.

The Scalable I/O Initiative (SIO) was created to address the balance that should exist between the performance demanded by today's applications and the I/O requirements to support these applications. Analyzing several I/O-intensive applications using tools such as PABLO, the SIO has achieved a ten-fold performance gain in applications codes by matching file operations to the application's I/O patterns.

www.hpcc.gov/blue97/software



Security and Privacy, and Electronic Commerce

DARPA, NSF, NASA, DOE, NIH, NSA, NIST, NOAA, EPA

Computerization and networking have brought great benefits of efficiency and new capabilities to national infrastructure functions, such as the financial system, healthcare, telecommunications, power generation and distribution, air traffic control, and defense. All now rely critically on the ad hoc assemblage of computer systems and internetworking technologies that serves as the nation's computing infrastructure. With these benefits, however, comes the prospect that failures in the computing infrastructure, whether due to physical damage, errors, or intentional attacks, could disrupt these vital activities. They are dangerously vulnerable. As a whole, the computing infrastructure, developed in a partially spontaneous fashion and managed under decentralized control, is fragile and has many poorly understood interactions. Making it resilient enough to support critical services safely requires security mechanisms to prevent break-ins or malicious misuse of systems and survivability technologies to ensure that critical capabilities are preserved when problems occur.

A number of improvements have been made in network protection technology. Firewalls, one of the primary mechanisms for protecting an organization's internal networks and computers from outsiders, have been combined with a strong Domain and Type Enforcement (DTE) security mechanism that can flexibly restrict what clients can use what applications and services over the network. Researchers have experimentally enhanced an industry standard firewall with the DTE mechanism and shown the ability to guard against theft of sensitive data over the network.

To enable widespread use of cryptography for authentication and authorization of access to network-based services, the concept of a generalized public key certificate infrastructure is being addressed by a National Laboratory effort. The Synergy research program is developing an open architecture based on security policy flexible operating system microkernels. The Secure Heterogeneous Application Runtime Environment (SHARE) is developing a high-bandwidth secure System Area Network for the embedded HPC market. SHARE will have dedicated cryptographic

hardware supporting packet-switched internode communication.

Mechanisms for secure transactions are needed to enable electronic commerce. Researchers have extended the Kerberos authentication system to allow use of public-key cryptography and digital signatures. In collaboration with leading industry consortia (CommerceNet, World Wide Web Consortium, Open Software Foundation, and Financial Services Technology Consortium), a U.S. company has initiated a cross-industry pilot to develop and demonstrate an electronic payment negotiation framework which has attracted interest from Netscape, MasterCard, Visa, Microsoft, and VeriFone. NetBill was developed as an electronic commerce system supporting economical, secure sales of goods costing as little as ten cents. (The expense of conventional credit card processing makes such small transactions impractical.)

Survivable systems must be able to find and react to attempted attacks. Researchers have developed a new specification-based intrusion detection technique able to detect even types of attacks that have never been seen before, since hackers constantly discover new ways to enter and interfere with a system, and a communications thumbprinting scheme to trace the activities of attackers who range widely over the network in the course of a break-in. Progress has also been made in another aspect of survivability. Multicast communications protocols have been developed for distributed systems that continue to provide reliable service even when a compromised processor behaves maliciously. These protocols take advantage of Fortezza's hardware-based encryption and authentication.



Visualization, Virtual Reality, and Human-Machine Interfaces

DARPA, NSF, NASA, DOE, NIH, NIST, ED, EPA

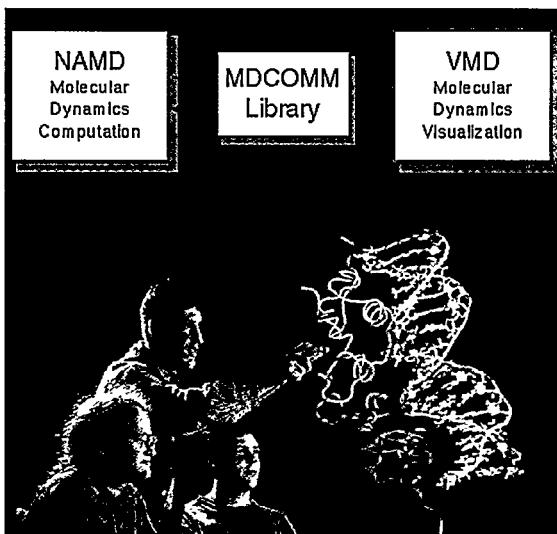
The advances in computers and communications have increased our ability to generate, collect, and manipulate vast amounts of data relating to our everyday lives: bank transactions and airline reservations, automation of manufacturing and commercial enterprises, sophisticated simulations of physical phenomena, state-of-the-art medical treatment, and education and learning. Research in visualization, virtual reality, and human-machine interface technology will enhance human capability to use and manage data, create information, and facilitate a symbiotic relation among humans, computers, and information.

Visualization techniques, developed with HPCC support, translate complex scientific data into images that help scientists and engineers uncover and understand underlying processes. Two recent examples are the application of computer-aided design and graphics to an ophthalmology project for modeling and visualization of the human eye cornea, and the custom design and fitting of contact lenses. Other efforts are pursuing "perfect realism" in computer generated images by including factors such as the effect of light on the human perception of a subject or a scene.

For remote access and use of large data sets, the Distributed Visualization Project has developed data compression algorithms that provide fast data browsing and interactive visualization capabilities. The 3-D Imaging of Complex Geologies project allows remote and rapid processing of terabyte data sets. Such capabilities have enormous potential as, for example, they allow a geologist to remotely access relevant data for seismic imaging of complex underground structures.

Virtual Reality (VR) is created by combining visualization, display and control technologies. Using computerized goggles and gloves, a user can create and be surrounded by an artificial environment that emulates real life, thus enabling selective probing and analysis of real processes. The research to enable such capabilities is led by several HPCC projects; the MDScope, shown right, is an application of VR.

Human-machine interface technology research on image understanding has led to the development of



MDScope allows scientists to explore the attributes of macromolecules in an immediate and visual way, and facilitates research into more complex systems that could not be readily understood using traditional methods.

a prototype "omnifocus" electronic camera, with potential applications (medical) such as endoscopy with extremely high depth of field. In addition, manufacturing (rapid visual inspection of 3-D objects, boresight applications, workcell monitoring, and robot visual guidance), construction, environmental, space, and other field operations involving surveillance of extended sites have been improved by this technology. Other efforts include the creation of remote access capabilities such as a remotely controlled, automatically focused, network-based video microscope, and the Microtele-operation-at-Atomic-Scale Nanomanipulator. Research on new algorithms for parts handling is applied in industry for testing and visualizing ways to assemble and disassemble a structure or a machine. Theoretical investigations of robot motion planning are applicable for wheeled or snake robots operating in space and underwater environments, and have led to establishment of a space robotics spinoff for aircraft in free flight.

www.hpcc.gov/blue97/viz



Digital Libraries

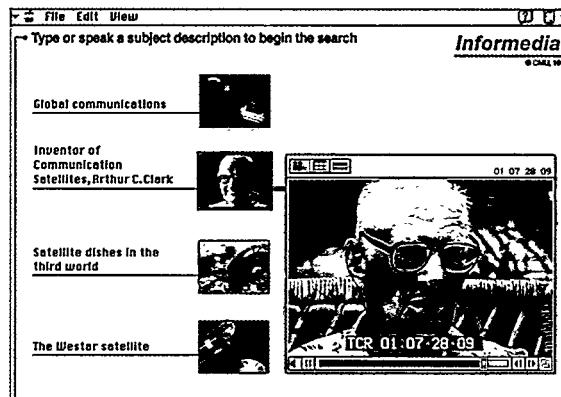
DARPA, NSF, NASA, DOE, NIH, NSA, NIST, NOAA, EPA

A multi-agency research initiative in digital libraries focuses on creating technologies required to manage future globally distributed information systems. Issues of primary interest are: interoperability, management of large and complex information spaces, and support for effective user interaction with large, distributed information systems. Research in digital libraries explores critical issues in the development and use of large scale networked knowledge repositories. The goals are to efficiently capture, store, organize, search, process, and retrieve knowledge from electronic collections containing text, images, maps, audio recordings, video and film clips, and the combinations thereof (multimedia).

A digital library is conceptually analogous to a traditional library in the diversity and complexity of its collection. A single digital library contains many terabytes of information, distributed throughout the world. Digital libraries are designed to be used by a broad spectrum of the population—not only scholars and researchers, but educators, students, and the general public. By developing digital libraries, many of the limitations and disadvantages associated with managing and using physical collections can be overcome.

One information intensive challenge of digital libraries is semantic interoperability—the ability to search for a topic across multiple sources that use different vocabularies. Such interoperability takes the form of “vocabulary switching,”—using the terminology of one source to search another and identifying the same concept expressed with different words. Scientists need this technology in order to find information outside their specialties effectively. In one of the largest information science calculations to date, ten days of dedicated supercomputer time were used to generate concept spaces for 10 million journal abstracts across one thousand science and engineering subject areas.

An important addition to the Alexandria Digital Library project (one of the six projects under the cross-agency research initiative) is a gazetteer that maps named earth features (such as towns or rivers) to their spatial footprints, a critical attribute of a geographic information system. This mapping was produced by merging the Geographic Names



The Informedmedia system is a Digital Video Library with voice recognition allowing users to enter semantic queries. The query results are displayed in rank order with a video snippet to provide context.

Information System Gazetteer (which characterizes an international set of features and is obtained from the Defense Mapping Agency) and the Board of Geographic Names gazetteers (which characterizes U.S. geographic features). The integration required error correction and rationalization of the different feature classifications in the two gazetteers. This new gazetteer is used to support content-based searches in which a user enters the name of a geographic feature, and the search engine retrieves the footprint of the feature, which is then matched with footprints of items (such as maps and images) in the collection.

Another digital library project provides on-line access to a large collection of California environmental data. An innovative “news-on-demand” service provides search and retrieval of radio and television broadcast news stories tailored to a user defined profile. This service delivers news items to a subscriber on a regular basis via the Internet.



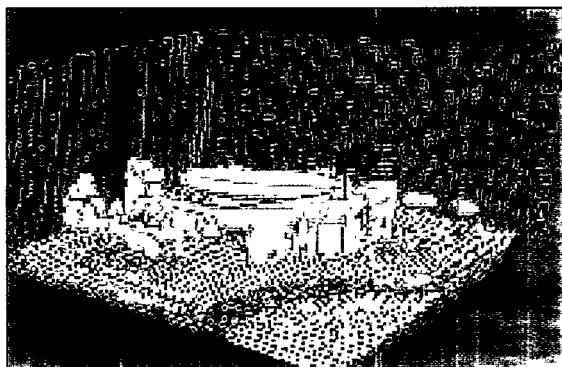
Weather Modeling

NSF, NASA, DOE, NOAA, EPA

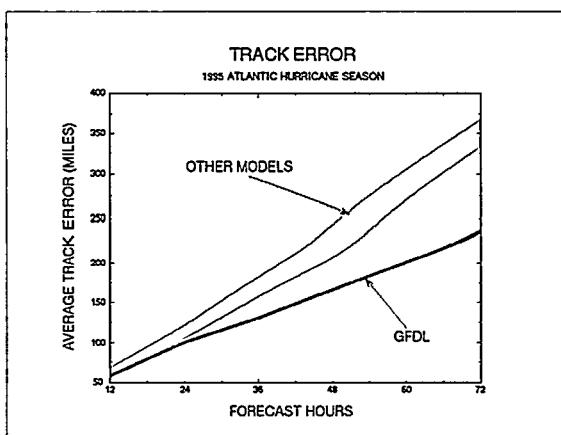
Weather forecasting, modeling, and research are advancing rapidly with the advent of high performance computing and communications systems. Research groups at several agencies are moving traditional models to scalable systems. These models are then used to test the validity of the current understanding of the physics of weather and to develop more detailed, robust models. When the models are sufficiently trustworthy, they are used for operational forecasting by the National Weather Service and to drive air quality assessments for environmental analysts.

Perhaps the most dramatic example of recent advances is the hurricane modeling system that became operational with the 1995 hurricane season as the culmination of a long development effort. The operational implementation of this model was made possible only by the use of the fastest computers available to the National Weather Service. When moved to a parallel operational setting, a speed-up of 18 over a serial implementation of the model was realized. Without this speed-up, forecasts could not be made soon enough to provide timely information. This is a clear example of the advantages of modern high performance computing applied to a problem which affects each of us every day. Further improvements are expected from better initialization procedures, more flexible grid design, and, of course, faster computers.

Our daily weather forecasts start as initial value problems on the National Weather Service supercomputer. Satellite, radar and balloon observations are assimilated and provide initial conditions for computer models of the atmosphere. Since these models have approximately a million degrees of freedom, and must be integrated with short time steps (minutes) for two weeks, it is essential to design the models in such a way that they can be efficiently implemented on Massively Parallel Computers (MPP). Since the current generation of models was designed for sequential vector architectures, transferring them to MPP or scalable parallel architectures is a challenging task.



Sample output from an MPP implementation of a weather model showing the improved level of detail obtainable through one of these models. The figure may show state-level detail not available in a traditional model.



The average forecast track error for the 1995 Atlantic hurricane season from the new forecast system developed by GFDL (red line) contrasted with the range of forecast track errors from earlier models.

Scientists at several National Laboratories and Supercomputer Centers, and the Center for Analysis and Prediction of Storms are investigating the application of scalable systems to the problems of mesoscale weather forecasting. These systems are exploring phenomena such as tornados, thunderstorms, and flash floods.

The National Centers for Environmental Prediction are examining two approaches to global weather prediction: ensemble forecasting for determining the robustness of forecasts and a "conformal expanded cube" that may be more efficient to parallelize than the traditional spectral models.



The Environment and Energy Management

NSF, NASA, DOE, NOAA, EPA



Simulation of nitrogen deposition to the Chesapeake Bay and surrounding area during a rain storm.

Global Climate Modeling

General global circulation models help us understand the processes controlling the earth's climate. Researchers compare models and evaluate the fidelity with which models simulate observations at different resolutions. Better resolution of the Atlantic ocean in an ocean circulation model accurately reproduced the path of the Gulf Stream into the open ocean north of Cape Hatteras, NC, whereas previous versions predicted the Gulf Stream would follow the U.S. coastline. Plans to develop high resolution global circulation models that couple atmosphere and oceans depend on continual performance improvements from HPCC technologies.

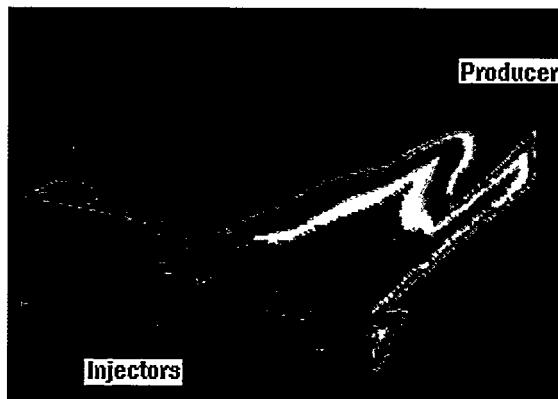
Four-dimensional data assimilation melds observation data from satellites and other platforms into earth system models to provide more accurate predictions of physical phenomena. The successful development of a parallel assimilation system for gases in the stratosphere is a major step toward providing consistent gridded data sets to the community studying global climate and air quality. Live access to climate data through the World Wide Web has eliminated many barriers to exchanging this data. Data repositories, like SALICIN, enable remote sensing data sets to be browsed, queried, processed, and downloaded over the Internet.

Environmental Modeling

Groundwater contamination is a serious environmental and economic problem, with detoxification costs for existing sites estimated in the hundreds of billions of dollars. Parallel algorithms for modeling flow in permeable media, used extensively by the petroleum industry to increase oil recovery, can also

be used to model movement of contaminants in groundwater. Collaborative efforts such as these involving industry, academia, and government researchers enable more realistic simulation of complex real world environmental problems. Research on nonlinear optimization and control techniques used to minimize groundwater clean up costs is being applied to traditional "pump and treat" methods and to more experimental bioremediation methods. Researchers are also exploring new approaches to solve broader regional-scale problems where groundwater flow impacts river tributaries and coastal aquifers.

Pollution in rivers, lakes, and estuaries can damage plant and animal life. Researchers use numerical simulations to estimate the impact of decades of agricultural pesticide application and to study the effect of major storms and floods on the distribution of toxic materials in lakes and streams. Recent modeling studies of PCBs in Lake Michigan demonstrate that resuspension of toxic sediments is affected by the amount of clean sediment overlaying the contaminated sediment.



Output of compositional simulator evaluating use of horizontal wells with vertical drainholes for tertiary oil recovery with carbon dioxide.

Concern about increased risk of cardiopulmonary disease from airborne particles accelerated the development of a regional particulate model (RPM) enabling researchers to assess the impact of several air quality control programs on fine particle concentrations in the air. A multi-pollutant modeling framework has been designed to facilitate the use of complex air quality models, like RPM, by state and industry environmental groups to support risk



assessments and predict consequences and costs of alternative control strategies.

Earthquake Damage Prevention

Advances in parallel computing technology enable computational ecology models typically developed for a specific area and time to be scaled up for useful analysis. For example, researchers are able to simulate the response of the San Fernando Basin to earthquake-induced ground motion. The ability to quantify the effects of earthquakes over an entire basin has practical implications for building code requirements and design of bridges, dams, and overpasses.

Energy Management

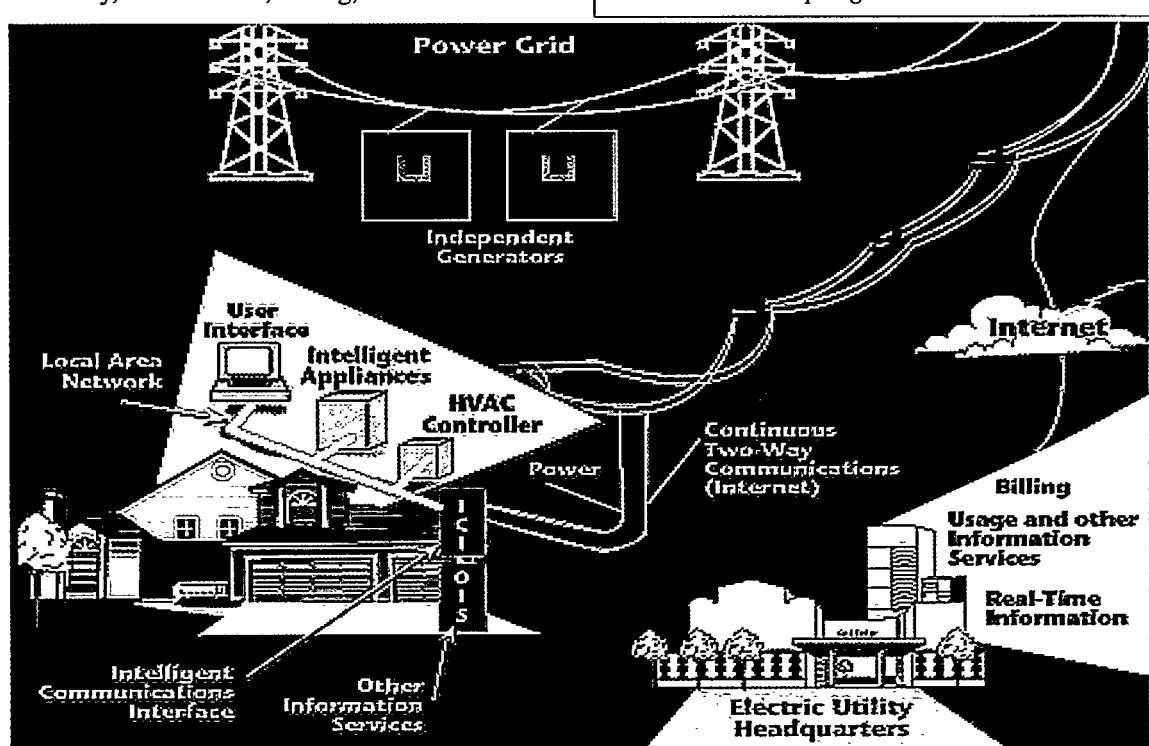
Deregulation in the electric utility industry is increasing the need for effective use of information technologies. To enhance the existing capabilities of utilities for real-time energy supply and demand management, new technology for distributed systems needs to be developed in areas such as interoperability, authentication, privacy control, and multicast data aggregation. Efforts that support the development and implementation of both wide-area-based and distributed network tools, services, and protocols that enable energy utilities to improve efficiency, conservation, billing, and customer ser-

vice, and that promote end-user interaction and control over their energy use are being supported.

The two projects currently under way:

- One group is developing a prototype system that permits the remote monitoring and control of the energy usage of multiple commercial buildings across the Internet from a single control center. The system will employ an Internet-to-building Energy Management Control System gateway configured to drive the development of a standards-based user interface.
- A private company is working to demonstrate a system for the remote acquisition, central analysis, and distribution of energy information at a level of detail not previously available, and to address fundamental electric utility issues that will improve the way the utilities interact with their customers. The project integrates "off-the-shelf" hardware and software into a system that will monitor power quality and energy usage devices placed at industrial sites and demonstrate services such as power quality event summaries, custom billing services, and on-line energy usage information reports.

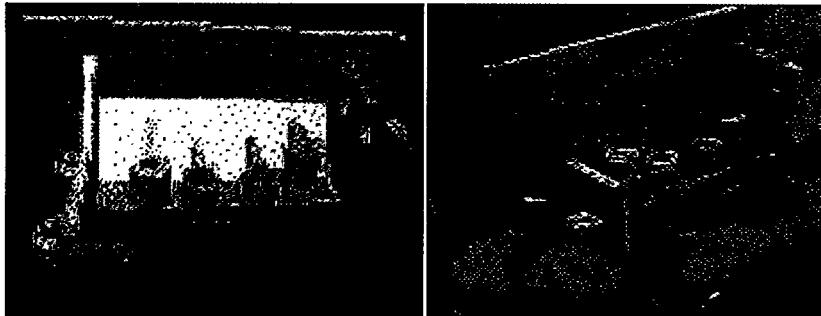
www.hpcc.gov/blue97/env





Manufacturing — Design, Processes, and Products

DARPA, NSF, NASA, DOE, NIST



A computer generated image of a metal flange [far left], part of a subway token collection machine, and a 3-D solid prototype [left] fabricated by a stereolithography device at the San Diego Supercomputing Center's Telemansufacturing Facility.

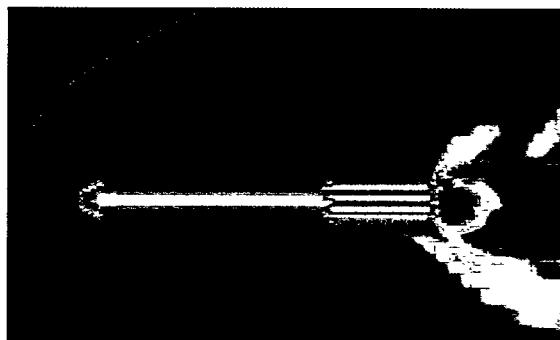
Modern manufacturing is increasingly driven by the need for both high flexibility and speed in process and product development. HPCC technologies enable high fidelity computational prototyping in a shorter time, resulting in products that are lighter, stronger, safer, and cheaper to manufacture and use. HPCC infrastructure can also provide more efficient and flexible production systems through integration of planning, design, production, quality control, marketing, and user services. Such needs span many industries, including aerospace, chemical, electronics, and materials processing.

The continuing need for high end computing in manufacturing is illustrated by semiconductor and integrated circuit design where feature sizes have reached the submicron level. Since the cost for a new semiconductor process exceeds a billion dollars, accurate device and process simulation is key to maintaining industry leadership. However, such simulations necessitate modeling complex nonlinear effects not present in current simplified models, and require both teraflops of computational power and terabytes of memory. Agency sponsorship includes several projects in ultra small device modeling including a joint effort by several universities and private companies.

Aerodynamics research also has severe computational requirements but high payoff. For example, researchers used a dedicated 512-node computer for a 4.5 million grid point simulation analysis of the flow about a Delta II rocket, explaining an anomaly that had sent a launch vehicle into the wrong orbit. In another project, a heterogeneous parallel instrumentation, data collection, and visualization facility for wind tunnel testing uncovered previously unknown noise sources in the DC-10.

Optimizing complex multidisciplinary designs is also beginning to be possible thanks to distributed environments and toolkits such as DAKOTA and FIDO. DAKOTA has been used to improve chemical vapor decomposition reactors used in integrated circuit manufacturing.

Integrating advanced design technologies with production, quality control, marketing, and customer service systems is the focus of the Advanced Manufacturing Systems and Networking Testbed (AMSANT), an experimental facility for testing integration of distributed virtual manufacturing enterprises. Data exchange protocols are a necessary part of this infrastructure. Recently, the PlantSTEP Consortium developed the first such protocol for the process plant industries.



The pressure field surrounding a Delta II rocket.



Health Care and Biomedical Imaging

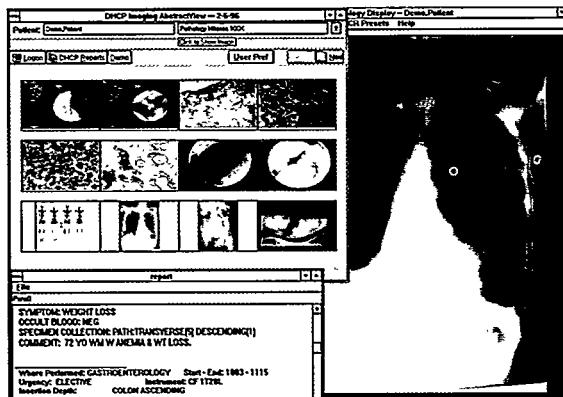
DARPA, NSF, NIH, VA, AHCPR

HPCC technology provides a potentially huge payoff in health care. From teleconsultations to availability of health related databases, the effect of using high performance computing and communications will enhance patient care, improve drug design, and broaden access to medical information.

Current work on Chagas disease — a parasitic illness — is an example of successful health care collaborations. Biochemists, computer scientists, and computational chemists combine rational drug design methods with virtual reality and video teleconferencing technology to evaluate potential drugs for Chagas disease and related parasitic illnesses. In another effort, a magnetic resonance imaging resource center recently made its instrument available through the Internet, supporting the concept of laboratories without walls (collaboratories). Real-time imaging is also a growing capability in many medical centers, allowing physicians to use scientific visualization techniques to address health care problems.

Distant clinical teleconsultation is possible through a project that uses wide area communications supported by an infrastructure that incorporates local and remote storage. Standard interfaces allow multiple systems to collaborate in image acquisition and storage.

The health care industry is actively engaged in using HPCC technology. The Medical Connections program provides "jump start" funding to academic medical centers, community hospitals, and other health care organizations to connect to the Internet.



On the workstation, this prototype visual medical chart integrates multidisciplinary medical images, including diagnostic-quality digital radiographs, with scanned documents and the patient's clinical data.

Over 170 health care organizations have been connected during the past three years. A program is available through three year contracts designed to help physicians practice better medicine through use of advanced computing and networking capabilities (telemedicine). Analytic studies and the development of computer-based systems are also studied to promote the widespread transmission and use of medical records and clinical images in patient care.

A current telemedicine activity involves developing the Imaging Science and Information System (ISIS), which has the potential to detect and classify tiny precancerous formations (microcalcifications) in the breast, and making this facility available to area and regional patients via telecommunications. Other work uses the Multisurface Method (MSM) based on linear programming and multiple separating hyperplanes to analyze fine needle aspiration for diagnosing breast cancer. MSM may achieve 97% accuracy without the need for an experienced oncologist.

The computer-based Medical Literature Analysis and Retrieval System (MEDLARS) was established to achieve rapid access to the National Library of Medicine's vast store of biomedical and health-related information. MEDLARS search services are available on-line to individuals and institutions through the World Wide Web using a web browser and the Grateful Med search engine. Another project supports the development of: (1) three-dimensional interfaces to medical knowledge, including several large databases; (2) prototype vehicles for making current medical knowledge available to the general public via community networks such as freenets; and (3) prototypes for smart software agents to interact between medical knowledge servers and clinical information systems or computerized patient record systems.

The Visible Human data sets, consisting of images from a male cadaver and a female cadaver, have been completed. The data are detailed atlases of human anatomy at unprecedented resolution, created from thousands of images of a human body collected with state-of-the-art radiographic and photographic techniques. The larger, long-term goal is to transparently link visual knowledge to symbolic knowledge formats. This activity creates a virtual environment for those studying and researching human anatomy.



Education, Lifelong Learning, and Access to Information

All HPCC Organizations

Educational applications address long term national needs that augment human resource skills and effectively apply high performance computing and communications technologies to education and training. They also respond to human resource needs in computer and computational sciences at the K-12, undergraduate, graduate, and postdoctoral levels.

The HPCC Program emphasizes public access to developments in technology and archived scientific information. The HorizonNet and another system of World Wide Web servers have been established to demonstrate that affordable computer infrastructures with direct Internet access are supportable by a school consortium which can have hundreds of thousands of users. HorizonNet is implemented through low cost direct Internet connections using voice-grade analog phone lines. This activity is scalable and recognizes that most schools will only be able to afford low cost Internet connectivity solutions. It is an enabling technology to assist school districts to meet the goal of full connectivity for their schools by the year 2000.

The Math and Science Gateway and the Gateway for Educators are good examples of educational use of the Internet. The former is a World Wide Web site for secondary school students and educators that provides an easy-to-use starting point for locating science and mathematics resources and is tailored to the needs of students in grades 9 through 12 in astronomy, biology, chemistry, computing, the environment, health, mathematics, and physics. The Gateway for Educators originated as a portion of the Math and Science Gateway and contains links to information on curriculum, lesson plans, software for the classroom, and information on setting up Web servers in the schools, as well as extensive resources in specific subject areas.

The Program is challenged to make educational, vocational, and cultural materials available to all citizens regardless of age, geographic location, or ability. Activities include education and training, curriculum development, education infrastructure development, interactive and virtual training environments, and digital libraries. The use of visualization as an important educational tool has already been established — for example, special projects



The images above are from a few of the many projects in various agencies that use HPCC technology to improve the classrooms in our nation.

like the "Whole Frog" provide for high school biology students to explore the anatomy of a frog.

The HPCC Program has enabled the design and deployment of eight multimedia modules that illustrate a new paradigm in interactive tutorials. This work is conducted through the Center for the New Engineer with materials consisting of texts, graphics, demonstrations, links to select World Wide Web resources, and workbenches. Originally designed to convey HPCC research to undergraduates, these modules are also used as powerful training aids.

Faculty at several universities, along with staff at Supercomputer Centers and Government agencies, and private consultants have cooperated to develop an outline to be included in an overall curriculum plan for users of advanced air quality models. The main subject areas include basic meteorology, atmospheric chemistry, emissions, and modeling techniques. Web servers are used to develop and test training delivery concepts and prototypes that can be used by organizations that may have limited computing or programming expertise.

www.hpcc.gov/blue97/edu

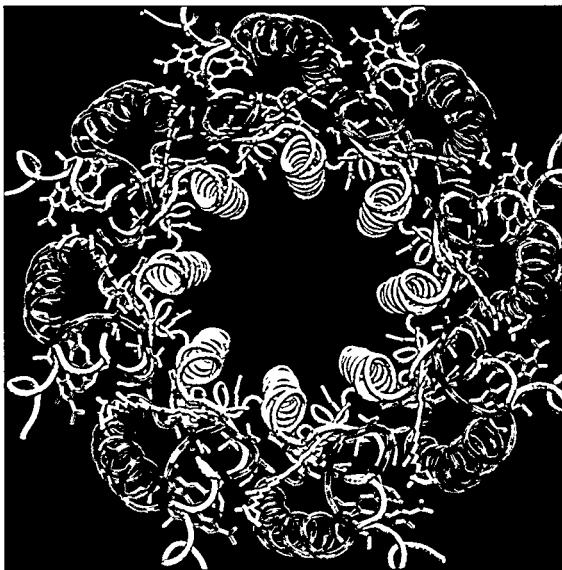


Biomedicine, Biomechanics, and Molecular Biology

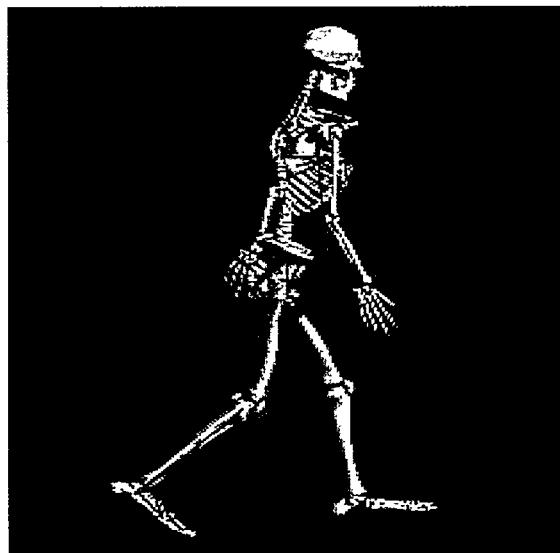
NSF, DOE, NASA, NIH, EPA

The impact of HPCC on all phases of biomedicine, ranging from basic research in molecular and cell biology to scientific and clinical visualization and modeling, has been impressive. The expanded capability of computer hardware and software means that scientific and clinical Grand Challenges not contemplated a few years ago can now be vigorously addressed. The examples summarized below represent only a small fraction of the recent accomplishments in research supported by various agencies.

Several models for protein folding have been developed that may become sufficiently robust to handle an arbitrary protein — an invaluable aid to the pharmaceutical industry in their quest for new and more disease specific drugs. A new method has been developed that combines computer modeling with x-ray scattering to determine protein structure. This approach yields an accurate structure for a light-harvesting protein. Three protein design teams are collaborating on the use of networking and real-time high performance computing to design *de-*



*Structure of the light harvesting complex II, a pigment-protein complex capturing sunlight in the photosynthetic membranes of the purple bacterium of *Rs. molischianum*. This is the first structure obtained through a combination of high performance computer modeling (to determine scattering phases) and x-ray diffraction (to determine scattering amplitudes). The approach promises to extend the applicability of x-ray scattering in structural biology and to simplify structure analysis.*



A 23-degrees-of-freedom model for simulating human motion on earth and in space. The algorithm for this model achieves nearly linear scaling on a MIMD parallel architecture.

novo proteins. Researchers are also investigating how damage to human DNA is recognized by repair enzymes and how repairs are made. Researchers have developed a three-dimensional description of processes that are templates for structure-based design of drugs capable of regulating the intricate intracellular signaling system. This work points to ways to attack human ailments ranging from birth defects to spinal cord injuries.

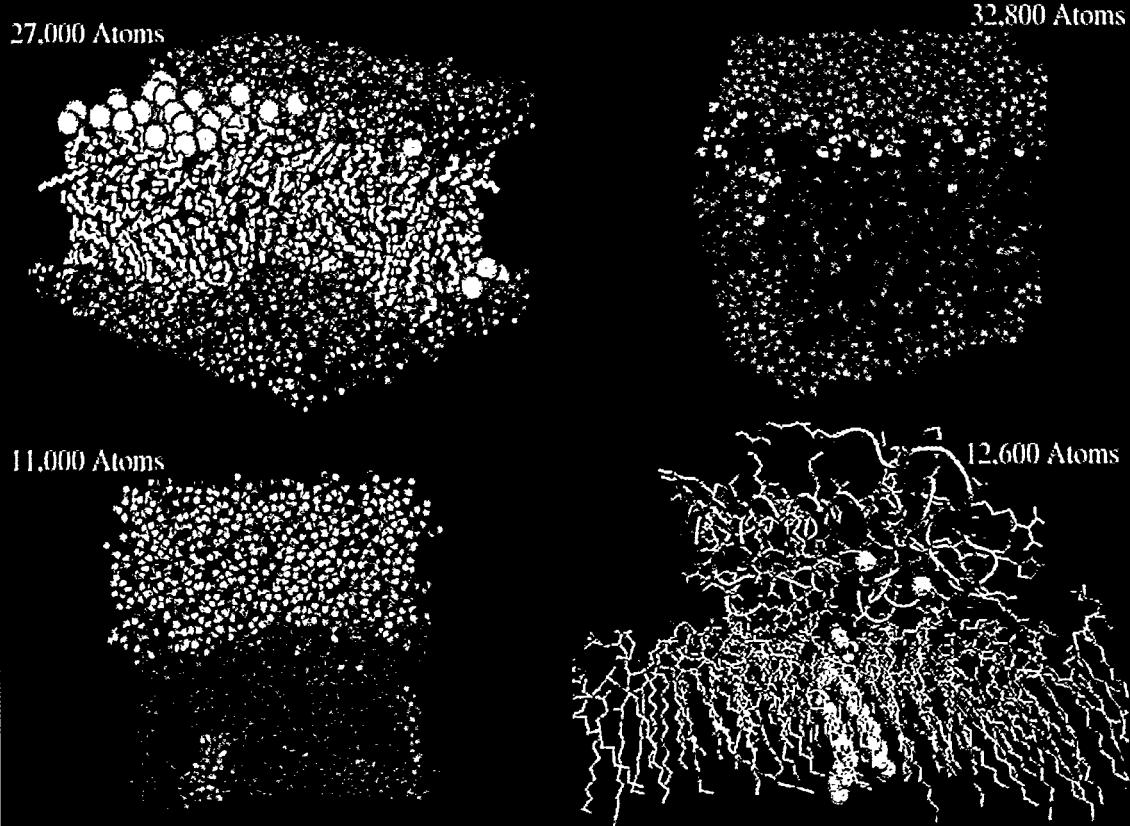
Other important high performance computing applications in structural biology include algorithm development to model and manipulate DNA, search and analyze protein sequence data, and determine energy profiles of important biological molecules such as enzymes.

In one research center, high performance computer systems that store and analyze knowledge about molecular biology, biochemistry, and genetics have been created. Specifically, parallelized search methods for sequence comparison have been developed that enable researchers to compare unknown sequences with every known sequence in under one minute, thereby facilitating the gene discovery process. This resource has over a half-million DNA sequences in its public database, which is accessed by over 20,000 different sites daily.

www.hpcc.gov/blue97/bio



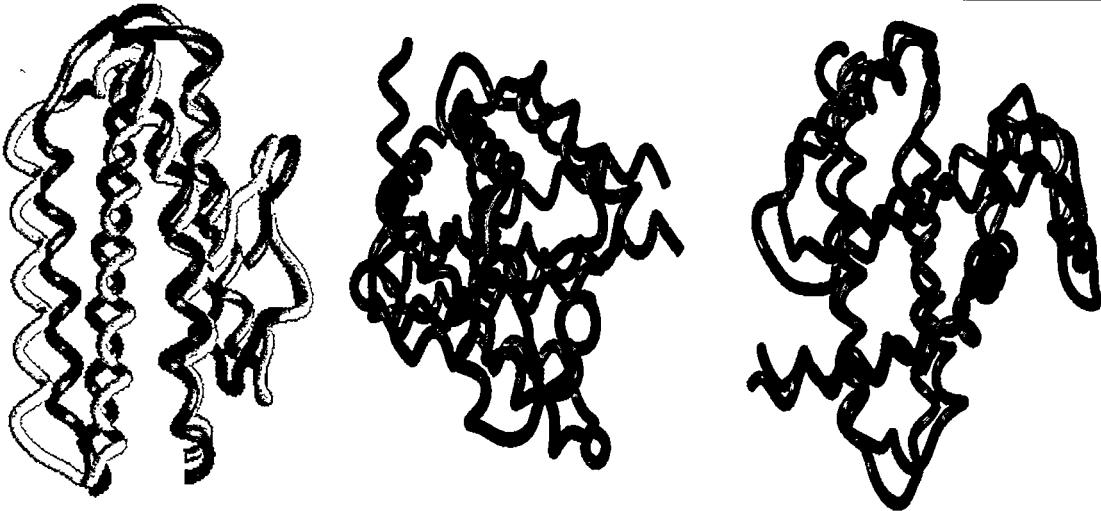
Advances in Membrane Simulations



Biological membranes can be viewed in their natural disordered state solely through computer simulations. The necessary large scale simulations, as illustrated here, have become feasible through the rapid advancement of hardware and software. A single layer of lipids grafted chemically to the surface of silicates (lower left in figure) contains about 11,000 atoms and barely can be modeled today with conventional hardware and software tools such as XPLOR. Using novel algorithms, a complete lipid bilayer with 27,000 atoms (upper left) was modeled on a 60 processor machine built in-house and required two years of runtime (1990-1992). More recently, advances in algorithm and processor technology have allowed the simulation of a 32,800-atom membrane (upper right) that was run on a workstation cluster in two months. This latter membrane model is a prerequisite to modeling the complex of one of its monolayers with the protein phospholipase A2 (lower right). These advances allow research today to contribute to the design of drugs that need to pass through membranes effectively. Proteins functioning in membranes or transporting bilayers of lipids through the bloodstream can soon be described in simulations involving 100,000 atoms. This latest accomplishment is from the Resource for Concurrent Biological Computing at the Beckman Institute, University of Illinois, with support from the National Center for Research Resources at NIH.



Progress in Predicting Protein Folding



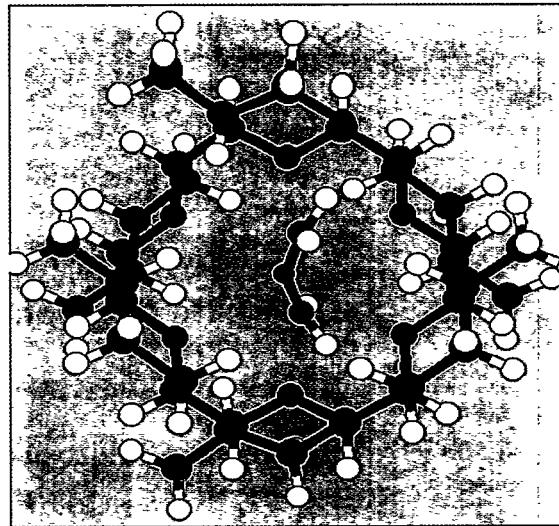
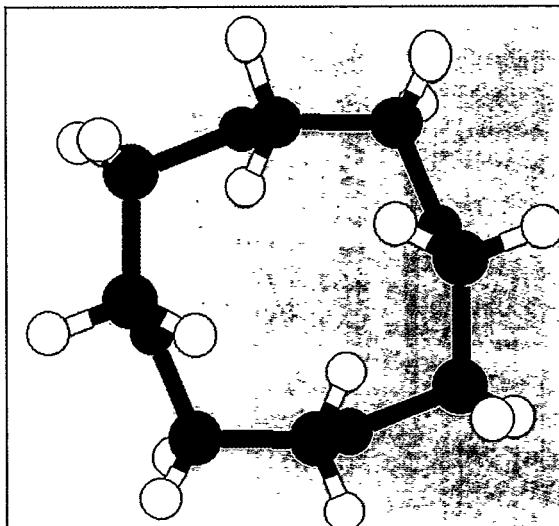
The process of a protein folding into its final (functional) structure is often classified in three stages: primary (protein amino acid sequence), secondary (helix bundles and beta sheets present in the protein), and tertiary (the final protein structure). The figures above illustrate the progress made during the last four years — enabled by HPC hardware and algorithm enhancements — in predicting tertiary structure from a specified secondary structure. The earliest effort, shown in the leftmost figure, shows good agreement (light strands are calculated and dark strands are experimental) for a relatively small four helix bundle protein, myoerythrin; this calculation took 30 minutes on an IBM RS/6000 workstation. The middle figure shows the results of calculations on a 16-node CM-5 parallel computer a few years later to predict the tertiary structure of a relatively large protein, myoglobin (red strands are calculated and blue strands are experimental). The novel algorithms used on the CM-5 were capable of examining approximately 10 billion structures, an effort that took 48 hours. Although promising, agreement between the calculated and experimental structures is not very good. The most recent simulations, shown in the right-most figure, used improved algorithms and refined potential functions to yield good agreement between calculated and experimental structures of myoglobin, especially in the core region where the helices are packed. A near term goal is to render this method capable of folding a large number of different proteins using information from existing databases.

The objective is to be able to predict protein structure directly from sequence information. In order to do this, the algorithms discussed here must be augmented by an accurate approach to secondary structure prediction. New methods to carry this out are currently being developed by the National Center for Research Resources at NIH through their support of the Center for Theoretical Simulation of Biological Systems at Columbia University.



Computational Chemistry

NSF, NASA, DOE, NIH, NSA, NIST, EPA



(Left) A view of 12-crown-4 ether, a model for binding alkali ions for waste separation;

(Right) A view of a more selective reagent, 18-crown-6 derivative, with a sodium ion and two waters of solvation.

Computational complexity effectively grows greater than the cube of the number of atoms.

Computational chemistry plays a critical role in developing molecular-level descriptions of physical, chemical, and biological processes in natural and industrial systems.

Grand Challenge applications in computational chemistry seek to develop new algorithms, software, and diagnostics to overcome limitations imposed by shortcomings in current computational chemistry software. The objective is to enable the use of high performance parallel computing platforms to solve problems in computational chemistry.

One of the most challenging problems ever undertaken in theoretical and computational chemistry is the modeling of the effect of anthropogenic (man-made) materials on the environment. To provide the required accuracy in a reasonable amount of time and to solve complex environmental impact problems, it is necessary to use the most advanced massively parallel computer systems. However, the current generation of computational chemistry software, often well optimized for serial scalar and/or vector supercomputers, is constrained by algorithmic, performance, memory, and I/O limitations on parallel computers. These problems are being addressed by a focused program of research and development in three related areas: (1) new algorithms that both scale to larger numbers of processors and use much less memory than traditional

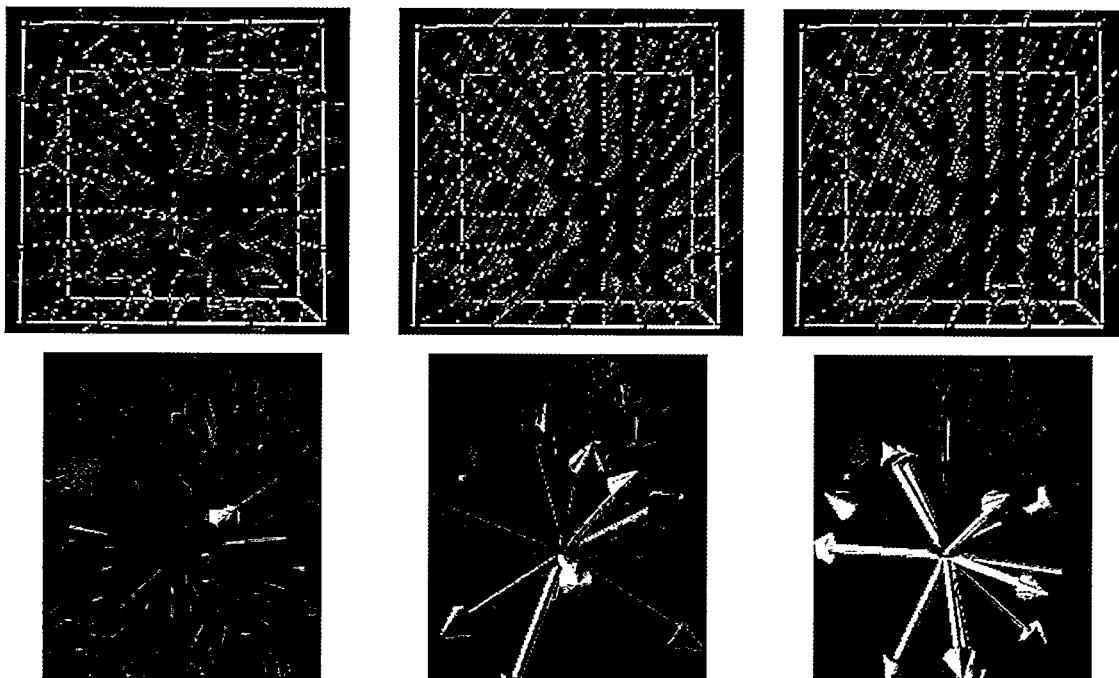
methods, (2) prototype codes to test these algorithms, and (3) software technology adaptations that will permit the efficient integration of these new algorithms into major existing computational chemistry applications. Examples of these new approaches to software problems include the environmental degradation of chlorinated hydrocarbons by natural and artificial processes, and characterization and processing of nuclear waste.

The tools designed to harness the computational power indirectly benefit other scientific areas such as the design of polymers, composite materials, and drugs. For example, one such effort brings together computer scientists, mathematicians, computational scientists, and theoretical chemists from several National Laboratories. An important component of the project is the participation of the pharmaceutical and chemical industries, thereby ensuring that the focus remains on real environmental and industrial problems including the minimization of waste streams and the optimal use of energy in industrial processes.



Materials Sciences

NSF, NASA, DOE, NIH, NSA, NIST, EPA



The illustrations show results of a calculation of the ground state ($T=0K$) magnetic structure of a large cell (256-atom) model of Fe65Ni35. In order not to prejudice the outcome, the magnetic moments associated with the Fe (big moment/big arrows) and Ni (small moment/small arrows) sites are initialized to point in random directions. As the calculation proceeds (left to right) most of the moments align ferromagnetically. However, some Fe moments align anti-ferromagnetically (anti-parallel) and some remain non-collinear. This is particularly apparent in the "rose" plots (lower frames) where the moments associated with all the Fe(Ni)-sites are projected onto a common Fe(Ni)-origin.

A project involving cooperative research among scientists at several National Laboratories uses high performance computing platforms such as the Intel Paragon XP/S 150. A massively parallel code for studying magnetism in metals and alloys has been developed based on the formulation of *ab initio* spin dynamics and a scalable method for performing large scale first-principles electronic structure calculations of solids (100-1000 atoms). The new method has the potential to allow investigation of a wide range of magnetic properties previously not accessible to *ab initio* theoretical study. Preliminary investigations of the ground state magnetic structure of large cell models (256 atoms) of disordered Fe65Ni35 alloys (the classic Invar system) are pointing to unusual magnetic behavior (see caption above) involving non-collinear arrangements of the local magnetic moments associated with Fe-rich clusters within the otherwise disordered alloy.

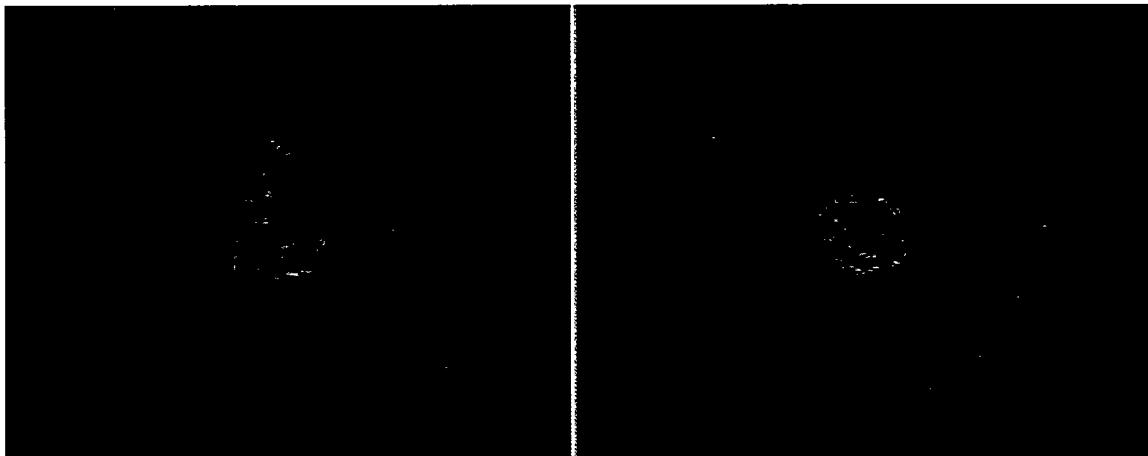
Magnetic materials represent a multi billion dollar industry. Despite this great driving force, a micro-

scopic understanding of metallic magnetism has proved to be elusive. Although much progress has been made in understanding magnetism in the 3-D transition metals using first-principles quantum mechanical methods based on the local spin density approximation, many significant scientific and technologically important problems remain unsolved. These range from a first-principles theory of the magnetic phase transition, to magnetism in inhomogeneous materials, to a microscopic theory of technologically important extrinsic properties such as permeability, coercivity and remanence. The methods that are being developed, coupled with the availability of massively parallel computers, will provide the necessary tools to perform realistic simulations on many of these problems, and have the potential to enable significant advances in magnetic materials research and development.



Computational Physics

NSF, NASA, DOE, NIH, NSA, NIST, EPA



Before-and-after frames from a 100-million-object simulation of a disk galaxy interacting with a smaller galaxy. The smaller galaxy passes through the disk of the larger galaxy, causing ring structures to form. The simulation is consistent with the Cold Dark Matter model of the universe and compatible with astronomical observations.

The field of computational physics has exploited HPCC technologies leading to new science, including new computational models in astronomy and astrophysics, charged plasma, and elementary particle physics.

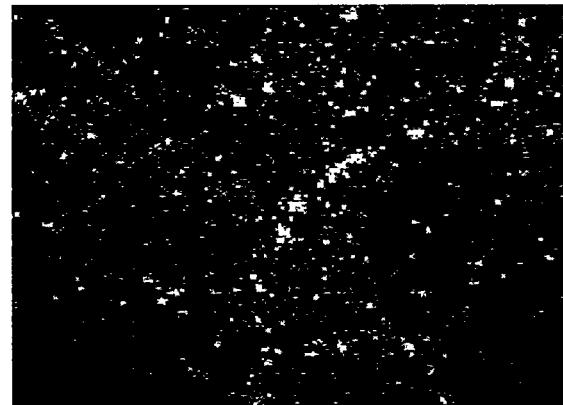
Astronomy and Astrophysics

Astrophysical models of “black holes,” galaxy formation, and solar dynamics help astronomers and astrophysicists better understand the forces and mechanisms in our Universe. One of the open questions in present day astronomy is the existence and behavior of black holes — astronomical objects so dense that not even light can escape — which are only inferred by indirect observations at this time. Because the equations governing the dynamics of black holes are relativistic, an understanding of a collision of two black holes requires simulations that were too difficult for the computing technology before recent advances made in the HPCC Program. A group of astrophysicists have formed the Binary Black Hole Alliance and have used HPCC technology to enable new ways for computational scientists and astrophysicists (distributed over five universities) to collaborate in modeling the collision of two black holes. The Alliance’s calculations led to the discovery of unsuspected qualitative features of the collision and to simple analytic models that reveal the underlying physics.

Advances in computational sciences are being used to better understand and simulate large scale galaxy formation and accretion astrophysics. Using information on the universe’s power spectrum from the

Cosmic Background Explorer and three cosmological models, scientists have simulated galaxy interactions using up to 46 million objects to illustrate their understanding that the Galactic Harassment model for the universe drives galactic morphological evolution. These simulations have been verified with observations from the Hubble Space Telescope.

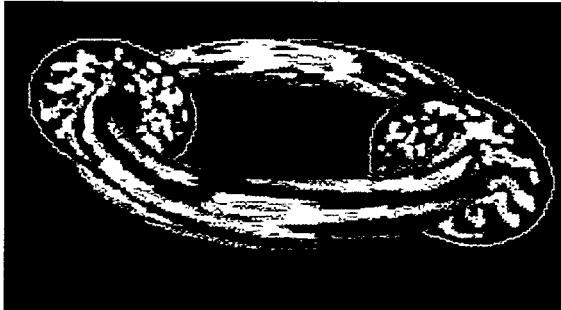
In other computational science efforts, a suite of three Computational Fluid Dynamics algorithms for massively parallel processor architectures is providing greater understanding of the mechanisms that control the behavior of the solar heliosphere. This research has already led to the discovery of the phe-



The simulation complements Hubble Space Telescope (HST) imaging of the formation of the central regions of galaxies. Following their evolution in clusters to the present epoch, this model ties together a vast range of data to create a unified model of structure formation and galaxy evolution.



nomenon of two orthogonal magnetic flux tubes passing through one another and then reconnecting.



A three-dimensional non-linear gyrokinetic plasma simulation showing turbulent fluctuations extending along the twisting magnetic field lines in a toroidal direction.

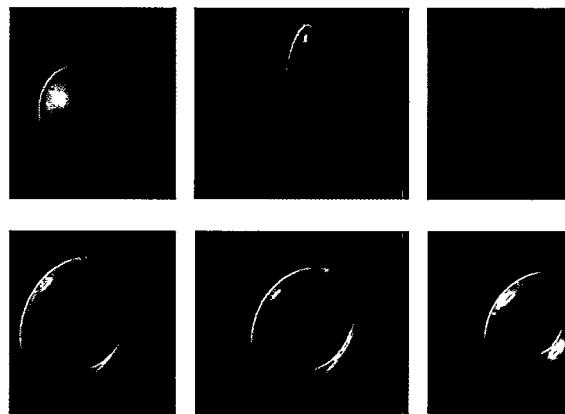
Charged Plasma Simulations

The knowledge of turbulent plasma transport in magnetically confined plasma fusion devices called tokamaks has significantly increased over the past four years through the use of HPCC technology in combination with advances in plasma theory. The most elusive problem of tokamak design and operation has been the anomalous loss of plasma particles and energy. Experiments observe turbulent fluctuations driven by several modes of oscillation and losses that always exceed the rates calculated for a quiescent plasma. These fluctuations enhance the transport of heat in a plasma in much the same way that fluid turbulence in ocean waves or eddies in the atmosphere enhance heat transport rates. The Numerical Tokamak Project is a multidisciplinary effort, involving several institutions, with the goal of realistic simulations of tokamak plasma turbulence needed to optimize performance of fusion devices. The near-term HPCC work concentrates on developing a predictive numerical description of the core plasma transport in tokamaks driven by low frequency collective fluctuations. The combination of emerging massively parallel processing hardware and improvements in algorithms is resulting in an estimated performance increase of 10^2 to 10^6 for the simulations.

Quantum Chromodynamics (QCD)

QCD is the fundamental theory that describes the interactions between quarks and gluons, the underlying constituents of protons, neutrons, and other so-called elementary particles. Lattice QCD is a discrete computational technique for determining the consequences of QCD theory from first principles. In the

HPCC Grand Challenge program, Lattice QCD simulations are being used by two U.S. collaborations to calculate the spectrum decay rates of hadrons by probing their quark and gluon wave functions. Lattice QCD simulations are used to determine the parameter that determines the decay of particles called B-mesons. B-mesons are the subject of intense study at experimental particle physics laboratories throughout the world and are the main focus of the B-Factory under construction in California.



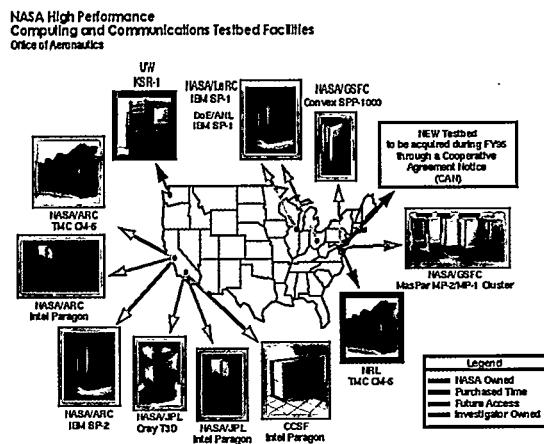
As a function moves along an interpolating path, the spacetime of the superstring theory undergoes a drastic topological jump.

Superstrings (SS)

Recent advances in the SS theory have shown how troubling singularities disappear when certain black holes, already known to occur as composite objects, are incorporated as fundamental particles. This result has the striking implication that SS theory ground states, which had been thought to be disjointed, are seen to fit together into a connected web with SS physics smoothly interpolating from one component to another. As one moves along an interpolating path, the spacetime of the SS theory undergoes a drastic topological change. The striking property is that the SS theory is perfectly smooth even at the places at which topological characteristic numbers jump.



III. HPCC Research Centers and Facilities



Computing, Information, and Communications Research Facilities

These facilities provide capabilities to (1) evaluate early prototype systems and provide valuable feedback to developers, (2) integrate visualization and virtual reality systems into existing high performance systems, (3) run full scale applications, such as Grand Challenges, on systems not otherwise available, and (4) develop parallel software using scaled down systems. In addition, these facilities provide access to innovations in network connectivity that allow large scale applications to run over remote connections to systems at HPCC facilities across the country, thus demonstrating future directions for advanced R&D in academia, industry, and the Federal government.

Researchers at these Centers rely on many enabling technologies — high speed networks, supercomputers, parallel architectures, massive data stores, virtual reality display devices — in order to succeed. Many groups in addition to the researchers contribute to this success, among them facility staff, hardware and software vendors, and industrial affiliates. HPCC funding is thus leveraged heavily through equipment and personnel from vendors, discipline-specific agency funds, as well as state and local funds, and industrial affiliate contributions. Industrial affiliation offers a low risk environment for exploring and ultimately exploiting HPCC technologies.

Applications software developers from the Centers access their resources over the Internet. The wide range of hardware and applications software that is available also makes these Centers ideal sites for benchmarking systems and applications and for providing feedback to hardware and software vendors.

All facilities provide extensive K-12 and undergraduate educational opportunities as well as training for researchers, graduate students, and faculty; they also provide for publication of articles in professional journals, annual reports, and newsletters.

The systems listed below are funded by the named agency and receive additional funding from other HPCC agencies. For example, funding for systems at NSF centers also comes from DARPA, NASA, and NIH.

NSF Supercomputer Centers

NSF funds four Supercomputer Centers, augments the computing facilities at NCAR (the National Center for Atmospheric Research), and funds Metacenter activities. The term Metacenter refers to the joint cooperative activities of these centers and others in naturally overlapping research and technology areas. A Metacenter facilitates collaboration, communication, technical progress, and interoperability among participating institutions. Metacenter Regional Alliances (MRAs), which were set up to augment national support activities, are also intended to complement, expand, and strengthen existing Metacenter activities at the regional, state, or local level.

Cornell Theory Center (CTC), Ithaca, NY

The primary computing resources at CTC are a 512 processor IBM SP-2, an SGI Power Onyx Array, consisting of two 8 processor Power Onyx systems, and a 16 processor SGI Power Challenge. One CTC focus area is a globally scalable computing environment, including mass storage, I/O capability, networking, archival storage, data processing, and graphics power.

National Center for Supercomputing Applications (NCSA), Urbana-Champaign, IL

Resources include:

- Convex C-3880 with 8 processors and 4 GB memory
- Convex Exemplar with 64 processors and 8 GB memory
- Silicon Graphics Power Challenge with 16 processors and 4 GB memory
- Thinking Machines CM-5 with 512 compute nodes and 16 GB memory

The three-tiered NCSA network consists of (1) Ethernet or FDDI to the desktop; (2) FDDI backbone between buildings, high end systems, and the Internet; and (3) HiPPI between high performance computing systems, mass storage, and high end peripherals.

NCSA is also involved in ATM research in (1) a local area network, (2) a trans-continental 155 Mb/s (SONET OC-3) national network, and (3) the BLANCA gigabit testbed at 622 Mb/s (SONET OC-12).

Pittsburgh Supercomputer Center (PSC), Pittsburgh, PA

Resources include:

- The first single-vendor heterogeneous system consisting of a Cray Research T3D (with 512 processors, each with



64 MB of memory) coupled to a C90 (with 16 processors and 4 GB of memory)

- 14-processor DEC Alpha workstation cluster
- HiPPI and FDDI network connecting these resources to each other and to storage devices

San Diego Supercomputer Center (SDSC), San Diego, CA

Resources include:

- Cray Research C90 with 8 vector processors
- Intel Paragon with 400 processors
- Thinking Machines CM-2 with 8,192 processors
- Eight workstation DEC Alpha cluster

National Center for Atmospheric Research (NCAR), Boulder, CO

NSF HPCC funds provided partial support for the acquisition of a 64 processor Cray Research T3D and an 8 processor IBM SP-1 for use in the global climate modeling Grand Challenge.

NSF Science and Technology Centers

Each of these four Centers addresses a particular research area. Common to all four is cross-disciplinary focus, knowledge transfer, links to the private sector, and education and outreach. The Centers are:

The Center for Research in Parallel Computation (CRPC) at Rice University

CRPC's aim is to make parallel computing systems as easy to use as conventional computing systems — efforts include HPF, PVM, MPI, and NHSE, HPC++, algorithms for physical simulation, algorithms using parallel optimization, and ScaLAPACK.

The Center for Computer Graphics and Scientific Visualization at the University of Utah

This Center is building and displaying models that are visually and measurably indistinguishable from real world entities.

The Center for Discrete Mathematics and Theoretical Computer Science at Rutgers University

This Center is applying discrete mathematics and theoretical computer science to solving fundamental problems in science and engineering.

The Center for Cognitive Science at the University of Pennsylvania

This Center studies the human mind through the interaction of disciplines such as psychology, philosophy, linguistics, logic, and computer science. Work in human cognition, perception, natural language processing, and parallel computing has applications in robotic and manufacturing systems,

human machine interfaces, and language teaching and translational tools.

NASA Testbeds

NASA maintains testbeds throughout the country to offer diversity in configuration and capability. The testbeds are:

Ames Research Center, Moffett Field, CA

Resources include:

- IBM SP-2 with 160 processors and 20 GB memory
- Intel Paragon with 208 compute nodes, 16 service nodes, and 7 GB memory
- Thinking Machines CM-5 with 128 compute nodes — each consisting of a SPARC processor and 4 vector processors and 4 GB memory, which is also used by the Naval Research Laboratory

Goddard Space Flight Center, Greenbelt, MD

Resources include:

- Convex SPP-1 with 8 processors
- MasPar MP-2 with 16,384 processors and 1 GB memory, SIMD (Single Instruction Multiple Data); MP-1 with 8,192 processors and 512 MB memory; 2 MP-1's with 4,096 processors and 256 MB memory. These four systems have all been connected by a 4-by-4 HiPPI switch. MasPar applications are being modified to take advantage of the HiPPI network in order to distribute the workload across the four MasPar systems. This cluster demonstrates combining SIMD and MIMD (Multiple Instruction Multiple Data) programming styles to enable MasPar to move beyond its current 16,384 processor ceiling.

Jet Propulsion Laboratory (JPL), Pasadena, CA

Resources include:

- Cray Research T3D with 256 processors and 16 GB memory,
- Intel Delta with 528 processors at Caltech
- Intel Paragon with 56 compute nodes, and a total of 1.8 GB of memory

Langley Research Center, Langley, VA

Resources include an Intel Paragon with 72 compute nodes.

Lewis Research Center, Cleveland, OH

Resources include an IBM SP-1 with 16 processors.

Other Resources

These include:

- IBM RS6000 workstation cluster with 32 nodes and 3 GB memory, that is part of a 128 node IBM SP-1 consortium at Argonne National Laboratory



- Kendall Square Research KSR-1 with 56 nodes and 1.8 GB memory, at the University of Washington

DOE Laboratories

National Energy Research Supercomputer Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA

The Supercomputer Access Program at NERSC provides production computing for investigators supported by the Office of Energy Research in the following areas: material sciences, chemistry, geosciences, biosciences, engineering, health and environmental research, high energy and nuclear physics, fusion energy, and applied mathematics and computational science. NERSC resources include:

- Cray Computer Cray-2 with 8 processors and 128 mega-words (millions of 64-bit words (Mw) memory)
- Cray Computer Cray-2 with 4 processors and 128 Mw memory
- Cray Research C90 with 16 processors and 256 Mw memory
- The National Education Supercomputer, a four processor Cray Research X-MP EL provided by Cray Research and available to high schools over the Internet

Los Alamos National Laboratory (LANL), NM, and Oak Ridge National Laboratory (ORNL), TN

These DOE HPCC Research Centers provide full scale high performance computing systems for work on Grand Challenge applications and use in scalability studies. These applications require large prototype systems — they cannot be scaled down without removing essential aspects of their physics.

LANL operates a Thinking Machines CM-5 with 1,024 compute nodes and 32 GB memory.

ORNL resources include:

- Intel Paragon XP/S 150 with 1,024 MP nodes (3,072 processors) and a total of 70 GB of memory
- Intel Paragon XP/S 35 with 512 GP nodes (1024 processors) and a total of 16 GB of memory
- Intel Paragon XP/S 14 with 96 MP nodes and a total of 8 GB of memory
- Kendall Square KSR 1 with 64 processors and a total of 2 GB of memory, used to study shared memory algorithms
- nCube2 with 8 processors and a total of 256 MB of memory, used by high school students in the Adventures in Supercomputing program

NIH Systems

The Division of Computer Research and Technology (DCRT) has a 56 processor IBM SP-2 and a 128 processor Intel

iPSC/860. Both systems are used by NIH staff in biomedical applications.

The National Cancer Institute's (NCI) Frederick Biomedical Supercomputing Center has an 8 processor Cray Y-MP and a MasPar MP-2 with 4,096 processors along with a comprehensive collection of biomedical software available to all scientists who use the facility.

The National Center for Research Resources (NCRR) supports various systems for biomedical research applications at its six High Performance Computing Resources Centers:

- Resource for Concurrent Biological Computing, Beckman Institute, University of Illinois;
- Supercomputing for Biomedical Research, Pittsburgh Supercomputing Center;
- Theoretical Simulation of Biological Systems, Columbia University;
- Parallel Computing Resource for Structural Biology, University of North Carolina, Chapel Hill;
- Biomedical Computation Resource, University of California, San Diego;
- Parallel Processing Resource for Biomedical Scientists, Cornell Theory Center, Cornell University;

and two Scientific Visualization Resource Centers:

- Interactive Graphics for Molecular Studies, University of North Carolina, Chapel Hill;
- Special Research Resource for Biomolecular Graphics, University of California, San Francisco.

NOAA Laboratories

The Forecast Systems Laboratory in Boulder, CO, has a 221 processor Intel Paragon, with 6.5 GB memory. This system is used to parallelize regional and mesoscale forecast models.

The Geophysical Fluid Dynamics Laboratory (GFDL) in Princeton, NJ, has acquired a Cray T-90 (PVP) and a T3E (SPP) and the National Centers for Environmental Prediction in Camp Springs, MD, will acquire a scalable computing system in FY 1997. These systems are used for the global climate modeling and weather forecasting Grand Challenges.

EPA Systems

EPA's National Environmental Supercomputing Center in Bay City, MI, has a Cray C-94 with 3 processors and 64 Mw memory, and a Cray T3D with 128 processing elements and 8 GB of memory. These systems are dedicated to environmental research, problem solving and related educational programs.

IV. CIC R&D Program Organization

The Computing, Information, and Communications R&D (CIC R&D) programs in FY 1997 encompass twelve Federal departments and agencies. The CIC R&D Subcommittee (formerly the High Performance Computing, Communications, and Information Technology (HPCCIT) Subcommittee) consists of representatives from each of the participating organizations. The CIC R&D Subcommittee reports to the Committee on Computing, Information, and Communications (CCIC). The CCIC is one of the nine National Science and Technology Council (NSTC) Committees.

The National Coordination Office

The National Coordination Office for Computing, Information, and Communications (CIC) R&D (formerly High Performance Computing and Communications) coordinates the activities of the twelve participating Federal departments and agencies through the CIC R&D Subcommittee of the CCIC. The NCO serves as a liaison to the U.S. Congress, state and local governments, international initiatives, industry, universities, and the public. The Director of the NCO reports jointly to the Director of the Office of Science and Technology Policy (OSTP), Executive Office of the President, and to the Chair of the CCIC.

Interaction with Government, Industry, and Academia

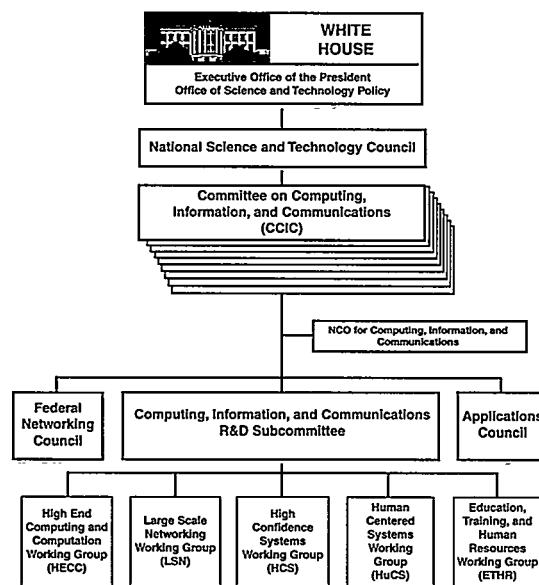
The agencies that participate in the CIC R&D programs work closely with other government agencies, academia, and industry in developing, supporting, and using the technologies developed through the programs. The NCO, the CIC R&D Subcommittee, and the participating agencies meet with representatives from the U.S. Congress, Federal, state and local organizations, industry, academia, professional societies, foreign governments, and others, to exchange information about technical and programmatic needs, issues, and trends.

The Program has a long history of successful government-industry interaction. In October 1995, for example, the NCO and the HPCCIT Subcommittee sponsored non-disclosure briefings for agency representatives by 12 high performance computing systems vendors. These vendors discussed hardware and software trends, R&D needs, standards, international issues, and the HPCC Program. Other examples include sponsorship of the PetaFLOPS Architecture Workshop, held April 22-25, 1996, in Oxnard, CA and the PetaSoft '96 Workshop, held June 1996 in Bodega Bay, CA. These workshops address technical approaches to meet future architecture and software needs of the most demanding applications.

Program Coordination

The CIC R&D Subcommittee and its Executive Committee coordinate program planning, budgeting, implementation, and review. Monthly meetings include information exchanges, development of multi-organization programs, and the review of the plans and budgets of participating organizations. Over the past year, six Task Forces participated in a planning process that contributed to the development of the Program Component Areas (PCAs) (described on pages 32-33). The Subcommittee has established Working Groups, one for each PCA, that will succeed and expand upon the activities of the HPCCIT Working Groups.

The CIC R&D Subcommittee is responsible for Federal networking R&D. The CIC R&D program and the Federal Networking Council (FNC), which





is chartered by NSF, work together to establish an effective interagency forum and long-term strategy to oversee the operation and evolution of the Federally funded portion of the Internet in support of science, research, and education. FNC members represent Federal agencies that need to operate and use increasingly advanced networking facilities, mainly for research and education, but also for administrative functions.

The CCIC has established an Applications Council to promote multi-agency leadership in the early application of advanced computing, information, and communications technologies, with special focus on projects that are widely applicable to Federal agency missions. Additionally, the Council will encourage pilot projects to assess technologies needed by these applications and will support the Administration's broad international goals to eliminate applications barriers. With member agencies throughout the Federal government, the Council provides a mechanism for those agencies to communicate with the organizations that participate in the CIC R&D programs. Council members provide user feedback and identify the research needed to enable their most demanding applications, while they are kept informed about the CIC R&D agenda, priorities, and results.

Program Evaluation

Federal and Federally chartered organizations, industrial and academic organizations, and professional societies provide critical analyses of the Program through conferences, workshops, and reports. These efforts strengthen program planning and management and help make program goals and accomplishments better understood.

As part of continuing internal program assessments during the past two years, the CIC R&D Subcommittee has conducted intensive evaluations of the changing role of information technology in Federal agency missions. This effort has led to the broader agenda with focused R&D programs that is now being coordinated by the Subcommittee and implemented across the participating Federal agencies and departments.

Buy American Report

The Congress has requested that it be informed about certain funding of non-U.S. high performance computing and communications activities.

In FY 1996, DARPA is the only HPCC agency that has entered into a grant, contract, cooperative agreement, or cooperative research and development agreement, for HPCC with either (1) a company other than a company that is either incorporated or located in the U.S., and that has majority ownership by individuals who are citizens of the U.S., or (2) an educational institution or nonprofit institution located outside the United States. That activity is with University College, London, for approximately \$120,000 per year. It is part of an effort developing the underlying technology to be used in pilot demonstrations of multi-country, multi-way, wide area multimedia services.

In FY 1996 no HPCC R&D procurement exceeds \$1 million for unmanufactured articles, materials, or supplies mined or produced outside the U.S., or for manufactured articles, materials, or supplies other than those manufactured in the U.S. substantially all from articles, materials, or supplies mined, produced, or manufactured in the U.S.

Reports about the CIC R&D Programs

The NCO provides electronic, print, and video materials to hundreds of media representatives in the U.S. and abroad and responds to thousands of requests for information from Congressional offices, industry, academia, and the public. The NCO has both Web and Gopher servers that contain the FY 1994, FY 1995, and FY 1996 annual reports, information about sources of funding for CIC R&D, links to the servers of participating agencies, and related information. In November 1995, users of over 16,000 computing systems from 70 countries accessed the Web server, and nearly 6,000 accessed the Gopher server.



V. Future Directions

The HPCC Program has been a driving force behind U.S. leadership in advanced computing, communications, and information technology for the past five years. The Program is often cited as a model for successful Federal multi-agency planning and program coordination. As a result of the successes of the Program and the changing role of information technology in agency mission applications, broader collaborative R&D investments in computing, information, and communications are needed. This effort has led the Computing, Information, and Communications (CIC) R&D Subcommittee to define a broader agenda with focused R&D programs. The CIC R&D programs reflect long term efforts in advanced computing, information, and communications technologies that build on the strong foundation of the Federal HPCC Program. The CIC R&D programs are organized into five Program Component Areas (PCAs), which evolved from the five original HPCC Program components (HPCS, NREN, ASTA, IITA, and BRHR, described in the Glossary) and will be used to guide R&D investments by the participating Federal agencies into the 21st century.

Federal agency mission applications are the driving force behind the development of this R&D agenda. Mission-driven applications include computation intensive science and engineering applications known as Grand Challenges and information intensive applications known as National Challenges.

The PCAs and CIC R&D planned activities for FY 1997 are presented below.

Program Component Areas

High End Computing and Computation (HECC)

HECC R&D is focused on continued U.S. leadership in high performance computing and computation. Investments concentrate on leading-edge innovations in hardware and software such as storage and data technologies for high end computing systems, experimentation with novel devices, development of system software technologies, advanced simulation techniques, and fast, efficient algorithms for simulation and modeling. In addition, HECC research supports exploration of advanced computing concepts in quantum, biological, and optical computing at both the hardware and software levels. At the high end, these technologies enable distributed, multidisciplinary, computation intensive, scientific and engineering applications. Scalable systems allow effective deployment of these technologies to the workplace, school, and home.

Large Scale Networking (LSN)

LSN R&D will assure U.S. technological leadership in communications through R&D that advances the leading edge of networking technologies and services. This includes advanced network components and technologies for engineering and management of large scale networks, both for scientific and engineering R&D and for other purposes. Areas of particular focus include (1) technologies and services that enable wireless, optical, mobile, and wireline communications; (2) networking software that enables information to be disseminated to individuals, multicast to select groups, or broadcast to an entire network; (3) software for efficient development and execution of scalable distributed applications; (4) software components for distributed applications, such as electronic commerce, digital libraries, and health care; and (5) infrastructure support and testbeds.

High Confidence Systems (HCS)

HCS R&D will provide users with high levels of security, protection, reliability, and restorability of information services. Such systems are resistant to system failure and malicious penetration or damage and readily respond to interference by adaptation or recovery. These systems include both physical components, wired and wireless technologies, the data they contain and transmit, and the software that manipulates these data. HCS R&D focuses on (1) system reliability (such as management of networks under load, failure, or intrusion; emergency response; firewalls; secure enclaves; and formal methods), (2) security and privacy (including personal identification, access control, authentication, encryption and other privacy assurance techniques, public key infrastructures, and trusted agents for secure distributed computing), and (3) testing and evaluation. Key applications include national security, law enforcement, life- and safety-critical requirements, personal privacy, and protection of critical elements of the National Information Infrastructure.

Human Centered Systems (HuCS)

HuCS R&D makes computing systems and communications networks more easily accessible to and useable by a wide range of user communities. These communities include scientists and engineers, educators and students, the workforce, and the general public. Technologies enabling such systems include: (1) "knowledge repositories" and "information agents" for managing, analyzing, and presenting massive amounts of multimedia and multi-source information; (2) "collaboratories" that provide access to knowledge repositories and that facilitate knowledge sharing, group authorship, and control of remote instruments; (3) systems that enable multi-modal human system interactions including speech, touch, and gesture recognition and synthesis; and (4) virtual reality environments and their application to fields including scientific research, health care, manufacturing, and training.

Education, Training, and Human Resources (ETHR)

The focus of ETHR R&D is on education and training technologies. The goals of this education and training are to produce (1) researchers and students in high performance computing, communications, information technologies, and their application, and (2) a citizenry with the skills to compete and prosper in the 21st century's information age. ETHR includes curriculum development, fellowships, and scholarships for computational, computer, and information sciences, and engineering. It includes the application of interdisciplinary research to learning technologies, and R&D in information-based learning tools, lifelong learning, and distance learning for people in remote locations. This PCA collaborates closely with the Committee on Education and Training (CET) on developing its programs and research agenda.

Planned FY 1997 CIC R&D Activities by Program Component Area

High End Computing and Computation	Large Scale Networking	High Confidence Systems	Human Centered Systems	Education, Training, & Human Resources
Scalable systems, software, and environments	Multi-gigabit network testbeds including wireline, wireless, and mobile technologies	Infrastructure protocols for secure and reliable networks	Intelligent systems and software	Integration of research and education technologies
Uniform memory access systems with low latency and high bandwidth	Very high speed backbones for network, scientific, and engineering applications	High speed cryptography for information security in virtual laboratories	Knowledge acquisition, fusion, aggregation, and summarization tools	Lifelong and distance learning technologies
Integrated circuit simulation for microsystems	Network reconfiguration tools	Interoperability standards	Virtual environments with remote instruments	Teacher development tools
Distributed operating systems	Global grid communications and inter-network protocols	Secure interface protocols	Multi-modal communication, speech understanding	K-12 education cooperative agreements
Compilers, debuggers, program development environments	Network management, cost accounting and authentication	Network security and restorability	Collaboratories, tools for group authoring, remote control of instruments	Summer school in high performance computation
Algorithms and software for Grand Challenge scientific and engineering applications	Network-centric computation for National Challenge applications	Reliability and security for mobile computing environments	Graphical user interfaces for medical imaging and patient record systems	Research fellowships for science, engineering, and health professionals
Supercomputer centers and centers for advanced research on experimental systems	Data mining, integration, and visualization for scientific data	Telemedicine testbed networks, remote surgical procedures	Clinical decision support systems, patient records, visible human database	Health care information infrastructure
Advanced manufacturing design and processes	Optoelectronic and optical media	Authentication and verification procedures	Remote visualization of environmental data in digital libraries	Technologies for ubiquitous access to the World Wide Web
Superconducting components	Convergence of computing and communications			Automatic document translation
Quantum and biological computing methods				

Agency HPCC Budgets by Program Component Area

FY 1996 Budget (Dollars in Millions)

<u>Agency</u>	<u>HECC</u>	<u>LSN</u>	<u>HCS</u>	<u>HuCS</u>	<u>ETHR</u>	<u>TOTAL</u>
DARPA	77.96	96.04	10.00	112.17	6.24	302.41
NSF	140.32	104.47		15.28	31.03	291.10
DOE	84.49	12.64		8.56	4.00	109.69
NASA	75.55	27.45			23.60	126.60
NIH	22.40	21.51	4.38	24.34	7.11	79.74
NSA	29.48	3.00	7.40		0.15	40.03
NIST	5.59	2.20	5.36	10.36		23.51
ED				11.40	17.53	28.93
VA	3.00	14.13	2.90	1.90		21.93
NOAA	3.30	2.70		0.50		6.50
EPA	8.70			0.60	0.08	9.38
AHCPR				3.20		3.20
TOTAL	450.79	284.14	30.04	188.31	89.74	1,043.02**

FY 1997 Budget Request (Dollars in Millions)

<u>Agency</u>	<u>HECC</u>	<u>LSN</u>	<u>HCS</u>	<u>HuCS</u>	<u>ETHR</u>	<u>TOTAL</u>
DARPA	75.05	115.93	10.00	117.03	7.51	325.52
NSF	129.17	72.26	1.21	57.76	19.11	279.51*
DOE	93.29	14.79		12.98	3.50	124.56
NASA	70.70	20.20	1.00	6.40	10.70	109.00
NIH	23.40	22.90	4.94	29.59	5.38	86.21
NSA	25.93	3.50	7.30			36.73
NIST	5.79	2.96	4.40	10.36		23.51
ED				11.40	6.61	18.01
VA	1.00	9.45	2.30	1.80		14.55
NOAA	6.30	2.70		0.50		9.50
EPA	6.58			0.60		7.18
AHCPR				4.20		4.20
TOTAL	437.21	264.69	31.15	252.62	52.81	1,038.48**

* Following the National Research Council and NSF Advisory Committee recommendations, the NSF FY 1997 allocation for HPCC excludes \$29.60 million for applications in cutting edge discipline-oriented computations, as well as for basic research in theoretical computer science that were previously allocated to HPCC. Other NSF activities in HPCC are increasing in FY 1997, resulting in a slight net reduction in the total.

** These totals vary slightly from the President's HPCC Budget. For example, funding for the Department of Transportation, a candidate HPCC agency, is not included.

VI. CIC R&D Summary

CIC R&D Goals

Assure continued U.S. leadership in computing, information, and communications technologies to meet Federal goals and to support U.S. 21st century industrial, academic, and defense interests

Accelerate deployment of advanced and experimental information technologies to maintain world leadership in science, engineering, and mathematics; improve the quality of life; promote long term economic growth; increase lifelong learning of citizens; protect the environment; harness information technology; and enhance national security

Advance U.S. productivity and industrial competitiveness through long term scientific and engineering research in computing, information, and communications technologies

CIC R&D Agencies

AHCPR – Agency for Health Care Policy and Research, Department of Health and Human Services

DARPA – Defense Advanced Research Projects Agency, Department of Defense

DOE – Department of Energy

ED – Department of Education

EPA – Environmental Protection Agency

NASA – National Aeronautics and Space Administration

NIH – National Institutes of Health, Department of Health and Human Services

NIST – National Institute of Standards and Technology, Department of Commerce

NOAA – National Oceanic and Atmospheric Administration, Department of Commerce

NSA – National Security Agency, Department of Defense

NSF – National Science Foundation

VA – Department of Veterans Affairs



Evaluation Criteria for CIC R&D Programs

Relevance/Contribution

The research must significantly contribute to the overall goals of the Federal Computing, Information, and Communications (CIC) R&D programs, which include the goals of the five Program Component Areas – High End Computing and Computation, Large Scale Networking, High Confidence Systems, Human Centered Systems, and Education, Training, and Human Resources – to enable solution of Grand Challenge- and National Challenge-class applications problems.

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 CIC R&D Program Component Areas.

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 CIC R&D resources (e.g., a balance among Program Component Areas), promote prospects for joint funding, and address long term resource implications.

Agency Approval

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



Committee on Computing, Information, and Communications

Dr. Anita K. Jones <i>Co-Chair</i>	<i>Director, Defense Research and Engineering</i> Department of Defense
Lionel S. Johns <i>Co-Chair</i>	<i>Associate Director for Technology</i> Office of Science & Technology Policy
Dr. Paul Young <i>Vice Chair</i>	<i>Assistant Director</i> Directorate for Computer and Information Science and Engineering National Science Foundation
Paul Regeon <i>Executive Secretary</i>	<i>CCIC Executive Secretary</i> Office of Science & Technology Policy

Members

Dr. Anita K. Jones <i>Director, Defense Research and Engineering</i> Department of Defense	Dr. Gilbert Weigand <i>Deputy Assistant Secretary for Strategic Computing & Simulations</i> Department of Energy
Lionel S. Johns <i>Associate Director for Technology</i> Office of Science & Technology Policy	Dr. Gary J. Foley <i>Director, National Exposure Research Lab</i> Environmental Protection Agency
Dr. Paul R. Young <i>Assistant Director</i> Directorate for Computer and Information Science and Engineering National Science Foundation	Robert Whitehead <i>Associate Administrator for Aeronautics</i> National Aeronautics and Space Administration
John Dahms <i>Associate Deputy Director for Administration for Information Services</i> Central Intelligence Agency	John C. Toole <i>Director</i> National Coordination Office for CIC
Larry Irving <i>Administrator, National Telecommunications and Information Administration</i> Department of Commerce	Thomas A. Kalil <i>Director to the National Economic Council</i> National Economic Council
Dr. Arati Prabhakar <i>Director, National Institute of Standards and Technology</i> Department of Commerce	Dr. Donald A.B. Lindberg <i>Director, National Library of Medicine</i> National Institutes of Health
Dr. Howard Frank <i>Director, Information Technology</i> Defense Advanced Research Projects Agency	George R. Cotter <i>Chief Scientist</i> National Security Agency
Dr. Linda G. Roberts <i>Special Advisor on Educational Technology</i> Department of Education	Eric Macris <i>Policy Analyst</i> Office of Management and Budget
	E. Fenton Carey <i>Research and Special Programs Administration (RSPA)</i> Office of Research, Technology & Analysis Department of Transportation

Alternate Members

Susannah Schiller <i>Program Analyst</i> Department of Commerce	Lee B. Holcomb <i>Director, Aviation Systems Technology Division</i> National Aeronautics and Space Administration
Dr. Elaine Bunten-Mines <i>Director, Program Office</i> Department of Commerce	Dr. Henry Kelly <i>Assistant Director, Technology</i> Office of Science & Technology Policy
Frank A. Schiermeier <i>Director, Atmospheric Modeling Division</i> Environmental Protection Agency	



Computing, Information, and Communications R&D Subcommittee

Computing, Information, Communications R&D Subcommittee

Chair: John C. Toole

DARPA

Representative

Howard Frank

Alternates

Robert Parker

Stephen L. Squires

NSF

Representative

Melvyn Ciment

Alternate

Robert Voigt

NASA

Representative

Anngienetta R. Johnson

Alternates

William Feiereisen

Philip L. Milstead

DOE

Representative

David B. Nelson

Alternates

Daniel A. Hitchcock

Norm Kreisman

Paul H. Smith (DP)

NIH

Representative

Robert L. Martino

Alternates

Michael J. Ackerman

Judith L. Vaitukaitis

NSA

Representative

George R. Cotter

Alternate

Norman S. Glick

NIST

Representative

Frederick C. Johnson

Alternate

R. J. (Jerry) Linn

VA

Representative

Daniel L. Maloney

Alternate

Rebecca L. Kelley

ED

Representative

Linda G. Roberts

Alternate

Alexis T. Poliakoff

NOAA

Representative

Thomas N. Pyke, Jr.

Alternates

William T. Turnbull

Ernest J. Daddio

EPA

Representative

Joan H. Novak

Alternate

Robin L. Dennis

AHCPR

Representative

J. Michael Fitzmaurice

Alternate

Luis Kun

OMB

Eric L. Macris

OSTP

Lionel S. Johns

Michael R. Nelson

Paul A. Regeon

Applications Council

Chair:

Melvyn Ciment, NSF

Federal Networking Council

Co-Chairs:

Tony Villasenor, NASA

George O. Strawn, NSF

PCA Working Groups:

High End Computing and Computation Working Group

Co-Chairs:

Lee B. Holcomb, NASA

Paul H. Smith, DOE

Large Scale Networking Working Group

Co-Chairs:

George O. Strawn, NSF

David B. Nelson, DOE

High Confidence Systems Working Group

Co-Chairs:

Howard Frank, DARPA

Brian Snow, NSA

Human Centered Systems Working Group

Chair:

Y.T. Chien, NSF

Vice Chairs:

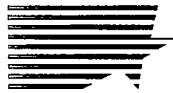
Michael J. Ackerman, NIH

David Gunning, DARPA

Education, Training, and Human Resources Working Group

Chair:

John Cherniavsky, NSF



VII. Glossary

ACTS

Advanced Communications Technology Satellite, a NASA sponsored program. A joint NASA/DARPA collaboration will demonstrate high speed ATM/SONET transmission over the ACTS satellite, and will provide interface and operations experience in mating high speed terrestrial communications systems with high speed satellite communications systems.

AHCPR

Agency for Health Care Policy and Research, part of the Public Health Service of the HHS.

Air Pollution Distance Learning Network

Used for advanced models training and information dissemination.

Algorithm

A procedure designed to solve a problem. Scientific computing programs implement algorithms.

AMSANT

Advanced Manufacturing Systems and Networking Testbed. Experimental facility for testing integration of distributed virtual manufacturing enterprises.

ARPANet

Advanced Research Projects Agency Network.

ASTA

Advanced Software Technology and Algorithms, a component of the original HPCC Program.

ATDNet

Advanced Technology Demonstration Network.

ATM

Asynchronous Transfer Mode, a telecommunications technology, also known as cell switching, which is based on 53-byte cells.

Backbone Network

A high capacity electronic trunk — for example the NSFNET backbone — connecting lower capacity networks.

Bandwidth

A measure of the capacity of a communications channel to transmit information; for example, million of bits per second or Mb/s.

BC

Biomolecular computing.

Benchmark

A point of reference (artifact) to compare an aspect of systems performance (for example, a well known set of programs). Also, to conduct and assess the computation (or transmission) capabilities of a system using a well known artifact.

Binary Black Hole Alliance

A group of astrophysicists who use HPCC technology to enable new ways to model the collision of two black holes.

Bit

An acronym for binary digit.

Bps, or B/s

An acronym for bytes per second.

bps, or b/s

An acronym for bits per second.

BRHR

Basic Research and Human Resources, a component of the original HPCC Program.

**Browser**

A system that provides access to and rendering of distributed information objects located at network-based repositories. That is, a service allowing a user to locate, access, and display information composed of text and still images; to animate moving images; and to play associated sound tracks. Mosaic is an example of a public domain browser.

Byte

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

C

C programming language.

C++

C++ programming language, an object-oriented descendant of the C language.

CAVE

A surround screen, surround sound, projection based virtual reality (VR) system.

CCIC

Committee on Computing, Information, and Communications (formerly the Committee on Information and Communications) of the NSTC.

CIC R&D

Computing, Information, and Communications R&D (formerly HPCCIT), part of CCIC.

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.

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.

CRPC

Center for Research in Parallel Computation.

CTC

Cornell Theory Center.

DAKOTA

Toolkit that aids in optimizing complex multidisciplinary designs. Has been used to improve chemical vapor decomposition reactors used in integrated circuit manufacturing.

DARPA

Defense Advanced Research Projects Agency, part of DOD. Formerly ARPA.

DCRT

Division of Computer Research and Technology, part of NIH.

DFS

Distributed File System.

Digital Libraries Initiative

Program that developed an interoperability protocol for interconnecting diverse libraries and services.

DII

Defense Information Infrastructure.

Distributed Visualization Project

Develops data compression algorithms that provide fast data browsing and interactive visualization capabilities.

DNA

Deoxyribonucleic Acid, a biomolecule from which genes are composed.



DOC	Flops
Department of Commerce.	Acronym for floating point operations per second. The term “floating point” refers to that format of numbers that 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.
DOD	
Department of Defense.	
DOE	
Department of Energy.	
DSM	FNC
Distributed Shared Memory.	Federal Networking Council.
D-System Project	Four Dimensional Data Assimilation
Supports development tools and debugging aids for data parallel programs.	Melds observation data from satellites and other platforms into earth system models to provide more accurate predictions of physical phenomena.
DTE	
Domain and Type Enforcement.	
ED	G, or Giga-
Department of Education.	A prefix denoting 10^9 , or a billion ... (for example, Gflops or gigaflops; gigabytes, gigabits).
EPA	Gateway
Environmental Protection Agency.	A system that interconnects networks (or applications) that communicate using different protocols, and bridges their differences by transforming one protocol (message) into another.
ESNet	
Energy Sciences Network.	
ESS	GB
Earth and Space Sciences.	An acronym for Gigabyte.
ETHR	Gb
Education, Training, and Human Resources.. One of the five Program Component Areas (see page 32).	An acronym for Gigabit.
Exa-	Gb/s
A prefix denoting 10^{18} , or a million trillion ... (for example, exabytes).	Gigabits per second.
FDDI	Gflops
Fiber Distributed Data Interface.	Gigaflops, billions of floating point operations per second.
FIDO	GII
Toolkit for optimizing complex multidisciplinary designs.	Global Information Infrastructure.
	GII Testbed
	System created for SC '95 to accelerate development of distributed computing tools to meet the Grand and National Challenges and beyond. Features interactive 2- and 3-D scientific visualization and virtual reality demonstrations.



Grand Challenge

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

HCS

High Confidence Systems. One of the five Program Component Areas (see page 32).

HECC

High End Computing and Computation. One of the five Program Component Areas (see page 31).

Heterogeneous system

A distributed system that contains more than one kind of computer.

HHS

Department of Health and Human Services.

High performance computing

The full range of advanced computing systems including workstations, networks of workstations with servers, scalable parallel systems, vector parallel systems, and more specialized systems. Scalable input/output interfaces, mass storage systems, and archival storage are components of these systems. Included also are system software and software development environments that enable users to view their workstations and the rest of their computing environments as a unified system.

HiPPI

High Performance Parallel Interface, 800 Mb/s.

HPCC

High Performance Computing and Communications.

HPCCI

High Performance Computing and Communications Initiative.

HPCCIT

High Performance Computing, Communications, and Information Technology Subcommittee, part of the CIC (currently CIC R&D Subcommittee).

HPCS

High Performance Computing Systems, a component of the original HPCC Program.

HPF

High Performance Fortran.

HuCS

Human Centered Systems. One of the five Program Component Areas (see page 32).

IITA

Information Infrastructure Technology and Applications, a component of the original HPCC Program.

ImmersaDesk

A more portable version of CAVE.

Internet

The global collection of interconnected, multiprotocol computer networks including Federal, private, and international networks.

Interoperability

The ability of any two computers that are interconnected to understand each other and perform mutually supportive tasks such as client/server computing.

I/O

Acronym for Input/Output.

IP

Internet Protocol.

ISIS

Imaging Science and Information System. Under development as a system having the potential to detect and classify tiny precancerous formations (microcalcifications) in the breast. May be made locally available via telecommunications.



I-WAY

An experimental, high performance network project at SC'95 based on ATM technology linking over a dozen high performance computers and advanced visualization machines by integrating existing high bandwidth networks with telephone systems.

K, or Kilo-

A prefix denoting 10^3 , or a thousand ... (for example, kilobits/second).

Kb/s

Kilobits per second or thousands of bits per second.

LANL

Los Alamos National Laboratory.

Lattice QCD

Discrete computational technique for determining the consequences of QCD theory from first principles.

LLNL

Lawrence Livermore National Laboratory.

LSN

Large Scale Networking. One of the five Program Component Areas (see page 31).

M, or Mega-

A prefix denoting 10^6 , or a million ... (for example, Mbps, or megabits per second; Mflops).

Math and Science Gateway

Internet link that connects students and educators to educational resources.

MB

An acronym for Megabyte.

Mb

An acronym for Megabit.

MBONE

Multicast backbone.

Mb/s

Megabits per second or millions of bits per second.

MDScope

Application that allows scientists to explore the attributes of macromolecules.

Medical Connections Program

Provides "jump start" funding to academic medical centers, community hospitals, and health care organizations to connect to the Internet.

MEDLARS

Medical Literature Analysis and Retrieval System. Provides rapid access to the National Library of Medicine's biomedical and health information.

Metacenter

A virtual supercomputing center composed of the networked high performance computing resources at several supercomputing centers.

Mflops

Megaflops, millions of floating point operations per second.

MIMD

Multiple Instruction, Multiple Data. An architectural paradigm for parallel computers.

Moore's Law

A rule of thumb stating that computing speed doubles every 18 months.

MosquitoNet

An experimental system developed to investigate issues such as seamless and transparent migration across different kinds of networks, variable bandwidth availability, data consistency issues due to frequent disconnections, and access to resources in the host environment.

MPI

Message Passing Interface.

MPP

Massively parallel processor.

MRAs

Metacenter Regional Alliance.

**MSM**

Multisurface Method. Uses linear programming and multiple separating hyperplanes to analyze fine needle aspiration for diagnosing breast cancer.

Multicast

A mode of communications that allows for simultaneous distribution of information to multiple designated recipients in a single transmission.

NASA

National Aeronautics and Space Administration.

National Challenge

A fundamental application that has broad and direct impact on the Nation's competitiveness and the well-being of its citizens, and that can benefit from the application of HPCC technology and resources.

National Information Infrastructure (NII)

The integration of hardware, software, and skills that will make it easy and affordable to connect people with each other, with computers, and with a vast array of services and information resources.

NCAR

National Center for Atmospheric Research.

NCBI

National Center for Biotechnology Information.

NCI

National Cancer Institute, part of NIH.

NCO

National Coordination Office for Computing, Information, and Communications.

NCRR

National Center for Research Resources, part of NIH.

NCSA

National Center for Supercomputing Applications, Urbana-Champaign, IL.

NERSC

National Energy Research Supercomputer Center.

Network

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

NHSE

National High performance Software Exchange

NIH

National Institutes of Health, part of HHS.

NIST

National Institute of Standards and Technology, part of DOC.

NLM

National Library of Medicine, part of NIH.

NOAA

National Oceanic and Atmospheric Administration, part of DOC.

NOWs

Network of Workstations.

NREN

National Research and Education Network, a component of the original HPCC Program.

NSA

National Security Agency, part of DOD.

NSF

National Science Foundation.

NSFNET

National Science Foundation computer network program.

NSTC

National Science and Technology Council.

Numerical Tokamak Project

Multidisciplinary effort to provide realistic simulations of tokamak plasma disturbances needed to optimize performance of fusion devices.

OMB

Office of Management and Budget.

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.

Optoelectronic

A descriptor of technology which combines optical and electronic components.

ORNL

Oak Ridge National Laboratory.

OS

Operating system.

OSTP

White House Office of Science and Technology Policy.

PABLO and Paradyne

Tools to monitor behavior of application software.

Parallel processing

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

PCA

Program Component Area. Structure of the new Computing, Information, and Communications R&D programs. Each PCA spans an area in which multiple agencies have activities. The five PCAs are: High End Computing and Computation (HECC); Large Scale Networking (LSN); High Confidence Systems (HCS); Human Centered Systems (HuCS); and Education, Training, and Human Resources (ETHR).

Peta-

A prefix denoting 10^{15} , or a thousand trillion ... (for example, petabits).

PGRT

Parallel Graph Reduction Tool. Tool for instrumentation and visualization of realtime embedded systems.

PlantSTEP

A consortium of U.S. process, engineering, and construction industry leaders, including Bechtel, Black & Veatch, DuPont, Eastman Chemical, Merck, H.B. Zachry, Autodesk, Bentley, CadCentre, ComputerVision, Dassault Systems of America, EA Systems, Intergraph, Jacobus, John Brown Systems, Rebis, and Shaw/Sunland.

PSC

Pittsburgh Supercomputer Center.

PVM

Parallel Virtual Machine.

QCD

Quantum Chromodynamics. Theory describing interactions between quarks and gluons.

R&D

Research and development.

RPM

Regional Particulate Model.

SALICIN

Climate data repository.

SC '95

SC 'XY (where XY is a year). General name of an annual conference on high performance computing and communications.

Scalable

A system is scalable if it can be made to have more (or less) computational power by configuring it with a larger (or smaller) number of processors, amount of memory, interconnection bandwidth, input/output bandwidth, and amount of mass storage.

ScaLAPACK

Scalable Linear Algebra Package. A library of equation solvers and linear algebra routines.



SDSC	vBNS
San Diego Supercomputer Center.	Very high speed Backbone Network Services.
SHARE	Visible Human Datasets
Secure Heterogeneous Application Runtime Environment.	Detailed atlases of human anatomy at unprecedented resolution.
SIMD	VLSI
Single Instruction Multiple Data.	Very Large Scale Integration. A type of computer chip.
SIO	VR
Scalable I/O Initiative.	Virtual reality.
SONET	Web
Synchronous Optical Network.	A reference to the World Wide Web, which is a subset of the Internet supported by a related set of protocols, services, and software tools including browsers.
SRTOS	Wireless technologies
Scalable Realtime Operating System.	Communications technologies that use radio, microwave, or satellite communications channels versus wire, coaxial or optical fiber.
SS	WWW
Superstrings.	World Wide Web.
SUIF	
A parallelizer using whole program analysis to achieve a goal.	
T, or Tera-	
A prefix denoting 10^{12} or a trillion ... (for example, terabits, teraflops).	
T1	
Network transmission of a digital signal at 1.5 Mb/s (DS-1).	
T3	
Network transmission of a digital signal at 45 Mb/s (DS-3).	
TCP/IP	
Transmission Control Protocol/Internet Protocol.	
TRP	
Technology Reinvestment Project. Supports ATM Interoperability Testbed for National Information Infrastructure and the SONET OC-192 ATM Self Healing Ring.	
VA	
Department of Veterans Affairs.	



VIII. Contacts

National Coordination Office for Computing, Information, and Communications R&D (NCO)

Suite 665
4201 Wilson Boulevard
Arlington, VA 22230
(703) 306-HPCC (4722)
FAX: (703) 306-4727
nco@hpcc.gov

Internet/World Wide Web Server:

<http://www.hpcc.gov/>

NCO

John C. Toole
Director
toole@hpcc.gov

Sally E. Howe, Ph.D.
Chief of Staff
howe@hpcc.gov

Donald M. Austin, Ph.D.
Assistant Director for Planning
austin@hpcc.gov

Margaret L. Simmons
Assistant Director for Program Integration
simmons@hpcc.gov

John J. Dumbleton
ComSci Fellow
jdumble@hpcc.gov

Riva R. Meade
Assistant to the Director
meade@hpcc.gov

Catherine McDonald
Senior Analyst
mcdonald@hpcc.gov

Terry L. Ponick, Ph.D.
Technical Writer
ponick@hpcc.gov

Ward Fenton
Systems Administrator
fenton@hpcc.gov

Vicki L. Harris
Administrative Assistant
harris@hpcc.gov

AHCPR

J. Michael Fitzmaurice, Ph.D.
Director, Center for Information Technology
Agency for Health Care Policy and Research
2101 East Jefferson Street, Suite 604
Rockville, MD 20852
(301) 594-1483
FAX: (301) 594-2333
mfitzmau@po5.ahcpr.gov

Luis G. Kun, Ph.D.
Senior Advisor, Center for Information
Technology
Agency for Health Care Policy and Research
2101 East Jefferson Street, Suite 602
Rockville, MD 20852
(301) 594-1483
FAX: (301) 594-2333
lkun@po5.ahcpr.gov

DARPA

Howard Frank, Ph.D.

*Director, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2228
FAX: (703) 696-2202
hfrank@darpa.mil*

Stephen L. Squires

*Special Assistant to the Director
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2226
FAX: (703) 696-2209
squires@darpa.mil*

Robert Parker, Ph.D.

*Deputy Director, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2255
FAX: (703) 696-2202
rparker@darpa.mil*

Robert Lucas, Ph.D.

*Assistant Director, Computation and
Networking, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2226
FAX: (703) 696-2202
rlucas@darpa.mil*

Kevin Mills, Ph.D.

*Intelligent Collaboration and Visualization
Program Manager, Information Technology
Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2227
FAX: (703) 696-2202
kmills@darpa.mil*

José L. Muñoz, Ph.D.

*Embeddable Systems Program Manager,
Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-4468
FAX: (703) 696-2202
jmunoz@darpa.mil*

Frederica Darema, Ph.D.

*Program Manager
Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2271
FAX: (703) 696-2202
fdarema@darpa.mil*

David Gunning

*Information Systems Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2218
FAX: (703) 696-2201
dgunning@darpa.mil*

DOE

CIC R&D Coordinator

*Office of Computational and Technology
Research (OCTR)
Department of Energy
OCTR, ER-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-5800
FAX: (301) 903-7774
hpcc@er.doe.gov*

DOE (continued)

David B. Nelson, Ph.D.
Associate Director, Office of Computational and Technology Research (OCTR)
Department of Energy
OCTR, ER-30
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-5800
FAX: (301) 903-7774
nelson@er.doe.gov

Daniel A. Hitchcock, Ph.D.
Acting Director, Mathematical, Information, and Computational Sciences (MICS) Division
Office of Computational and Technology Research (OCTR)
Department of Energy
OCTR/MICS, ER-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-6767
FAX: (301) 903-7774
hitchcock@er.doe.gov

Thomas A. Kitchens, Ph.D.
Program Manager, MICS Division,
Office of Computational and Technology Research (OCTR)
Department of Energy
OCTR/MICS, ER-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-5152
FAX: (301) 903-7774
kitchens@er.doe.gov

John S. Cavallini
MICS Division, Office of Computational and Technology Research (OCTR)
Department of Energy
OCTR/MICS, ER-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-5800
FAX: (301) 903-7774
cavallini@nersc.gov

Norman H. Kreisman
Advisor, International Technology
Department of Energy, ER-5
Mailstop 3H049-FORS
1000 Independence Avenue, S.W.
Washington, DC 20585
(202) 586-9746
FAX: (202) 586-7152
kreisman@er.doe.gov

Frederick A. Howes, Ph.D.
Program Manager, MICS Division,
Office of Computational and Technology Research (OCTR)
Department of Energy
OCTR/MICS, ER-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-3166
FAX: (301) 903-7774
howes@er.doe.gov

Mary Anne Scott, Ph.D.
Program Manager, MICS Division,
Office of Computational and Technology Research (OCTR)
Department of Energy
OCTR/MICS, ER-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-6368
FAX: (301) 903-7774
scott@er.doe.gov

Rod Oldehoeft, Ph.D.
Program Manager, MICS Division,
Office of Computational and Technology Research (OCTR)
Department of Energy
OCTR/MICS, ER-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-6904
FAX: (301) 903-7774
rro@er.doe.gov

DOE (continued)

George R. Seweryniak
*Program Manager, MICS Division,
Office of Computational and Technology
Research (OCTR)
Department of Energy
OCTR/MICS, ER-31
19901 Germantown Road
Germantown, MD 20874-1290
(301)903-0071
FAX: (301) 903-7774
seweryni@er.doe.gov*

Gilbert G. Weigand, Ph.D.
*Deputy Assistant Secretary, Strategic Computing
and Simulation
Department of Energy
1000 Independence Avenue S.W.
Washington, DC 20585
(202) 586-0568
FAX: (202) 586-7754
weigand@dp.doe.gov*

Paul H. Smith, Ph.D.
*Special Assistant, Advanced Computing
Technology, Defense Programs Office
Department of Energy
DP-50
1000 Independence Avenue, S.W.
Washington, DC 20585
(202) 586-0992
FAX: (202) 586-7754
phsmith@dp.doe.gov*

ED

Linda G. Roberts, Ed. D.
*Special Advisor on Educational Technology,
Office of the Deputy Secretary
Department of Education
FB10
600 Independence Avenue, S.W.
Washington, DC 20202
(202) 401-1000
FAX: (202) 401-3093
linda_roberts@ed.gov*

Alexis T. Poliakoff
*Network Planning Staff
Department of Education
OHRA/IRG
600 Independence Avenue, S.W.
Washington, DC 20202
(202) 708-5210
FAX: (202) 260-5669
alex_poliakoff@ed.gov*

EPA

Joan H. Novak
*HPCC Program Manager, MD-80
Environmental Protection Agency
Research Triangle Park, NC 27711
(919) 541-4545
FAX: (919) 541-1379
novak.joan@epamail.epa.gov*

Robin L. Dennis, Ph.D.
*Senior Science Program Manager, MD-80
Environmental Protection Agency
Research Triangle Park, NC 27711
(919) 541-2870
FAX: (919) 541-1379
rdennis@hpcc.epa.gov*

EPA FedEx address:

79 Alexander Drive
Building 4201
Research Triangle Park, NC 27709

NASA

Lee B. Holcomb
*Director, HPCC Office
Office of Aeronautics
National Aeronautics and Space Administration
Code RC
300 E Street, S.W.
Washington, DC 20546
(202) 358-2747
FAX: (202) 358-3557
l_holcomb@aeromail.hq.nasa.gov*



NASA (continued)

Anngienetta R. Johnson
HPCC Program Manager
National Aeronautics and Space Administration
Code RC
300 E Street, S.W.
Washington, DC 20546
(202) 358-4717
FAX: (202) 358-3557
ajohnson@aeromail.hq.nasa.gov

William J. Feiereisen
*Project Manager, Computational Aerosciences,
High Performance Computing and
Communications Program*
National Aeronautics and Space Administration
Ames Research Center
Moffett Field, CA 94035-1000
(415) 604-4225
FAX: (415) 604-7324
feieris@ames.arc.nasa.gov

Phillip L. Milstead
Information Technology Program Coordinator
National Aeronautics and Space Administration
Code RC
300 E Street, S.W.
Washington, DC 20546
(202) 358-4619
FAX: (202) 358-3550
p_milstead@aeromail.hq.nasa.gov

NASA express/commercial/courier address:

NASA High Performance Computing and
Communications Office
Mail Code RC
Attn: Receiving & Inspection (Rear of
Building)
300 E Street, S.W.
Washington, DC 20024-3210

NIH

NIH HPCC Coordinator and DCRT contact:

Robert L. Martino, Ph.D.
*Chief, Computational Bioscience and
Engineering Laboratory, Division of
Computer Research and Technology*
National Institutes of Health
9000 Rockville Pike
Building 12A, Room 2033
Bethesda, MD 20892
(301) 496-1111
FAX: (301) 402-2867
martino@alw.nih.gov

DCRT program:

William L. Risso
*Deputy Director, Division of Computer
Research and Technology*
National Institutes of Health
9000 Rockville Pike
Building 12A, Room 3033
Bethesda, MD 20892
(301) 496-8277
FAX: (301) 402-1754
risso@nih.gov

NCI program:

Jacob V. Maizel, Ph.D.
*Biomedical Supercomputer Center,
National Cancer Institute,
Frederick Cancer Research and
Development Center*
National Institutes of Health
P.O. Box B, Building 469, Room 151
Frederick, MD 21702-1201
(301) 846-5532
FAX: (301) 846-5598
jmaizel@ncifcrf.gov

NIH (continued)

Cherie Nichols

*Chief, Planning Evaluation and Analysis
Branch, Office of Program Operations and
Planning, National Cancer Institute
National Institutes of Health
Building 31, Room 11A21
31 Center Drive, MSC 2590
Bethesda, MD 20892
(301) 496-5515
FAX: (301) 402-1225
nicholsc@od.nci.nih.gov*

NCRR program:

Judith L. Vaitukaitis, M.D.
*Director, National Center for Research
Resources
National Institutes of Health
9000 Rockville Pike
Building 12A, Room 4011
Bethesda, MD 20892
(301) 496-5793
FAX: (301) 402-0006
vaitukaitis@nih.gov*

Richard M. DuBois, Ph.D.
*Head, Computer Technology
Section
National Institutes of Health
One Rockledge Center
6705 Rockledge Drive, Room 6164
Bethesda, MD 20892-7965
(301) 435-0758
FAX: (301) 480-3654
richarddd@ep.ncrr.nih.gov*

NLM program:

Donald A.B. Lindberg, M.D.
*Director, National Library of Medicine
National Institutes of Health
Building 38, Room 2E17
8600 Rockville Pike
Bethesda, MD 20894
(301) 496-6221
FAX: (301) 496-4450
lindberg@nlm.nih.gov*

Michael J. Ackerman, Ph.D
*Assistant Director for High Performance
Computing and Communications,
National Library of Medicine
National Institutes of Health
Building 38A, Room B1N30
8600 Rockville Pike
Bethesda, MD 20894
(301) 402-4100
FAX: (301) 402-4080
ackerman@nlm.nih.gov*

NIST

Frederick C. Johnson, Ph.D.
*Associate Director for Computing, Computing
and Applied Mathematics Laboratory
National Institute of Standards and Technology
Building 820, Room 672
Gaithersburg, MD 20899
(301) 975-2700
FAX: (301) 216-2075
fjohnson@nist.gov*

R. J. (Jerry) Linn
*Associate Director for Program Implementation,
Computer Systems Laboratory
National Institute of Standards and Technology
Building 820, Room 635
Gaithersburg, MD 20899
(301) 975-3624
FAX: (301) 948-1784
linnrj@nist.gov*

Ronald F. Boisvert, Ph.D.
*Leader, Mathematical Software Group
National Institute of Standards and Technology
Gaithersburg, MD 20899
301-975-3812
Fax: 301-990-4127
boisvert@nist.gov*



NOAA

Thomas N. Pyke, Jr.
Director for HPCC
National Oceanic and Atmospheric Administration
Room 15300
1315 East-West Highway
Silver Spring, MD 20910
(301) 713-3573
FAX: (301) 713-4040
tpyke@hpcc.noaa.gov

William T. Turnbull
Deputy Director, HPCC
National Oceanic and Atmospheric Administration
Room 15300
1315 East-West Highway
Silver Spring, MD 20910
(301) 713-3573
FAX: (301) 713-4040
wturnbull@hpcc.noaa.gov

Ernest Daddio
Program Officer
National Oceanic and Atmospheric Administration
Room 15400
1315 East-West Highway
Silver Spring, MD 20910
(301) 713-1262
FAX: (301) 713-1249
edaddio@hpcc.noaa.gov

Directors of high performance computing centers:

Jerry Mahlman, Ph.D.
Director, Geophysical Fluid Dynamics Laboratory
National Oceanic and Atmospheric Administration
Forrestal Campus, U.S. Route 1
P.O. Box 308
Princeton, NJ 08542
(609) 452-6502
FAX: (609) 987-5070
jm@gfdl.gov

Ron McPherson, Ph.D.
Director, National Centers for Environmental Prediction
National Oceanic and Atmospheric Administration
5200 Auth Road, Room 101
Camp Springs, MD 20746
(301) 763-8016
FAX: (301) 763-8434
ronmcp@nic.fb4.noaa.gov

Sandy MacDonald, Ph.D.
Director, Forecasting Systems Laboratory
National Oceanic and Atmospheric Administration
325 Broadway
Boulder, CO 80303
(303) 497-6378
FAX: (303) 497-6821
macdonald@fsl.noaa.gov

Internet Use:

Carl Staton
Manager, NOAA Network Information Center
National Oceanic and Atmospheric Administration
Federal Building 4
Suitland and Silver Hill Roads
Suitland, MD 20746
(301) 457-5165
FAX: (301) 457-5199
cstaton@nesdis.noaa.gov

NSA

George R. Cotter
Chief Scientist
National Security Agency
9800 Savage Road
Fort Meade, MD 20755-6000
(301) 688-6434
FAX: (301) 688-4980
grcotter@afterlife.ncsc.mil

NSA (continued)

Norman S. Glick
Senior Computer Scientist
National Security Agency
9800 Savage Road
Fort Meade, MD 20755-6000
(301) 688-8448
FAX: (301) 688-4980
nsglick@afterlife.ncsc.mil

Brian D. Snow
ISSO Technical Director
National Security Agency
9800 Savage Road, Suite 6577
Fort Meade, MD 20755-6577
(301) 688-8112
FAX: (301) 688-3090
bsnow@dockmaster.ncsc.mil

NSF

Paul Young, Ph.D.
Assistant Director, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1105
Arlington, VA 22230
(703) 306-1900
FAX: (703) 306-0577
pyoung@nsf.gov

Melvyn Ciment, Ph.D.
Deputy Assistant Director, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1105
Arlington, VA 22230
(703) 306-1900
FAX: (703) 306-0577
mciment@nsf.gov

Merrell Patrick, Ph.D.
Acting Executive Officer, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1105
Arlington, VA 22230
(703) 306-1900
FAX: (703) 306-0577
mpatrick@nsf.gov

Robert G. Voigt, Ph.D.
HPCC Coordinator, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1105
Arlington, VA 22230
(703) 306-1900
FAX: (703) 306-0577
rvoigt@nsf.gov

George O. Strawn, Ph.D.
Division Director, Networking and Communications Research and Infrastructure
National Science Foundation
4201 Wilson Blvd, Room 1175
Arlington VA 22230
(703) 306-1950
FAX: (703) 306-0621
gstrawn@nsf.gov

Y.T. Chien, Ph.D.
Division Director, Information, Robotics, and Intelligent Systems
National Science Foundation
4201 Wilson Boulevard, Room 1115
Arlington, VA 22230
(703) 306-1930
FAX: (703) 306-0599
ytchien@nsf.gov

John Cherniavsky, Ph.D.
Head, Office of Cross-Disciplinary Activities
National Science Foundation
4201 Wilson Blvd, Room 1160
Arlington VA 22230
(703) 306 1980
FAX: (703) 306-0589
jchernia@nsf.gov



VA

Daniel L. Maloney
*Director, Veteran's Health Administration
Technology Service*
Department of Veterans Affairs
8403 Colesville Road, Suite 200
Silver Spring, MD 20910
(301) 427-3700
FAX: (301) 427-3711
maloney.dan@forum.va.gov

Rebecca L. Kelley
*Management Analyst, Veterans Health
Administration Technology Service*
Department of Veterans Affairs
8403 Colesville Road, Suite 200
Silver Spring, MD 20910
(301) 427-3750
FAX: (301) 427-3711
kelley.rebecca@forum.va.gov

OMB

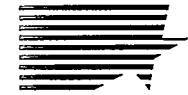
Eric L. Macris
Program Assistant
Office of Management and Budget
Executive Office of the President
New Executive Office Building, Room 8002
725 17th Street, N.W.
Washington, DC 20503
(202) 395-3404
FAX: (202) 395-4817
macris_e@al.eop.gov

OSTP

Lionel S. Johns
*Associate Director for Technology,
Office of Science and Technology Policy*
Executive Office of the President
Old Executive Office Building, Room 423
17th Street and Pennsylvania Avenue, N.W.
Washington, DC 20502
(202) 456-6046
FAX: (202) 456-6023
ljohns@ostp.eop.gov

Michael R. Nelson, Ph.D.
*Special Assistant for Information Technology,
Office of Science and Technology Policy*
Executive Office of the President
Old Executive Office Building, Room 423
17th Street and Pennsylvania Avenue, N.W.
Washington, DC 20502
(202) 456-6046
FAX: (202) 456-6023
mnelson@ostp.eop.gov

Paul A. Regeon
*CCIC Executive Secretary, Office of Science and
Technology Policy*
Executive Office of the President
Old Executive Office Building, Room 423
17th Street and Pennsylvania Avenue, N.W.
Washington, D.C. 20502
(202) 456-6042
FAX: (202) 456-6023
pregeon@ostp.eop.gov



FY 1997 Editorial Group

Editor

Margaret L. Simmons
National Coordination Office

Editorial Assistant

Terry L. Ponick
National Coordination Office

Writing Group

Donald M. Austin, NCO
Ronald F. Boisvert, NIST
Frederica Darema, NSF/DARPA
Richard Dubois, NIH
Norman S. Glick, NSA
Rebecca L. Kelley, VA
Thomas A. Kitchens, DOE
Luis G. Kun, AHCPR
Phillip L. Milstead, NASA
José Muñoz, DARPA
Joan H. Novak, EPA
Alexis T. Poliakoff, ED
William T. Turnbull, NOAA

Acknowledgments

Many people contributed to this book, and we thank them for their efforts. This especially includes the researchers who provided descriptions of their work. Thanks also go to Joe Fitzgerald of the National Library of Medicine for the cover design. We thank the staff of the National Coordination Office for their hard work and assistance.