Computing, Information, and Communications

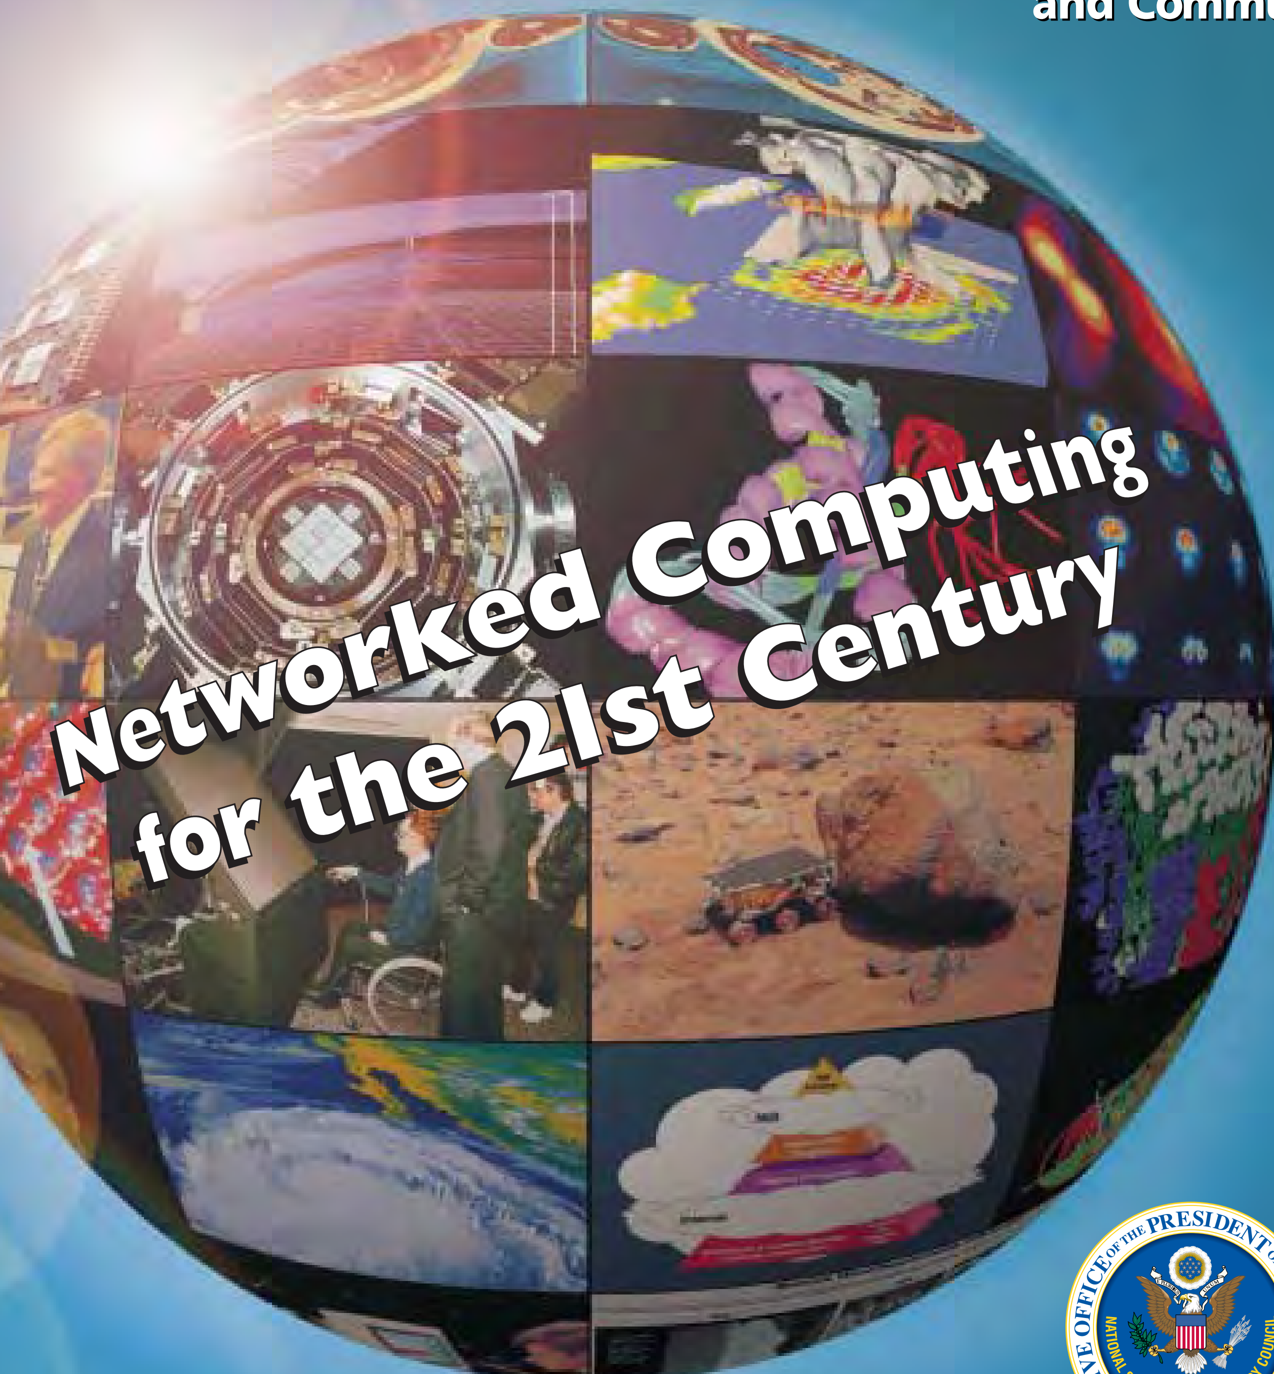

Networked Computing for the 21st Century

Supplement to the President's FY 1999 Budget

National Science and Technology Council
Committee on Technology
Subcommittee on Computing, Information, and Communications R&D

About the National Science and Technology Council:

President Clinton established the National Science and Technology Council (NSTC) by Executive Order on November 23, 1993. This cabinet-level council is the principal means for the President to coordinate science, space, and technology policies across the Federal Government. NSTC acts as a "virtual agency" for science and technology to coordinate the diverse parts of the Federal research and development enterprise. The NSTC is chaired by the President. Membership consists of the Vice President, Assistant to the President for Science and Technology, Cabinet Secretaries, Agency Heads with significant science and technology responsibilities, and other White House officials.

An important objective of the NSTC is the establishment of clear national goals for Federal science and technology investments in areas ranging from information technologies and health research, to improving transportation systems and strengthening fundamental research. The Council prepares research and development strategies that are coordinated across Federal agencies to form an investment package that is aimed at accomplishing multiple national goals.

To obtain additional information regarding the NSTC, contact the NSTC Executive Secretariat at (202) 456-6100.

About the Office of Science and Technology Policy:

The Office of Science and Technology Policy (OSTP) was established by the National Science and Technology Policy, Organization and Priorities Act of 1976. OSTP's responsibilities include advising the President in policy formulation and budget development on all questions in which science and technology are important elements; articulating the President's science and technology policies and programs; and fostering strong partnerships among Federal, State, and local governments, and the scientific communities in industry and academe. The Assistant to the President for Science and Technology serves as the Director of the OSTP and directs the NSTC on behalf of the President. OSTP Associate Directors co-chair the various committees of the NSTC.

To obtain additional information about OSTP, please call (202) 456-7116.

Cover Images

This year's cover is a collage of images gathered from throughout this document. The cover itself highlights the increasing importance of leading-edge technologies to the U.S. and to the world—including technologies, like the Internet, the Next Generation Internet, and advanced computing technologies that have been made possible by CIC R&D efforts.

COMPUTING, INFORMATION, AND COMMUNICATIONS

NETWORKED COMPUTING FOR THE 21ST CENTURY

A Report by the Subcommittee on Computing, Information, and Communications R&D

Committee on Technology

National Science and Technology Council

OSTP LETTER

[FPO-Letter to be supplied]

Committee on Technology

Dr. Mortimer L. Downey
Chair
Deputy Secretary
Department of Transportation

Dr. Duncan T. Moore
White House Co-Chair
Associate Director for Technology
Office of Science and Technology Policy

Gary R. Bachula
Vice Chair
Acting Under Secretary for Technology
Department of Commerce

Vacant
Vice Chair
Director, Defense Research and Engineering
Department of Defense

E. Fenton Carey
Executive Secretary
Research and Special Programs Administration
Department of Transportation

Lori A. Perine
White House Liaison
Senior Policy Advisor, Computing, Information, and Communications
Office of Science and Technology Policy

Members

Dr. Joseph Bordogna
Acting Deputy Director
National Science Foundation

David Boyd
Director, Office of Standards and Technology, National Institute of Justice
Department of Justice

Russell E. Dressell
Director of Research and Development, DS&T
Central Intelligence Agency

David M. Gardiner
Assistant Administrator for Policy, Planning, and Evaluation
Environmental Protection Agency

Dr. I. Miley Gonzalez
Under Secretary for the Research, Education, and Economics Mission Area
Department of Agriculture

Dr. Juris Hartmanis
Assistant Director, Directorate for Computer and Information Science and Engineering
National Science Foundation

Thomas A. Kalil
Senior Director to the National Economic Council
National Economic Council

Dr. Henry C. Kelly
Assistant Director for Technology
Office of Science and Technology Policy

Dr. Ruth L. Kirchstein
Deputy Director
National Institutes of Health

Eric L. Macris
Policy Analyst
Office of Management and Budget

Stan Ponce
Director for Research, Bureau of Reclamation
Department of the Interior

Dan W. Reicher
Assistant Secretary for Energy Efficiency and Renewable Energy
Department of Energy

Dr. Linda G. Roberts
Director, Educational Technology Office
Department of Education

Dr. Linda Rosenstock
Director, National Institute for Occupational Safety and Health
Department of Health and Human Services

Samuel L. Venneri
Chief Technologist
National Aeronautics and Space Administration

James W. Vollman
Administrator for One-Stop Career Centers/Labor Market Information
Department of Labor

Subcommittee on Computing, Information, and Communications R&D

NCO

Chair
Kay Howell

DARPA

Representative
David Tennenhouse
Alternate
Ronald L. Larsen

NSF

Representative
Melvyn Ciment
Alternate
Gary W. Strong

NASA

Representative
Lee B. Holcomb
Alternates
William Feiereisen
Philip L. Milstead

DOE

Representative
Daniel A. Hitchcock
Alternates
Norm Kreisman
Paul H. Smith

NIH

Representative
Robert L. Martino
Alternates
Michael J. Ackerman
Judith L. Vaitukaitis

NSA

Representative
George R. Cotter
Alternate
Norman S. Glick

NIST

Representative
R. J. (Jerry) Linn
Alternate
Frederick C. Johnson

VA

Representative
Daniel L. Maloney

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

Henry A. Kelly
Lori A. Perine

High End Computing and Computation Working Group

Co-Chairs
Lee B. Holcomb, NASA
Paul H. Smith, DOE

Large Scale Networking (LSN) Working Group

Co-Chairs
George O. Strawn, NSF
David B. Nelson, NASA

NGI Applications Team

Chair
William T. Turnbull, NOAA

Joint Engineering Team

Co-Chairs
Phil Dykstra, DoD
Javad Boroumand, NSF

Networking Research Team

Co-Chairs
Tatsuya Suda, NSF
Mari Maeda, DARPA

High Confidence Systems Working Group

Co-Chairs
Teresa Lunt, DARPA
Andrew Arenth, NSA

Human Centered Systems Working Group

Chair
Michael Lesk, NSF

Education, Training, and Human Resources Working Group

Chair
John Cherniavsky, NSF

Federal Information Services Applications Council

Co-Chairs
Melvyn Ciment, NSF
Marty Wagner, GSA

Table of Contents

Table of Contents

Table of Contents

Executive Summary

Overview

U.S. research and development (R&D) in computing, communications, and information technologies has enabled unprecedented scientific and engineering advances, transforming workplace products and processes, and benefiting society and individuals. Today's benefits are the result of investments made decades ago; the seeds of the Internet, for example, were planted in the 1960s. Continued investments will result in further dramatic advances, some of which may be as unpredictable and compelling as the World Wide Web. These investments include the development of increasingly powerful high performance computing systems, global-scale networking technologies with advanced capabilities, advances in software development technologies and applications software, improved reliability and safety, advances in managing and accessing vast distributed knowledge repositories, and human interface technologies that let people work — and work together — more effectively.

Computing, Information, and Communications (CIC) Program Component Areas (PCAs)

Federally-funded R&D in these technologies is coordinated within the Computing, Information, and Communications (CIC) R&D programs, which are successors to the highly successful, Congressionally-chartered High Performance Computing and Communications (HPCC) Program. The application of computing, communications, and information technologies beyond their decades-long concentration on science, engineering, and education has placed additional demands on the CIC programs. Over the past two years, these programs have redefined themselves and are in the process of reorganizing their coordinated activities into five Program Component Areas (PCAs):

HECC High End Computing and Computation
LSN Large Scale Networking, including the Next Generation Internet initiative
HCS High Confidence Systems
HuCS Human Centered Systems
ETHR Education, Training, and Human Resources

plus the Federal Information Services and Applications Council (FISAC).

The first step of the reorganization is reflected in this FY 1999 Supplement to the President's Budget. The Federal investments in LSN R&D continue to be conducted, essentially, by organizations that participate in the coordinated CIC programs, and this is true, though to a lesser degree, for HECC R&D.

Accordingly, the President’s proposed FY 1999 budget includes the crosscuts for the HECC and LSN PCAs. However, this President’s Budget does not include HCS, HuCS, or ETHR crosscuts, in recognition of the possibly multi-year process of reorganizing their R&D to include work being conducted both by CIC organizations — but outside their historical CIC budget crosscuts — and by other Federal organizations.

This document describes the HECC and LSN FY 1998 anticipated accomplishments and FY 1999 plans in substantial detail. It also includes descriptions of some representative HCS, HuCS, and ETHR activities and plans, primarily focusing on work by organizations that have been in the CIC programs.

The CIC R&D reporting structure changed in FY 1998. In December 1997, the White House National Science and Technology Council (NSTC) was reorganized and streamlined. Three NSTC Committees — the Committee on Computing, Information, and Communications (CCIC), the Committee on Technological Innovation (CTI), and the Committee on Transportation R&D (CTRD) were replaced by the new Committee on Technology (CT). The Subcommittee on CIC R&D of the CCIC now reports to the CT. The five PCA Working Groups continue to report to the Subcommittee on CIC R&D, as does the FISAC, which had reported to the CCIC.

High End Computing and Computation (HECC)

HECC R&D investments provide the foundation for U.S. leadership in high end computing and promote the use of high end computing and computation for government, academia, industry, and broad societal applications. HECC research explores algorithms for physical, chemical, and biological modeling and simulation in complex systems; information-intensive science and engineering applications; and the advanced concepts in quantum, biological, and optical computing that will keep the U.S. in the forefront of computing breakthroughs for years to come.

Short-term HECC development — with anticipated payoffs in three to five years — addresses needs for systems software for teraflops (10^{12} floating point operations per second) by means of investments in operating systems, languages and compilers, programming environments and libraries, debugging and performance tools, scientific visualization, data management, and developments leading to a common framework and infrastructure. FY 1998 and FY 1999 R&D areas include ease of use, performance measurement, provability, and resource reservation and guarantees.

Long-term HECC R&D — to be useful in 10 to 15 years — has helped establish scalable parallel processing as a standard for high performance computing, and has enabled the technology base for the $2 billion per year mid-range computing market. The next major long-term HECC milestone is a reliable, robust implementation of petaflops — 10^{15} flops level performance, and exabyte — 10^{18} bytes — storage capability. Applications requiring this level of computational power include simulating complex systems in biology, business, the Earth and the environment, materials science and manufacturing, and the physical sciences. FY 1998 and FY 1999 research topics include embedded systems, intelligent systems, optical storage, optoelectronics, quantum computing, scalability, and survivability.

HECC R&D includes incorporating HECC technologies in agency mission applications. In FY 1998 and FY 1999 these applications include acoustics, aerodynamics, air and water quality modeling, biology (cellular, molecular, and biochemistry), Earth and space systems modeling, numerical analysis software, pharmaceuticals, physics, reconstructing three-dimensional information from two-dimensional biomedical images, regional-scale weather modeling, and simulating hurricane structure.

HECC investments include the infrastructure for HECC R&D, which includes the Department of Energy's (DOE) national laboratories, the National Aeronautics and Space Administration (NASA) centers, and facilities at the Environmental Protection Agency (EPA), the National Institutes of Health (NIH), the National Oceanic and Atmospheric Administration (NOAA), and the National Security Agency (NSA). In FY 1998, the National Science Foundation's (NSF) Partnership for Advanced Computational Infrastructure (PACI) became a part of this infrastructure, succeeding the NSF supercomputing centers. PACI's goal is to prototype a national-scale computational grid for use in research and education. Currently, there are more than 101 academic and industrial PACI partners in 30 states.

Large Scale Networking (LSN)

FY 1999 LSN R&D, including the Next Generation Internet (NGI) initiative, will help assure U.S. technological leadership in high performance network communications through research that advances the leading edge of networking technologies, services, and performance. Early Federal investments in networking R&D helped build the technological foundation of today's global Internet. Development by Federal research laboratories, academia, and industry helped deploy prototype networking capabilities on a national scale and produced popular applications — like email and World Wide Web browsers — that transformed the way people use computer networks, paving the way for our Nation's leadership in the multi-billion dollar information technology industry. Key research areas today include advanced network components and technologies for engineering and management of large scale networks of the future.

LSN activities include coordinating the operation and select peering of advanced Federal networks, including the multiagency Washington, D.C. area Advanced Technology Demonstration network (ATDnet), the Department of Defense's (DoD) Defense Research and Engineering Network (DREN), DOE's Energy Sciences network (ESnet), the NASA Research and Education Network (NREN), and NSF's very high performance Backbone Network Service (vBNS). Areas of research and support in FY 1998 and FY 1999 include global-scale communications, networking security, satellite technologies, special purpose connectivity programs, and network-based applications.

Next Generation Internet (NGI)

The primary focus of LSN activities in FY 1998 has been the Presidential NGI initiative. The investments made by this initiative and associated academic and industrial investments are creating the foundation for the networks of the 21st century — networks that will be much more powerful and versatile than today's Internet. The NGI, in partnership with these other investment sectors, will keep the U.S. at the cutting edge of communications and information technologies. The NGI will also stimulate the introduction of new multimedia applications in our homes, schools, and businesses as the technologies designed and developed as part of the NGI are incorporated into products and services that are subsequently made available to the general public. The NGI initiative is essential to sustain U.S. technological leadership in computing and communications and to enhance U.S. economic competitiveness and commercial eminence. NGI activities are leveraged off of and tightly coupled with the base LSN network research and infrastructure support.

The NGI goals are:

- ❒ To conduct R&D in advanced end-to-end networking technologies, including differentiated services (including multicast and audio/video), network management (including allocation and sharing of bandwidth), reliability, robustness, and security.
- ❒ To prototype high performance network testbeds for systems scale testing of advanced technologies and services and for developing and testing advanced applications. One testbed will connect at least 100 sites with end-to-end performance at least 100 times faster than today's Internet; it will be built on the Federal networks in cooperation with academic campus and regional networks. The other testbed will connect more than 10 sites with end-to-end performance at least 1,000 times faster than today's Internet. This network will build on the Defense Advanced Research Projects Agency's (DARPA) advanced networks.
- ❒ To develop revolutionary applications. These include enabling applications technologies such as collaboration technologies, digital libraries, distributed computing, privacy and security, and remote operation and simulation, as well as disciplinary applications in basic science, crisis management, education, the environment, Federal information services, health care, and manufacturing.

For FY 1998, the NGI agencies are DARPA, NSF, NASA, National Institute of Standards and Technology (NIST), and NIH (National Library of Medicine [NLM], and National Center for Research Resources [NCRR]). The proposed FY 1999 NGI agencies will additionally include DOE.

Netamorphosis demonstrations

Seventeen NGI applications, developed by Federal LSN programs including the NGI initiative, were demonstrated at Netamorphosis, held March 11-13, 1998, at Highway1 near the U.S. Capitol in Washington, D.C. These included:

Advanced Regional Prediction System (ARPS):
Precision storm forecasting

Automating the construction site: A leap in capabilities

Cave5D: A tool for collaborative immersive visualization of environmental data

Collaboratory for structure-based drug design

Earth data from satellite to desktop

Echocardiography over the NGI

Exploring the Earth system on the "Second Web"

Informedia News-on-Demand

Interactive video dialogues

Magic: Viewing large geographic areas in 3-D

Materials Microcharacterization Collaboratory (MMC): Studying state-of-the-art materials

Nanotechnology research: Controlling atoms from a distance

Octahedral hexapod: An Information Age machine tool

Real-time functional MRI: Watching the brain in action

Security technology for the Next Generation Internet

SF Express: Advanced battle simulation

Visible Human Project

High Confidence Systems (HCS)

HCS R&D focuses on the critical technologies necessary to achieve high levels of availability, reliability, restorability, protection, and security of information services. Systems that employ these technologies will be resistant to component failure and malicious manipulation and will respond to damage or perceived threat by adaptation or reconfiguration. Applications include banking, law enforcement, life- and safety-critical systems, medicine and health care, national security, power generation and distribution, telecommunications, and transportation.

FY 1998 and FY 1999 HCS R&D includes work in assurance technologies, information security, information survivability, protecting the privacy of medical records, and secure programming languages for Internet-based applications.

In August 1997 and March 1998, as part of the process of developing new CIC R&D agendas that extend beyond high levels of performance, the HCS Working Group held workshops with CIC and other agencies to identify agency needs requiring HCS research and to develop a proposed HCS R&D agenda. These agencies included DARPA, DOE, the Department of the Treasury's Internal Revenue Service (IRS), Federal Aviation Administration (FAA), Federal Railroad Administration (FRA), Federal Transportation Administration (FTA), Food and Drug Administration (FDA), NASA, NIST, NSA, and the Nuclear Regulatory Commission (NRC).

Human Centered Systems (HuCS)

HuCS R&D leads to increased accessibility and usability of computing systems and communications networks. Scientists, engineers, educators, students, the workforce, and the general public are all potential beneficiaries of HuCS technologies and applications.

A new Digital Libraries Phase II initiative will begin in FY 1998. This is a joint effort of NSF, DARPA, NASA, NIH/NLM, and several non-CIC agencies — the Library of Congress and the National Endowment for the Humanities in partnership with the National Archives and Records Administration and the Smithsonian Institution. The new initiative will emphasize human and societal dimensions, including digital interoperability; electronic information life cycles; integration of information, computing, and communications technologies with human needs; and new types of content and collections.

HuCS collaboratories will allow researchers to conduct large-scale modeling and simulation, access appropriate information, share access and operation of remote facilities, and work within virtual environments to visualize scientific data and configure and control experiments, regardless of geographic and temporal separation among individual members. Other HuCS R&D includes active visualization, disability and rehabilitation research, educational technologies, finding and tracking information, knowledge networks, manufacturing applications, virtual reality, and the Visible Human project.

The February 12, 1998, issue of *Business Week* magazine highlighted leading-edge R&D in speech technologies funded by DARPA, NIST, and NSF, including automatic speech recognition, speech synthesis, speaker identification and authentication, and natural language understanding. Additional NSF activities include STIMULATE (the Speech, Text, Image, and MULtimedia Advanced Technology Effort), which supports fundamental research in understanding multimodal human communication and its application to computer technology.

Education, Training, and Human Resources (ETHR)

ETHR R&D supports computer and communications-related research to advance education and training technologies at all levels including K-12, community college, technical school, trade school, university undergraduate and graduate, and lifelong learning. The complex and technically challenging applications flowing from leading edge R&D in HECC and LSN make it increasingly important for today's students and professionals to update their education and training on an ongoing basis in order to exploit the latest technological advances. ETHR technologies improve the quality of today's science and engineering education and lead to more knowledgeable and productive citizens and Federal employees.

FY 1998 and FY 1999 activities include new NSF centers for developing innovative learning technologies, NIH/NLM and NIH/NCRR training grants, and NASA's Web-based classroom training. Beginning in FY 1999, NSF's Knowledge and Distributed Intelligence (KDI) initiative aims to achieve the next generation of human capability to generate, gather, model, and represent complex and cross-disciplinary scientific data and transform this information into knowledge by combining, classifying, and analyzing it in new ways. DoD and its Air Force Office of Scientific Research participate in ETHR through learner-centered education and automated training activities, although neither agency is part of the ETHR budget crosscut. NSF, the Department of Education (ED), and Department of Labor (not currently a CIC agency) are addressing the need for training the workforce in information technology.

CIC R&D programs

The five Program Component Areas of the CIC R&D programs are a driving force in information technologies, computing, and communications and are a key 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. The estimated FY 1998 HECC and LSN budgets for the participating Federal organizations was $717.4 million. For FY 1999, the President requested $860.9 million for these organizations.

Federal Information Services and Applications Council (FISAC)

FISAC, which reports to the Subcommittee on CIC R&D, succeeds the Applications Council of the CCIC. FISAC fosters the migration of technology from the information technologies R&D community to government application missions and information services communities and identifies challenges from applications to the information technologies R&D community. Multiagency teams focused on Crisis Management, Federal Statistics, Next Generation Internet, Universal Access, and Information Security report to the FISAC.

Presidential Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet

The Presidential Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet (PAC) and its Subcommittees have met eight different times between February 1997, when the Advisory Committee was established, and May 1998. The PAC held a Town Hall meeting at SC97 in San Jose in November 1997. It plans to report its findings in a report to the President, through the Director of OSTP.

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

The NCO facilitates multiagency CIC R&D activities, such as the development of inter-agency CIC programs and the preparation of planning, budget, and assessment documents, and supports the activities of the Presidential Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet. The NCO Director, who reports to the Director of the Office of Science and Technology Policy (OSTP), Executive Office of the President, serves as the chair of the Subcommittee on CIC R&D.

This report highlights many of the vital, ongoing efforts of the CIC R&D programs, focusing on representative FY 1998 accomplishments, key FY 1999 research and development areas, and the budget crosscut.

This report, all other CIC R&D publications, and links to related Web sites can be found at *http://www.ccic.gov/* and *http://www.ngi.gov/*.

High End Computing and Computation

Overview

HECC research and development (R&D) investments provide the foundation for U.S. leadership in high end computing and promote the use of high end computing and computation in government, academia, industry, and in broad societal applications. HECC research explores algorithms for physical, chemical, and biological modeling and simulation in complex systems; information-intensive science and engineering applications; and the advanced concepts in quantum, biological, and optical computing that will keep the U.S. in the forefront of computing breakthroughs for years to come.

HECC R&D supports critical Federal government mission needs, including:

- ❐ National security
- ❐ Environment/climate/weather
- ❐ Aeronautics and space exploration
- ❐ Energy research (solar, combustion, fusion)

HECC efforts also promote broad societal applications, including:

- ❐ Healthcare
- ❐ Crisis management/natural disaster warnings
- ❐ Long-term environment and energy management
- ❐ Education and lifelong learning

Major accomplishments of the Federal investment in high end computing include:

- ❐ Establishing scalable parallel processing as the commercial standard for high performance computing
- ❐ Creating massively parallel systems as effective high end computing devices
- ❐ Creating the scientific base for high end computing, including trained scientists and engineers, new architectural approaches and next-generation technologies

❒ Enabling the technology base for the $2 billion middle range high performance market that expanded access to high performance computing while reducing costs to the Government

❒ Enabling the near-term computing technology for the Department of Energy's (DOE) Accelerated Strategic Computing Initiative (ASCI)

The next major advance is anticipated to be a sustained rate of one petaflops (pflops — 10^{15} floating-point operations per second). With appropriate funding, emphasis, and technological attention this milestone may be achieved by 2007, assuming that certain key technologies continue to progress at current rates.

Federal investments in four HECC areas, or thrusts, will enable development of distributed, computation-intensive applications to support future U.S. science and engineering research, national security priorities, and economic competitiveness.

Thrust 1: System software technology

The aim of this thrust is to improve the usability and effectiveness of teraflops-scale systems across a wide range of government, industry, and academic applications, concentrating on medium term technology development (three to five years). Thrust 1 activities address high end architectures, including symmetric multiprocessor systems (SMPs), clusters of SMPs, and a computational grid of distributed homogeneous and heterogeneous clustered systems. Longer-term activities will focus on the system software technology requirements of future generations of high end system architectures.

Thrust 1 efforts recognize that Government investments are required, since the market size for high end computing is not large enough for system vendors and independent software vendors (ISVs) to make significant investments in the Federal mission-critical system software technology required for computationally intensive applications. It is appropriate that as Government-developed system software and tools mature they become the property of non-government entities and shared resources throughout appropriate segments of the research community and industry.

Thrust 1 investment focus areas are:

❒ Operating systems and input/output (I/O)

❒ Languages and compilers

❒ Programming interfaces and libraries

❒ Debugging and performance tools

❒ Scientific visualization

❒ Data management

❒ Common framework and infrastructure

Thrust 2: Leading-edge research for future generations of computing

Driven by Federal, academic, and commercial applications, Thrust 2 focuses on long range research and technology development to achieve a sustained pflops (10^{15} floating-point operations per second) computational rate and exabyte (10^{18} byte) storage.

Thrust 2 activities will support U.S. global leadership in high end computing (HEC), ensuring that scientists and engineers, especially those working on Federal missions, will continue to have access to the most powerful computers, and assuring that the research and technology necessary for HEC systems will be available to U.S. industry. Activities requiring petaflops speed and exabyte storage include:

Biology:

- ❐ Simulations of complex biological systems (membranes, synthesis machinery, and DNA) and "post genome" connection of genome information to biological function

Business:

- ❐ Modeling of complex transportation, communication, and economic systems

The Earth and the environment:

- ❐ Modeling of integrated Earth systems (ocean, atmosphere, bio-geosphere)
- ❐ Comprehensive modeling of ground water and oil reservoirs for contamination and management
- ❐ Design of new chemical compounds and synthesis pathways for environmental safety and cost improvements
- ❐ Data assimilation and data fusion capabilities applied to remote sensing and environmental models for 4-D/6-D integration of information

Materials science and manufacturing:

- ❐ Materials simulations that bridge the gap between microscale and macroscale (bulk materials)
- ❐ Coupled electro-mechanical simulations of nano-scale structures (dynamics and mechanics of micro-machines)
- ❐ Full plant optimization for complex processes (chemical, manufacturing, and assembly problems)
- ❐ Complete engine simulation combining high-resolution reacting flow problems (combustion, chemical mixing, and multiphase flow) with mechanical and material properties

Physics:

- ❐ Multidisciplinary optimization problems in combining structures, fluids, and geometry
- ❐ Simulation of plasma fusion devices and basic physics for controlled fusion (to optimize the design of future reactors) combining all length and time scales from electrons to macroscopic turbulence
- ❐ Total design of new experimental facilities in high-energy physics — from beams to magnets to detectors to tunnels
- ❐ Modeling and simulation of complex physical processes related to DOE's ASCI program, to include areas such as the simulation of gas turbine engines, shock waves, astrophysical thermonuclear flashes, accidental fires and explosions, and solid propellant rockets

Thrust 3: Incorporating technology into applications

Many HECC agencies support mission-driven scientific applications requiring large scale computation-intensive and/or data-intensive operations spanning the full space-time spectrum of scientific problems. Thrust 3 R&D will incorporate the first use of HECC methods into agency applications, and encourage the use of computational science algorithms to solve problems requiring high performance computational facilities, ensuring that key applications execute at full potential.

Advances are needed in fast, efficient algorithms for computational science techniques that address problems in very large, sparse matrices, searching, sorting, and pattern matching. Research on algorithms with large amounts of concurrency, fault tolerance, and latency-hiding is crucial for high end computational systems of the future.

Thrust 4: Infrastructure for research in HECC

The Thrust 4 objective is to ensure a balanced R&D infrastructure with maximum computational strength and bandwidth. Interdependent with LSN activities, this thrust supports the research facilities built on large scale test systems and on large scale, high performance computational grids and networks.

Prudent agency investments in coordinated R&D areas will enable development of the distributed, computation-intensive, and data-intensive applications needed to assure future scientific, engineering, and economic competitiveness, and fulfill national security requirements. Select agency R&D efforts, organized by thrust area, include the following:

Thrust 1

System software technology — FY 1998 accomplishments and FY 1999 plans

Distributed computing

National Science Foundation (NSF) -supported work in distributed multimedia systems has as objectives the development of:

- ❐ Methods for providing provable Quality of Service (QoS)
- ❐ QoS guarantees for multimedia applications over heterogeneous networks

❒ Algorithms and protocols optimized for transmitting live and stored audio and video streams

❒ Routing and resource reservation algorithms

❒ Techniques for integrating traffic management

Time-variant QoS management combines research from real-time systems modeling, computer networks, and distributed systems. Brokerage procedures to negotiate and renegotiate required timing constraints, and distributed control mechanisms to satisfy time-variant QoS provision and adaptation during the transmission are being developed. This NSF-supported work is being performed at the University of Maryland, the University of Texas at Austin, and the University of Illinois at Urbana-Champaign.

Proof carrying code

NSF-funded researchers at Carnegie-Mellon University have developed methods and implementations that allow highly secure execution of code downloaded from untrusted sources. The problem of intrusion via imported programs has become significant in the Java/Web-based environment of much of today's computing. The new technique is based on downloading with the code a proof that the code satisfies desirable safety characteristics. Theorem-proving software on the receiving site checks the given proof against the code, assuring that the code indeed conforms. If either the code or the proof has been tampered with or otherwise corrupted, the proof will fail, making this method highly resistant to attack. This work capitalizes on research in programming language semantics, theorem proving, compiler techniques, and formal methods. Early products of this research are relatively efficient proof-checking mechanisms and a "certifying" compiler for a significant subset of C++, a compiler that annotates the compiled code with predicates that it guarantees are satisfied by the object.

Scalable software and ease of use

Defense Advanced Research Projects Agency (DARPA) -supported projects develop scalable software geared toward easing the use of systems by application programmers. In FY 1998, DARPA will:

❒ Release HPC++ languages and runtime services supporting both task and data parallelism

❒ Develop languages and runtime services supporting parallel applications such as advanced distributed simulation

❒ Demonstrate experimental versions of new iterative solvers for radar cross-section modeling

In FY 1999, DARPA plans to:

❒ Release a prototype subsystem supporting adaptive resource allocation and consumption in response to changing workload and resource availability

❒ Demonstrate an experimental scalable image processing application using embedded systems

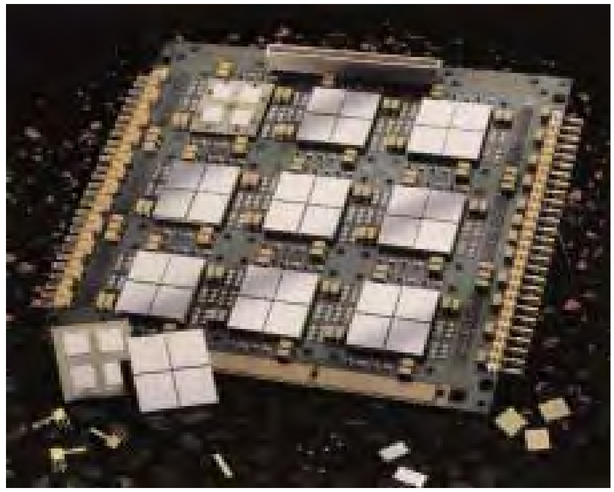

The SOLITAIRE board contains a complete SGI/CRAY J90 System Element, which includes four vector processors, network crossbar, I/O, and 1/2 Gbyte of dynamic random access memory (DRAM). There are nine diamond based die cast multichip modules (MCMs) mounted on one side of the board and eight die cast MCMs mounted on the reverse side. SOLITAIRE represents a new achievement in double sided board column grid attachment, featuring over 25,000 connections between the modules and the board. This repackaging implementation of the System Element is 80 percent smaller and 75 percent lighter than the commercial product, and concentrates 500 Watts onto a 10 inch square board.

Operating system extensions

In FY 1999, DARPA plans to design instruction set extensions and storage components to allow Defense applications to specify whether operations are executed in the central processor or in logic circuits embedded in the memory hierarchy.

High performance software tools

DOE computer science and software tools work focuses on software to facilitate the use of high performance systems for solving scientific problems. Efforts range from efficient operating systems and I/O for massively parallel processing (MPP), to frameworks for isolating application codes from the underlying hardware details, to tools for monitoring the performance of scientific applications, and to efforts to improve the management, visualization, and understanding of the results of high end computations.

DOE projects include software tools for program development environments and performance monitoring for Advanced Computational Testing and Simulation (ACTS). The goal is to develop an integrated set of algorithms, software tools, and infrastructure that will enable computer simulation to be used in place of experiments when real experiments are too dangerous, expensive, inaccessible, or not politically feasible.

In FY 1998, DOE will:

- ❒ Develop scripting interface languages for object-oriented numerics
- ❒ Complete a parallel program analysis framework for the ACTS Toolkit
- ❒ Develop a high performance software component architecture for the Scientific Template Library (SciTL)
- ❒ Extend SciTL with a shared-memory access programming model

In FY 1999, DOE will extend SciTL for wide area execution to support computational grid applications.

Earth and Space Sciences (ESS)

The FY 1998 goal of NASA's ESS scalable testbed is 50 gigaflops sustained. For this effort, NASA will develop pre-competitive prototype systems software that provides high availability and portability demonstrated in a large scale production environment with the objective of eventual commercial availability. In FY 1998, the goal of NASA's 100-250 gigaflops sustained scalable testbed is 100 gigaflops sustained.

MARQUISE

In FY 1999, NSA will flight test SOLITAIRE, the embedded high performance computer prototype developed under the MARQUISE program, on an Air Force aircraft. The agency will continue research on miniaturized, spray cooled, and embedded diamond power supplies. In FY 1999, NSA expects to demonstrate a 48-Volt to 2.8-Volt @ 200 Amps high-efficiency power converter with a power density of 2 Kwatts/inch^3.

NSA continues to conduct research in embedded scalable nodes (the follow-on architecture to MARQUISE) for mission scenarios, and has selected three potential repackaging implementations that could make a 100 percent software-compatible, miniaturized high performance computer available six months after the introduction of the commercial computer.

Microelectronics in high speed computing

Related NSA research areas include synthetic diamond packaging technology, all-optical switching, and optoelectronic integrated circuit packaging technology. Other activities include:

- ❒ Point-of-use power conversion for power reduction
- ❒ Area array I/O design studies for low power implementations of high performance MCM
- ❒ Very high level programmable accelerator plug-ins for standard architectures

Pictured to the right is a 128x128 superconductive crossbar switch. The switch matrix is mounted on the 4"-diameter MCM and maintained at 4 degrees Kelvin temperature. It is connected to its room-temperature environment by an eight multichannel electrical cable.

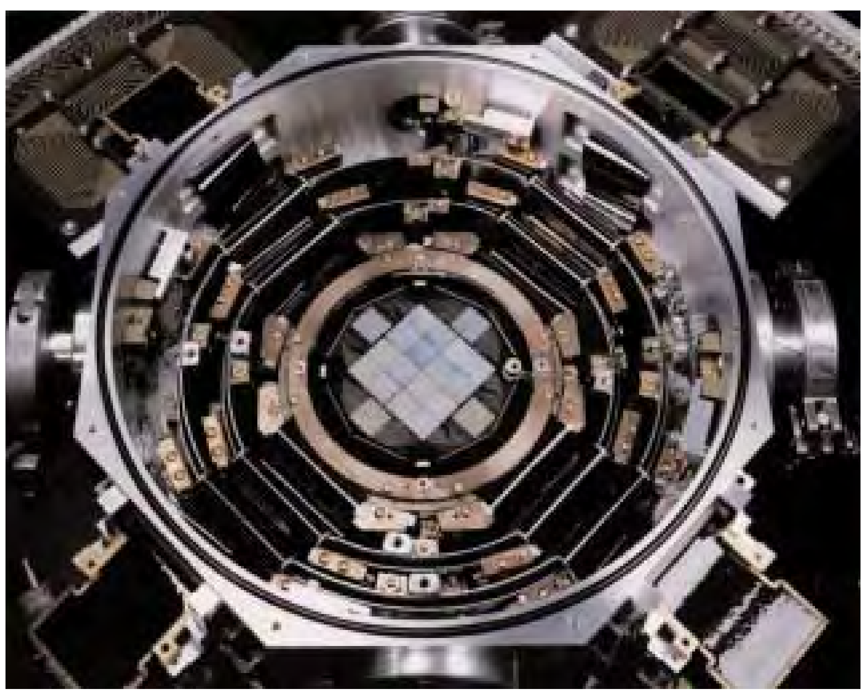

A spray cooled dc/dc converter can supply an impressive 200 W/in^3. This magnitude of miniaturization allows tight integration of the power converter and electronics in a shared environment. Point of load conversion eliminates heavy bus bars and rigid cables by distributing power at high voltage and low current. Spray cooling results in relaxed device specifications and improved reliability due to reduced operating junction temperatures and thermal gradient. A 28 vdc/4 vdc spray cooled converter is pictured to the right.

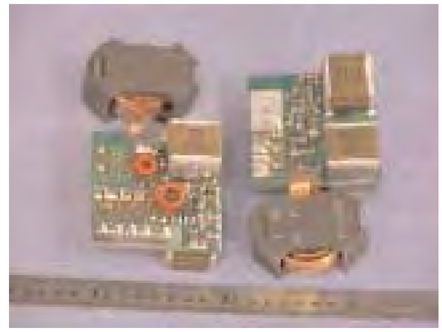

UPC

NSA's Center for Computing Sciences (CCS), in partnership with the University of California at Berkeley and Lawrence Livermore National Laboratory, is developing UPC, a language that combines features of the CCS-developed AC with features of Split-C and PCP, two other parallel C languages. UPC represents an attempt to develop a base language that could become a standard for explicitly-parallel C. UPC extends AC in three ways:

- ❒ It adds keywords that enable programmers to control memory consistency to balance safety and performance.
- ❒ It enables blocking of arrays across processors to improve the performance and ease of programming for select applications.
- ❒ It introduces a "forall" concept to handle load balancing in loops that can be executed in parallel.

Quantum computing

In FY 1998 and 1999, NSA will continue fundamental exploratory studies in quantum computation theory and experimentation, including algorithms and complexity theory, quantum error correction techniques, quantum decoherence processes, and small scale laboratory implementations.

Performance measurement and testing technologies

NIST continues developing application-level measurement techniques for local area network (LAN) -clustered workstations, including physical performance monitors and software tools. These include the acquisition of performance information from the operating system, communication libraries, and their associated daemons, all of which have an impact on concurrent applications, but today provide little, if any, feedback to the user. Since the kernel and daemons operate in parallel with the application, data acquisition and resolution are more difficult. NIST will investigate the Linux kernel and message passing libraries such as MPI to provide performance information and develop techniques that allow users to acquire information selectively. Other work includes investigations into optical tape subsystem designs and testing methods for digital video disk (DVD) media.

Parallelizing environmental modeling software

A framework for parallelizing a broad range of environmental modeling codes has been developed by EPA-supported researchers at Iowa State University. The key feature is class-specific automatic parallelization where high-level knowledge about the specific class is used to arrive at efficient parallel codes. This approach combines expert system and compiler technologies to automate tedious and time-consuming aspects of hand coding. The current system (PA-1) can parallelize 3-D time matching explicit finite difference codes. One person used PA-1 to parallelize the Mesoscale Meteorology model (MM5) in two weeks, as compared to three years for manual parallelization. A scalable 2-D Regional Acid Deposition Model was parallelized in three weeks versus 10 months for manual parallelization of a simpler one-dimensional model.

HPC modernization

Although not formally part of the CIC R&D crosscut, the Department of Defense (DoD) HPC Modernization Program, which is represented on the HECC Working Group, will provide advanced hardware, computing tools, and training to DoD researchers using the latest technology to support the warfighter. Incorporating high performance computing into the system design process will allow the U.S. to maintain its technological supremacy in weapons systems design into the foreseeable future, and the use of high performance computing in the early stages of system acquisition will decrease the total life-cycle costs of new warfighting support systems.

Common High Performance Computing Software Support Initiative (CHSSI)

Also part of the DoD HPC Modernization Program, CHSSI is an application software development program designed to provide DoD computational scientists and engineers with software for scalable computing systems. These products will improve DoD science and technology and developmental test and evaluation computation. Core applications focus on improved 21st century military capability.

Thrust 2

Leading edge research for future generations computing — FY 1998 accomplishments and FY 1999 plans

Computational models

In computational models, DARPA will conduct R&D in new classes of computing technologies that may offer performance/cost/size/weight/power improvements beyond the limits of today's semiconductor-based computing.

Intelligent systems technology

The focus of DARPA's intelligent systems technology effort is on advanced techniques for knowledge representation, reasoning, and machine learning, which enable computer understanding of spoken and written language and images. Intelligent systems technology research includes:

- Human computer interaction technology focusing on design methods and enabling technology for more natural interaction between people and computers
- Prototype robust spoken and text language technologies with emphasis on affordable dialog grammars and understanding (FY 1998)
- Human-computer dialogue interaction for crisis planning and automatic transcription of conversational speech over battlefield radio (FY 1999)
- Microelectronic science focusing on circuitry and software that will enable highly configurable computational and storage elements
- Advanced methods for planning, scheduling, and resource allocation

Scalable computing

This R&D focuses on hardware and software technologies leading to a secure scalable computing and communications technology base for systems configured over a wide performance range, from mobile handheld devices, to desktop workstations, to large scale distributed systems. In FY 1998, DARPA will demonstrate highly efficient, parallel nodes; auto-parallelization of file I/O for scalable systems; first node-level performance of ultra-low-power systems; performance of novel backplane networks supporting security; and hardware-accelerated, distributed, shared-memory performance on workstation clusters.

Ultrascale computing

In FY 1998, DARPA will:

- ❒ Design, model, and assess both the quantum-to-Si hardware and software interface and language for expressing amorphous algorithmic computations
- ❒ Demonstrate a 256-component addressed array of molecular computational mechanisms and evaluate surface patterning mechanisms for culturing neural components on silicon
- ❒ Investigate computational models suitable for implementation using ultra scale computing techniques

In FY 1999 DARPA will:

- ❒ Demonstrate and validate ultra scale computing models, with emphasis on DNA-based logic operations, cell-based computation, and novel communication pathways, and the scalability of these techniques in Defense applications
- ❒ Conduct a system-level design and simulation study of a computation model based on large amorphous arrays, demonstrate a prototype array with more than 1,000 elements, and establish the role of nuclear magnetic resonance technologies in developing a quantum computing research medium

Global scale system software

In FY 1998, DARPA will:

- ❒ Develop quality-of-service negotiation protocols and adaptive resource discovery protocols
- ❒ Demonstrate coarse- and fine-grained performance-driven resource allocation mechanisms, achieving performance between 30 and 50 percent of optimal

In FY 1999, DARPA plans to focus on integrating a multi-attribute quality of service specification language architecture and demonstrating path-based propagation of quality of service constraints across layer and network boundaries.

Information survivability

DARPA research in information survivability will focus on:

- ❒ Developing an architecture for low-power configurable computational elements

- ❒ Demonstrating real-time adaptive control and resource management
- ❒ Releasing Defense-critical software based on scalable library technologies

In FY 1999, DARPA plans to:

- ❒ Investigate novel control mechanisms for self-organizing and autonomous systems
- ❒ Validate low-power configurable architecture, develop supporting software, and demonstrate automated mapping of 500,000 elements

Multi-threaded architecture (MTA)

In FY 1998, as part of its work in evaluating prototype computing systems, DOE will evaluate the performance of a suite of algorithms and DOE-critical applications on the first Tera Computer system, which is used in evaluating the MTA that is critical to all proposed petaflops architectures. In FY 1999, DOE will evaluate software requirements for petaflops systems and new architectural developments.

Non-binary optical storage systems

Today's optical disks represent data in a binary format — zeroes and ones. NSF-funded researchers at the Georgia Institute of Technology are developing optical systems to store data in "M-ary" format, where M can be larger than two. Working with Optex Communication Corporation, researchers helped design the signal processing components of an optical memory in which each stored symbol can take on one of M=6 different values, resulting in greater information density per square inch in optical recording. This research is funded by a Presidential Early Career Award for Science and Engineering (PECASE).

Multi-user detection in CDMA communication systems

Code division multiple access (CDMA) permits multiple users to be assigned the same frequency in a digital cellular communication system: each user gets a unique signature sequence that can be recognized even when there is interference from other users. In this NSF-funded work, multi-user detection is a new approach to implementing CDMA that exploits the fact that the base station knows the signature sequences for all users and can use that knowledge to better combat interference.

Optimal multi-user detection is prohibitively complex to implement today. However, NSF-funded researchers at the New Jersey Institute of Technology, North Carolina State University, Ohio State University, Rice University, and the University of Colorado are developing low-complexity algorithms and associated hardware. The benefit will be the ability to handle more users per cell at higher data rates.

Knowledge Networking

NSF will continue to support a broad academic research program in computing systems, including "Knowledge Networking," an initiative focused on the next generation of interconnected networks and associated database and collaborative technologies.

HTMT

NSA's Hybrid Technology Multithreaded (HTMT) architecture effort is a collaborative project among a dozen research groups (Caltech/JPL, University of Delaware, SUNY Stony Brook, University of Notre Dame, Princeton University plus an association of other government and industry labs). The objective is to define a very high performance computer system that can reach a petaflop level of performance in significantly less time than a conventional approach. HTMT will be based on a unique multithreaded execution model and will use a combination of leading-edge technologies, including superconducting technology, high-speed VLSI semiconductor technology, optical interconnect, and storage technology. The system is expected to be used in several critical high performance applications of national strategic importance over the next decade.

An initial study of the original HTMT concept was funded by NSF and NASA in 1996-1997. A comprehensive design study of HTMT is now sponsored by DARPA, NSA, and NASA, and is managed through Caltech and JPL. The research being performed at University of Delaware is focused on the program execution and architecture model of HTMT.

High speed circuits

NSA is conducting research to determine whether electronics can handle the 100 Gb/s serial data rates of fiber optic links. Uses include Tb/sec data transfer.

RES, JACKAL, and LOTS

RES, a system developed in 1991-1992 with NSA HPCC funds, was designed to harness "idle" workstations. It is seeing increasingly widespread use and is the subject of continuing enhancement.

JACKAL is a cliche-based software reverse-engineering research system (where a "cliche" is an "interesting" piece of code). Input to JACKAL is a high-level abstract language translation of source code (when available) or decompiled code. JACKAL uses tree-matching and string-matching algorithms to identify cliches.

Archival store for I/O is a major element in all computing systems. NSA is developing LOTS, an optical tape system with ten times the storage per tape and at least ten times the input data rate compared with conventional large stores.

Thrust 3

Incorporation of HECC technologies into real applications — FY 1998 accomplishments and FY 1999 plans

High performance simulation of multiphysics problems

NSF will support the simulation of physically realizable systems, giving physical scientists the ability to predict properties and function from fundamental structure. R&D in this area could lead, for example, to new lightweight, high strength biomimetic materials, or could enable the design of chemicals to immobilize viruses.

This NSF-funded collaborative effort involves computer science and computational mathematics researchers at the University of Colorado-Boulder and the University of Colorado-Denver. The engineering activities are driven by the simulation of problems encountered in aerospace, civil, and mechanical engineering. The problem areas selected include at least two mechanical fields; a common basis in modeling, problem decomposition, and parallel simulation methods; a need for extensive scientific visualization; and the requirement for a common overarching software/hardware infrastructure. The research in computational mathematics will focus on the theoretical foundation of multiphysics simulation and the design of new algorithms in interface conditions, fast nonlinear solvers, and staggered multi-field timestepping. The computer science research will focus on the development of a Problem Solving Environment (PSE) for applications problems that is suitable for the parallel implementation of large scale multiphysics simulations. The primary goal is to provide transparent, portable, and scalable access to all of the resources needed for executing large parallel applications on heterogeneous and/or distributed systems, including interprocessor communication, visualization, debugging, and performance monitoring.

Multi-model multi-domain computational methods

A collaboration among Old Dominion University, University of Colorado-Boulder, Notre Dame University, New York University (NYU)-Courant Institute, Boeing Company, and Argonne National Laboratory, this NSF-funded project will develop a software laboratory for aerodynamics and acoustics problems through domain decomposition methods of task-parallel (involving multiple models, such as a viscous boundary layer and an inviscid free stream) and data-parallel (involving multiple subdomains, for additional concurrency) types. The project unifies several pairwise collaborations among numerical analysts, parallel algorithmicists, engineers, and computer scientists into a vertically integrated demonstration of a Grand Challenge computational philosophy; namely, that the solution to Grand Challenge problems is most feasible and cost-effective when software permits users to conveniently adapt the fidelity of the computational model and the strength of the local solver from region to region within an overall implicit algorithmic framework.

Smart antennas

The explosion of the cellular communications industry between 1995 and 1997 prompted communications companies to explore the use of antenna arrays at cellular base stations as a means to improve both the channel capacity and quality of wireless communications — cellular telephony, personal communications services (PCS), and wireless local area networks — by exploiting the enhanced interference rejection capability they provide. It is predicted that by the year 2005, all cellular and PCS base stations will employ some form of smart antenna technology, and many companies will demonstrate working systems. An example is ArrayComm's IntelliCell Base Station Array, which demonstrated uplink and downlink capacity improvements. NSF funded the development of this array and ultimately the creation of this company. ArrayComm provides hundreds of such arrays for Japan's personal communication systems.

The world market for base station antennas for digital wireless applications is projected to grow at an unprecedented rate. Cellular carriers in the U.S. may need 15,000 new cell sites over the next decade to upgrade their services and meet anticipated demand. PCS services may require an additional 100,000 sites.

NSF continues to fund work in this area at Brigham Young University, the University of Texas, and Purdue University. Researchers have built smart antenna systems to test uplink/downlink scenarios in real environments. Further experiments include channel measurements in troublesome

In these snapshots of simulated neutron star mergers, the lighter color signifies higher matter density and the darker color, lower density. Embodying some gravitational effects, simulations show that gravitational radiation losses influence instabilities during the merger. Only full relativistic calculations, now underway through NASA-supported R&D, can predict the gravitational wave signal and probe the possibility of resulting black holes.

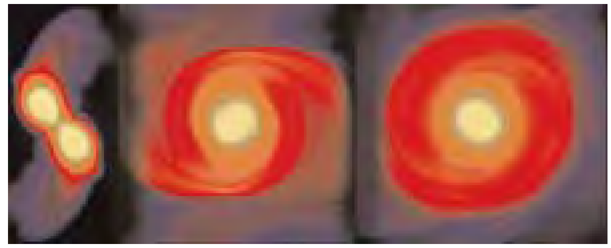

environments, such as highly reverberant areas, in order to determine when and where fades occur, and to determine the means to combat them electronically. The goal is to uncover performance gaps between existing commercial systems and theoretical optimal performance and improve compression algorithms.

Geometric shape analysis applied to molecular biology

The "alpha-shape" approach is used to describe geometric surfaces and volumes. In an NSF-funded multi-year computational geometry project at the University of Illinois, a molecular biology research group has applied the alpha-shape approach to study building blocks of light harvesting proteins. The approach is proving uniquely able to identify and simplify protein shapes without losing the topological characteristics of surfaces involved in aggregation.

NASA advanced supercomputing applications

NASA supports R&D in advanced supercomputing applications. Research areas include the following:

❒ Three-dimensional spherical simulations of the Earth's core and mantle dynamics, with Johns Hopkins University. This effort will simulate the chaotic processes that drive the evolution of the planet's interior, and in turn shape its surface, over timescales ranging from hundreds to millions of years.

❒ Advanced computing technology applications to Synthetic Aperture Radar (SAR) interferometry and imaging science, with Jet Propulsion Laboratory (JPL). Using multiple supercomputers to process and visualize satellite-collected SAR data will allow close monitoring of regional changes in alpine glaciers, plate tectonics, and rain forests.

❒ Four dimensional data assimilation — investigation of high performance computing and current algorithms, conducted at NASA's Goddard Space Flight Center and the University of Maryland. The focus is on melding observations and climate model prediction into a robust data analysis scheme for NASA's Earth Observing System to provide accurate pictures of the atmosphere through space and time.

❒ Development of an Earth system model — atmosphere/ocean dynamics and tracer chemistry, with the University of California, Los Angeles. Aimed at realistic portrayal of the Earth's climate, this effort will develop and couple four highly complex models with high spatial resolutions: atmospheric general circulation, oceanic general circulation, atmospheric chemistry, and oceanic chemistry.

- ❒ Scalable parallel finite-element computations of Rayleigh-Benard-Marangoni problems in a microgravity environment, with the University of Texas, Austin. Modeling of fluid flows in low gravity environments will test the effectiveness of manufacturing higher quality thin films and coating processes in space and the functioning of the space station’s life support and safety systems.
- ❒ Turbulent convection and dynamos in stars, with the University of Chicago. This research will probe fundamental and little-understood turbulent processes in the interior of stars like the sun.
- ❒ Understanding solar activity and heliospheric dynamics, with the Naval Research Laboratory. Using NASA observations, the tangled three-dimensional structures that develop in the magnetic field of the Sun’s corona, or outermost layer, will be modeled.
- ❒ Parallel adaptive methods for multiscale modeling of the heliosphere, a multipurpose three-dimensional code for relativistic astrophysics and gravitational wave astronomy, with the University of Michigan. From the corona to the free-streaming interstellar medium, computational studies will be used to understand the interaction of the solar wind with galactic gases and plasmas, as well as with magnetized and unmagnetized bodies in the solar system.
- ❒ A multipurpose three-dimensional code for relativistic astrophysics and gravitational wave astronomy — application to coalescing, or merging, neutron stars, with the University of Illinois, Urbana-Champaign. This project will combine fluid dynamics and general relativity to develop computational methods to study the merger of two neutron stars, a process that encompasses many aspects of relativistic astrophysics and thus provides a basis for studying similar phenomena such as black holes and supernovae.

Computational Aerosciences (CAS)

CAS supports NASA aeronautics programs and is driven by the needs of the aeronautics industry. Its goal is to accelerate the development and availability of high performance computing technology to the U.S. aerospace community, to facilitate the adoption and use of this technology by the U.S. aerospace industry, and to hasten the emergence of commercial hardware and software markets.

The CAS project is developing computational techniques and applications to solve Grand Challenge problems, which include multidisciplinary simulations of complete aircraft throughout their flight envelopes. These efforts will be structured with near-term, intermediate, and long-term

NASA is involved in ongoing research on high speed civil transport and high performance aircraft. Images generated by high performance computing technology, such as the one pictured to the right, illustrate the wake systems associated with each blade of an aircraft engine. Such illustrations help researchers to better understand blade-vortex interaction.

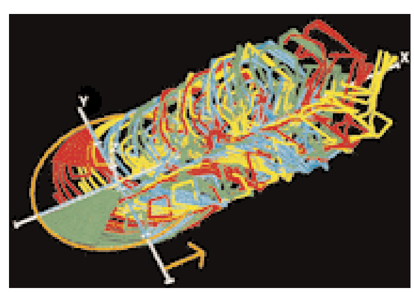

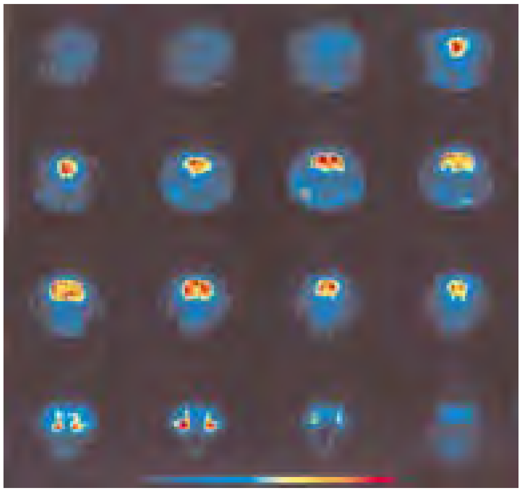

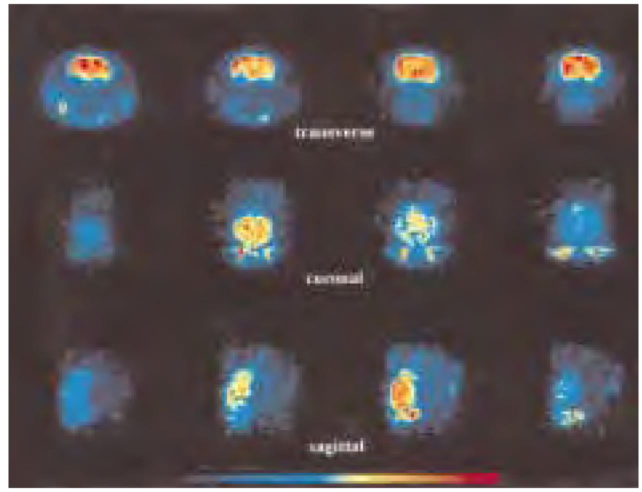

The image on the left consists of transverse sections of a 3-D positron emission tomography (PET) image reconstructed using a maximum likelihood (ML) algorithm. In this NIH study, a rat was injected with the fluorine-18 analogue of glucose (18F-2-deoxy-D-glucose) and subjected to visual stimuli. High-resolution cortical and subcortical activations are visible. The images on the right are transverse, coronal, and sagittal cross sections of a 3-D PET image reconstructed using an ML algorithm.

milestones that, at each stage, provide end-user benefits. These milestones will be pursued in close cooperation with academia and the aerospace industry.

CAS works with the Numerical Aerodynamic Simulation (NAS) program and the NASA Research and Education Network (NREN) to ensure that the aerospace community has access to the NREN and the CAS testbeds. The CAS project is also encouraging joint development of new metrics that address affordability and cycle time. Research on high speed civil transport (airframe and engine) and high performance aircraft (airframe and simple engine) is ongoing.

DOE Grand Challenge applications

In support of its Grand Challenge applications, DOE R&D covers topics that include raw computational efficiency in quantum chromodynamics software, massive data requirements of relativistic heavy ion collider experiments, Accelerated Strategic Computing Initiative (ASCI) applications, and visualization for global change modeling. DOE also conducts R&D in collaborative tools and information surety techniques. In FY 1998, DOE will coordinate its DOE 2000, Grand Challenge, ASCI, and ACTS R&D.

Molecular structure prediction and simulations

In FY 1998, NIH organizations National Library of Medicine (NLM), National Center for Research Resources (NCRR), Division of Computer Research and Technology (DCRT), National Cancer Institute (NCI), and National Institute of General Medical Sciences (NIGMS) will refine new methods for ab initio structure prediction for use in the pharmaceutical industry. The Institutes will improve computational technologies for larger simulations of protein, DNA, and membrane complexes in water environments. (*Note:* As of February, 1998, DCRT's activities became part of NIH's Center for Information Technology— [CIT].)

NCRR will provide parallel computing applications in biochemistry, molecular biology, and cellular biology with access to supercomputing systems, large databases, and other resources through Web-based browsers. Thus, complex simulations related to receptor sites and other drug design R&D can be conducted over the Internet. NCRR will integrate 3-D graphics software, software tools for magnetic resonance spectroscopy data analysis and molecular structure determination, and other software tools to provide new capabilities for structure-based drug design. CIT will improve methods for predicting protein-drug binding energies, and will develop high performance computing methods for biomedical applications.

NCI provides state of the art capabilities in a fully integrated high performance computing center, applies high performance parallel computing and communication methods to biomedical applications, and evaluates new scalable parallel architectures for biomedical applications.

Reconstruction of 3-D positron emission tomography images

In positron emission tomography (PET), images are reconstructed from a set of projected measurements in what is classically known as the inverse problem. While advances in detector technology have enabled researchers to use these techniques to study small animals, the data exhibit poor counting statistics. The quality of the reconstructed image can be improved by incorporating a statistical model of the scanning process into the reconstruction algorithm. In particular, reconstruction based on the maximum likelihood (ML) criterion show both reduced noise and improved resolution.

These methods are being used to study the phenotypical consequences of genomic manipulation in mice and rats. The Nuclear Medicine Department at NIH is developing a high resolution small animal PET imaging system using new scintillation crystal technology. Image sensitivity is being improved both through acquiring 3-D data and using ML-based reconstruction. CIT has developed 3-D algorithms for these scanners. These algorithms exploit the inherent sparsity and symmetries in the matrix that

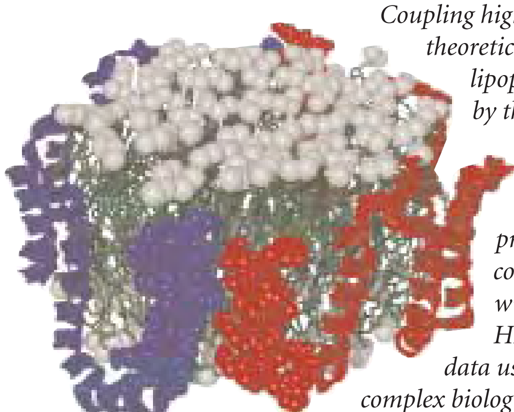

Coupling high performance computing systems with visualization tools permits theoretical investigations of complex biomedical systems. The figure shows a lipoprotein complex model with the lipids (green) in the center surrounded by the proteins (blue and red). High-density lipoprotein (HDL) circulates in the bloodstream, extracting cholesterol from body tissues and transporting it to the liver for excretion or recycling. Increased levels of HDL have been correlated with a decreased risk of arteriosclerosis, a primary cause of cardiovascular disease. NIH-supported researchers combined experimental evidence about the structure of HDL particles with protein structure predictions to produce a preliminary model of HDL. Predicting 3-D protein structure from one-dimensional sequence data uses high performance parallel computing techniques. Understanding complex biological functions requires determining the structure and properties of large molecular complexes.

Pictured at the right is a conformation of prion protein as determined through magnetic resonance spectroscopy. Prions are a novel class of "infectious" pathogens distinct from viruses with respect to both their structure and the neurodegenerative diseases that they cause. Prion diseases are manifested as sporadic, inherited, and infectious disorders including scrapie and bovine spongiform encephalopathy (mad cow disease) of animals as well as kuru, Creutzfeldt-Jakob Disease (CJD), and fatal familial insomnia of humans.

Prion protein (PrP) is the major, if not the only, component of prions. PrP exists in two isoforms: the normal cellular form (PrPC) shown here and the abnormal disease (scrapie)-related form (PrPSc).

models the scanner. The reconstruction software also decomposes the problem in parallel and takes advantage of high bandwidth I/O in CIT's IBM SP2 parallel computing system. Early visual stimulus studies revealed unprecedented cortical and subcortical activity in mouse and rat brains.

CIT plans to complete this work in FY 1998. The Nuclear Medicine Department will attempt to achieve the once unattainable goal of 0.5-mm resolution with less than 10 percent coefficient of variation for this new-generation small animal scanning system.

In FY 1999, CIT will continue developing methods and algorithms for biomedical applications that can benefit from computational speedup, such as image processing of electron micrographs, radiation treatment planning, medical imaging, and protein and nucleic acid sequence analysis.

Image management and communications systems

CIT, NCI, and the NIH Clinical Center are implementing a prototype high-speed telemedicine network to support multimedia communication for medical research and education. The prototype will include two Asynchronous Transfer Mode (ATM) switches, with 16 ports each, to support multimedia communication at 155 Mb/s. Washington University in St. Louis has developed an ATM-based multimedia consultation environment that includes prototype medical workstations. This network environment will initially support high performance radiation therapy planning.

Radiology Consultation Workstation (RCWS)

RCWS is a specialized imaging workstation for real-time telemedicine consultation and for medical education. Designed to provide high-resolution display of medical images, the RCWS allows remote consultations between radiologists and other medical specialists over the ATM network described above. A shared cursor permits collaborators to use outlining to identify organs and lesions. The identified regions of interest are transmitted in real time to all other RCWS sites participating in the consultation to plan a radiotherapy treatment. The CIT IBM SP2 parallel computing system will be used for the computationally intensive radiation therapy planning calculations.

RCWS goals include connecting the NIH prototype ATM telemedicine network to Washington, DC-area military medical centers and to Washington University in St. Louis. To this end, NIH is connected to the Advanced Technology Development Network (ATDNet) — begun by DARPA and now funded by several Federal agencies — that circles the Washington, DC metropolitan area. NIH will use NASA's Advanced Communication Technology Satellite (ACTS), which connects to ATDNet at the NASA Goddard Space Flight Center (GSFC), and provides connectivity to Kansas City, MO.

In FY 1998, additional RCWS nodes will be installed at Walter Reed Army Medical Center and the National Naval Medical Center, to allow multimedia telemedicine consultations between those institutions and NCI. Image communication software will be developed to allow the interchange of computerized tomography (CT) and magnetic resonance images (MRI) between the RCWS and select commercial radiotherapy treatment planning systems. In FY 1998, cardiology applications will be the second medical specialty for which the RCWS is adapted.

Environmental protection

The goal of EPA's program is to accelerate the evolution of HECC technologies to meet environmental protection mission objectives. EPA R&D in problem solving environments for multi-discipline environmental modeling builds on earlier achievements in technology development and integration for an air quality modeling framework.

EPA HECC R&D develops scalable computational algorithms and integrates HECC technologies developed by other agencies to explore the feasibility of new paradigms for community problem solving environments and for multi-discipline, multi-pollutant, multi-scale environmental exposure assessment. EPA fosters interagency collaboration on development and testing of HECC technologies to provide a foundation upon which the scientific and HPC communities can build, component by component, complex high performance environmental management tools. A collaborative approach can leverage scientific and technology advancements of other Federal agencies, academia, and research institutions to more rapidly evolve a unified, comprehensive approach to multi-discipline environmental modeling.

EPA's scalable software libraries

EPA-supported projects develop scalable software libraries, cross-platform distributed data management and computing methods, multivariate visualization techniques including high performance geospatial analysis, and modular software for an integrated problem solving framework to simplify community development and use of complex multidiscipline ecosystem and human exposure models. The framework technology facilitates interdisciplinary coordination, data and code sharing, and rapid infusion of new technology and science into integrated environmental problem solving tools.

In FY 1998, EPA will:

- ❐ Modularize 3-D hydrodynamics and water quality models for parallel solvers and reusable environmental process modules

New developments in EPA-supported analytical modeling are extending regional scale representations of porous media aquifers. The model to the right illustrates axisymmetric flow to a partially penetrating well in a stratified aquifer.

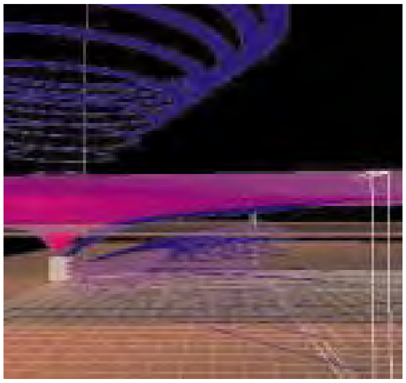

- ❒ Explore component-based model construction of tree/forest growth and succession suitable for parallel processing
- ❒ Develop parallel algorithms for clustering raster-based data and performing arithmetic operations on the clusters
- ❒ Demonstrate a parallelization agent to support legacy environmental modeling codes

Environmental modeling

EPA's Models-3 is a flexible software system for developing and using environmental assessment and decision support tools. The Community Multiscale Air Quality model has been integrated into Models-3 for urban to regional scale air quality simulation of ground level ozone, acid deposition, visibility, and fine particulates. Components of Models-3 assist in design and preparation of source emission inventories compatible with a variety of air quality modeling capabilities. Models-3 is a community framework for continual advancement and use of environmental assessment tools.

The Models-3 framework provides an interface between the user and operational models, between the scientist and developing models, and between the hardware and software, thus enhancing the user's ability to perform a wide range of environmental management tasks, from regulatory and policy analysis to understanding the interactions of atmospheric chemistry and physics. The framework uses specialized object libraries and a layered design that isolates critical system components, thus minimizing the impact of hardware and software upgrades. A client-server architecture in conjunction with a standardized data interface and object-oriented database containing metadata enables transparent use of scalable computing platforms and access to data across the network. The object data base also contains application-specific globally shared data, such as model domain, map projections, grid resolution, and chemical species, that enable the interchange of science codes while maintaining user control of model specifics. A library-based graphical user interface facilitates easy model executions and access to a variety of visualization and analysis packages.

In FY 1999, EPA will:

- ❒ Develop a framework supervisor to synchronize the exchange of instructions and data among parallel tasks with different spatial and temporal resolutions

- ❒ Develop techniques for visually comparing multispectral and multisource satellite data with model prediction data
- ❒ Develop a prototype temporal geographic information system for multidimensional, multiscale representation of space/time environmental data
- ❒ Demonstrate the use of ad hoc parallel algorithms for processing intersections associated with adaptive grid techniques
- ❒ Demonstrate a parallel high performance spatial analysis kernel
- ❒ Develop parallel data assimilation linking meteorology and hydrology models

Aquifer modeling

In FY 1998, new developments in analytical modeling are extending regional scale representations of porous media aquifers. This EPA-funded work is based on the superposition of many closed form analytic solutions representing hydrologic features, such as point sinks for wells, line-sinks for rivers, and areal elements for variable recharge. The solutions contain degrees of freedom that allow them to be chosen and combined so that boundary conditions are met along certain internal boundaries. Cooperative research with Indiana University exploring the implementation of 3-D well elements in parallel processing computers has resulted in public domain analytic element research code (ModAEM) available for use on scalable parallel computers and personal computers. The University of Minnesota is implementing variable density flow into analytic element representations of coastal aquifers. A case study is simulating the fresh water to salt water transition in the Upper Chesapeake aquifer.

Large scale environmental flow and transport

EPA-funded University of Kentucky researchers have simulated turbulent flows by combining the least-squares finite element method (LSFEM) with large eddy simulation (LES). This has been demonstrated in 3-D shear flow problems and natural convection problems, and has reduced the single-processor computational burden by 75 percent for large shear flow problems containing up to one million finite elements with over 7 million unknowns. This robust method enables more efficient solution of larger fluid flow and transport problems in turbulence prediction and control, meteorology, climatic changes, air pollution, water pollution, and soil contamination.

Online reference data for computational science

In FY 1998, NIST will evolve mathematical software repositories into problem-solving environments, complete object-oriented libraries for basic linear algebra and related capabilities, and demonstrate capabilities on distributed systems.

NIST continues to develop network-based resources that will bring more of the process of doing science online. For example, the Matrix Market (*http://math.nist.gov/MatrixMarket/*) provides researchers and software developers with a collection of large sparse matrices from industrial applications that can be used in comparative studies of linear algebra algorithms and software. Such problems form the core of many large scale applications. NIST is revising and updating the classic NBS (National Bureau of Standards, NIST's predecessor) Handbook of Mathematical Functions to be the basis for an interactive online library of mathematical reference data. It will include extensive formulas, graphics, and tables suitable for downloading into computer algebra systems and word processors.

High performance linear algebra software

NIST continues work in software design strategies that promote usability and maintainability, as well as high performance, for core numerical linear algebra operations. For example, NIST is working with the Basic Linear Algebra Subprograms (BLAS) Technical Forum, an industrial/government/academic consortium, to specify a standard interface to sparse matrix kernels. Such standards permit manufacturers and software developers to develop specially tuned high performance kernels with a portable interface. NIST is also demonstrating the use of object-oriented software design strategies in developing numerical toolkits in C++ and Java.

Algorithm development at NOAA

In FY 1998 NOAA plans to continue algorithm development on a scalable system to achieve five- to ten-kilometer resolution in mesoscale atmospheric models. The agency will explore the design of next-generation environmental observing systems to test data assimilation needs for optimizing future forecast systems and will develop software tools to facilitate conversion from traditional shared-memory machines to scalable systems. NOAA will enhance scientific experiments running on high performance computing systems at its Princeton, NJ, Geophysical Fluid Dynamics Laboratory (GFDL), and will evaluate the performance of the ETA model at various grid resolutions and assess its potential for operational forecasting.

Advancing the science of weather forecasting

The Nation's meteorological observations are a billion-dollar resource. Because it has not been possible to design the observing system from a whole-system view, the availability of new technology has driven observing systems development. The goal of NOAA's North American Atmospheric Observing System (NAOS) Program is to exploit existing observational resources more fully and guide the development of future observing systems more purposefully. The tests and evaluations of meteorological observing systems that NAOS advocates require high performance computing hardware and techniques because both incorporating observations into prediction models and running the prediction models themselves are computationally demanding.

In order to enable more rapid development of the models used in the NAOS program, NOAA's Forecast Systems Laboratory (FSL) in Boulder, CO, has developed the Scalable Modeling System (SMS) software toolset. Built on the industry-standard Message Passing Interface (MPI), SMS is designed to ease the transfer of existing numerical weather prediction models to highly parallel computing systems. A key feature is that models parallelized using SMS are portable to a wide variety of high performance computing systems. In FY 1999, FSL plans to acquire a high performance computing system, continue enhancing its national weather prediction (NWP) models to take advantage of the new hardware, and use the new hardware and higher resolution models to conduct NAOS experiments.

Geophysical Fluid Dynamics Laboratory (GFDL)

The strongest hurricanes experienced on Earth may be upstaged by even more intense hurricanes over the next century as the Earth's climate is warmed by increasing levels of greenhouse gases in the atmosphere. Most

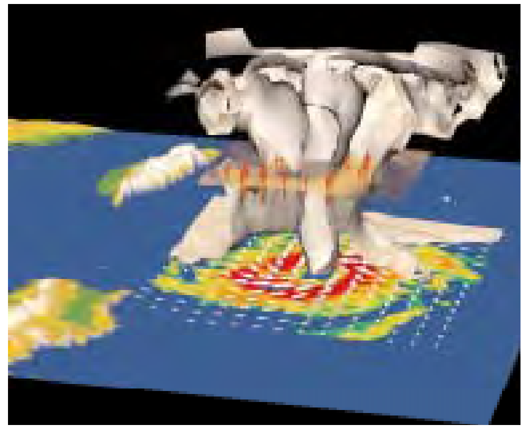

NOAA simulated samples of hurricanes from both the present-day climate and from a greenhouse-gas warmed climate.

hurricanes do not reach their maximum potential intensity before weakening over land or cooler ocean regions. However, those storms that do approach their upper-limit intensity are expected to be stronger in the warmer climate due to the higher sea surface temperatures. According to a new simulation by GFDL scientists, a 5 to 12 percent increase in wind speeds for the strongest hurricanes is projected if tropical sea surfaces warm by a little over 2˚C.

Although such an increase in the upper-limit intensity of hurricanes due to global warming was suggested on theoretical grounds a decade ago, this investigation is the first to examine the question using a hurricane prediction model that is being used operationally to simulate realistic hurricane structures. GFDL simulated samples of hurricanes from both the present-day climate and from a greenhouse-gas warmed climate by linking information from its global climate model into its high-resolution hurricane prediction model. The latter has been successfully used by NOAA's National Centers for Environmental Prediction to predict tropical storm tracks for the last several hurricane seasons. The climate model is a leading model used by climate researchers to project possible effects of greenhouse gases on future climate.

These accomplishments illustrate how research in two distinct areas that require high performance computing — climate change and hurricane prediction — can be combined to provide new information about the potential impact of global climate change upon future weather systems.

Thrust 4

Infrastructure for research in HECC — FY 1998 accomplishments and FY 1999 plans

Partnerships for Advanced Computational Infrastructure (PACI)

In FY 1998, NSF initiated the Partnerships for Advanced Computational Infrastructure (PACI) to provide access to high performance computing for the academic research community at a performance level two orders of magnitude greater than that available at a typical major research university. (Special PACI section begins on page 35.)

DOE High Performance Computing Resource Providers (HPCRPs)

The National Energy Research Scientific Computing Center (NERSC) supports the computational needs of DOE's energy science investigators. DOE supports High Performance Computing Resource Providers (HPCRPs) at Argonne, Los Alamos, and Oak Ridge National Laboratories. Systems at these sites include Intel Paragon, IB M SP2, Cray T3E and C90, and SGI

Origin 2000. DOE's Energy Sciences Network (ESnet) is part of the initial national computational grid. In FY 1998, DOE will evaluate the success of the HPCRPs, including productivity and usability.

DOE Advanced Computational Testing and Simulation Research (ACTS) Toolkit

FY 1998 ACTS research includes Sandia National Laboratory's computational biology work that focuses on developing and implementing algorithms for genome sequencing and protein folding. FY 1999 plans include prototype deployment of software tools designed around ACTS interface definitions.

DOE's Academic Strategic Alliances Program (ASAP)

DOE's Academic Strategic Alliances Program (ASAP), part of DOE's ASCI program, works with the CIC R&D program sin areas of common interest. Currently, ASAP works with U.S. universities to advance its high performance simulation capabilities:

- ❐ Stanford University — the Center for Integrated Turbulence Simulations (CITS) will develop simulation technology for designing gas turbine engines. The 1999 goal of this project is single-stage unsteady simulations of rotor-stator flow.
- ❐ California Institute of Technology — the Computational Facility for Simulating the Dynamic Response of Materials will investigate the effect of shock waves induced by high explosives on materials in different phases (i.e., gas, liquid, solid). Its 1999 goal is developing empirical reaction rates based on gas-phase detonation.
- ❐ University of Chicago — Center for Astrophysical Thermonuclear Flashes will study astrophysical thermonuclear flashes through simulation and analysis. Its 1999 goal is to develop adaptive grid methods for solving the FLASH problems on massively parallel machines.
- ❐ University of Utah/Salt Lake — the Center for Simulation of Accidental Fires and Explosions will provide state-of-the-art science-based tools for numerical simulation of accidental fires and explosions, especially in handling and storing highly flammable materials. Its 1999 goal is to conduct firespread scenario simulations.
- ❐ University of Illinois at Urbana/Champaign — The Center for Simulation of Advanced Rockets will focus on detailed whole-system simulation of solid propellant rockets under normal and abnormal operating conditions. Its 1999 goal is the preliminary investigation of abnormal scenarios.

In their initial five-year relationships, each university will interact with DOE's Los Alamos, Lawrence Livermore, and Sandia National Laboratories, and will use some of their high end computing systems.

Computing, Information, and Communications Research Facilities

To foster technical research and demonstrate future directions of CIC R&D, Federal agencies operate numerous CIC research facilities across the country. These facilities provide capabilities to (1) evaluate early prototype systems and provide feedback to developers, (2) integrate visualization and virtual reality systems into existing high performance systems, (3) run full-scale applications on systems not otherwise available, and (4) develop parallel software using scaled-down systems. These facilities provide access to innovations in network connectivity that allow large-scale applications to run over remote connections.

The success of these research centers is attributed to many enabling technologies (such as high speed networks, supercomputing systems, parallel architectures, massive data stores, and virtual reality display devices), and to the dedication of researchers, facility staff, hardware and software vendors, and industrial affiliates. All facilities provide extensive K-12 and undergraduate educational opportunities as well as training for researchers, graduate students, and faculty, and are optimal benchmarking sites for both systems and applications.

Funding for these facilities, primarily provided by CIC R&D agencies, is leveraged heavily through equipment and personnel from hardware and software vendors, discipline-specific agency funds, state and local funds, and industrial affiliate contributions. These facilities provide a low risk, collaborative environment for exploring and ultimately exploiting CIC technologies. Below is a list of CIC R&D research facilities categorized by the primary funding agency.

NSF

Supercomputing Centers (FY 1985 through FY 1998)

Cornell Theory Center (CTC), Ithaca, NY

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

Pittsburgh Supercomputing Center (PSC), Pittsburgh, PA

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

Computing research facilities at the National Center for Atmospheric Research (NCAR), Boulder, CO

PACI Centers

Begun in FY 1998, PACI succeeds the NSF Supercomputing Centers and takes advantage of newly emerging opportunities in high performance computing and communications.

National Computational Science Alliance (NCSA), Urbana-Champaign, IL

National Partnership for Advanced Computational Infrastructure (NPACI), San Diego, CA

Science & Technology Centers

Center for Cognitive Science, University of Pennsylvania

Center for Computer Graphics and Scientific Visualization, University of Utah

Center for Research in Parallel Computation (CRPC), Rice University

NASA

Testbeds

Ames Research Center, Moffett Field, CA

Goddard Space Flight Center, Greenbelt, MD

Jet Propulsion Laboratory, Pasadena, CA

Langley Research Center, Langley, VA

Lewis Research Center, Cleveland, OH

DOE

Laboratories

Argonne National Laboratory, Chicago, IL

Los Alamos National Laboratory, Los Alamos, NM

National Energy Research Supercomputer Center at Lawrence Berkeley National Laboratory, Berkeley, CA

Oak Ridge National Laboratory, Oak Ridge, TN

NIH

Systems

Frederick Biomedical Supercomputing Center at the National Cancer Institute

Supercomputing resources at the Division of Computer Research and Technology (now part of the CIT)

National Center for Research Resources' High Performance Computing Resource Centers

Biomedical Computation Resource, University of California, San Diego

Parallel Computing Resource for Structural Biology, University of North Carolina, Chapel Hill

Parallel Processing Resource for Biomedical Scientists, Cornell Theory Center, Cornell University

Resource for Concurrent Biological Computing, Beckman Institute, University of Illinois

Supercomputing for Biomedical Research, Pittsburgh Supercomputing Center

Theoretical Simulation of Biological Systems, Columbia University

National Center for Research Resources' 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

Forecast Systems Laboratory, Boulder, CO

Geophysical Fluid Dynamics Laboratory, Princeton, NJ

National Centers for Environmental Prediction, Camp Springs, MD

EPA

Systems

National Environmental Supercomputing Center, Bay City, MI

Partnerships for Advanced Computational Infrastructure (PACI)

Overview

NSF's PACI program, led by NCSA — headquartered at the University of Illinois at Urbana-Champaign — and NPACI — located at the University of California, San Diego — will enable the U.S. to maintain its lead in computational science, furthering the use of computing systems in all disciplines of research and providing new educational opportunities for students ranging from K-12 through advanced degree programs.

NCSA envisions a distributed environment whose goal is to prototype a national information infrastructure that enables the best computational research in the country. The Alliance is organized into four major groups: Application Technologies Teams that drive technology development; Enabling Technologies Teams that convert computer science research into usable tools and infrastructure; Regional Partners with advanced and mid-level computing resources that help distribute the technology to sites throughout the U.S.; and Education, Outreach, and Training Teams that will educate the population and promote the use of the technology to various sectors of society. The leading-edge site at the University of Illinois at Urbana-Champaign will support a variety of high end machines and architectures that will enable high end computation for scientists and engineers across the country.

NPACI includes a national-scale metacomputing environment with diverse hardware and high end sites. In addition to supporting the computational needs of high end scientists and engineers across the country via a variety of leading-edge machines and architectures at the University of California-San Diego, NPACI will foster the transfer of technologies and tools developed by applications and computer scientists for use by these high end users. Major focus areas are data-intensive computing, digital libraries, and large data set manipulation across multiple disciplines, including engineering and the social sciences, supported by partners across the country.

National Computational Science Alliance (NCSA)

NCSA is a partnership among computational scientists, computer scientists, and professionals in education, outreach, and training at more than 50 U.S. universities and research institutions. NCSA's national technology grid — a prototype of the computational and information infrastructure of the 21st century — will consist of a broad range of high end parallel computing systems located at the University of Illinois at Urbana-Champaign and other leading-edge facilities within the Alliance.

Alliance partners

Most NCSA partners are university and Government laboratories. Researchers at these institutions will build the technology for the grid and disseminate it to user communities across the Nation. Federal and industrial partners and strategic vendors test new technologies and provide feedback on industrial and commercial potential.

Research partners

Alabama Supercomputer Authority
American Indian Higher Education Consortium
Argonne National Laboratory
Arizona State University
BBN Systems and Technologies
Boston University
California State University - San Diego
California Institute of Technology
Center for Advanced Research in Biotechnology (CARB)
Committee on Institutional Cooperation
Computing Research Association
EPSCoR Foundation
Georgia Institute of Technology
Indiana University
Jet Propulsion Laboratory (JPL)
Kitt Peak National Observatory
Krell Institute
Lawrence Berkeley National Laboratory (LBNL)/National Energy Research Supercomputer Center
Lawrence Livermore National Laboratory (LLNL)
Los Alamos National Laboratory (LANL)
Maryland Virtual High School of Science and Math
Massachusetts Institute of Technology
Montana State University
National Radio Astronomy Observatory
North Carolina State University
North Carolina Supercomputing Center
North Central Regional Education Laboratory
Northwestern University
Ohio State University
Ohio Supercomputer Center
Old Dominion University
Oregon State University
Pacific Northwest National Laboratory (PNL)/Environmental Molecular Sciences Laboratory
Princeton University
Purdue University
Rice University
Rutgers University
Salk Institute
Shodor Education Foundation
South Dakota State University
Southeastern Universities Research Association
Stanford University
State University of New York - Albany
Syracuse University
The Scripps Research Institute
University of Alabama - Huntsville
University of Arizona
University of California - Berkeley
University of California - Davis
University of California - Irvine

University of California - Los Angeles
University of California - Santa Barbara
University of California - Santa Cruz
University of California - San Diego
University of California - San Francisco
University of Houston
University of Illinois - Chicago
University of Illinois - Urbana-Champaign
University of Iowa
University of Kansas
University of Kentucky
University of Maryland
University of Massachusetts
University of Michigan
University of Minnesota
University of New Mexico/Sevilleta Long-Term Ecological Research Program/
Maui High Performance Computing Center
University of North Carolina - Chapel Hill
University of Oregon
University of Pennsylvania
University of Pittsburgh
University of Southern California
University of Tennessee
University of Texas - Austin
University of Utah
University of Virginia
University of Washington
University of Wisconsin - Madison
Virginia Technical University
Washington University in St. Louis

Industrial partners

Fifteen Fortune 500 corporations are partnering with NCSA to preview emerging technologies:

Corporation	Fortune 500 Industrial Sector
Allstate Insurance Company*	Insurance
Boeing Company*	Aerospace
Caterpillar Inc.*	Industrial and farm equipment
Eastman Kodak Company*	Scientific and photographic equipment
Eli Lilly and Company*	Pharmaceuticals
FMC Corporation	Agricultural chemicals/diversified
Ford Motor Company	Automobiles
J.P. Morgan*	Banking/finance
Motorola, Inc.*	Electronics
Phillips Petroleum Co.	Petroleum and support
SABRE Group, Inc.	Transportation
Schlumberger Ltd.*	Petroleum and support
Sears, Roebuck and Co.*	General merchandise
Shell Oil Company*	Petroleum and support
Tribune Company*	Information/entertainment

**Have signed letters of intent to join the Alliance industrial partnership*

Strategic vendors

Commercial vendors include major computing, communications, and software companies that work closely with Alliance partners to develop and evaluate early prototypes of hardware, software, and services. The partnerships are between individual institutions rather than with the Alliance.

Ameritech
Computer Associates
Hewlett-Packard Company/Convex Technology Center
IBM Corporation
Microsoft Corporation
Platform Computing Corporation
The Portland Group, Inc.
Pyramid Systems, Inc.
Silicon Graphics, Inc./Cray Research
Sun Microsystems

Federal partners

NCSA and Federal agencies are developing and applying advanced technologies in Government. Federal partners include:

Department of Defense High Performance Computing Modernization Program
World Wide Web Federal Consortium
Department of Energy Accelerated Strategic Computing Initiative

The National Partnership for Advanced Computational Infrastructure (NPACI)

NPACI is creating a ubiquitous, continuous, and pervasive computational infrastructure to support interdisciplinary research by the national community that will help maintain U.S. world leadership in scientific advancement and economic progress. NPACI's aim is to make this infrastructure — consisting of disparate platforms linked into progressively higher performance systems — transparent, easy to use, and available everywhere.

Led by the University of California - San Diego (UCSD) and building on the foundation of the San Diego Supercomputing Center (SDSC), NPACI includes 37 of the Nation's academic and research institutions, located in 18 states. Driven by applications needs, NPACI is developing a software infrastructure to link high performance computers, data servers, and archival storage systems to aggregate computing power. Development work is focused in HECC thrust areas that leverage separately funded research projects to ensure rapid deployment and robustness of the resulting infrastructure. This work is complemented by education, outreach, and training programs and collaborations with industry.

Resource partners

NPACI consists of resource and research and education partners as well as associate and industrial partners:

UCSD/SDSC
California Institute of Technology
University of California - Berkeley
University of California - Davis
University of California - Los Angeles
University of California - Santa Barbara
University of Houston/Keck Observatory (with Baylor and Rice Universities)
University of Maryland

University of Michigan
University of Texas
Washington University in St. Louis

Research and education partners

California State University
Center for Research on Parallel Computation
LTER/University of New Mexico
Oregon State University
Rutgers University
Salk Institute
Stanford University
The Scripps Research Institute
University of California - Santa Cruz
University of Kansas
University of Southern California
University of Tennessee
University of Virginia
University of Wisconsin

Associate partners

NPACI associate partners provide additional linkages or share infrastructure requirements:

CARB
EPSCoR Foundation
JPL
Kitt Peak National Observatory (KPNO)
LANL
LBNL/NERSC
LLNL
Montana State University
PNL/Environmental Molecular Sciences Laboratory (EMSL)
University of California - Irvine
University of California - San Francisco
University of Massachusetts
University of Pennsylvania

NPACI industrial partners

NPACI industrial partners help determine the research agenda of NPACI and provide substantial cost sharing. Their involvement builds on the following collaborations:

- ❐ SGI/Cray and UCSD are integrating Cray's GigaRing technology into a heterogeneous computing environment.
- ❐ Hewlett-Packard/Convex and Caltech are enhancing Convex's system software to support large cache-coherent domains.
- ❐ IBM and UCSD are continuing their collaboration to integrate object-relational database systems with mass storage systems and develop a Massive Data Analysis Testbed.
- ❐ Sun Microsystems is continuing its collaborations with UC Berkeley on Network of Workstations (NOW) technology and with UCSD on security. It has initiated a new collaboration with UCSD to demonstrate the effectiveness of Sun technology for data serving.

❒ DEC and UCSD continue to explore using the NT operating system as an alternative to Unix in a scientific setting.

❒ Seagate, other leading disk manufacturers, and UCSD will study disk reliability using UCSD's multi-TByte storage resources as a testbed.

Large Scale Networking

Goals and focus areas

LSN R&D will help assure U.S. technological leadership in high performance network communications through research that advances the leading edge of networking technologies, services, and performance. Early Federal networking R&D investments helped build the technological foundation of today's global Internet. Development by Federal research laboratories, academia, and industry helped deploy prototype networking capabilities on a national scale and produced popular applications — such as email and World Wide Web browsers — that transformed the way people use computer networks, paving the way for our Nation's leadership in the multi-billion dollar information technology industry.

Key research areas today include advanced network components and technologies for engineering and managing large scale networks of the future. LSN activities will:

- ❒ Foster Federal research to address requirements in LSN technologies
- ❒ Facilitate interagency collaborations in LSN R&D
- ❒ Provide mechanisms for cooperation in LSN R&D among Federal agencies, Government laboratories, academia, and industry.

Beginning in FY 1998 and building on base LSN research, the Next Generation Internet (NGI) initiative is the dominant focus of LSN R&D. This section describes FY 1998 accomplishments and FY 1999 plans in LSN R&D and in the NGI initiative.

Global Grid Communications and Global Mobile Information Systems

DARPA's Global Grid Communications effort develops and demonstrates advanced communications technologies needed for Defense and intelligence operations for the 21st century. This program will show that advanced optical components developed in the program can be integrated with commercial communications resources and technologies as well as DoD tactical and satellite technologies developed elsewhere. In FY 1998, DARPA will demonstrate multi-wavelength network management and control in local area testbeds, 40 billion bit per second cross-connect switching, and a 32 channel transceiver chip.

DARPA's Global Mobile Information Systems effort will enable mobile users to access and use the full range of services available in the Defense Information Infrastructure (DII) by developing nomadic technologies and techniques at the applications, networking, and wireless link/node levels. In FY 1999, DARPA plans to develop a mobile wireless network incorporating software radio technology and demonstrate application support for distributed computing in mobile environments, continuous multi-tier networking across wireless domains, and integrated high data-rate untethered nodes.

Advanced Networking Infrastructure and Research (ANIR)

NSF's ANIR efforts support the enhancement of networks connecting U.S. researchers and educators to information resources, computational resources, and special facilities. ANIR's centerpiece is the very high performance Backbone Network Service (vBNS), which is targeted to provide very high performance connections to approximately 100 leading universities and provide interconnections to other Federal research networks. Responding to the goals of the NGI initiative (page 48), these connections and advanced facilities serve as a testbed for advanced communications and networking technologies that might eventually be deployed in support of the greater Internet, and for developing revolutionary applications and information services. Related components of the ANIR program provide for network growth engineering, development support for network-based applications, and campus engineering design.

ANIR supports basic research and experimental projects that focus on networking and communications systems. Research topics include wireless access and networks, collaboration technology — especially where enabled by active networks — and the convergence of computing, communications, and information. The projects typically are conducted by small groups of researchers in networking or communications and, as appropriate, from other areas of computer science and engineering, such as operating systems, data bases, software environments, and architecture; or from the social sciences, such as economics, psychology, or sociology.

National Scalable Cluster Project (NSCP)

ANIR and NSF's Cross-Disciplinary Program will support the acquisition of equipment for the NSCP. This project combines supercomputing systems, high-speed networking protocols, and custom-designed software to extend the boundaries of applications in data mining, data archiving, parallel processing, and other research projects that require large scale computational engines to support knowledge discovery.

Chesapeake Bay Virtual Environment (CBVE)

NSF's Chesapeake Bay Virtual Environment fuses 3-D visualizations of numerically-generated data into a large scale interactive virtual world. This framework will incorporate runtime computational steering, interactive visualization, data ensonification, and wide-area information dissemination. CBVE will support technology transfer and serve as a medium for information exchange by taking advantage of Internet-based media and complete World Wide Web functionality. The use of a virtual world as a new paradigm of information exchange coupled with a Web-based information architecture will allow multidisciplinary, multi-institutional groups of scientists, educators, students, and administrators to apply emerging high performance computing and communications technologies to studying the country's largest estuarine ecosystem. CBVE will initially couple a working 3-D circulation model of the Chesapeake Bay and adjoining shelf with working biological models for larvae of two commercially important species in Chesapeake Bay, the Atlantic menhaden (*Brevoortia tryannuas*) and the blue Crab (*Callinectes sapidus*). Visualization of the entire model output in an immersive, interactive environment is the goal.

Experimental Program to Stimulate Competitive Research (EPSCoR)

EPSCoR's goal is to help states identify, develop, and use their academic science and technology resources to support a more productive and fulfilling way of life for their citizens. EPSCoR increases the R&D competitiveness of a state by developing and using the science and technology (S&T) resources residing in its major research universities. EPSCoR achieves its objective by:

- ❒ Stimulating sustainable S&T infrastructure improvements at the state and institutional levels that significantly increase the ability of EPSCoR researchers to compete for Federal and private sector R&D funding
- ❒ Accelerating the movement of EPSCoR researchers and institutions into the mainstream of Federal and private sector R&D support

EPSCoR currently operates in 18 states: Alabama, Arkansas, Idaho, Kansas, Kentucky, Louisiana, Maine, Mississippi, Montana, Nebraska, Nevada, North Dakota, Oklahoma, South Carolina, South Dakota, Vermont, West Virginia, and Wyoming, as well as the Commonwealth of Puerto Rico.

NSF satellite communications

In FY 1998, NSF will support NOAA's GOES-3 satellite for Internet communications to South Pole Station, with continued development of the link to support remote operation of the instruments in its astronomy and astrophysics programs. NSF will also collaborate with NASA on its experimental satellite communications link to the South Pole for higher speed Internet connectivity and high speed file transfer.

Energy Sciences Network (ESnet)

DOE conducts network research, advanced network deployment, and advanced application support for over 20,000 users of dozens of DOE experimental facilities and high performance computing resources. DOE's LSN program includes its core network and security research program in addition to the ESnet, its advanced production network. ESnet has played a

The Energy Sciences Network (ESNet) is the DOE Office of Energy Research nationwide network that supports open scientific research.

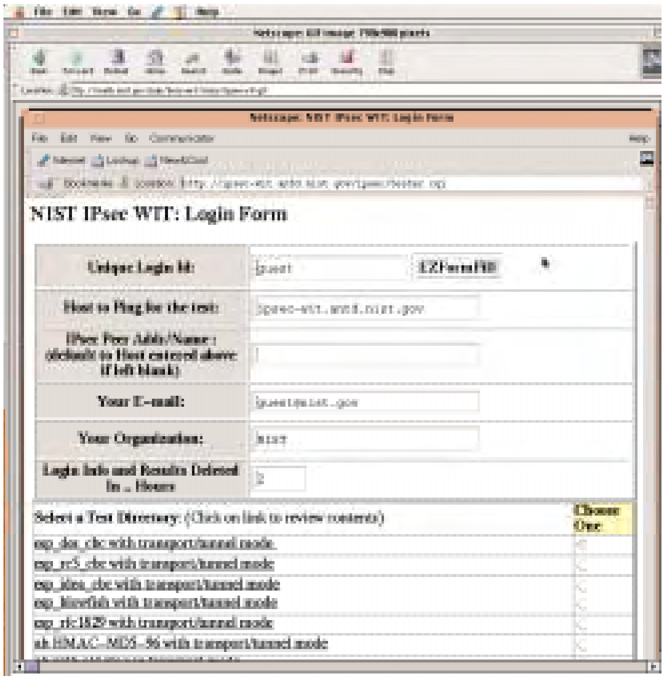

Software developers select interoperability test scripts through NIST's Web-based Interoperability Tester (WIT) Web interface.

major role in the development of the worldwide Internet and will continue to contribute to its evolution through participation in the LSN programs.

ESnet is a service-oriented production network chartered to support mission-oriented DOE science. It provides advanced Internet Protocol (IP) and ATM services to 30 DOE sites, including national laboratories, universities, and international partners. ESnet will adopt and integrate leading-edge technologies to support DOE's mission applications, and will continue to implement and enhance advanced interconnection and peering arrangements with other Federal research networks as well as university networks and aggregation points (e.g., gigapops) to support DOE/university collaborations on DOE mission programs.

Information technology metrology, testing, and applications

NIST produces tools to evaluate the software needed for next generation networks and to evaluate the quality of service that those networks deliver. In FY 1998, NIST built a laboratory testbed including many vendor implementations of the resource reservation protocol and the real time transport protocol. NIST also developed the Integrated Services Performance Instrument (ISPI) to evaluate the quality of service these protocols deliver, and the NIST Network Emulation Tool (NIST NET) that allows an application developer to artificially impair network performance to evaluate how much quality of service the application needs.

Simulating protocols can save time and money by uncovering problems before systems are developed. In FY 1998 NIST completed simulations of the Hybrid Fiber Coax (HFC) Media Access Control (MAC) protocols and of the performance of IP over ATM in an HFC environment. These protocols are being evaluated for high-speed residential access over the Next Generation Internet.

NIST is developing methods to speed the interoperability testing cycle. In FY 1998 a reference system for the IP Security (IPSEC) protocol was developed and is now available for Unix PCs. Because every vendor can test against the same reference, the number of incompatibilities between systems can be reduced. As an extension of this idea, NIST has developed a Web-based Interoperability Tester (WIT) for IPSEC. This allows developers to test against the NIST reference system using the Web, eliminating the need to download and install the reference system itself. WIT will be extended to other test systems in FY 1999.

Public Key Infrastructure

In cooperation with a dozen industry partners, NIST has developed and issued a Minimum Interoperability Specification for Public Key Infrastructure Components. This specification helps ensure that PKI components, such as certificate authorities, from multiple vendors will interoperate across entire networks and the Internet.

Internetworking security

NIST contributes to the development of advanced security technologies needed to ensure high levels of confidentiality, integrity, and availability of network systems and data. This is a critical need for the full realization of the potential of electronic commerce. With the support of the LSN and the Government Information Technology Services (GITS) Innovation Fund, NIST conducted Collaborations in Internet Security (CIS), a project involving multiagency collaborative research in the use of advanced network security mechanisms such as Kerberos, security smart cards, secure messaging, and PKI components.

NIST has begun identifying, evaluating, and establishing the advanced encryption standard, intended to replace the existing Data Encryption Standard (DES) as the standard algorithm for symmetric key encryption and to provide high-quality encryption capability that will be useful well into the 21st century.

FedCIRC

As Federal agencies expand their use of the Internet both to provide citizen access to government services and to support agency needs, there is an increasing threat of outsider attacks. Thus, there is a need for a 24-hour network threat monitoring and security incident response capability. NIST has collaborated with the Computer Emergency Response Team (CERT), DOE's Computer Incident Advisory Capability (CIAC), and other organizations to develop methods for monitoring and responding to these threats. To ensure that all agencies have access to such resources, NIST has established the Federal Computer Incident Response Capability (FedCIRC) to conduct the dual roles of threat monitoring and vulnerability analysis and incident notification and response support.

NLM's Medical Connections program

At NIH, the NLM's Medical Connections program provides "jump start" funding to academic medical centers, community hospitals, and other healthcare organizations to connect to the Internet. The overall goal is to provide Internet connectivity to the top 3,000 healthcare institutions in the U.S. More than 250 U.S. medical schools and healthcare facilities have been connected during the past five years.

Approximately 50 grants are being awarded in FY 1998. Special emphasis is given to linking medical libraries with healthcare delivery organizations and data base servers to support timely, accurate clinical decision making. The program also supports the creation of regional consortia of healthcare institutions for sharing medical information and distributing Internet capability within an institution.

NLM's National Center for Biotechnology Information (NCBI)

NCBI creates automated systems for storing and analyzing a vast and growing volume of molecular biology, biochemistry, and genetics data — a field known as bioinformatics. Through its use of analytical and predictive methods to identify key molecular patterns associated with health and disease, bioinformatics is an essential component of genome research, protein engineering, and drug design. Within a distributed database architecture, NCBI collects sequence data from researchers worldwide and incorporates them into GenBank, the NIH DNA sequence databank, a key data resource of the Human Genome Project. NCBI has produced an integrated database system consisting of GenBank, the genetic scientific literature in Medline, taxonomy, and 3-D molecular structures. This database is accessed daily through the Internet by more than 35,000 sites. Basic research on efficient data analysis techniques and large scale genome analysis is conducted within NCBI's Computational Biology Branch and has been a key factor in gene discovery.

Combined high throughput sequencing efforts by NIH-funded centers and by commercial organizations has led to doubling of the size of the database each year. User demands have steadily increased largely due to availability of the World Wide Web and the diffusion of sequencing technology across multiple biological disciplines. The human gene sequence and mapping data are represented in the Web database of the Human Gene Map, a continuing collaborative project with *Science* magazine and 15 leading international laboratories.

NLM's Integrated Academic Information Management Systems (IAIMS)

Academic medical centers are the backbone of the American biomedical research enterprise. These 120-plus institutions include health profession schools and their associated teaching and research hospitals, clinics, and laboratories. Information about patient care, research, education, and administration is the life blood of these complex centers. This information — databases of bibliographic and factual information, molecular databases, patient records, laboratory and clinical data — is in electronic form, but the electronic information sources are typically disconnected and isolated from one another. Communications among the various computerized systems in academic medical centers is often primitive or non-existent.

The focus of the IAIMS program, initiated in 1984, is first to develop the technical and organizational infrastructure to link and retrieve conceptually related information from disparate sources within the medical center, and then to link the medical centers. The administrative, clinical, educational, and research databases should be able to communicate and to appear as one database to the user. The goal of the program is to develop, implement, and test generalizable systems of information flow management within university health science centers and major teaching hospitals. The expected outcomes are greater research productivity, improved access to patient data for technology assessment and health outcomes research, and more efficient patient care leading to increased efficiency in the use of healthcare resources. The work is expected to benefit all health delivery organizations, including community hospitals and outpatient services.

Agency for Health Care Policy and Research (AHCPR)

In FY 1999, AHCPR will begin to evaluate quality, cost, and medical effectiveness of home care in rural, urban, and suburban areas for the elderly who are afflicted with chronic disease.

Joint Engineering Team (JET)

The JET helps coordinate large scale Federal networking activities to support research and engineering. The JET is the successor to the Engineering and Operations Working Group of the Federal Networking Council (FNC), which was incorporated into the LSN Working Group in FY 1997. The JET is tasked with coordinating NGI Goals 2.1 and 2.2 (pages 52-55) as well as other LSN network research operational issues. Participants in JET meetings include Federal agency representatives; representatives from the university-based Internet2 (I2) community; and representatives from commercial network vendors for whose services the Federal networks contract.

The JET coordinates networking activities and operations among multiple Federal agency networks including:

- ❐ DoD's DREN
- ❐ DOE's ESnet
- ❐ NASA's NREN
- ❐ NSF's vBNS
- ❐ DARPA's SuperNet

In addition, the JET coordinates with Internet2 (I2) and the Abilene Network. Coordinated projects include network access points (NAPs; I2 gigapops; STAR TAP; interagency network peering points such as ESnet-vBNS; etc.), international connections, traffic monitoring, performance measurement, multicast distribution, and new technology deployment. (STAR TAP is the Science, Technology And Research Transit Access Point, an international networking meeting point in Chicago.) In FY 1998, the JET supported the FY 1998 NGI demonstrations at Highway 1 in Washington DC, and is coordinating DREN's and NREN's links to Alaska. In FY 1999 the JET plans to coordinate Federal networking for NGI demonstrations at SC98 in Orlando, FL.

Networking Research Team (NRT)

The NRT coordinates the networking research of the LSN agencies, including activities associated with Goal 1 of the NGI initiative. The NRT's goals are to enhance overall Government R&D in networking research and to ensure LSN agencies jointly implement a comprehensive program in critical research areas such as differentiation of service, multicasting, network management and operations, and privacy and security. These goals will be accomplished through keeping NRT members informed about each participating agency's activities and research results, using Federal/academic/industry workshops and other forums to identify research needs, and planning and implementing a balanced program of individual agency and multiagency activities.

Next Generation Internet

The Federal NGI initiative, together with the country's other investment sectors, will create the foundation for the networks of the 21st century, setting the stage for networks that are much more powerful and versatile than the current Internet. The NGI will foster partnerships among academia, industry, and governments that will keep the U.S. at the cutting-edge of information and communications technologies. The NGI will also stimulate the introduction of new multimedia applications in our homes, schools, and businesses as the technologies designed and developed as part of the NGI are incorporated into products and services that are subsequently made available to the general public. The NGI initiative is essential to sustain U.S. technological leadership in computing and communications and enhance U.S. economic competitiveness and commercial eminence.

The NGI initiative is a key component of the ongoing multiagency R&D of the LSN Working Group. NGI activities will be leveraged off of and tightly coupled with the basic network research and infrastructure support that are funded under the LSN budget.

NGI Goals

The NGI initiative has three goals:

1. To advance research, development, and experimentation in the next generation of networking technologies to add functionality and improve performance in:

   - ❒ Reliability
   - ❒ Security
   - ❒ Robustness
   - ❒ Differentiated services including multicast and audio/video — also known as Quality of Service (QoS) and Class of Service (CoS)
   - ❒ Network management including allocation and sharing of bandwidth

2. To develop NGI testbeds for system-scale testing of advanced technologies and services to:

   - ❒ Connect at least 100 NGI sites with end-to-end performance at least 100 times faster than today's Internet
   - ❒ Connect about 10 sites with end-to-end performance at least 1,000 times faster than today's Internet
   - ❒ Develop and test advanced applications

3. To develop and demonstrate revolutionary applications in enabling applications technologies such as:

   - ❒ Collaboration technologies
   - ❒ Digital libraries
   - ❒ Distributed computing
   - ❒ Privacy and security
   - ❒ Remote operation and simulation

and disciplinary applications in:

- ❒ Basic science
- ❒ Crisis management
- ❒ Education
- ❒ The environment
- ❒ Federal information services
- ❒ Healthcare
- ❒ Manufacturing

The Federal agencies participating in the NGI in FY 1998 are DARPA, NSF, NASA, NIST, and NIH (NLM and NCRR). The NGI initiative is managed by the participating agencies and coordinated by the LSN Working Group. Experts from academia, industry, and Federal laboratories will be asked to provide input. While DOE is not a participant in the NGI in FY 1998, the Administration proposes to add the agency beginning in FY 1999. This proposed participation in the NGI is described on pages 58-60.

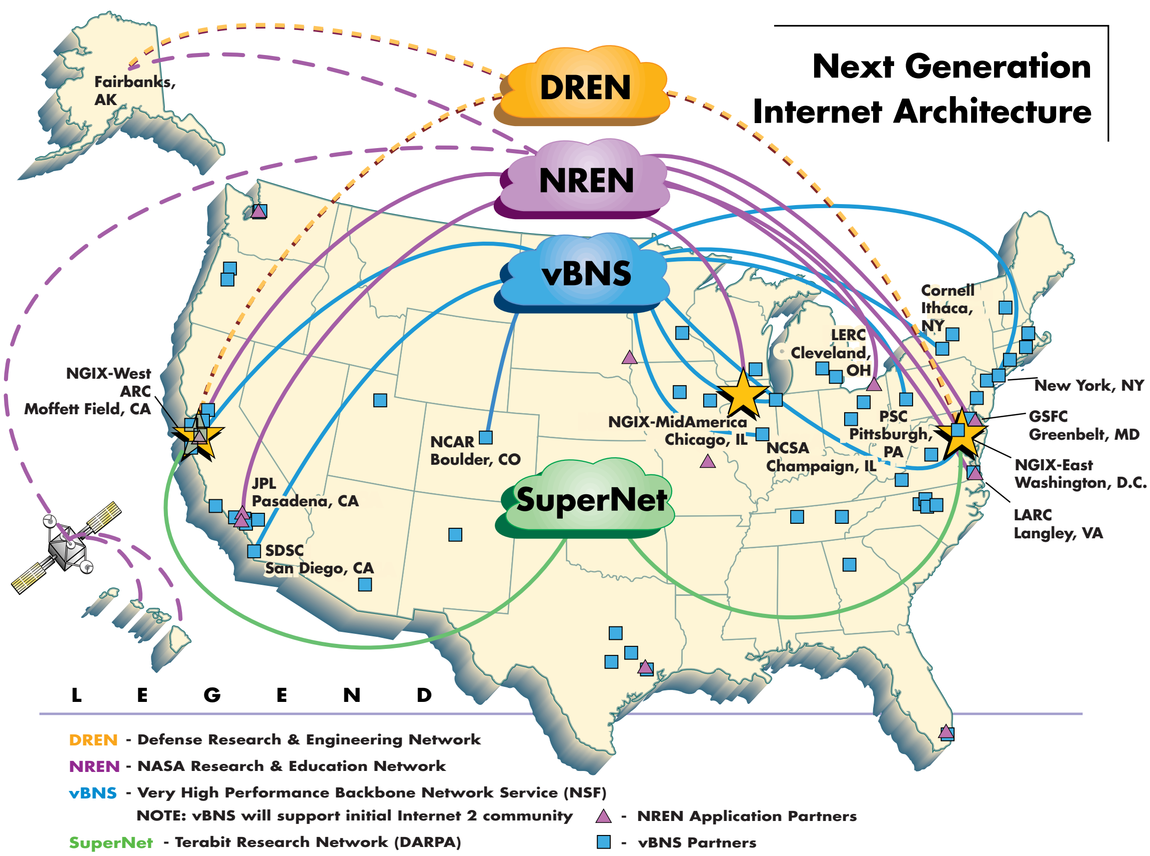

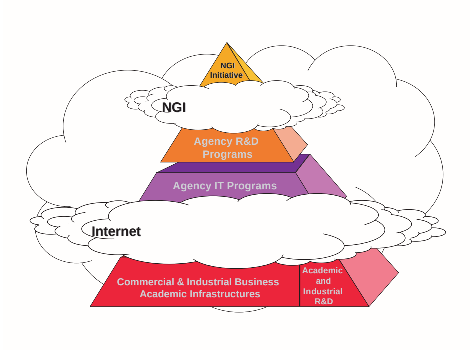

The NGI will foster partnerships among academia, industry, and governments that will keep the U.S. at the cutting edge of information and communications technologies.

NGI GOAL 1: Experimental research for advanced network technologies

Goal 1 activities focus on research, development, and testbed deployment and demonstration of the technologies necessary to permit the effective, robust, and secure management, provisioning, and end-to-end delivery of differentiated classes of service. These activities cluster into three major tasks: (1) network growth engineering, (2) end-to-end QoS, and (3) security. These technologies along with increased bandwidth will help meet the needs for advanced functionality and for the real time demands of emerging and next generation applications technologies such as collaboration, distributed computing, and teleoperation and telecontrol.

This is a multiagency effort with DARPA as the lead; participation by NSF, NASA, and NIST; and contributions by non-NGI-funded agencies. Notice of funding opportunities is through normal agency mechanisms such as solicitations and broad area announcements. Cross-agency participation in review panels and coordination by program managers help ensure that efforts are not unnecessarily duplicated and lead to an integrated solution — a robust, scalable, shared infrastructure supporting lead agency users, other government agencies, the research community, and a large number of commercial users.

DARPA, in partnership with industry, will develop network management and QoS technologies. Portions of DARPA's existing Quorum program in global distributed system technologies will form the basis for end-to-end QoS. Other agency programs will complement and leverage the Quorum program in developing advanced network services, QoS, and security technologies.

NSF's FY 1998 NGI plans include providing integrated IP and ATM monitoring and analysis tools at its Goal 2.1 sites (page 52), and implementing caching and native multicast at appropriate sites. Anticipated FY 1999 activities include the implementation of security at appropriate sites.

NASA's NGI Goal 1 plan is to sponsor R&D in new networking technologies and services to support high performance applications. NASA's NREN is the basis of this plan. NASA will continue to be an early adopter of emerging networking technologies.

NASA-sponsored research will focus on network performance measurement, network interoperability scaling, management, QoS, and network security. NASA will fund and manage research in advanced network technologies that are richer in features, higher in performance, and deliverable at a reasonable cost. For example, NASA will enable real-time networking, group collaborations, and a seamless interface for space-to-ground communications.

NASA will deploy a suite of advanced networking services to enable high performance applications. By partnering with industry and academia on R&D in internetworking technologies to achieve an interoperable high performance network testbed, NASA will help deliver advanced networking technologies to the aerospace community and ultimately to the public.

In FY 1998, NASA's activities include developing standard simulation models to "grow" internetworks/intranetworks; developing baseline simulation statistics; and demonstrating the interconnection of NASA's NGI infrastructure with other agency NGI networks. FY 1999 activities will include NGI requirement analysis and configuration management and the deployment of ATM debugging, monitoring, and analysis tools.

NIST is refocusing part of its on-going research program in advanced networking technologies, computer security, and conformance testing to support NGI goals. Through its traditional focus on measurement, standards, and test methods, NIST will help ensure that research in these areas results in standardized, interoperable commercial technologies.

NIST will support the U.S. information technology industry by fostering the rapid commercialization and deployment of enabling and infrastructural networking technologies developed through the NGI. The complexities of, and interdependencies among, future network control systems and services (for example, multilayer QoS signaling, routing, flow control, and security) and their global scaling will require more than simple analysis.

NIST is conducting R&D in techniques and tools to test and evaluate new networking technologies at all stages of development and deployment. Advances in measurement and testing technology will enable the rapid evaluation of research designs and prototypes and will facilitate the transfer of new technology to the NGI Goal 2 testbeds and the communications industry. The goal is for test and instrumentation technologies to become

NREN implemented a consistent, wide bandwidth data pipeline between the Jet Propulsion Laboratory (JPL) and the NASA Ames Intelligent Mechanisms Group that provided the Mars Pathfinder mission with MarsMap — a photo-realistic virtual reality model of the Mars environment that assisted with rover data archiving and operational planning. Raw image data from JPL was relayed via NREN to the NASA Ames Intelligent Mechanisms Group for processing, and 3-D images were returned from Ames to JPL for display. Here, the rover is seen exploring the Mars rock dubbed "Yogi."

part of the protocol design and specification process, and be integrated into, and evolve with, the implementation and deployment of the network itself.

NGI GOAL 2: NGI testbeds

The networking testbeds developed under the NGI initiative will connect at least 100 sites — universities, Federal research institutions, and other research partners — at speeds 100 times faster end-to-end than those of today's Internet, and will connect on the order of 10 sites at speeds 1,000 times faster end-to-end than the current Internet. This end-to-end connectivity (such as between two workstations) will be at speeds from 100+ million bits per second (Mbps) to 1+ billion bits per second (Gigabits per second or Gbps). Some networks have already achieved OC-12 (Optical Carrier) speeds (622 Mbps) on their backbone links, and some experimental links are running at 1+ Gbps, but end-to-end usable connectivity is typically limited to less than 10 Mbps because of bottlenecks or incompatibilities in switches, routers, local area networks, and workstations. Goal 2 addresses these shortcomings through development and demonstration involving two subgoals, described below. Goal 2 testbed networks will aggressively incorporate Goal 1 technologies.

GOAL 2.1: High performance connectivity

The high performance connectivity testbed is a distributed laboratory delivering, at a minimum, 100 times faster speeds than current Internet performance on an end-to-end basis to at least 100 interconnected NGI participating universities, national laboratories, and Federal research sites conducting networking and applications research. This testbed will be large enough to provide full system proof-of-concept for hardware, software, protocols, security, and network management that will be required in the future commercial Internet. It will include easily accessible sites, remote sites, and sites in EPSCoR (Experimental Program to Stimulate Competitive Research) states. Experiments to help researchers go beyond the current Internet infrastructure are planned. Goal 2.1 is a multiagency effort led by NSF, NASA, and DOE (beginning in FY 1999). Participants include DoD and other agencies.

NSF has a two-phase strategy to build its Goal 2.1 testbed. The first task is to significantly expand and enhance NSF's existing program for high performance connections to its vBNS network to serve over 100 leading universities and to link them to their research partners by improving the interconnections among the vBNS and other Federal research networks. This work builds on existing high performance connections to the NSF supercomputing centers and their partners (PACI); the NSF-sponsored National Laboratory for Applied Networking Research (NLANR); dozens of funded individual investigators in university and industry laboratories; and ongoing funded research with investigators in DARPA, NASA, and other agencies.

In particular, the NLANR applications team will help users maximize performance of their applications, solve network problems, maintain information and links about applications, and provide training to network and applications engineers. Additionally, NSF's partnership with Internet2 will focus the collective efforts of more than 100 leading universities on next generation networking technologies and associated issues of deployment, management, and testing for NGI Goal 1.

NSF's second task is to test and deploy Goal 1 technologies and Goal 3 applications. At the same time, NSF will begin forming a national organization of universities to plan and coordinate their ongoing role in the NGI and associated efforts. In FY 1999, NSF will work with other agencies to design and implement a more unified Federal research network that can better serve their entire research community.

DARPA's FY 1998 network research objectives include establishing a Peer Network-to-Network Interface (P-NNI) hierarchy across ATM domains, network performance measurement, congestion management, IP and ATM address resolution mechanisms, and ATM signaling behavior across multiple providers. Another way to add new ACTS ATM International (AAI) nodes is to establish gateway agreements with other providers, such as through NSF for select vBNS-attached collaborator organizations. (ACTS is NASA's Advanced Communications Technology Satellite program.)

NASA will leverage NREN in meeting its NGI goals. NASA will provide both a high performance network application testbed and a network research testbed for the NASA community and its partners. These testbeds exist at the various NASA centers now and can be interconnected via NREN, thus

Via the NREN, remote, interactive simulations were conducted between the NASA Ames Vertical Motion Simulator in California (pictured at the right) and engineers located in Houston, Texas.

providing virtual testbeds and harnessing the expertise distributed throughout NASA. NASA will focus on delivering a leading-edge application environment. Therefore, in FY 1998 and FY 1999, NASA will enable next generation application demonstrations across the network; internetwork with other Federal agencies and academic and industry partners at both the IP and ATM service level; and deploy advanced networking services such as IPv6, multicast, QoS, security, and network management tools.

GOAL 2.2: NGI technologies and ultrahigh performance connectivity

Goal 2.2 addresses the development of ultrahigh speed switching and transmission technologies, and the demonstration of end-to-end network connectivity at 1+ Gbps. DARPA's SuperNet is the Goal 2.2 testbed. Because of its high risk and pioneering nature, this testbed will initially involve about 10 NGI sites and a limited number of applications. There will be some overlap of Goal 2.1 and Goal 2.2 nodes.

Attaining this goal, together with the technologies developed in Goal 1, will be the pathway to terabit-per-second (Tbps) networks, operated with the appropriate network management and control and guaranteed end-to-end QoS. Working in partnership with industry is the key to a shared infrastructure that can be profitably used to support high end scientific users and large numbers of commercial users. This is a multiagency effort with DARPA as the lead, with participation by NASA, NSF, DOE (beginning in FY 1999), and other Federal agencies.

Wide Area Broadband Core

DARPA's Broadband Information Technology (BIT) program has developed basic wavelength division multiplexing (WDM) transmission capabilities and will demonstrate a metropolitan network of five nodes, with link transmission capacities of 20 Gbps. DARPA will extend these technologies and deploy them in more complex, mesh-like topologies that involve long-distance links.

The metropolitan testbed will be expanded into a wide area network using WDM with about 10 nodes. This wide area backbone will have sufficient aggregate transmission and switching resources to support hundreds of users at Gbps rates. This network will use existing network fibers (usually at different wavelengths) that also provide general Internet services.

Terabit-per-second technologies

DARPA will develop the generation-after-next multiplexing, switching, and routing technologies that will bridge the gap among packet-based Gbps tributaries and the WDM-based optical core. This task will also lay the groundwork for the direct optical support of packet-based communication. A major component of this task will be to investigate statistically sound techniques for performing "space-division"-like spreading of the resultant time division multiplexing (TDM) traffic across a set of wavelengths. A second component will be the design and demonstration of a highly parallel and distributed switching fabric.

Taken together, these efforts will enable the development of a highly distributed approach to Tbps switching, based on a combination of optical and electronic technologies, with many-to-many multicast capability.

Broadband local trunking

The need to provide select sites with orders-of-magnitude-above-average access to the network core has been a recurring source of delay in commissioning advanced research facilities. This task will explore new and cost-effective approaches to delivering broadband access to select sites within a geographically restricted area. DARPA will examine the terrestrial extension of its SuperNet-rate facilities to the building and explore the effectiveness of high capacity (>150 Mbps) radio frequency (RF)-based trunking. In addition, DARPA will address wireless broadband local access.

DARPA demonstrations and field trials

Most of the technologies to be developed by the previous tasks are associated with the physical, link, and networking layers. DARPA plans to demonstrate the newly developed capabilities through collaboration with some of its application-oriented activities, such as the Human Computer Interaction, Information Management, and Intelligent Collaboration and Visualization programs.

Multiagency partnerships

NASA will partner with DARPA's ultrahigh speed testbeds through the participation of at least two NASA sites. NASA will investigate the feasibility and performance of engineering application demonstrations across these testbeds. The goal is to achieve an end-to-end high speed hybrid network capable of supporting both wireless and bounded-media applications.

NSF's activities include select connections to Goal 2.2 networks and the funding of competitive academic research proposals. NSF will participate with DARPA and other agencies in ultrahigh speed networking links and technologies through NSF's PACI partners. The focus will be on technologies and protocols for advanced distributed computing.

End-to-end technologies

NSF will:

- ❒ Select and tune PACI applications for high speed research
- ❒ Study and tune ultrahigh speed performance using future-generation tools
- ❒ Connect to applications at select PACI partner universities
- ❒ Adapt Goal 1 results to Goal 2.2 networks
- ❒ Coordinate these activities through NLANR and PACI, as well as through awards to individual investigators

NGI GOAL 3: Revolutionary applications

Applications are the ultimate metric for NGI success. Faster and more advanced networks will enable a new generation of applications that include crisis response, distance education, environmental monitoring, healthcare, national security, and scientific research.

To achieve this goal, agencies will leverage NGI investments with other major application investments. Agencies will demonstrate new applications, as well as enhance and enable current mission applications that address national goals. Each demonstration will partner the advanced networking technologies developed in Goals 1 and 2 with modern applications technologies.

Each community will bring its knowledge, skills, and methods to the partnership. The applications partner will provide the bulk of the resources and support needed to implement its applications, working within the framework of the NGI initiative to develop and demonstrate its applications over the high performance networking infrastructure by using advanced network technologies provided by other parts of the NGI. The applications demonstrations will primarily be proof-of-concept demonstrations intended to suggest new ways for the application partners to meet their mission needs. As part of an ongoing research effort, these applications will initially be built on less-than-fully-robust technologies and be operating in less-than-optimal networking environments.

Many agencies have critical signature applications that will benefit from advanced networking services and capabilities. Both the Federal government's information technology services and the Federally supported R&D community have networking requirements that cannot be met with today's technologies. Higher speed networks with more advanced services and functionality will enable a new generation of applications that support fundamental governmental interests. For example, NLM is working to define NGI capabilities needed for routine use of NGI technologies in healthcare, public health and health education, and biomedical, clinical, and health services research. NLM will fund testbed projects demonstrating the use of these technologies by the healthcare community. These demonstrations will be designed to improve understanding of the impact of NGI capabilities on the nation's healthcare, health education, and health research systems in such areas as cost, efficacy, quality, security, and usability.

As the NGI initiative develops new capabilities such as adaptive networking, collaboration technologies, medical data privacy, network management technologies, nomadicity, QoS, and security, advanced demonstration applications will take advantage of the new services that these capabilities enable. It is expected that agencies not in the NGI budget crosscut will participate in these applications. For example, EPA and NOAA have identified key applications requiring NGI speed and services. The education community is also putting significant effort into connecting K-12 schools to the current generation Internet. Advanced education applications such as distance learning are expected to be key components as the NGI matures. NGI applications prototypes will test these new capabilities to ensure that the protocols developed in Goals 1 and 2 are complete, robust, and useful in real applications and to provide a road map useful in future governmental and commercial services.

Selection process

The participating Federal agencies have established a coordinated selection process to identify NGI applications. These applications require the advanced networking capabilities of Goals 1 and 2. Agencies will be asked to adapt their applications as these capabilities are developed.

The selection process will be used to ensure that applications tested and demonstrated on the NGI testbeds provide robust, realistic, complete tests of technologies that are extensible and adaptable to other applications. Selection will be an iterative process with Federal, academic, and industry participation. Applications will be derived from the Federally-focused applications in appropriate technology classes such as digital libraries, remote operation of instruments, environmental monitoring, crisis management, manufacturing, basic sciences, and Federal information services.

Prototype Radiation Treatment Consultation Network

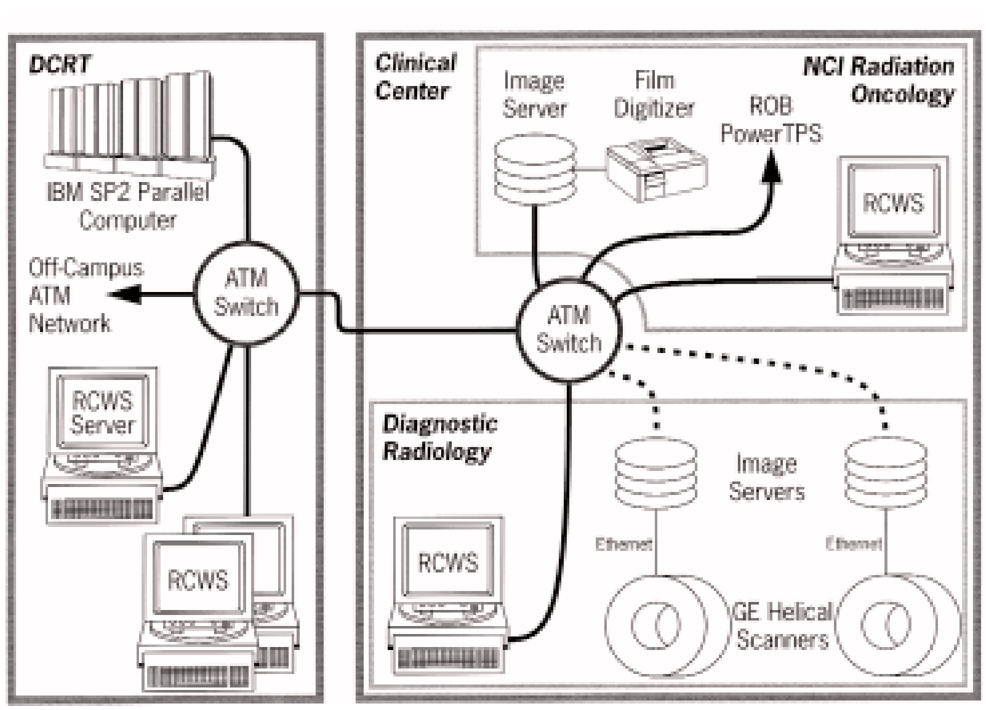

Block diagram of the Prototype Radiation Treatment Consultation Network being implemented on the NIH campus. Four RCWS nodes are shown above, as well as the conference server and CIT's IBM SP2 supercomputer, and connections within the NCI Radiation Oncology Branch to their PowerTPS treatment planning system node and a sheet film digitizer node. This is the type of leading edge health services delivery network application that will benefit greatly from advances made through the NGI initiative.

Coordination

This multiagency effort will be coordinated by the participating agencies through the LSN Working Group. Since most of the funding for applications will come from the applications themselves, leadership will be provided by domain-specific affinity groups. Participation will be encouraged from a broad range of agencies with demanding networking applications. Applications will also be solicited from other interested research entities within academia and industry.

Applications

Seventeen NGI applications were demonstrated at Netamorphosis and are described beginning on page 61.

NGI and I2: Complementary and interdependent

Specific ways in which the Federal NGI initiative and the university-led Internet2 work together include:

- ❒ NSF has approved 92 institutions for High Performance Connections to its vBNS. Connections are based on merit. Seventy seven of the institutions are Internet2 universities. vBNS connectivity is a key part of NSF's NGI program.
- ❒ Internet2 universities are establishing gigapops (gigabit per second points of presence) that provide regional connectivity among universities and other organizations. Through the gigapops, universities may connect to NGI testbeds and other advanced Federal networks, including the vBNS, NREN, DREN, and the ESnet. The NGI and Internet2 will help ensure that advanced networking services are available on interoperable backbone, regional, and local networks that are competitively provided by multiple vendors.
- ❒ Researchers at Internet2 universities are developing a wide range of applications that require advanced networking. Many of these applications are funded by Federal initiatives including the NGI.

DOE FY 1999 NGI plans

In FY1999, DOE proposes to join in the NGI initiative. DOE will leverage its current core programs in network and application research, as well as its system integration expertise, to enhance the Department's ability to satisfy mission requirements through the development and prototyping of advanced technologies and to further the NGI goals.

DOE's proposed FY 1999 NGI program has three major components:

- ❒ Core network research focuses on developing new technologies and capabilities to be integrated into the network infrastructure to support DOE application requirements.
- ❒ Enhanced DOE application control of network technologies focuses on intelligent middleware — the software between an application and the network transport or operating system. The aim is to provide easy-to-use interfaces and software to DOE's applications, as well as any supporting network transport layer capabilities, so that users can ascertain the status of the network and intelligently and dynamically make the best use of that infrastructure to support their application.

❐ A new DOE-university partnership program that enhances the collaborative application environment through joint technology development and deployment. This partnership will focus on jointly developing NGI technologies, accelerating end-to-end lab-to-campus network and testbed infrastructures, and adapting DOE application codes at both labs and universities to support DOE missions.

These components contribute to all three NGI goals.

In order to enhance DOE and university collaboration on DOE mission critical applications and NGI research, DOE will support joint DOE-university network research to develop the capabilities and tools required by the applications and infrastructure administrators at DOE labs and select universities, as well as deploy DOE2000 tools and capabilities to support critical DOE mission applications. DOE will support the researchers at both the labs and the universities, to enable them to adapt their DOE application codes to make use of these new technologies as they are being developed, and to work with the network researchers to ensure that the new technologies are responsive to their application requirements.

DOE will also support enhancements of certain critical path infrastructure elements such as ESnet, aggregation and interconnection points (for example, gigapops), and local networks and services, to implement and support these new technologies to provide the appropriate level of end-to-end services to the application. Network management and analysis tools that function across networks and administrative boundaries at interconnection points and support these new capabilities will be developed.

DOE FY 1999 research focus

In FY 1999 DOE will experiment with and integrate network and applications research technologies on LSN and NGI testbeds, which include DOE laboratories, universities, and other Federal research centers.

DOE's goal 1 activities will focus on providing easy to use middleware that is network aware and able to interact with the network infrastructure, including system software, tools, and libraries to support DOE applications. The objective is to enable applications to exercise more efficient and smarter control and use of network resources and to support greater end-to-end capabilities required by DOE's applications.

In support of Goal 2, DOE will coordinate interconnection and peering arrangements and mechanisms with NGI networks to satisfy agency applications needs and provide programmatically-justified access to DOE's on-line facilities. This will include interconnecting with the other Federal agency and university networks at various speeds, locations, and media types (that is, IP, ATM, and WDM). It will also include developing and supporting advanced tools and infrastructure to address the technical and business issues associated with managing and supporting cross-domain internetwork interconnections and peerings.

DOE Goal 3 applications that require NGI-type technologies and infrastructure are largely part of the DOE2000 initiative. Pilot projects include the Materials MicroCharacterization Collaboratory and the Diesel

Combustion Collaboratory. The latter focuses on accelerating the iterative process between research and deployment by creating an environment that provides tools for archiving, sharing, discussing, and protecting scientific and proprietary information for participants in diesel combustion research, which includes major U.S. diesel engine manufacturers. One of the first steps — currently in progress — is establishing a public key infrastructure for protecting proprietary information centralized in an image library that will be accessed remotely by all partners.

CIC R&D Highlights

Netamorphosis

On March 11-13, 1998, Congressional representatives, administration officials, and the general public had an opportunity to view the future of the Internet firsthand at "Netamorphosis," a demonstration of technologies and applications being developed by the Federal NGI R&D initiative. In cooperation with Highway 1 and the High Performance Computing and Communications Consortium, representatives from the White House, seven Federal agencies, academia, and industry showed members of Congress how further development of Internet technologies will lead to advancements in healthcare, the environment, manufacturing, defense, and education. The demonstrations were mounted by the agency members of the LSN Working Group — DARPA, DOE, NASA, NIH, NIST, NOAA, NSA, NSF, and VA. Support and technical contributions were also provided by Bell Atlantic, Cisco, Highway1, IBM, Internet2/University Corporation for Advanced Internet Development (UCAID), and MCI.

Netamorphosis demonstrations were made possible by a state-of-the-art networking systems including NSF's vBNS and DARPA's ATDnet. However, all demonstrated applications require further advancements in networking technologies to make them fully functional, widely available, and affordable. Demonstrations included:

Advanced Regional Prediction System (ARPS)

Researchers at the University of Oklahoma, with Federal funding from NSF, NOAA, and the Federal Aviation Administration (FAA), demonstrated a prototype numerical weather prediction system developed to forecast intense small-scale spring and winter storms up to several hours in advance. In some cases, the system can pinpoint their locations within one or two counties.

"This new technology has the potential to save lives and millions of dollars in weather-related revenues lost each year in commercial aviation, agriculture, surface transportation, power and communications utilities, and recreation," said Kelvin Droegemeier, Director of the University of Oklahoma Center for Analysis and Prediction of Storms (CAPS), and one of the demonstration's participants. "In forecasting, speed is the name of the game. Unfortunately, current limited network capacities prevent CAPS from providing all available data to forecasters and other end-users in a timely manner. We need about a

Demonstration of virtual reality simulation of wide area environments such as the Chesapeake Bay environment using the Cave5D ImmersaDesk .

factor of 100 more bandwidth, and it is just not available right now. This is where the NGI Initiative will provide a valuable service," said Droegemeier.

Other Netamorphosis demonstrations included:

Automating the construction site — a leap in capabilities

TETRA, a versatile robotic crane, is one of several experimental technologies being investigated at NIST's new National Construction Automation Testbed (NCAT). These technologies will be used in developing standards for wireless tracking, positioning, and control of machines at construction sites and for supporting interactive construction management systems. At NCAT, researchers are developing an infrastructure for modeling, simulating, and automating dangerous or error-prone construction tasks; remote site-management; and up-to-the-moment access to all site-related information, from architectural designs, to subcontractor schedules, to the status of machinery. Anticipated benefits include improvements in safety, productivity, and quality.

❒ NGI research needs include the transmission and processing of huge volumes of position data and other information originating from global positioning satellites, hard-hat-mounted displays, and innovative surveying systems for wireless tracking of machines and other construction-site components.

Cave5D: A tool for collaborative immersive visualization of environmental data

Cave5D uses atmospheric and oceanographic data to create interactive virtual reality simulations of wide area environments such as the Chesapeake Bay. Simulations are viewed in 3-D on the ImmersaDesk. Researchers study the effects of physical phenomena such as wind, rain, and chemical pollutants on ecosystem behavior. For example, simulations can help scientists and environmental and fisheries managers forecast the likely path of an oil spill and analyze how winds and tides affect distributions of larval fish.

❒ NGI research needs include an advanced networking infrastructure linking ImmersaDesk sites to allow scientists in many locations to collaborate on environmental research.

Collaboratory for structure-based drug design

A prototype collaborative environment for carrying out interactive 3-D studies of molecular structure among scientists at distant locations. The collaboratory is used for drug design, protein engineering, biomaterials design and fabrication, and bioremediation.

- ❒ NGI research needs include advancements in high performance network access, collaborative molecular modeling software, and desktop videoconferencing.

Earth data from satellite to desktop

With the Distributed Image SpreadSheet (DISS), scientists working at their own desks visualize, manipulate, and analyze massive amounts of data collected by Earth Observing System satellites and stored at specialized archive centers around the country. The current version of the DISS, available to researchers for evaluation, enables the analysis of large amounts of data produced by next generation satellite systems.

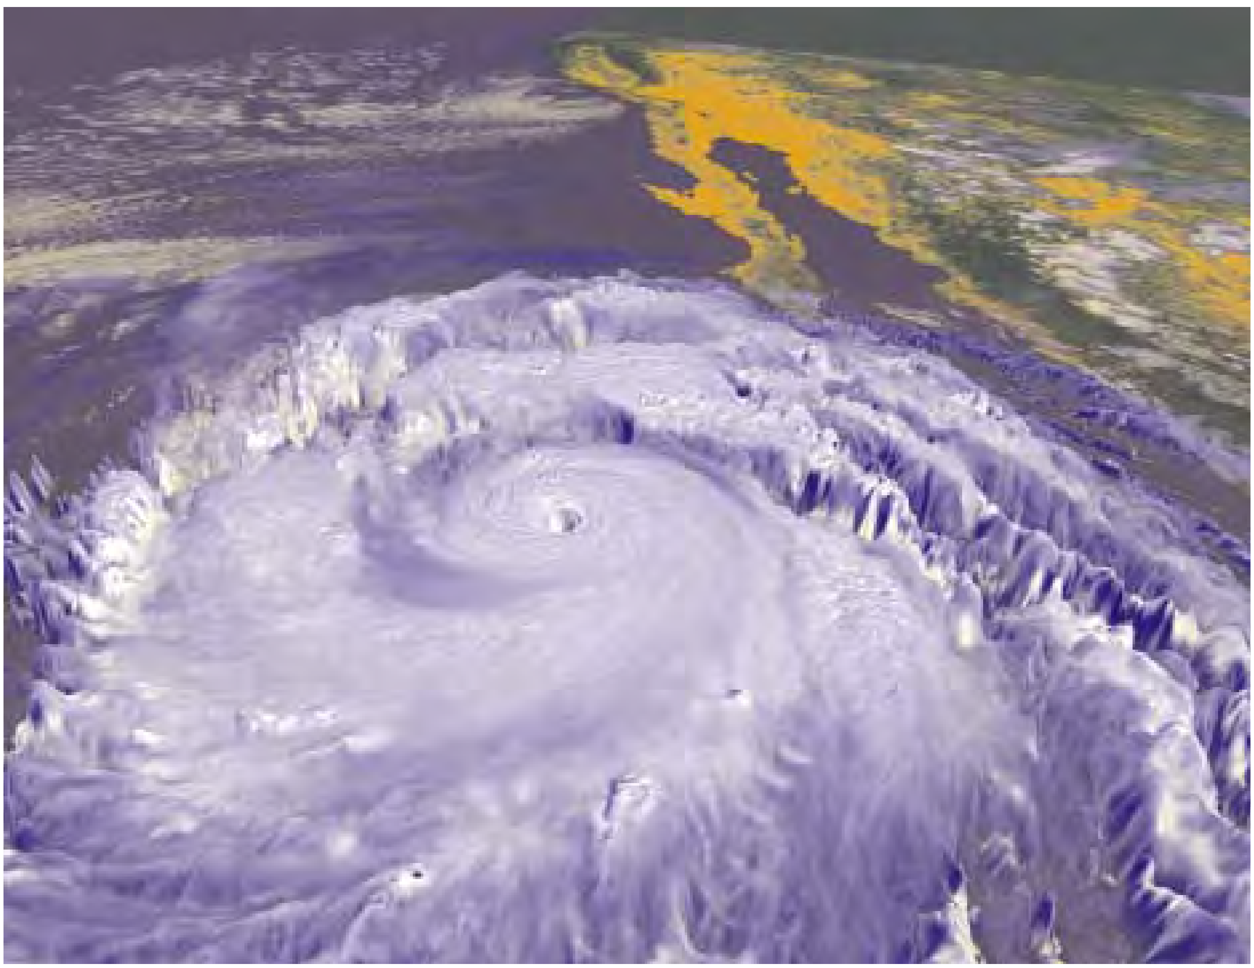

High performance connections between data repositories and the scientists' desktops are critical to remote visualization. NREN and the Distributed Image SpreadSheet allowed scientists to visualize, analyze, and manipulate years of Earth Observing Satellite data at their desktops. Pictured is an image of Hurricane Linda.

Senator Bob Kerrey (D-Nebraska) discusses the Netamorphosis demonstrations with Presidential Advisory Committee Co-Chairs Bill Joy (Sun Microsystems) and Ken Kennedy (Rice University).

- ❐ Currently using OC-3 bandwidth, a significant improvement over T-3, the DISS requires very low delay times and steady picture quality. NGI research needs include advancements toward OC-12 or higher bandwidths on distributed file systems and optimization of input/output performance.

Echocardiography over the NGI

Interactive echocardiography generates full-motion video of cardiac structure and cardiovascular blood flow and delivers these images in real-time to physicians in remote locations. It is being used at Cleveland Clinic Foundation, satellite facilities in Ohio and Florida, and the clinic's outpatient labs. Echocardiographs are transmitted from the cardiac operating room to locations where cardiologists provide guidance even when they are not physically present in the operating room. Echocardiography images are also relayed from satellite facilities to the main facility for diagnosis. Remote echocardiography will be critical for the future international space station, in battlefield conditions, and in medically underserved areas around the world.

- ❐ NGI research needs include increased guaranteed bandwidth and end-to-end differentiation of service to distribute 30 megabytes/second of full-screen, full-motion image data over wide-area networks.

Exploring the Earth system on the "Second Web"

This environment for exploring the Earth system uses new 3-D Web technologies to teleport viewers into high-resolution, stereo/3-D explorations of tropical storms, forest fires, clear air turbulence, cyclones, and El Niño. Researchers and educators create and share Earth system data to study patterns and behaviors behind naturally-occurring and sometimes dangerous phenomena.

- ❐ NGI research needs include high-bandwidth, wide-area networks with end-to-end differentiation of service, allowing multiple remote users to explore science together in virtual 3-D worlds.

Informedia News-on-Demand

Informedia News-on-Demand provides “full content” search and retrieval from broadcast TV and radio news. Present applications include digital video library access for K-12 education and evaluation prototypes at DARPA and NSA.

- ❒ NGI research needs include networking advancements in end-to-end differentiation of service, dynamic bandwidth adaptation, and secure and shared data annotation capabilities.

Interactive video dialogues

These multimedia voice-controlled scenarios engage users and video characters in face-to-face dialogues concerning realistic, dramatic situations using the IDI (Interactive Drama Inc.) Defense Language Institute’s Conversim interface software. Interactive video dialogues have a wide range of training and education uses, from combat casualty triage training, to information kiosks, to language training. For example, military linguists located anywhere in the world can sustain language proficiency by routinely talking in virtual dialogue with native speakers.

- ❒ NGI research needs include enhanced network capacity to accommodate large quantities of full-motion broadcast-quality video in an interactive format.

MAGIC: Viewing large geographic areas in 3-D

The MAGIC information system allows interactive viewing of large geographic areas in 3-D by retrieving and merging massive volumes of remotely located data, including aerial and satellite imagery, and data describing topography, weather, buildings, and other objects. MAGIC is applicable to military personnel planning a mission or assessing battle damage, emergency teams fighting a forest fire or organizing hurricane relief efforts, and environmental workers evaluating flood conditions.

- ❒ NGI research needs include widespread access to high-speed networks, distributed processing, and techniques for controlling access to computing and data resources. Combined with MAGIC’s application- and network-based storage system, these capabilities will enable people in disparate locations to respond quickly and collaboratively to time-sensitive events.

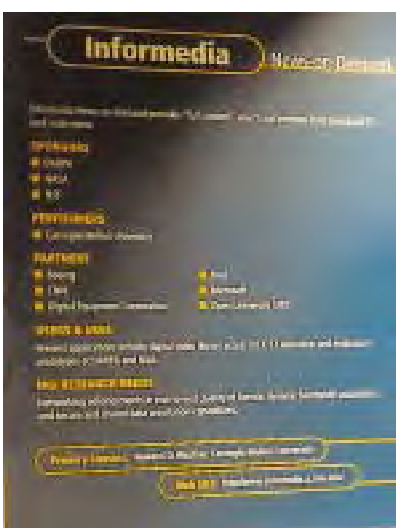

Explanatory posters (left) were displayed for each demo, highlighting the application, its NGI technology requirements, and sponsoring institutions or agencies. The MAGIC information system (right) demonstrated interactive viewing of large geographic areas in 3-D by retrieving and merging massive volumes of remotely located data.

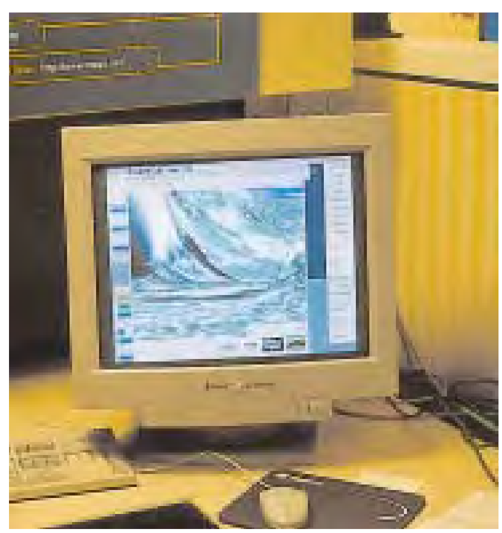

Setting up the camera for the MMC collaboratory demonstration. The MMC allows scientists and researchers to use state-of-the-art instruments to study the structure and composition of materials.

Materials Microcharacterization Collaboratory (MMC): Studying state-of-the-art materials

The MMC, linking five national research facilities, allows scientists and researchers to use state-of-the-art instruments to study the structure and composition of metals, ceramics, and alloys, and to conduct basic and applied research in materials science — especially strength and fracture toughness, corrosion and wear resistance, critical fields and transport current in superconductors, electronic and magnetic properties of metallic alloys, ceramics, semiconductors, artificially layered superlattices, and the properties of catalysts and engineered materials.

❒ NGI research needs include security, privacy, reliable connectivity, real-time sustained high-speed data delivery, video/audio conferencing capabilities, and end-to-end differentiation of service assurances that unite the five research facilities into a single online interactive Internet-accessible laboratory.

Nanotechnology research: Controlling atoms from a distance

The Scanning Tunneling Microscope (STM) can measure and manipulate atomic structures (measured in nanometers, or billionths of a meter), whose images have been magnified to a workable human scale. The Field Ion Microscope is used to shape the STM's probe, or tip, so it can accurately measure these structures. Manufacturing researchers use the STM and its close relative, the Atomic Force Microscope, as quality control tools for developing standard measurements of small-scale products and their component parts, such as computer chips and their circuitry.

❒ NGI research needs include a high-speed, secure, reliable network and simultaneous voice, video, and data transmission to make these microscopes, located at NIST'S Gaithersburg, MD, campus, accessible by remote control to companies and universities nationwide.

Octahedral hexapod: An Information Age machine tool

The hexapod is an innovative experimental metal-cutting machine tool with the potential to deliver an unprecedented combination of versatility, speed, accuracy, and portability. Industry and university researchers are working with NIST to investigate the hexapod's potential performance advantages, from lower production costs to faster methods for making parts, molds, and dies.

NIST exhibitors demonstrated real-time control and interaction of the octahedral hexapod.

❐ NGI research needs include real-time, full-motion video plus 15 Mbps bandwidth in a dedicated, secure environment — technology not widely available and currently prohibitively expensive for most organizations that contribute to and benefit from network-enabled manufacturing research collaborations.

Real-time functional MRI: Watching the brain in action

The "brain in action" allows remote viewing of brain activity while a patient is performing cognitive or sensory-motor tasks. Neurosurgeons, neurologists, psychiatrists, and brain scientists investigate brain function to diagnose and treat brain diseases. This application will enable neurosurgeons to develop surgical plans for removing a tumor based on an understanding of the cognitive and sensory-motor abilities located near the tumor site.

❐ NGI research needs include improvements in available capacity, interactive real-time capability, security, privacy, and integration with advanced computing to ensure high performance, widespread availability, online visualization, and patient confidentiality.

Security technology for the Next Generation Internet

A NIST prototype of Internet security protocol technology created a secure virtual private network, over the commercial Internet, for the telerobotic control data being transmitted between the Highway 1 demonstration site and the NIST campus. More than 60 commercial vendors of Internet technology, Internet service providers, universities, and Federal laboratories

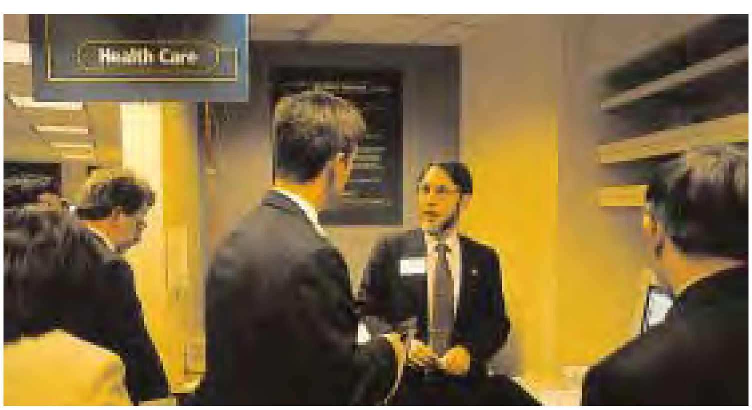

Michael Ackerman of NIH/NLM outlines how biomedical research applications and healthcare will benefit from developing NGI technologies.

The Visible Human Project has produced a highly-detailed digital image library of normal male and female anatomy. Images can be accessed for testing medical imaging algorithms, and as an atlas of the human body by healthcare researchers and users. Such uses require the high bandwidth and data security and integrity that are a major focus of NGI R&D efforts.

use this technology as a basis for further research into Internet security and as a reference from which to build and test commercial products.

- ❒ NGI research needs include further work in security technology to address issues such as scalability, support for multicast, and automated management of cryptographic keys.

SF Express: Advanced battle simulation

SF Express is a military training tool that integrates data on weather; terrain; land, sea, and air assets; and other strategic factors to create simulated battle environments. DoD personnel will create realistic distributed synthetic environments for training and mission rehearsal.

- ❒ NGI research needs include high bandwidth, wide-area networks accommodating multiple users and real-time interactivity requirements.

Visible Human Project

The project has produced a digital image library of normal male and female human anatomy, combining CT, MR, and digital anatomical images that allow the creation of 3-D renderings of any part of the human body. This library serves as a common reference point for studying human anatomy, as a source for common public domain data for testing medical imaging algorithms, and as a testbed and model for constructing other network-accessible image libraries. Physicians, educators, artists, mathematicians, engineers, physicists, and industry researchers incorporate Visible Human data into applications ranging from virtual reality surgical simulations and crash simulation models to kindergarten curricula and multimedia artwork.

- ❒ NGI research needs include high bandwidth on demand, end-to-end differentiation of service, and security for patient data.

Telehealth

Research in telehealth

A major goal of Federal CIC R&D investments in telehealth, or remotely-provided healthcare, is to close the gaps in healthcare expertise that exist between major urban medical centers and remote rural healthcare facilities. NLM, VA, and AHCPR support numerous telehealth projects. Cardiology, dermatology, ophthalmology, pathology, psychology, radiology, and surgery (endoscopic/laparoscopic) are some of the application areas where telehealth has been tested and evaluation has begun.

Computer-based patient records (CPR)

The primary function of the CPR is to support the delivery of medical care to patients by facilitating communication among healthcare providers and documenting the reasoning behind clinical decisions. CPRs can also be used to build information repositories about effective medical treatments that could improve population-based care.

CPR data can be entered from keyboards, by dictation and transcription, voice recognition and interpretation, light pens, touch screen terminals, and personal digital assistants. Data can also be input automatically via electronic patient monitors and bedside terminals, nursing stations, EKG systems, laboratory autoanalyzers, and clinical imaging digital systems.

Patient care data collected by a CPR system can be stored centrally or distributed, and can be retrieved by authorized users, who can then view the data as text, tables, graphs, sound, images, or full-motion video.

Clinical decision support systems (CDSS)

As a component of clinical decision support systems (CDSSs) — software designed to aid clinical decision making — CPRs can provide physicians with medical knowledge pertinent to patient care, such as diagnostic suggestions, testing prompts, therapeutic protocols, practice guidelines, alerts of potential drug-drug and drug-food reactions, and treatment suggestions. The link between the CPR and CDSS is the "knowledge server" that acquires the necessary information for the decision maker.

Knowledge sources range from the medical literature, which can be searched for review articles and specific subjects using NLM's Internet-based Grateful Med program to explore the Medline literature data base; to evidence-based practice reports sponsored by AHCPR; the Physicians Data Query program at NCI; other consensus panel guidelines sponsored by the NIH; guidelines developed by medical and specialty societies; and internal development and approval by a hospital's staff.

Software can provide rule-based alerts, reminders, and suggestions for the care provider when and where health service is delivered. This software can be developed as independent, portable modules. NLM's Unified Medical Language System (UMLS) provides a uniform medical nomenclature consistent with scientific literature, which is needed for widespread use of CDSSs.

Research databases

CPRs are useful in developing research data bases, medical knowledge, and quality assurance information. Outcomes research supported by AHCPR, such as evidence-based studies, can provide consumers, providers, administrators, and insurers with cost and medical effectiveness information. National repositories make systematic reviews such as statistical analyses possible. High performance computing and communications systems are needed to perform these population studies, retrieve data, produce information and knowledge, and send them to authorized users promptly and safely.

CPR in telehealth

Using CPR in telehealth permits the exchange of information between clinics and healthcare providers' offices in cities and rural areas, filling in the gaps where specialized expertise may not be available, thus allowing remote patients to receive proficient medical care. For example, digital images, such as MRI, should appear no different whether viewed at a patient's site of care or by a radiologist hundreds of miles away from the patient — assuming compression and decompression algorithms that reduce transmission time and cost do not cause discernible loss of quality. Digital X-rays, on the other hand, are produced by scanning conventional film, making it necessary for radiologists to have training and specialized experience. In both examples, digital information can be assembled, shared, and discussed with appropriate specialists via telehealth communications links.

High Confidence Systems

Overview

HCS R&D focuses on the critical technologies necessary to achieve high levels of availability, reliability, security, protection, and restorability of information services. Systems that employ these technologies will be resistant to component failure and malicious manipulation and will respond to damage or perceived threat by adaptation or reconfiguration. HCS R&D supports interagency collaborations for Federal high confidence systems.

Applications requiring HCS technologies include national security, law enforcement, life- and safety-critical systems, personal privacy, and the protection of critical elements of the National Information Infrastructure. Systems for power generation and distribution, banking, telecommunications, medical implants, automated surgical assistants, and transportation also need reliable computing and telecommunication technologies. This section highlights some recent accomplishments in HCS R&D.

Information survivability

DARPA's Information Survivability program is developing technologies that can be used to create survivable systems. These technologies will create strong barriers to attack, will detect malicious and suspicious activity, will isolate and repel such activity, and can be used to guarantee minimum essential continued operation of critical system functions in the face of concerted information warfare attacks. The program aims to create affordable, verifiable, scalable technologies for a robust and secure Defense infrastructure — technologies that will enable the construction of secure enclaves and allow distributed computing to span such enclaves.

This program is creating advanced technologies that can be used to protect DoD's mission-critical capabilities as well as critical national infrastructures against electronic attack upon or through their supporting computing infrastructure. The technologies developed by the program will provide the strength needed for DoD while retaining the cost savings resulting from the use of commercial technologies. Following are a few highlights:

- ❐ Under DARPA funding, the Boeing Corporation has developed and successfully demonstrated the Intruder Detection and Isolation Protocol. This involves a cooperative exchange of information about intrusion attacks by network components in order to isolate and cut off an attack. In a successful demonstration, ten attacks on the demonstration environment were detected and isolated.

- DARPA has funded Reliable Software Technologies, Inc. to develop a tool for whitebox analysis of security vulnerabilities in source code. This tool has been used to identify potential security flaws in one of the most common servers in use on the Internet. Analysis revealed how a dozen simple perturbations and three buffer overflow flaws could compromise the security of the FTP server's system, allowing unauthorized retrievals of sensitive files. The results demonstrate that automated tools can be used to detect potential areas of weakness, enabling developers to take preventive actions to fortify their products against malicious attacks.
- DARPA has funded Secure Computing Corporation (SCC) to develop an approach for protecting users browsing the Web from malicious attacks on their systems. The approach allows users to run their favorite browser in a confined manner so that any actions initiated by the browser or its children are restricted to only those files that the user has specifically permitted. In particular, rogue Java applets, JavaScripts, Netscape plug-ins, and Microsoft ActiveX components cannot access portions of the system that the user has declared off-limits to them. The code used to implement the protection is called a kernel hypervisor. To illustrate the approach, SCC has prototyped a Netscape Navigator kernel hypervisor. The prototype was developed for a Linux system, but the approach is applicable to many other modern operating systems, including Sun's Solaris and Windows NT. The kernel hypervisor is not bypassable, since it runs in the kernel. Also, since it is a loadable module, it does not require the kernel to be modified or even rebooted when it is loaded. SCC is currently investigating other potential applications.

Information security

NSA's Information Security (INFOSEC) Research Program continues to deliver a broad range of security technology solutions. Fundamental mathematical work in cryptography, including elliptic curve technology, has produced more secure and efficient algorithms for privacy protection and authentication, while analytic work in electronic cash technology has provided valuable guidance to the financial and legal communities. NSA has provided demonstrations and standards developments to ease the integration of security services into commercial products and services. Engineering breakthroughs in high speed/low power electronics and in optical encryption technology will provide the foundation for emerging high performance communication systems. Improved biometric authentication techniques are finding widespread acceptance for improving government and commercial access control systems. Security enhancements for next generation operating systems and for object technology have been developed and transferred to the R&D community. New visualization and risk assessment tools have been developed and applied to assessing system security. Finally, NSA has established cooperation across the INFOSEC research community to address network security.

NSA has developed a technology forecast and set of challenge problems that focus on the development of a high assurance computing platform, technology for secure internetworking, and technologies needed for a high assurance security management infrastructure. The technology needs and gaps of these challenge problems will direct the bulk of NSA's INFOSEC research resources. Problem areas that need to be addressed include the development of system security engineering methods to specify and design security characteristics into a system; the management of network security and the development of an infrastructure to support that management; tools and techniques to detect and respond to local and national level attacks on critical information systems and infrastructure components; the development of strong mechanisms to allow the controlled sharing of information among disparate communities; and improved assurance technology for increasing the level of trust in the secure operation of system hardware, software, and procedures. Following are some highlights:

- ❒ NSA is conducting research on technologies for high-speed encryption. In 1997, NSA engineers, in collaboration with DOE's Sandia National Laboratory, completed the design for an ATM encryptor that operates at 10 Gbps. In 1998 they evaluate the vulnerability of the chip, which will be a follow-on to the FASTLANE ATM encryptor. This design is capable of context (key, algorithm, mode) agility rather than simply key agility.

- ❒ NSA scientists have invented a process for fabricating etched mirrors on semiconductor laser surfaces, a process necessary to cascade optical logic gates on a single chip, and an advancement needed for the development of an all-optical encryptor.

- ❒ NSA has authored the Internet Security Association and Key Management Protocol (ISAKMP), an Internet draft standard. NSA researchers are modeling ISAKMP and analyzing it for completeness and security. The model contains more than 100 different data flow diagrams, state transition diagrams, and mini-specifications. Included within the more than 50 state transition diagrams are 250 states and over 600 state transitions. Future modeling will include interaction of ISAKMP with other protocols.

- ❒ NSA contracted with Virtual Motion Inc. to deliver a driver with a protocol structure that enables a laptop user to encrypt network data independent of the particular vendor's wired or wireless PC card. This improves the original idea of modifying each vendor's device driver to operate with a Fortezza encryption card. A wireless local-area network was also developed that uses the Fortezza card for security.

Assurance technologies

NASA is developing several technologies to help achieve high confidence in system safety. Following are some highlights:

- ❒ NASA has funded Odyssey Research Associates (ORA) and Honeywell Air Transport Systems Division to study the incorporation of formal methods into the company's software development processes. In particular, ORA developed TableWise, a prototype tool to analyze the characteristics of decision tables. A decision table is a tabular format for defining the rules that choose a particular action to perform based on the values of certain parameters. TableWise generalizes binary decision diagrams to determine if a particular table is exclusive (for

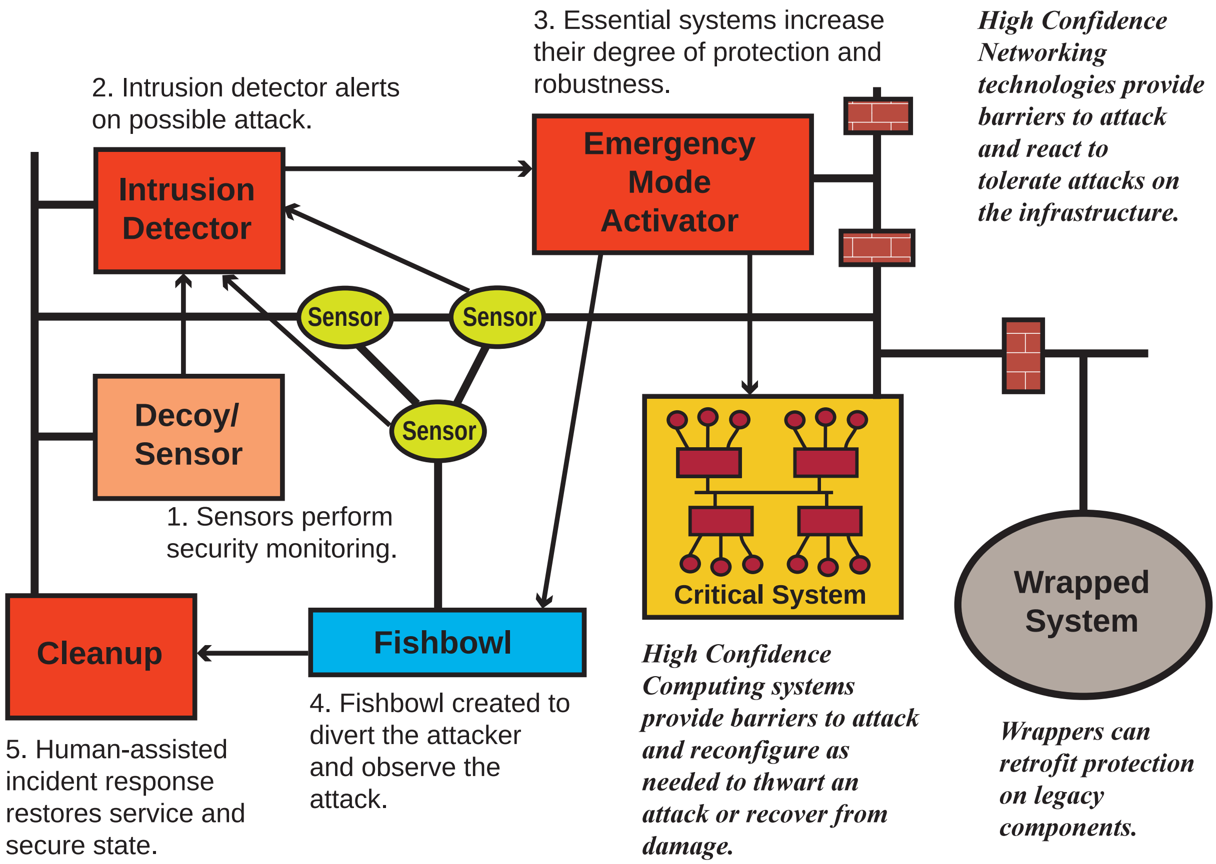

DARPA information survivability research focuses on technology that will guarantee that critical information systems continue to function adequately in the face of attack, even when the precise type of attack has not been anticipated.

every combination of parameter values, at most one action can be chosen) and exhaustive (for every combination of parameter values, at least one action can be chosen). The tool can automatically generate documentation and Ada code from a decision table. Honeywell is evaluating TableWise in developing their commercial autopilots, where the bulk of the creative effort is in developing and validating the mode selection logic.

❒ NASA has initiated a joint project with Rockwell Collins to investigate the use of formal methods to reduce mode confusion in the cockpit. The project focuses on a composite flight guidance system appropriate for business jets. Using a formal model to drive an animation of the user interface while at the same time displaying the system's behavior is intended to help users and designers develop an architecture that matches the mental model of the flight crew. The investigation will look to formal methods to model the behavior of the system, visualize the internal states of the model, drive an animation of the user interface directly from the formal model, and analyze the models for desired properties such as consistency, completeness, and safety.

Protecting privacy for medical records

In FY 1998 and FY 1999, NLM and AHCPR will continue to support research in technologies for storing and transmitting patients' medical records while protecting the accuracy and privacy of those records. Projects will promote the application of HCS technologies to healthcare, telemedicine evaluation, and the testing of methods for protecting the authenticity, integrity, confidentiality, and privacy of electronic health data.

Secure Internet programming

NSF is supporting a secure Internet programming project at Princeton University that focuses on the security of mobile code systems such as Java, JavaScript, and ActiveX. Software-based protection can allow for more extensible security models that improve performance over hardware-based solutions. Extensible security mechanisms can protect subsystems and implement policies created after the original system has been shipped. This project has identified and analyzed different software-based security schemes and has popularized the extended stack inspection model.

National Information Assurance Partnership (NIAP) program and Role-Based Access Control (RBAC)

Systems and networks are trusted to perform their intended functions with a high degree of confidence. Systems and networks performing mission-critical functions or managing high-value assets or embedded systems require unprecedented levels of reliability and quality. Two NIST programs focus directly on these needs.

Under the NIAP program, NIST has partnered with NSA to establish a center to foster the development of formal laboratories to test and certify security products against published formal specifications. This program will help ensure that both vendors and users can cite third-party assurance of the functionality and quality of security products and systems.

In the complex information technology environment, the careful and correct specification of rules to control access to online documents, capabilities, or systems has become critical — and increasingly difficult. While traditional access control methods focus on individual users, files, or other system objects, management of access in the real world is more often based on the role that a user assumes. NIST has pioneered the new RBAC model that better meets the needs of user organizations and is implementing it in environments, including a Web-based application.

FAA high confidence systems activities

Beginning in FY 1997, the FAA has participated in the coordinated CIC R&D process, particularly through the HCS Working Group. The FAA is interested in high confidence systems from two viewpoints:

❒ Achieving a higher degree of system dependability

❒ Assuring the integrity and security of aviation information systems.

The FAA has established "Streamlining Software Aspects of Certification," a program to obtain faster certification at reduced cost while maintaining high confidence and reliability. Other FAA-funded research is investigating how to rigorously define architecture constraints for protecting safety-critical processes (performed by SRI International), and approaches for structural test coverage analysis (performed by NASA and Boeing).

Future HCS R&D

The high confidence community is in the process of defining a future multiagency research program that would address the following research needs:

- ❐ Architectural approaches to isolate safety-critical system components from non-safety-critical components
- ❐ Approaches to achieving survivability for collections of autonomous entities
- ❐ Approaches and techniques for testing and verifying integrity. For example, automated testing could lower the cost while increasing the level of assurance of high confidence systems.
- ❐ Measuring confidence in integrity, availability, and overall security of highly complex systems. Metrics are needed to identify whether confidence has been increased or decreased over time.
- ❐ Evaluation methods that would shorten the security product evaluation cycle. Such methods should allow evaluation to be completed within a fraction of the half-life of the product.
- ❐ Strategies for securing a system of systems, including intrusion detection and monitoring techniques. Examples include the use of intelligent agents for security administration and monitoring; one-time sign-on for large systems of systems; verified secure communication protocols; key management and key distribution for very large systems; and techniques to identify and contain denial of service attacks in high-speed networks.
- ❐ Non-intrusive access control technology to protect individual privacy
- ❐ Technology for performing information warfare situation assessment, in order to discern whether an attack against the Nation's critical infrastructure may be in progress.

Human Centered Systems

HuCS technologies

HuCS R&D leads to increased accessibility and usability of computing systems and communications networks. Scientists, engineers, educators, students, the workforce, and the general public are all potential beneficiaries of HuCS technologies, which include:

- ❒ Collaboratories
- ❒ Information agents for collecting and analyzing data
- ❒ Interdisciplinary research in human and distributed cognition applied to environments
- ❒ Knowledge repositories for information access, management, and applications
- ❒ Missions and applications
- ❒ Multilingual technology for document translation and understanding
- ❒ Multimodal interactions between humans and computer systems
- ❒ Universal access
- ❒ Virtual reality environments
- ❒ Visualization systems and tools/multimodal presentation

HuCS FY 1998 and FY 1999 R&D will address the following areas:

Digital Libraries Phase 2

Building on the successful interagency Digital Libraries initiative, the new Digital Libraries Phase II initiative will begin in FY 1998. It is a joint effort of NSF, DARPA, NASA, NLM, the Library of Congress (LoC), and the National Endowment for the Humanities (NEH) in partnership with the National Archives and Records Administration and the Smithsonian Institution. The new initiative will place a greater emphasis on the human and societal dimensions of digital libraries, including research in electronic information life cycles; digital interoperability; integration of information, computing, and communication technologies with human needs; and new types of content and collections.

Since digital libraries can serve as intellectual infrastructure, this initiative will stimulate the partnerships necessary for creating next generation operational systems in such areas as education, engineering and design, Earth and space sciences, biosciences, geography, economics, and the arts and humanities. It will address the entire digital libraries life cycle from information creation, access, and use, to archiving and preservation.

The goal of the initiative is to provide leadership in fundamental research, to advance the usability of globally distributed, networked information resources, and encourage a focus on innovative applications. Digital Libraries Phase II plans to:

- ❒ Selectively build on and extend research and testbed activities in promising digital libraries areas
- ❒ Accelerate development, management, and accessibility of digital content and collections
- ❒ Create new capabilities and opportunities for digital libraries to serve existing and new user communities, including all levels of education
- ❒ Encourage the study of interactions between humans and digital libraries in various social and organizational contexts

Computer and Information Science and Engineering (CISE)

HuCS R&D programs in NSF's Computer and Information Science and Engineering (CISE) Directorate focus on improving the fundamental understanding of computing and information processing by enhancing the training of scientists and engineers. Special attention is given to the computing and communications technologies — including software — employed to manage these processes and to select application areas. CISE FY 1998 HuCS R&D focuses on human centered systems and knowledge networks.

Knowledge and Distributed Intelligence (KDI) and Knowledge Networking (KN)

As a result of technological advances in computing power and connectivity, there are now unprecedented opportunities for providing rapid and efficient access to enormous amounts of knowledge and information; for studying vastly more complex systems than was hitherto possible; and for advancing in fundamental ways our understanding of learning and intelligence in living and engineered systems. NSF's Knowledge and Distributed Intelligence (KDI) theme is an NSF-wide effort to realize these opportunities. KDI's Knowledge Networking (KN) research will help facilitate the evolution from today's distributed information access environment to new technical and human capabilities for interactive knowledge creation and use. Through this evolution, interdisciplinary communities can be joined in sharing data and building knowledge to address complex problems that have been explored strictly within traditional disciplinary boundaries.

Finding and tracking information

NSF-supported research at Columbia University focuses on developing technologies to help people find and track the information they need to keep current in their jobs and lives. The Columbia Digital News System provides up-to-the-minute briefings on news of interest, linking users into an integrated collection of related multimedia documents. Depending on the user's profile or query, events can be tracked over time, and a summary of the most recent, breaking developments, including a representative set of images or videos, can be generated automatically. A user can follow up with multimedia queries to obtain more details and further information.

Active visualization

NSF-supported active visualization R&D is being undertaken at New York University. Human vision is an active process involving four principal eye movements: saccades — rapid shifts in the fovea, a small depression at the back of the retina that forms the point of sharpest vision; pursuit movements; vestibular (entrance) systems; and vergence (borders). Biological vision uses "foveated images" in which the resolution is higher at the fovea than at the periphery. To compensate for the loss of peripheral vision in images, active visualization technology allows the user interactive control of several parameters of vision, including rapid saccades and adjustable resolution. Active visualization techniques attempt to present images that match some of these eye motions.

Active visualization is particularly suitable for thin-wire communication of images. Since foveated images require a much smaller bandwidth for transmission, they are ideal for distributed computing, collaborative computing, and client-server applications over the Internet.

Speech technologies

Business Week magazine (February 12, 1998) highlighted the increasing importance of leading-edge speech technologies in computing. Continuing DARPA- and NIST-supported research has made significant contributions to this field, which is a HuCS R&D focus area. For example, the DARPA-supported Jupiter effort, cited in the *Business Week* article, is an experimental system that allows individuals to dial an 800 number to obtain weather information on 500 cities around the world. Jupiter will respond to spoken requests for weather information and answer these queries in its own voice in real time.

DARPA's Speech Recognition Technology Integration effort will provide an information navigation and exploration toolkit that can provide military decision makers with rapid access to multisource, multimedia information. In FY 1998, DARPA-led researchers will design, develop, integrate, and demonstrate key Human Computer Interaction (HCI), speech recognition, "News on Demand" (demonstrated at Netamorphosis), and video paragraphing technologies with various applications. Additionally, they will create a full command vocabulary for U.S. Atlantic Command operations and deliver a speech-enabled version of the Joint Readiness Automated Management System for use in the command center.

The Space and Naval Warfare Center is supporting research in the Command Center of the Future (CCOF). In FY 1998, commercial off-the-shelf (COTS) and DARPA speech engines capable of running in server mode will be installed onto either the Command Center's systems or other reachable systems. The core speech technologies comprising the speech kernel include automatic speech recognition, text-to-speech, speaker identification and authentication, and natural language understanding.

Collaboration among contributing DARPA researchers via integration and testing in the CCOF's speech environment will provide a feedback loop, *An investigator, with a head-mounted display and a force-feedback manipulator controlling the tip of a scanning tunneling microscope (STM) probe, explores the surface of a sample of material using NANOSCAPE, a virtual-reality interface for controlling an atomic force microscope and for observing its output as a three-dimensional image in real time. The image projected on the wall shows what the investigator sees and changes as his head position and orientation change. The region of the surface shown is 25 Angstroms on a side. The investigator can change the viewpoint, "feel" the contours of the surface, and control the microscope as it modifies the surface.*

preparing systems for installation aboard the USS Coronado, part of the sea-based Battle Lab Initiative. (The USS Coronado is an information warfare command ship, where an entire war can be planned, monitored, and directed.) One of the goals of this R&D is a consistent, integrated, and standardized speech applications programming interface (API) that can be made available to developers to encourage software reuse and to leverage both COTS and DARPA technologies.

Healthcare data

NLM's computer-based Medical Literature Analysis and Retrieval System (MEDLARS) was established to achieve rapid access to the Library's vast store of biomedical and health-related information. Through worldwide communications networks, MEDLARS search services are available online to individuals and institutions, and the system is accessed more than 18,000 times a day.

To make searching easier and to provide a user-friendly way to use MEDLARS, NLM developed Internet Grateful Med and PubMed, free services accessed via standard Web browsers. The availability of Internet Grateful Med has resulted in an upsurge of usage by health professionals as well as the general public. In FY 1998 NLM will continue developing and deploying new capabilities for automatic source selection and for retrieving and sorting information from multiple databases.

With the large and rapidly growing number of computerized database resources and services offering bibliographic, full text, and factual data via the Internet, the user must deal with structured retrieval languages that differ from one database service to the next. NLM's Unified Medical Language System (UMLS) compensates for the dissimilarity in the ways related information is classified in different automated systems. Intelligent-agent-mediated gateways provide users with a single point of access to information and free the user from having to know the peculiarities of the various information sources. Under continuing development and refinement in FY 1998 and FY 1999, the UMLS is functioning as an electronic Rosetta Stone, making the thicket of medical classifications invisible to the user and enabling retrieval of related biomedical information from many sources.

NCRR activities in collaboratories and virtual reality

NCRR-supported research resource centers can provide remote control of complex scientific instruments such as a 400 kilovolt (kv) electron microscope and a sophisticated nuclear magnetic resonance imaging system. In FY 1998, NCRR will grant up to four awards to establish, demonstrate, and evaluate collaboratory testbeds to interact with these resources. The goal is to demonstrate the efficiency and efficacy of conducting research by interactively linking people, instruments, databases, and sophisticated software.

In FY 1998, NCRR is supporting two resource centers that apply virtual reality/environments to enhance the capabilities of instruments such as the Nanomanipulator (described in the 1998 Supplement to the President's Budget) to enable visual and tactile interaction with large biomedical molecules such as proteins (Protein Interactive Theater) and to create controllable virtual environments for brain cognition research. Plans include establishing a resource center that will employ virtual and augmented reality for additional brain research and guided surgery.

The Visible Human Project

In FY 1998 and FY 1999, NLM is continuing to develop and evaluate a 55 gigabyte digital image library of male and female anatomy. Full use and understanding of the anatomical structures depicted in such a library requires the integration of high performance computing and communications technologies with technologies used in medical imaging systems, including CT and MR imaging. (Illustrated on page 82.) Combining this library with efficient rendering algorithms will provide new educational tools for researchers, healthcare providers, students, and the general public. NLM is working with academia and industry to encourage the development of interoperable methods for representing and communicating such electronic images.

More than 900 non-financial license agreements to access these data sets have been signed by governmental, commercial, and academic organizations. Applications using the data sets include multimedia materials for all levels of education; virtual reality programs including surgical simulators; and modeling applications such as surgical and radiation treatment planning, radiation absorption, ergometrics, and crash testing.

Systems Integration for Manufacturing Applications (SIMA)

Scientists and engineers using advanced software to develop and produce manufactured goods face obstacles that prevent seamless communication and interpretation of information among manufacturing applications as well as between people and the applications themselves. HuCS R&D at NIST

Section through Visible Human Male — head; including cerebellum, cerebral cortex, brain stem, nasal passages (from head data subset)

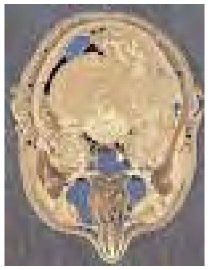

Section through Visible Human Male — thorax, including heart (with muscular left ventricle), lungs, spinal column, major vessels, musculature (from thorax data subset)

Section through Visible Human Male — abdomen, including large and small intestines, spinal column, musculature, subcutaneous fat (from abdomen data subset)

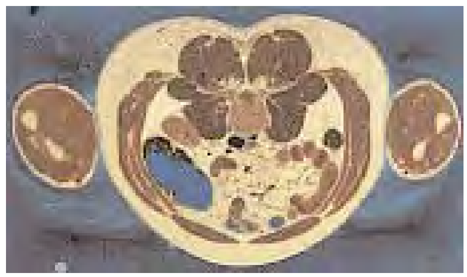

focuses on developing and testing information interfaces for manufacturing applications and on the application of HuCS technologies in support of these efforts.

In projects such as the Process Specification Language, NIST researchers work with collaborators from around the world to develop a unified process representation scheme. This will enable improved communication of manufacturing process information among different systems. As part of this effort, NIST is using Internet collaboration technologies that enable remote researchers to participate in virtual discussions and round table events.

NIST's Virtual Environments and Visualization for Manufacturing project is leveraging existing standards developments, such as VRML (Virtual Reality Modeling Language), to enable the integration of VRML environments from multiple application sources. A number of VRML exporters have been developed, one of which has recently been incorporated into a commercial product.

Other NIST efforts are making scientific and engineering data available via the Internet in support of engineering and manufacturing applications. The NIST Chemistry WebBook brings evaluated data directly to a scientist's desktop, simplifying the job of finding reliable property data for thousands of chemicals. The Engineering Statistics Handbook, being undertaken jointly by NIST and SEMATECH, will provide an Internet-accessible means for engineers to classify statistical problems, identify statistical solution methods, and be linked to software that can be used to execute those methods and analyze the results.

National Institute on Disability and Rehabilitation Research (NIDRR)

In FY 1998, the Department of Education's (ED) NIDRR will support four Rehabilitation Engineering Research Centers for Information Technology Access, Communication Enhancement, Universal Telecommunications Access, and Telerehabilitation. NIDRR also supports jointly with NSF a World Wide Web accessibility initiative at MIT. The major focus of all these activities is to learn how individuals with various disabilities might use advanced computational and communications tools. In such cases, the interface between the user and the system is critical, because not all sensory or motor modalities may be usable by persons with disabilities. For example, if an individual cannot see, a graphical user interface is not useful, and an alternative must be available.

NIDRR also supports ABLEDATA, a database on assistive technology devices for individuals with disabilities, and the National Rehabilitation Information Center that maintains REHABDATA, a rehabilitation bibliography database.

Regional Technology in Education Consortia (RTEC)

The ED RTECs provide state and local education agencies with technical assistance in developing comprehensive educational technology plans, drawing upon existing resources and anticipating future technological needs and innovations. In FY 1998, the RTECs will continue to provide assistance in integrating technology with teaching and learning, including the acquisition, maintenance, and effective use of new technologies. The RTECs will also focus on making the most valuable resources on technology more accessible to educators through a variety of sources including the Internet and telecommunications technologies. A new focus will be assisting schools in making the best use of Internet technologies.

Regional educational laboratories

ED's regional laboratories assist educators and policy makers in implementing effective school reform strategies. A major focus of the

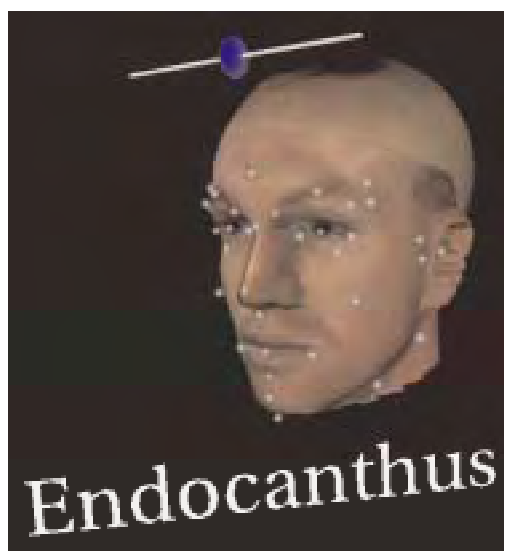

NIST-supported research developed this visual glossary of landmarks of the head. The 3-D face is overlaid with markers identifying anthropometric landmarks. By moving the cursor over the markers, the name of the landmark appears.

laboratories is to improve access to the research and resources of the labs. The labs will work with the Educational Resources Information Center (ERIC) and the RTECs to develop a strategy for providing access to decentralized full-text online materials, using a search engine that would include resources from all labs and other Department-funded institutions. Specific labs will continue to provide telecommunications services for outlying areas without modern technology; a vast network for teachers to share ideas and lessons and increase technological literacy; and tools to train teachers in effective uses of technology, and to model exemplary teaching and learning through video and CD-ROM.

Virtual reality

The ability to record and play back collaborative experiences will be critical to the future development of teleimmersion environments. The Voyager media server developed at Argonne National Lab enables researchers to experiment with scalable, multistream, multimedia recording and playback. Future extensions will enable the capture and playback of virtual reality-based teleimmersive experiences.

ManyWorlds, also developed at Argonne, is a scalable prototype system for constructing shared virtual reality environments. It provides the communications and sharing services needed to build large-scale teleimmersion systems that can be shared among multiple users and locations, and is being used as an experimental vehicle for exploring human performance issues in teleimmersion. ManyWorlds provides an environment that allows users to collaborate on data exploration, interactive steering of computational simulations, cooperative model construction, and access to remote instruments.

Electronic Notebook project

Electronic notebooks are beginning to be used by scientists in collaboratories to share ideas, data, experiments, and research programs. The goal of DOE’s Electronic Notebook Project is to design a modular, extensible architecture for a notebook that is easy to use, secure, cross platform, and interoperable. DOE’s electronic notebook is required to be a legally binding record that can be maintained in archives for as long as 25 years. This notebook will share components for input and display of sketches, text, equations, images, graphs, and other data types as well as provide tools for authentication, notarization, and other services.

STIMULATE

Overview

The Speech, Text, Image, and MULtimedia Advanced Technology Effort (STIMULATE) supports fundamental research devoted to understanding multimodal human communication and its application to computer technology. STIMULATE is a multiagency collaboration that includes researchers at NSF, NSA, and DARPA. The aim is to accelerate the progress of information technology by supporting new directions in R&D for understanding human communications in multiple languages and modalities, such as text, images, video, gestures, facial expressions, handwriting, and other means by which humans communicate. STIMULATE projects include:

Synergistic modalities for human/machine communication

Natural communication with machines is a crucial factor in bringing the benefits of networked computers to mass markets, and the sensory dimensions of sight, sound, and touch are, at present, comfortable and convenient modalities for the human user to interact with these computers. However, new technologies are now emerging that can support human/machine communication with features that emulate face-to-face interaction. Because speech is a preferred means for human information exchange, conversational interaction with machines will play a central role in collaborative knowledge work performed with networked computers. Using speech in combination with simultaneous visual gestures and physical signaling requires software agents that can fuse error-susceptible sensory information into reliable interpretations that are responsive to and anticipatory of human user intentions. A Rutgers University project seeks to design methods and evaluation metrics for providing human users the benefits of natural communication with computers.

Information retrieval

The Center for Intelligent Information Retrieval (CIIR) at the University of Massachusetts is investigating how to retrieve information from general image databases. Retrieval is based on image content and any associated text. Given an image database, images similar to an example image are retrieved. Similarity is evaluated on the basis of appearance (that is, the shape of the gray-level intensity surface), color, and texture. The CIIR's approach is not limited to specific image types, nor does it depend upon learning.

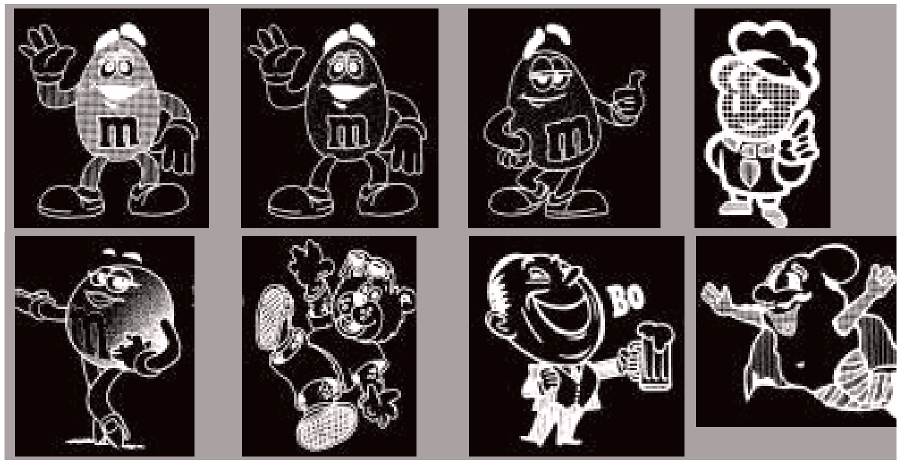

This figure is an example of a retrieval.performed at the CIIR at the University of Massachusetts. The queried database is a set of 2048 trademark images from the U.S. Patent and Trademark Office. The figure shows the retrieved images ranked as most similar to the first image (the query).

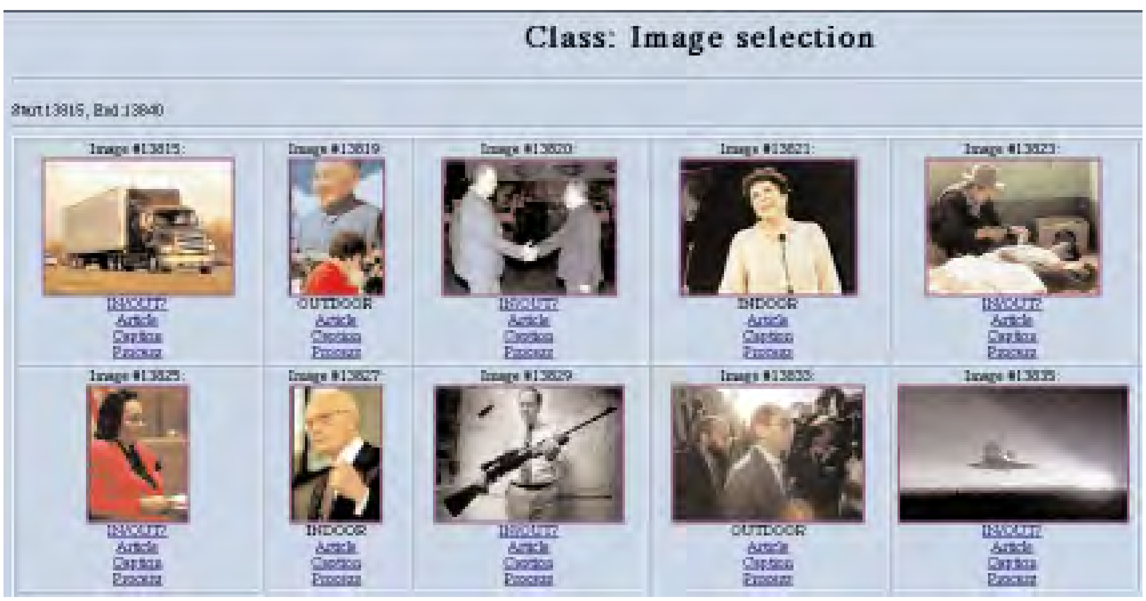

The figure above depicts the results of a sample query made to the Columbia Digital News System. After examining the retrieved images, the user can link to articles, captions, and other information associated with them.

Search over live multimedia information

Research at Columbia University focuses on developing technologies to help people find and track the information they need to keep current in their jobs. The Columbia Digital News System provides up-to-the-minute news briefings on topics of interest, linking the user into an integrated collection of related multimedia documents. Depending on the user's profile or query, events will be tracked over time and an automatically generated summary of the most recent developments provided. A representative set of images or videos can be incorporated into the summary, and the user can follow up with multimedia queries for more details and further information. In order to allow searches over images, a major component of this research is categorizing images through coordinated use of image features and associated text.

First annual STIMULATE workshop

The first annual STIMULATE workshop was held at NSF on March 5-6, 1998. Numerous Government agency representatives attended this workshop, which featured fifteen project presentations focusing on topics such as "Video Scene Segmentation and Classification Using Motion and Audio Information" and "Human-Computer Communication and Collaboration."

Collaboratories

Collaboratories and collaborative technologies

In the future, science and education will operate in a highly distributed but fully connected environment where every element is aware of and can interact with all other elements. Research teams will be able to work together effectively regardless of geographic and temporal separation among the individual members. This collaboratory approach will allow researchers to conduct large scale modeling and simulation, quickly access appropriate information, share access to and remotely operate scarce and expensive facilities and instruments, and work within virtual environments to better visualize scientific data and configure and control experiments.

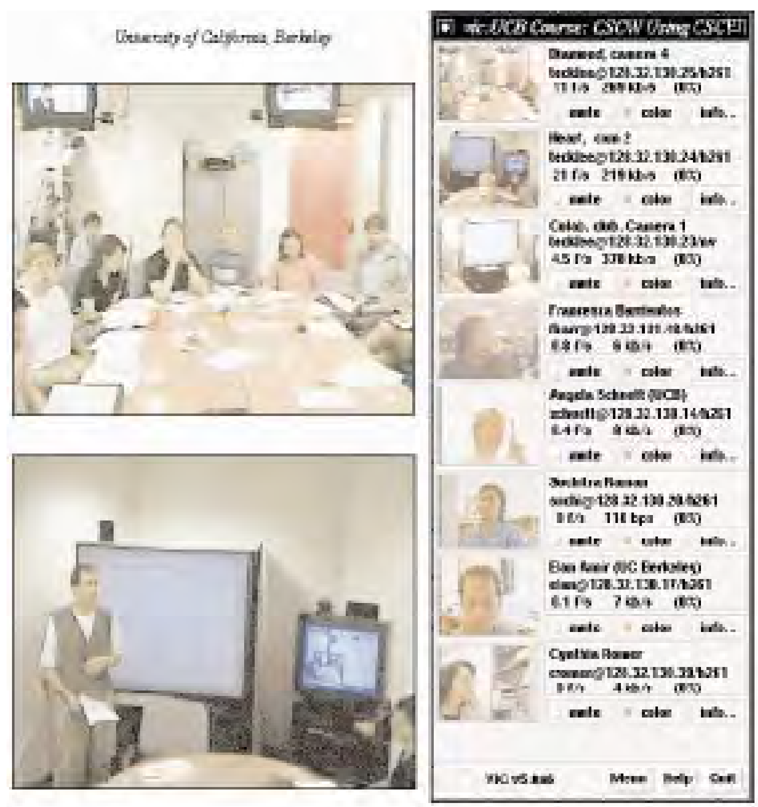

As shown here in a distance learning application, MASH enables collaboration among groups in multimedia rooms with large screen displays and individuals at remote sites, connected via various wired and wireless bandwidths, who access the collaboration using desktop, notebook, or handheld devices.

From a live MMC multisite session between collaboratory members operating a state-of-the-art field emission scanning electron microscope at Oak Ridge National Laboratory.

Achieving tomorrow's vision requires research efforts today, including R&D in advanced technologies for collaborative environments. In programs such as DARPA's Intelligent Collaboration and Visualization program and the Department of Energy's DOE 2000 initiative for national collaboratories, Federal agencies are cooperating to develop new collaborative technologies and prototype tools and to apply and evaluate those prototypes. In FY 1998, tools developed by integrating research from networking, software engineering, and distributed systems will become available to Government and commercial entities and will be deployed in several scientific collaboratories operated by DOE.

Multimedia Architecture that Scales across Heterogeneous environments (MASH)

MASH is an object-oriented framework for scalable, multimedia collaboration being developed at the University of California, Berkeley. MASH will enable developers of collaborative applications to manage data type heterogeneity, coordinate distributed resources in collaborative sessions, archive and retrieve the contents of multimedia collaborative sessions, and compose new collaborative applications from a toolkit of components. It is being validated in a collaboratory testbed focused on distance learning across the Internet.

MASH will provide next-generation Internet multimedia conferencing tools, replacing the MBONE (multicast backbone) videoconferencing tools developed at the Lawrence Berkeley National Laboratory. MASH integrates the real-time image and video "transcoding proxies" developed for DARPA's global mobile information systems program. These proxies adapt image representations and video streams across bandwidth-constrained wireless communications links. In MASH, the transcoding proxies, coupled with a consensus-based resource allocation protocol, enable fixed-session bandwidth to be allocated automatically among collaborators by monitoring the level of interest the group has in data streams flowing during a session.

Materials Microcharacterization Collaboratory

Scientific research teams, working in pilot collaboratories, are providing validation and feedback that will guide the design of these and other tools intended to provide shared collaborative access to unique computational and experimental resources. The teams and pilot collaboratories will also demonstrate how advanced collaboration technology can affect the conduct of science. The Materials MicroCharacterization Collaboratory (MMC — demonstrated at the Netamorphosis NGI exhibit) brings together five

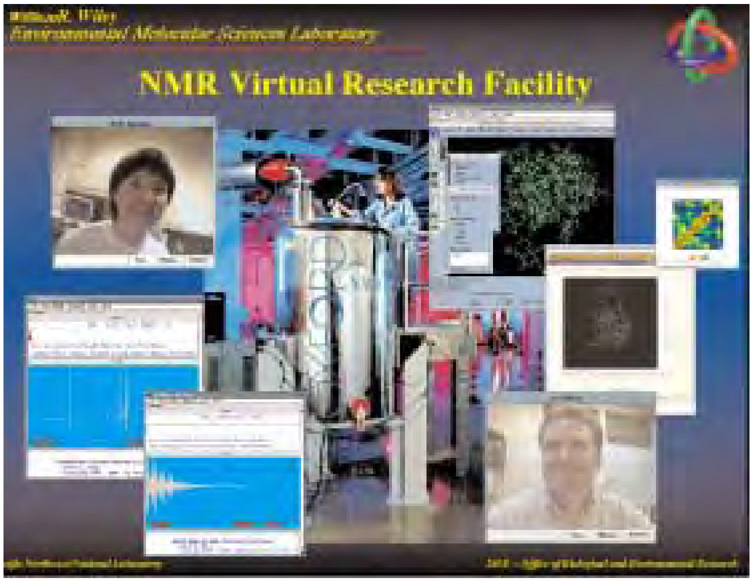

CORE2000, illustrated in this composite, will provide an application programming interface for on-line instruments and collaborative applications.

facilities in an interactive electronic laboratory. Participants include Argonne, Lawrence Berkeley, and Oak Ridge National Laboratories, NIST, and the University of Illinois. These facilities bring together virtually every characterization technique that employs electrons, ions, photons, x-rays, neutrons, mechanical and/or electromagnetic radiation to elucidate the microstructure of any material. Through the MMC, distributed teams can use and share these complementary facilities. One early problem uncovered by the MMC was that latency, introduced by remote access across networks, leads to difficulty in controlling delicate instruments. To deal with this problem, the concept of visual servoing was developed. It hides network latency and simplifies the use of scientific imaging instruments.

CORE 2000

Problems in scheduling and managing distributed resources during real-time collaborative sessions are being addressed through CORE2000, a real-time collaboration engine being developed at DOE's Environmental Molecular Sciences Laboratory (EMSL) to handle membership, authentication, event distribution, and invocation. The CORE2000 interface specification allows incompatible collaboration tools to be integrated. The CORE2000 session manager provides users with integrated access to tools such as Habanero, a framework developed by NCSA for sharing Java objects. CORE2000 will operate on a variety of platforms and will provide an application programming interface (API) for on-line instruments and collaborative applications such as DOE's WindTunnel and NLM's Visible Human project.

NIST's manufacturing collaboratory

A manufacturing collaboratory is being deployed by NIST, the University of Michigan, and industrial participants. This effort will provide a technology environment to implement and assess the effects of distributed integrated manufacturing.

Education, Training, and Human Resources

Overview

ETHR R&D supports computer- and communications-related research to advance education and training technologies at all levels including K-12, community college, technical school, trade school, university undergraduate and graduate, and lifelong learning. The complex and technically challenging applications flowing from leading-edge R&D in HECC and LSN make it increasingly important for today's students and professionals to update their education and training on an ongoing basis in order to adopt the latest technological advances. ETHR technologies improve the quality of today's science and engineering education and lead to more knowledgeable and productive citizens. ETHR R&D will:

- ❒ Encourage collaborations in research, development, and evaluation of education, training, and human resources among Federal agencies, academia, and industry
- ❒ Advance technologies for high-quality, affordable software learning tools
- ❒ Encourage development of information-based models of educational systems and learning productivity
- ❒ Support research on information technology applied to learning and cognitive processes
- ❒ Demonstrate innovative technologies and networking applications.

The training of the next generation of citizens skilled in developing and using information technologies is critical to the ongoing effort of maintaining U.S. competitiveness in today's highly aggressive international market environment. FY 1998 and anticipated FY 1999 ETHR R&D investments will continue to support this effort.

Centers for Learning Technologies

NSF has recently funded three national centers for learning to develop a new generation of researchers and computer tools focused on learning technologies:

- ❒ The Center for Interdisciplinary Research in Constructive Learning Environments, a partnership of the University of Pittsburgh and Carnegie-Mellon University, aims to develop a new generation of computer tutoring systems that adds advanced planning and natural language components to existing intelligent tutoring systems. The Center will test the new tutors in collaborating schools and will serve as a technology transfer center enabling developers, researchers, and educators to integrate these advances into existing educational practices and upgrade tutors already in use by thousands of students.
- ❒ The Center for Learning Technology in Urban Schools, a partnership of the Chicago and Detroit public schools with researchers at Northwestern University and the University of Michigan, will assist teams of teachers and researchers in learning how to make new technologies work in urban and non-urban communities across the Nation. The Center will encourage new technologies and improved school organizations that are better structured to adopt new technologies and rigorous new science curricula.
- ❒ The Center for Innovative Learning Technologies, a partnership of SRI International, University of California-Berkeley, Vanderbilt University, and the Concord Consortium, will investigate virtual learning communities, visualization and modeling, uses of low-cost ubiquitous computing, and learning assessment. The Center will provide the infrastructure for synthesizing learning technology R&D and implementing lessons across projects, stimulating rapid innovation, multidisciplinary collaboration, and support for mathematics and science learning. The Center will train postdoctoral scholars from multiple disciplines to provide future leadership in learning technologies.

Knowledge and Distributed Intelligence (KDI) initiative

In FY 1999, NSF will continue investing in both Collaborative Research in Learning Technologies and Learning and Intelligent Systems (LIS) through its KDI initiative.

NLM provides educational grants to cultivate cross-trained biomedical/ computer/ telecommunications professionals.

NSF's LIS program will study cognition, neuroscience, information technologies, and related disciplines. Supported work focuses on "learning about learning" by emphasizing the integration of theory with experiments that ground, test, and advance basic understanding of learning and intelligent behavior. Learning technology projects funded under LIS include research on learning applied to reading proficiency, learning applied to navigation problems, learning using intelligent agents, and the development of multiple-user environments for teacher training. As part of this investment, NSF plans to fund an educational digital library under a new Digital Library solicitation and will continue to support work in educational object environments and advanced distributed learning.

Integrative Graduate Education and Research Training (IGERT)

Meeting the challenge of educating scientists, mathematicians, and engineers for the 21st century will require a new paradigm in graduate training. To address this need, NSF is establishing the IGERT program in FY 1998. A Foundation-wide, multidisciplinary, graduate training effort, IGERT will develop innovative, research-based, graduate education and training activities that will produce a diverse group of new scientists and engineers well-prepared for a broad spectrum of career opportunities in both academic and industrial environments. Supported projects must be based upon a multidisciplinary research theme and organized around a diverse group of investigators from U.S. Ph.D.-granting institutions with appropriate research and teaching interests and expertise.

Two additional programs are eligible for IGERT support:

- ❐ KDI — to understand how educational and scientific communities can educate new generations of learners
- ❐ Graduate education — to integrate education and research

NLM HPCC training grants

There is a shortage of biomedical professionals trained in the use of modern computer and telecommunications systems. Current U.S. needs include both biomedical professionals cross-trained in informatics, and professionals from computer and information sciences and engineering who have had doctoral or post doctoral training in the application of these technologies to health problems. To help address these needs, NLM is expanding its successful predoctoral and postdoctoral grants program for career training in medical informatics. Grants will be made in research and application, and in HPCC-in-medicine fellowship training support.

National Center for Research Resources (NCRR) training

NCRR supports high performance computing training at six high performance computer resource centers and at two graphics/visualization resource centers. These centers provide training for students (undergraduate, graduate, and postdoctoral) attending the host institutions, and for scientists and students (including postdoctoral) outside these institutions. Current participants include biologists and scientists in related disciplines who need to learn about high performance computing, and computer scientists who need to better understand the capabilities and needs of high performance computing as applied to biomedical research.

NASA learning technologies projects

Through the Internet, NASA is reaching into America's classrooms to make a difference in the quality and content of mathematics, science, and technology education. Using the Web as a primary medium, NASA hosts virtual electronic field trips, collaborative science projects, and distance learning activities, and allows students to interact with scientists in real time.

NASA's cooperative agreement, "Public Use of Earth and Space Science Data over the Internet," demonstrated mature K-12 education products and innovative digital library technologies. Current activities available through

LTP is NASA communicating science — letting students and teachers know that while imagination starts at the ground floor, the sky is not the limit.

NASA Regional Outreach Centers and education cooperative agreements will continue to make an impact on education and information delivery over the Internet.

NASA's Learning Technologies Project (LTP) is now developing a digital audio network testbed that will initially support 5,000 schools nationwide. This audio-based Internet infrastructure will be used to communicate NASA science and make distance learning technologies more widely available. Interaction and training will be provided through NASA's Learning Technologies Channel.

Presidential Technology Initiative and Training Research for Automated Instruction

Two Federal agencies — DoD and its Air Force Office of Scientific Research (AFOSR) — that are not part of the CIC budget crosscuts, participate in ETHR activities. Among them, the Department of Defense Educational Program (DoDEA) is a learner-centered educational organization providing its students with the knowledge and skills required for high levels of achievement.

DoDEA's Presidential Technology Initiative encourages software developers to use their institution's Web sites for sharing and interacting with the DoDEA teachers who will be their partners in finishing software projects and rapidly making them available to DoDEA students.

The AFOSR-supported Training Research for Automated Instruction (TRAIN) project is a large scale, long-term effort designed to investigate instructional approaches for automated education and training. These include a human performance taxonomy, benchmarks, and experiments. This work is being made available to extramural researchers. A computer laboratory has been built at Lackland (CA) Air Force Base to support TRAIN goals.

CIC R&D Programs

Committee on Technology

In FY 1998, the National Science and Technology Council (NSTC) realigned its structure into five new NSTC Committees, including the Committee on Technology (CT). The CT oversees the activities of seven technology R&D Subcommittees, including the Subcommittee on Computing, Information, and Communications (CIC) R&D.

The purpose of the CT is to advise and assist the NSTC to increase the overall effectiveness and productivity of Federal technology R&D efforts. The CT — which encompasses the activities of the former Committee on Computing, Information, and Communications; the Committee on Technological Innovation; and the Committee on Transportation R&D — will address national policy matters that cut across agency boundaries and provide a formal mechanism for interagency policy coordination and development of balanced and comprehensive technology R&D programs. The Committee

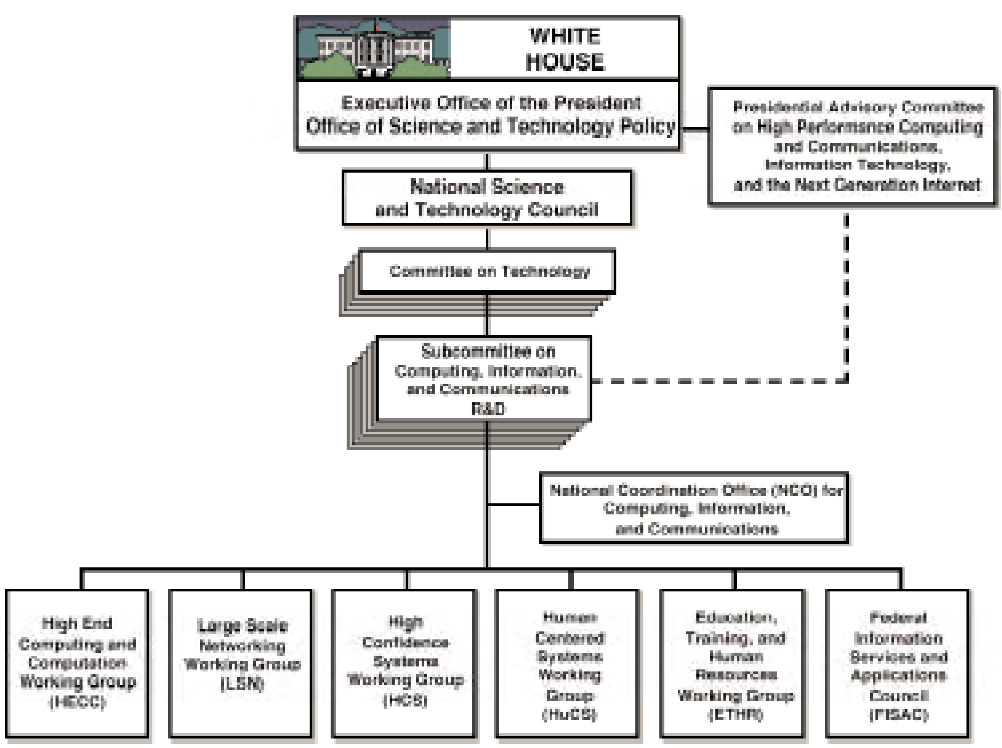

will advise the Directors of OSTP and OMB on technology R&D budget crosscuts and priorities.

Subcommittee on Computing, Information, and Communications R&D

The Subcommittee on CIC R&D consists of representatives from each of the twelve agencies that participate in the Federal CIC R&D programs. The Subcommittee and its Executive Committee work with the agencies to plan, budget, implement, and review the multiagency CIC programs, helping to shape the Nation's 21st century information infrastructure.

The Subcommittee has five Working Groups that meet regularly to coordinate activities, propose new initiatives, and address specific programmatic objectives. These Working Groups are associated with the five Program Component Areas (PCAs) of the CIC R&D programs: High End Computing and Computation; Large Scale Networking; Human Centered Systems; High Confidence Systems; and Education, Training, and Human Resources. The FY 1998 accomplishments and FY 1999 plans for these PCA Working Groups are highlighted in this book.

National Coordination Office for Computing, Information, and Communications (NCO/CIC)

The NCO facilitates multiagency CIC R&D activities, such as the preparation of planning, budget, and assessment documents and the development of interagency CIC programs, and supports the activities of the Presidential Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet. The NCO Director, who reports to the Director of the Office of Science and Technology Policy (OSTP), Executive Office of the President, serves as the chair of the Subcommittee on CIC R&D.

The NCO and the Subcommittee on CIC R&D meet frequently with representatives from Congress; Federal, state and local organizations; academia; industry; professional societies; foreign organizations; and others to exchange technical and programmatic information about CIC needs, issues, and trends. For example, during FY 1998, the NCO and the Subcommittee's LSN Working Group sponsored "Netamorphosis," a three-day demonstration to help educate members of Congress and the public about Federal CIC R&D programs such as the Next Generation Internet initiative. Also, in FY 1997, a delegation from the U.S., including the NCO Director and representatives from CIC agencies, visited St. Petersburg, Russia, and briefed Russian officials on CIC R&D.

Federal Information Services and Applications Council (FISAC)

The FISAC was created to 1) stimulate and foster the migration of technologies from the information technologies R&D community to Government application missions and information services communities, and 2) identify challenges from applications to the information technologies R&D community. FISAC activities are conducted through interagency working groups in the areas of Crisis Management, Federal Statistics, Next Generation Internet, Universal Access, and Information Security. With membership across the Federal agencies and in collaboration with the Government Information Technology Services (GITS) Board, the FISAC initiates multiagency studies and workshops, encourages pilot projects, and provides an information exchange among members.

Buy American Report

Congress requires annual reporting of non-U.S. high performance computing and communications funding activities.

In FY 1998, DARPA was the only CIC R&D agency that entered into grants, contracts, cooperative agreements, or cooperative research and development agreements for CIC R&D 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 U.S. DARPA funded an award related to HECC and LSN to the University of Warwick (United Kingdom) in the amount of $165,796.

In FY 1998, no CIC 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.

Outreach

The NCO responds to thousands of information requests from Congressional offices, industry, academia, and the public each year. To fulfill these requests, the NCO provides print and video materials and maintains Web servers that include all NCO publications since 1994; information on the Subcommittee on CIC R&D and the organizations that report to it; the Presidential Advisory Committee; the Next Generation Internet initiative; Congressional testimony; links to the servers of participating agencies; and other related information.

In FY 1998, the NCO produced a publicly-available video entitled "Advanced Networking: Connecting to the Future" in cooperation with the University Corporation for Advanced Internet Development. The video highlights recent and potential advancements in large scale networking through Federal and non-Federal projects such as the Next Generation Internet initiative and Internet2.

Several CIC R&D agencies and the NCO also participated in SC97, a national conference focusing on networking, distributed computing, data-intensive applications, supercomputing, scalable computing, and other emerging technologies that push the frontiers of computing, communications, and computational science. The NCO sponsored a research exhibit to help educate the public about Federal investments in CIC R&D, participated in panel discussions, and helped coordinate a Town Hall meeting of the Presidential Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet. At SC98, the CIC R&D agencies and the NCO plan to demonstrate Next Generation Internet applications in the conference's research exhibit area and participate in panel discussions of Federal CIC R&D activities.

For the most up-to-date information on CIC programs, the reader is invited to visit *http://www.ccic.gov/*.

The Presidential Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet

Established by President Clinton on February 11, 1997, the Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet will provide the President, the Office of Science and Technology Policy, and the Federal agencies involved in CIC R&D with guidance and advice on all areas of high performance computing, communications, and information technologies. Representing the research, education, and library communities and including network providers and representatives from critical industries, the Committee members will help guide the Administration's efforts to accelerate development and adoption of information technologies vital for American prosperity in the 21st century.

On October 31, 1997, President Clinton appointed four new members to the Committee, bringing the total number of members to twenty-five. The new members are David Dorman, John Miller, Joe Thompson, and Irving Wladawsky-Berger. Their affiliations are given on the next page.

As part of the Executive Order creating the Committee, President Clinton asked for an independent assessment of the Federal government's role in high performance computing and communications, including:

- ❐ Progress made in implementing the High Performance Computing and Communications Program
- ❐ Progress in designing and implementing the NGI initiative
- ❐ The need to revise the HPCC Program
- ❐ Balance among components of the HPCC Program
- ❐ Whether the research and development undertaken pursuant to the HPCC Program is helping to maintain U.S. leadership in advanced computing and communications technologies and their applications
- ❐ Other issues as specified by the Director of the Office of Science and Technology Policy

During FY 1998, the Committee held public meetings December 9-10, 1997; March 11, 1998; and May 19, 1998. At the December meeting, the CIC R&D Program Component Area chairs presented budget briefings and summaries of the CIC R&D agenda. The Committee was also briefed on critical infrastructure protection; societal issues; and the regulation, deregulation, and growth of the Internet. In March, 1998, the Committee's Highend and Broadbased Subcommittees presented reports on their activities and findings, and the Committee announced its intention to produce a report by the end of 1998. The Committee also received an update from the FCC, an update on the NGI initiative, and a report on "Connecting All Americans."

At SC97 in San Jose, CA, in November, 1997, several Advisory Committee members participated in a Town Hall meeting. They briefed conference participants on the Committee's activities and answered questions about the future of U.S. information technology.

In March, 1998, the Committee also attended "Netamorphosis," a demonstration of NGI technologies and applications, held at Highway 1 in Washington, DC. At the Netamorphosis kick-off reception, Committee co-chairs Ken Kennedy and Bill Joy spoke about the success of the Federal NGI initiative and the need for further Federal networking R&D.

Committee Members

Committee Co-Chairs

Ken Kennedy is Director of the Center for Research on Parallel Computation and is Ann and John Doerr Professor of Computer Science at Rice University.

Bill Joy is co-founder and Vice President of Research at Sun Microsystems.

Committee Members

Eric A. Benhamou is President, Chairman, and CEO of 3Com Corporation.

Vinton Cerf is Senior Vice President of Internet Architecture and Engineering at MCI Communications.

Ching-chih Chen is a Professor in the Graduate School of Library and Information Science at Simmons College.

David Cooper is Associate Director of Computation at the Lawrence Livermore National Laboratory.

Steven D. Dorfman is Vice Chairman of Hughes Electronics Corporation, Chairman of Hughes Telecommunications and Space Company, and a member of Hughes Electronics Office of the Chairman.

David W. Dorman is President and CEO of PointCast.

Robert Ewald is Executive Vice President for Computer Systems at Silicon Graphics, Inc.

David J. Farber is Alfred Fitler Moore Professor of Telecommunications at the University of Pennsylvania.

Sherrilynne S. Fuller is Director of the Health Sciences Libraries and Information Center; Acting Director, Informatics, School of Medicine; and Director of the National Network of Libraries of Medicine, Pacific Northwest Region at the University of Washington.

Hector Garcia-Molina is Leonard Bosack and Sandra Lerner Professor in the Departments of Computer Science and Electrical Engineering at Stanford University.

Susan Graham is Chancellor's Professor of Computer Science in the Department of Electrical Engineering and Computer Science at the University of California, Berkeley.

James N. Gray is a senior researcher in Microsoft's Scalable Servers Research Group and manager of Microsoft's Bay Area Research Center.

W. Daniel Hillis is a Vice President and Disney Fellow at Walt Disney Imagineering, Research and Development, Inc.

John P. Miller is Director of the Center for Computational Biology at Montana State University, Bozeman.

David C. Nagel is President of AT&T Labs.

Raj Reddy is Dean of the School of Computer Science and Herbert A. Simon University Professor of Computer Science and Robotics at Carnegie Mellon University.

Edward H. Shortliffe is Associate Dean for Information Resources and Technology, Professor of Medicine, and Professor of Computer Science at Stanford University School of Medicine.

Larry Smarr is Director of the National Computational Science Alliance and Professor of Physics and Astrophysics at the University of Illinois at Urbana-Champaign.

Joe F. Thompson is the William L. Giles Distinguished Professor of Aerospace Engineering in the Department of Aerospace Engineering at Mississippi State University.

Leslie Vadasz is Senior Vice President and Director of Corporate Business Development at Intel Corporation.

Andrew J. Viterbi is a co-founder of QUALCOMM Incorporated and Vice Chairman of its Board of Directors.

Steven J. Wallach is Advisor to CenterPoint Ventures.

Irving Wladawsky-Berger is General Manager, IBM Internet Division at IBM Corporation.

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 academic, defense, and industrial 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; 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 (HECC), Large Scale Networking (LSN), High Confidence Systems (HCS), Human Centered Systems (HuCS), and Education, Training, and Human Resources (ETHR) – 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.

Agency CIC R&D Budgets by Program Component Area

FY 1998 Budget (Dollars in Millions)

Agency	HECC	LSN	HCS	HuCS	ETHR	TOTAL
DARPA	84.80	89.23	9.40	137.87		**321.30**
NSF	132.90	69.20	0.90	60.17	20.96	**284.13**
DOE	90.53	25.79		9.94	3.00	**129.26**
NASA	90.10	25.00	2.80	2.20	8.30	**128.40**
NIH	23.74	28.19	4.13	29.28	6.38	**91.72**
NSA	26.42	2.18	7.20			**35.80**
NIST	3.99	5.46	3.40	13.66		**26.51**
VA		7.45	5.35	9.20		**22.00**
ED				12.00		**12.00**
NOAA	4.30	2.70		0.50		**7.50**
EPA	5.38					**5.38**
AHCPR				5.50		**5.50**
TOTAL	**462.16**	**255.20**	**33.18**	**280.32**	**38.64**	**1069.50**

FY 1999 Budget Request (Dollars in Millions)

Agency	HECC	LSN	TOTAL
DARPA	57.50	102.70	**160.20**
NSF	235.70	74.20	**309.90**
DOE	88.00	40.20	**128.20**
NASA	71.60	19.80	**91.40**
NIH	42.33	67.67	**110.00**
NSA	24.00	3.00	**27.00**
NIST	3.50	5.20	**8.70**
VA		1.90	**1.90**
ED			
NOAA	10.30	2.70	**13.00**
EPA	5.00		**5.00**
AHCPR	3.50	2.10	**5.60**
TOTAL	**541.43**	**319.47***	**860.90***

NGI totals are included under LSN.

HCS, HuCS, and ETHR crosscuts are not included in the President's FY 1999 budget.

* These totals vary from the President's FY 1999 Budget in that the Department of Veterans Administration (VA) LSN funding was not included in that Budget.

Glossary

AAI
ACTS ATM International.

AC
A parallel C language.

ACTS
NASA's Advanced Communication Technology Satellite.

ACTS
Advanced Computational Testing and Simulation.

AFOSR
Air Force Office of Scientific Research.

AHCPR
Agency for Health Care Policy and Research, part of the Public Health Service of the Department of Health and Human Services.

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

ANIR
NSF's Advanced Networking Infrastructure and Research.

API
Applications programming interface.

ARPS
Advanced Regional Prediction System.

ASAP
DOE's Academic Strategic Alliances Program.

ASCI
DOE's Accelerated Strategic Computing Initiative.

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 NSF vBNS backbone — connecting lower capacity networks.

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

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.

BIT
DARPA's Broadband Information Technology program.

Bit
An acronym for binary digit.

BLAS
Basic Linear Algebra Subprograms.

Bps, or B/s
An acronym for bytes per second.

bps, or b/s
An acronym for bits per second.

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.

Caltech
California Institute of Technology.

CAPS
Center for the Analysis and Prediction of Storms.

CAS
NASA's Computational Aerosciences Project.

CAVE
Cave Automatic Virtual Environment. A surround screen, surround sound, projection-based virtual reality (VR) system.

CBVE
Chesapeake Bay Virtual Environment.

CCIC
Former Committee on Computing, Information, and Communications of the NSTC; now merged into the Committee on Technology.

CDMA
Code division multiple access.

CERT
Computer Emergency Response Team.

CIAC
Computer Incident Advisory Capability.

CIC
Computing, Information, and Communications.

CIC R&D, Subcommittee on
Subcommittee on Computing, Information, and Communications R&D, which reports to the Committee on Technology.

CIS
Collaborations in Internet Security.

CISE
NSF's Directorate for Computer and Information Science and Engineering.

CIT
NIH's Center for Information Technology.

CORE2000
A real-time collaboration engine being developed at DOE's Environmental Molecular Sciences Laboratory (EMSL) to handle membership, authentication, event distribution, and invocation.

CoS
Class of Service.

COTS
Commercial off-the-shelf. Describes hardware and software that are readily available commercially.

CPR
Computer-based patient records.

CRPC
Center for Research in Parallel Computation.

CT
Committee on Technology of the NSTC.

DARPA
Defense Advanced Research Projects Agency, part of DoD. Formerly ARPA.

DCRT
NIH's Division of Computer Research and Technology. In February 1998, the Center for Information Technology (CIT) combined the functions of DCRT, the Office of Information Resources Management (OIRM), and the Telecommunications Branch (TCB) of the Office of Research Services. *(See CIT.)*

DES
Data Encryption Standard.

DII
Defense Information Infrastructure.

DISS
Distributed Image Spread Sheet.

DNA
Deoxyribonucleic Acid, a biomolecule from which genes are composed.

DoD
Department of Defense.

DoDEA
Department of Defense Education Program, the public school system for dependents of US personnel worldwide.

DoD HPC Modernization Program
Department of Defense High Performance Computing Modernization Program.

DOE
Department of Energy.

DOE 2000
Department of Energy program focusing on solving DOE's complex scientific problems.

DRAM
Dynamic Random Access Memory.

DREN
DoD's Defense Research and Engineering Network.

ED
Department of Education.

EMSL
DOE's Environmental Molecular Sciences Laboratory.

EPA
Environmental Protection Agency.

EPSCoR
Experimental Program to Stimulate Competitive Research.

ERIC
ED's Educational Resources Information Center.

ESNet
DOE's Energy Sciences Network.

ESS
NASA's Earth and Space Sciences.

ETA
A model currently being tested at the National Centers for Environmental Prediction Environmental Modeling Center (EMC) Mesoscale Modeling Branch.

ETHR
Education, Training, and Human Resources. One of the five CIC R&D Program Component Areas.

Exa-
A prefix denoting 10^{18}, or a million trillion. (For example, exabytes).

FAA
Federal Aviation Administration.

FDA
Food and Drug Administration.

FedCIRC
Federal Computer Incident Response Capability.

FISAC
Federal Information Services and Applications Council, formerly the Applications Council, which reports to the Subcommittee on CIC R&D.

Flops
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.

FRA
Federal Railroad Administration.

FSL
NOAA's Forecast Systems Laboratory in Boulder, CO.

FTA
Federal Transit Administration.

FTP
File transfer protocol.

FY
Fiscal Year.

G, or Giga-
A prefix denoting 10^9, or a billion. (For example, Gflops or gigaflops; gigabytes, gigabits).

Gateway
A system that interconnects networks (or applications) that communicate using different protocols, and bridges their differences by transforming one protocol into another.

GB
An acronym for Gigabyte.

Gb
An acronym for Gigabit.

Gb/s or Gbps
Gigabits per second.

GFDL
NOAA's Geophysical Fluid Dynamics Laboratory.

Gflops
Gigaflops, billions of floating point operations per second.

Gigapops
Gigabit per second points of presence.

GITS
Government Information Technology Services.

GOES
Geostationary Operational Environmental Satellite.

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 and/or communications techniques and resources.

GSA
General Services Administration.

GSFC
NASA Goddard Space Flight Center.

HCS
High Confidence Systems. One of the five CIC R&D Program Component Areas.

HECC
High End Computing and Computation. One of the five CIC R&D Program Component Areas.

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

HFC
Hybrid fiber coax.

HPC
High performance computing.

HPCC
High Performance Computing and Communications, a Federal R&D Program that is the predecessor to the CIC R&D programs.

HPCRP
High Performance Computing Resource Provider.

HuCS
Human Centered Systems. One of the five CIC R&D Program Component Areas.

HTMT
Hybrid technology multithreaded architecture.

IAIMS
NIH's Academic Medical Centers.

IGERT
NSF's Integrative Graduate Education and Research Training.

INFOSEC
Information Systems Security.

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

I2
Internet2.

I/O
Input/Output.

IP
Internet Protocol.

IPSEC
IP Security Protocol.

IRS
Internal Revenue Service, part of the Department of Treasury.

ISV
Independent Software Vendor.

JET
Joint Engineering Team, which reports to the LSN Working Group.

JPL
Jet Propulsion Laboratory.

JTO
Joint Technology Office.

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.

KDI
NSF's Knowledge and Distributed Intelligence Program.

kv
Kilovolt.

LAN
Local area network.

LANL
DOE's Los Alamos National Laboratory.

LBNL
DOE's Lawrence Berkeley National Laboratory.

LIS
NSF's Learning and Intelligent Systems Program.

LLNL
DOE's Lawrence Livermore National Laboratory.

LoC
Library of Congress.

LOTS
An optical tape system with ten times the storage per tape and at least ten times the input data rate compared with conventional large stores.

LSN
Large Scale Networking. One of the five CIC R&D Program Component Areas.

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

MAGIC
An information system that allows interactive viewing of large geographic areas in 3-D by retrieving and merging massive volumes of remotely located data.

MASH
Multi-media architecture that scales across heterogeneous environments.

MB
An acronym for Megabyte.

Mb
An acronym for Megabit.

MBONE
Multicast backbone.

Mb/s or Mbps
Megabits per second or millions of bits per second.

MCM
Multichip module.

MEDLARS
Medical Literature Analysis and Retrieval System. Provides rapid access to NLM's biomedical and health information.

Mflops
Megaflops, millions of floating point operations per second.

Mips
Millions of instructions per second.

MIT
Massachusetts Institute of Technology.

ML
Maximum likelihood.

MM5
Mesoscale Meteorological model.

MMC
Materials Microcharacterization Collaboratory.

MPI
Message Passing Interface.

MPP
Massively parallel processing.

MRI
Magnetic Resonance Imaging.

NAPs
Network access points.

NAS
Numerical Aerodynamic Simulation.

NAS
National Academy of Sciences.

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 CIC technology and resources.

NBS
National Bureau of Standards, NIST's predecessor.

NCAR
National Center for Atmospheric Research.

NCBI
NLM's 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 Computational Science Alliance, Urbana-Champaign, IL, successor to the National Center for Supercomputing Applications.

NEH
National Endowment for the Humanities.

NERSC
DOE's National Energy Research Supercomputer Center.

Network
Computer communications technologies that link multiple computers for sharing information and resources across geographically dispersed locations.

NGI
Next Generation Internet, a Presidential CIC R&D initiative.

NIDRR
ED's National Institute on Disability and Rehabilitation Research.

NIGMS
National Institute of General Medical Sciences, part of NIH.

NIH
National Institutes of Health, part of HHS.

NIST
National Institute of Standards and Technology, part of the Department of Commerce.

NLANR
National Laboratory for Applied Networking Research, sponsored by NSF.

NLM
National Library of Medicine, part of NIH.

NOAA
National Oceanic and Atmospheric Administration, part of the Department of Commerce.

NOW
Network of workstations.

NPACI
National Partnership for Advanced Computational Infrastructure.

NRC
Nuclear Regulatory Commission.

NREN
NASA's Research and Education Network.

NRT
Networking Research Team, which reports to the LSN Working Group.

NSA
National Security Agency, part of DoD.

NSF
National Science Foundation.

NSFNET
National Science Foundation computer network program.

NSTC
National Science and Technology Council.

OC-nn
Optical carrier. "nn" is a number, with a higher number indicating a faster data transmission rate.

OMB
White House Office of Management and Budget.

ORNL
DOE's Oak Ridge National Laboratory.

OS
Operating system.

OSTP
White House Office of Science and Technology Policy.

PACI
Partnership for Advanced Computational Infrastructure.

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

PCA
Program Component Area. Structure of the 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).

PCP
A Parallel C language.

PCS
Personal communication services.

PECASE
Presidential Early Career Award for Science and Engineering.

PET
Positron emission tomography.

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

pflop/s
Petaflops, 10^{15} flops.

PKI
Public key infrastructure.

Prions
Prions are a novel class of "infectious" pathogens distinct from viruses with respect to both their structure and the neurodegenerative diseases, such as Creutzfeldt-Jakob Disease, that they cause.

PrP
Prion protein. It is the major, if not the only, component of prions, the protein implicated in Creutzfeldt-Jakob Disease.

PrPC
The normal cellular form of PrP *(see above)*.

PrPSc
The abnormal, disease (scrapie)-related form of PrP *(see above)*.

PSC
Pittsburgh Supercomputer Center.

PSE
Problem Solving Environment.

QoS
Quality of Service.

R&D
Research and development.

RCWS
Radiology Consultation Workstation.

RES
A system developed to harness "idle" workstations.

RTEC
Regional Technology in Education Consortium.

SAR
Synthetic aperture radar.

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.

SciTL
Scientific Template Library. The initial projects undertaken under the Advanced Computational Testing and Simulation portion of DOE2000.

SDSC
San Diego Supercomputer Center.

SEMATECH
A national consortium dedicated to conducting semiconductor research.

SGI
Silicon Graphics, Inc.

SIMA
NIST's Systems Integration for Manufacturing Applications.

SMPs
Symmetric multiprocessor systems.

SONET
Synchronous Optical Network.

Split-C
A parallel C language.

S&T
Science and technology.

STAR TAP
Science, Technology And Research Transit Access Point, an international transit networking meeting point in Chicago.

STIMULATE
Speech, Text, Image, and Multimedia Advanced Technology Effort.

STM
Scanning Tunneling Microscope.

T, or Tera-
A prefix denoting 10^{12} or a trillion. (For example, terabits, teraflops).

Tbps
Terabit-per-second.

TCP/IP
Transmission Control Protocol/Internet Protocol.

TDM
Time division multiplexing.

TETRA
A versatile robotic crane, one of several experimental technologies being investigated at NIST's National Construction Automation Testbed.

THz
Trillions of cycles per second

TTS
Text-to-Speech, a core speech technology comprising the speech kernel in DARPA's natural language understanding technology.

UCAID
University Corporation for Advanced Internet Development.

UC-Berkeley
University of California at Berkeley.

UCLA
University of California-Los Angeles.

UCSD
University of California-San Diego.

UIUC
University of Illinois at Urbana-Champaign.

UMLS
Unified Medical Language System.

UPC
A programming language that combines features of AC with features of Split-C and PCP, two other parallel C languages.

USC
University of Southern California.

VA
Department of Veterans Affairs.

vBNS
NSF's very high performance Backbone Network Services.

Vis5D
System for visualizing the output of atmospheric and ocean models with composite images from the Geostationary Operational Environmental Satellite (GOES).

VLSI
Very Large Scale Integration. A type of computer chip.

VRML
Virtual Reality Modeling Language.

WAN
Wide area network.

WDM
Wavelength division multiplexing.

Web
A reference to the World Wide Web.

Wireless technologies
Communications technologies that use radio, microwave, or satellite communications channels versus wire, coaxial, or optical fiber.

WIT
Web-based Interoperability Tester.

WWW
World Wide Web.

Contacts

National Coordination Office for Computing, Information, and Communications (NCO/CIC)

Suite 690
4201 Wilson Boulevard
Arlington, VA 22230
(703) 306-4722
FAX: (703) 306-4727
nco@ccic.gov

Internet/World Wide Web Servers:

http://www.ccic.gov/
http://www.ngi.gov/

NCO

Kay Howell
Director

Sally E. Howe, Ph.D.
Associate Director

Yolanda Comedy, Ph.D.
Technical Staff (Presidential Advisory Committee Liaison)

Ward Fenton
Systems Administrator

Kelly Goins
Administrative Officer

Vicki L. Harris
Administrative Assistant

Kristin Janger
Director's Assistant

Catherine W. McDonald
Technical Staff

Betty McDonough
Assistant System Administrator

Grant Miller, Ph.D.
Technical Staff (LSN Liaison)

Krish Namboodiri, Ph.D.
Technical Staff (Implementation Plan Liaison)

Terrence L. Ponick, Ph.D.
Technical Writer

Denis Roark
Technical Staff (FISAC Liaison)

Robert I. Winner, Ph.D.
Technical Staff (Presidential Advisory Committee Liaison)

AHCPR

J. Michael Fitzmaurice, Ph.D.
Director, 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

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

DARPA

David L. Tennenhouse, 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

Ronald L. Larsen, Ph.D.
Assistant Director, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2227
FAX: (703) 696-6416

Bertram Hui, Ph.D.
Program Manager, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2239
FAX: (703) 696-2202

Teresa F. Lunt
Program Manager, Information Survivability Program
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-4469
FAX: (703) 696-2202

Mari W. Maeda, Ph.D.
Program Manager, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2255
FAX: (703) 696-2202

Hilarie Orman
Program Manager, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2234
FAX: (703) 696-2202

Jay Allen Sears, Ph.D.
Program Manager, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
Phone: please use email
FAX: (703) 696-2202

Stephen L. Squires
Special Assistant for Information Technology
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2226
FAX: (703) 696-2209

DOE

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

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) 903-5800
FAX: (301) 903-7774

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

Edgar W. Lewis
Office of Strategic Computing and Simulation, Defense Programs Office
Department of Energy
DP-50
Mail Stop 4A-028/FORR, Room 1F-020/FORR
1000 Independence Avenue, S.W.
Washington, DC 20585-1290
(202) 586-5092
FAX: (202) 586-2168

Thomas A. Kitchens, Ph.D.
Program 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) 903-5152
FAX: (301) 903-7774

Norman H. Kreisman
Advisor, International Technology
Department of Energy, ER5
Mailstop 3H049-FORS
1000 Independence Avenue, S.W.
Washington, DC 20585
(202) 586-9746
FAX: (202) 586-7152

Mary Anne Scott, Ph.D.
Program Manager, 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) 903-6368
FAX: (301) 903-7774

George R. Seweryniak
Program Manager, 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) 903-0071
FAX: (301) 903-7774

ED

Linda G. Roberts, Ed.D.
Director of Educational Technology, Office of the Deputy Secretary
Department of Education
Room 5164
600 Independence Avenue, S.W.
Washington, DC 20202
(202) 401-1444
FAX: (202) 401-3941

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

EPA

Joan H. Novak
HPCC Program Manager, MD-80
Environmental Protection Agency
Research Triangle Park, NC 27711
(919) 541-4545
FAX: (919) 541-1379

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

EPA FedEx address:

79 T.W. Alexander Drive
Building 4201
Research Triangle Park, NC 27709

NASA

Lee B. Holcomb
Chief Information Officer
National Aeronautics and Space Administration
Code AO
300 E Street, S.W.
Washington, DC 20546
(202) 358-1824
FAX: (202) 358-2810

David B. Nelson, Ph.D.
Acting Deputy Chief Information Officer
National Aeronautics and Space Administration
Code AO
300 E Street, S.W.
Washington, DC 20546
(202) 358-1817
FAX: (202) 358-3063

William J. Feiereisen, Ph.D.
Program Manager, High Performance Computing and Communications Program
National Aeronautics and Space Administration
Mailstop 269-3
Ames Research Center
Moffett Field, CA 94035-1000
(650) 604-4225
FAX: (650) 604-7324

Betsy Edwards
Chief Information Officer, Office of Aeronautics and Space Transportation Technology
National Aeronautics and Space Administration
Code RT
300 E Street, S.W.
Washington, DC 20546
(202) 358-4639
FAX: (202) 358-3550

Phillip L. Milstead
Information Technology Research Coordinator, Office of Aeronautics and Space Transportation Technology
National Aeronautics and Space Administration
Code RC
300 E Street, S.W.
Washington, DC 20546
(202) 358-4619
FAX: (202) 358-3550

Kul B. Bhasin
Acting Branch Chief, Satellite Networks and Architectures Branch
National Aeronautics and Space Administration
NASA Lewis Research Center
21000 Brookpark Road
Mail Stop 54-2
Cleveland, OH 44135
(216) 433-3676
FAX: (216) 433-8705

Malcom V. Phelps, Ph.D.
Assistant Director, Programs, Education Division, Office of Human Resources and Education
National Aeronautics and Space Administration
Code FE
300 E Street, S.W.
Washington, DC 20546
(202) 358-1536
FAX: (202) 358-3550

Wayne H. Bryant
Assistant Chief, Flight Electronics Technology Division, NASA Langley Research Center
National Aeronautics and Space Administration
Mail Stop 150
Hampton, VA 23681
(757) 864-1690
FAX: (757) 864-8821

Christine M. Falsetti
Project Manager, NASA HPCC NREN/NGI, NASA Ames Research Center
National Aeronautics and Space Administration
Mail Stop 233-10
Moffett Field, CA 94035-1000
(650) 604-6935
FAX: (650) 604-6999

James R. Fischer
Project Manager, NASA HPCC/ESS Project, NASA Goddard Space Flight Center
National Aeronautics and Space Administration
Code 930
Greenbelt, MD 20771
(301) 286-3465
FAX: (301) 286-1634

Kenneth M. Ford, Ph.D.
Associate Director for Information Technology,
NASA Ames Research Center
National Aeronautics and Space Administration
Mail Stop 200-3
Moffett Field, CA 94035-1000
(650) 604-1847
FAX: (650) 604-3786

Mark Leon
Project Manager, Learning Technology,
NASA Ames Research Center
National Aeronautics and Space Administration
Mail Stop N269-3
Moffett Field, CA 94035-1000
(650) 604-6498
FAX: (650) 604-4036

Stephen Scott Santiago
Chief Information Officer, NASA Ames Research Center
National Aeronautics and Space Administration
Moffett Field, CA 94035-1000
(650) 604-5015
FAX: (650) 604-6999

William J. Campbell
Head, Applied Information Sciences Branch,
NASA Goddard Space Flight Center
National Aeronautics and Space Administration
Code 935 - NASA/GSFC
Greenbelt, MD 20771
(301) 286-8785
FAX: (301) 286-1776

Milt Halem, Ph.D.
Chief, Earth and Space Data Computing Division,
NASA Goddard Space Flight Center
National Aeronautics and Space Administration
8800 Greenbelt Road
Code 930/Building 28, Room W230
Greenbelt, MD 20771
(301) 286-8834
FAX: (301) 286-1777

NIH

NIH HPCC Coordinator and CIT contact:

Robert L. Martino, Ph.D.
Associate Director, Office of Computational Bioscience
Chief, Computational Bioscience and Engineering
Laboratory, Center for Information Technology (CIT)
National Institutes of Health
12 South Drive, MSC 5624
Building 12A, Room 2033
Bethesda, MD 20892-5624
(301) 496-1112
FAX: (301) 402-2867

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

Cherie Nichols
Chief, Planning, Evaluation, and Analysis Branch,
Office of Science Policy, Office of the Director,
National Cancer Institute
National Institutes of Health
Federal Building, MSC 9010
7550 Wisconsin Avenue, Room 312
Bethesda, MD 20892
(301) 496-5515
FAX: (301) 435-3876

NCRR program:

Judith L. Vaitukaitis, M.D.
Director, National Center for Research Resources
National Institutes of Health
31B Center Drive
Building 31, Room 3B11
Bethesda, MD 20892-2128
(301) 496-5793
FAX: (301) 402-0006

Richard M. DuBois, Ph.D.
Head, Computer Technology Section, BT
National Center for Research Resources
National Institutes of Health
One Rockledge Center
6705 Rockledge Drive, Room 6148
Bethesda, MD 20892-7965
(301) 435-0758
FAX: (301) 480-3659

Martin B. Blumsack
Extramural Science Administrator
Research Infrastructure, National Center for Research Resources
National Institutes of Health
One Rockledge Center, Room 6144
6705 Rockledge Drive, MSC 7965
Bethesda, MD 20892-7965
(301) 435-0769
FAX: (301) 480-3770

NLM program:

Donald A.B. Lindberg, M.D.
Director, National Library of Medicine
National Institutes of Health
Building 38, Room 2E17B
8600 Rockville Pike
Bethesda, MD 20894
(301) 496-6221
FAX: (301) 496-4450

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

NIGMS program:

James C. Cassatt
Director, Division of Cell Biology and Biophysics,
National Institute of General Medical Sciences
National Institutes of Health
4500 Center Drive, MSC 4500
Bethesda, MD 20892
(301) 594-0828
FAX: (301) 480-2004

NIST

R. J. (Jerry) Linn
Associate Director for Program Implementation,
Information Technology Laboratory
National Institute of Standards and Technology
Building 820, Room 601
Gaithersburg, MD 20899-0001
(301) 975-3624
FAX: (301) 216-2075

Frederick C. Johnson, Ph.D.
Associate Director for Computing,
Information Technology Laboratory
National Institute of Standards and Technology
Building 820, Room 601
Gaithersburg, MD 20899-0001
(301) 975-2700
FAX: (301) 216-2075

Ronald F. Boisvert, Ph.D.
Leader, Mathematical Software Group,
Information Technology Laboratory
National Institute of Standards and Technology
Building 820, Room 364
Gaithersburg, MD 20899-0001
(301) 975-3812
FAX: (301) 990-4127

James Fowler
SIMA Program Manager
National Institute of Standards and Technology
Building 304, Room 04
Gaithersburg, MD 20899-0001
(301) 975-3180
FAX: (301) 926-3842

Craig W. Hunt
Acting Chief, Advanced Network Technologies Division
National Institute of Standards and Technology
Building 820, Room 445
Gaithersburg, MD 20899-0001
(301) 975-3600
FAX: (301) 590-0932

Steven R. Ray
Chief, Manufacturing Systems Integration Division
National Institute of Standards and Technology
Building 220, Room A127
Gaithersburg, MD 20899-0001
(301) 975-3524
FAX: (301) 258-9749

Shukri A. Wakid, Ph.D.
Director, Information Technology Laboratory
National Institute of Standards and Technology
Technology Building 225, Room B260
Gaithersburg, MD 20899-0001
(301) 975-2904
FAX: (301) 840-1357

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

William T. Turnbull
Deputy 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

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

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

Ronald D. 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

Alexander E. MacDonald, Ph.D.
Director, Forecast Systems Laboratory
National Oceanic and Atmospheric Administration
325 Broadway
Boulder, CO 80303
(303) 497-6378
FAX: (303) 497-6821

Carl Staton
Manager, NOAA Network Information Center
National Oceanic and Atmospheric Administration
4700 Silver Hill Road
MS 9909
Washington, DC 20233-9909
(301) 457-5165
FAX: (301) 457-5199

NSA

George R. Cotter
Chief Scientist
National Security Agency
9800 Savage Road, Suite 6217
Fort Meade, MD 20755-6217
(301) 688-6434
FAX: (301) 688-4980

Norman S. Glick
Senior Computer Scientist
National Security Agency
9800 Savage Road, Suite 6217
Fort Meade, MD 20755-6217
(301) 688-8448
FAX: (301) 688-4980

Andrew J. Arenth
Chief,
INFOSEC Research and Technology Office(R2)
National Security Agency
Fort Meade, MD 20755
(301) 688-0840
FAX: (301) 688-0255

Bruce B. Bottomley
Technical Director, Information Technology Group
National Security Agency
9800 Savage Road, Suite 6643
Fort Meade, MD 20755-6643
(301) 688-9935
FAX: (301) 688-9169

Robert V. Meushaw
Technical Director,
Office of INFOSEC Research and Technology
National Security Agency
9800 Savage Road, Suite 6529
Fort Meade, MD 20755-6529
(301) 688-0840
FAX: (301) 688-0255

Richard L. Bloom
Special Assistant to the Director of the National Computer Security Center
National Security Agency
9800 Savage Road, Suite 6765
Fort Meade, MD 20755-6765
(410) 859-4372
FAX: (410) 859-4375

NSF

Juris Hartmanis, 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

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

Robert Borchers, Ph.D.
Division Director, Division of Advanced Scientific Computing
National Science Foundation
4201 Wilson Boulevard, Suite 1122
Arlington, VA 22230
(703) 306-1970
FAX: (703) 306-0632

John Cherniavsky, Ph.D.
Head, Office of Cross-Disciplinary Activities
National Science Foundation
4201 Wilson Boulevard, Suite 1160
Arlington, VA 22230
(703) 306-1980
FAX: (703) 306-0589

Michael Lesk, Ph.D.
Division Director, Information and Intelligent Systems
National Science Foundation
4201 Wilson Boulevard, Suite 1115
Arlington, VA 22230
(703) 306-1930
FAX: (703) 306-0599

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

Gary W. Strong, Ph.D.
Program Director, Human Computer Interaction Information and Intelligent Systems
National Science Foundation
4201 Wilson Boulevard, Suite 1115
Arlington, VA 22230
(703) 306-1928
FAX: (703) 306-0599

Javad Boroumand
Associate Program Director, Advanced Network Infrastructure Program
National Science Foundation
4201 Wilson Blvd., Suite 1175
Arlington, VA 22230
(703) 306-1949
FAX: (703) 306-0621

Tatsuya Suda, Ph.D.
Program Director, Networking Research Program, Division of Advanced
Networking Infrastructure and Research
National Science Foundation
4201 Wilson Blvd.
Arlington, VA 22230
(703) 306-1950
FAX: (703) 306-0621

VA

Daniel L. Maloney
Director, Veterans Health Administration, Emerging Technology
Department of Veterans Affairs
8403 Colesville Road, Suite 200
Silver Spring, MD 20910
(301) 427-3700
FAX: (301) 427-3711

OMB

Eric L. Macris
Program Examiner
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) 395-4706
FAX: (202) 395-4817

OSTP

Duncan T. Moore, Ph.D.
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

Henry C. Kelly, Ph.D.
Assistant 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) 456-6033
FAX: (202) 456-6023

Lori A. Perine
Senior Policy Advisor, Computing, Information, and Communications, 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) 456-6039
FAX: (202) 456-6023

FY 1999 Editorial Group

Executive Editor
Sally E. Howe
National Coordination Office

Editor
Terrence L. Ponick
National Coordination Office

Editorial Assistant
Kristin Janger
National Coordination Office

Writing Group

Michael J. Ackerman, NIH
Martin Blumsack, NIH
Ronald F. Boisvert, NIST
Richard Dubois, NIH
Christine M. Falsetti, NASA
Norman S. Glick, NSA
Daniel A. Hitchcock, DOE
Luis G. Kun, AHCPR
Robert Jaeger, ED
Mark Leon, NASA
Teresa F. Lunt, NSA
Robert L. Martino, NIH
Phillip L. Milstead, NASA
Joan H. Novak, EPA
Mary Anne Scott, DOE
Gary W. Strong, NSF
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. We thank the staff of the National Coordination Office for their hard work and assistance.

Abstract

The Federal Computing, Information, and Communications (CIC) programs are investing in long-term research and development (R&D) to advance computing, information, and communications in the United States, and to help Federal departments and agencies fulfill their evolving missions in the 21st century. This report summarizes the goals and objectives of the CIC R&D programs by Program Component Area (PCA) and presents a condensed view of some of the programs' FY 1998 accomplishments and FY 1999 plans.

The CIC PCAs are:

HECC High-End Computing and Computation

LSN Large Scale Networking, including the Next Generation Internet initiative

HCS High Confidence Systems

HuCS Human-Centered Systems

ETHR Education, Training, and Human Resources.

Special sections are included on the Partnership for Advanced Computational Infrastructure (PACI) and Netamorphosis — the demonstration of leading-edge Next Generation Internet applications held in Washington in March, 1998. The document includes a comprehensive list of Government personnel who are involved in the CIC programs.

For additional copies or further information please contact:

National Coordination Office for Computing, Information, and Communications
4201 Wilson Boulevard, Suite 690
Arlington, VA 22230
Voice: (703) 306-4722
Fax: (703) 306-4727

NCO:
Web: *http://www.ccic.gov/*
E-mail: nco@ccic.gov
NGI:
Web: *http://www.ngi.gov/*
E-mail: ngi@ccic.gov

Also available on the NSTC Home Page via a link from the OSTP Home Page at:

http:/www.whitehouse.gov/WH/EOP/OSTP/html/OSTP_Home.html

National Science and Technology Council

Committee on Technology

Subcommittee on Computing, Information, and Communications R&D

Suite 690
4201 Wilson Blvd.
Arlington, VA 22230
(703) 306-4722

nco@ccic.gov

http://www.ccic.gov
http://www.ngi.gov