High Performance Computing and Communications

Information Technology Frontiers for a New Millennium

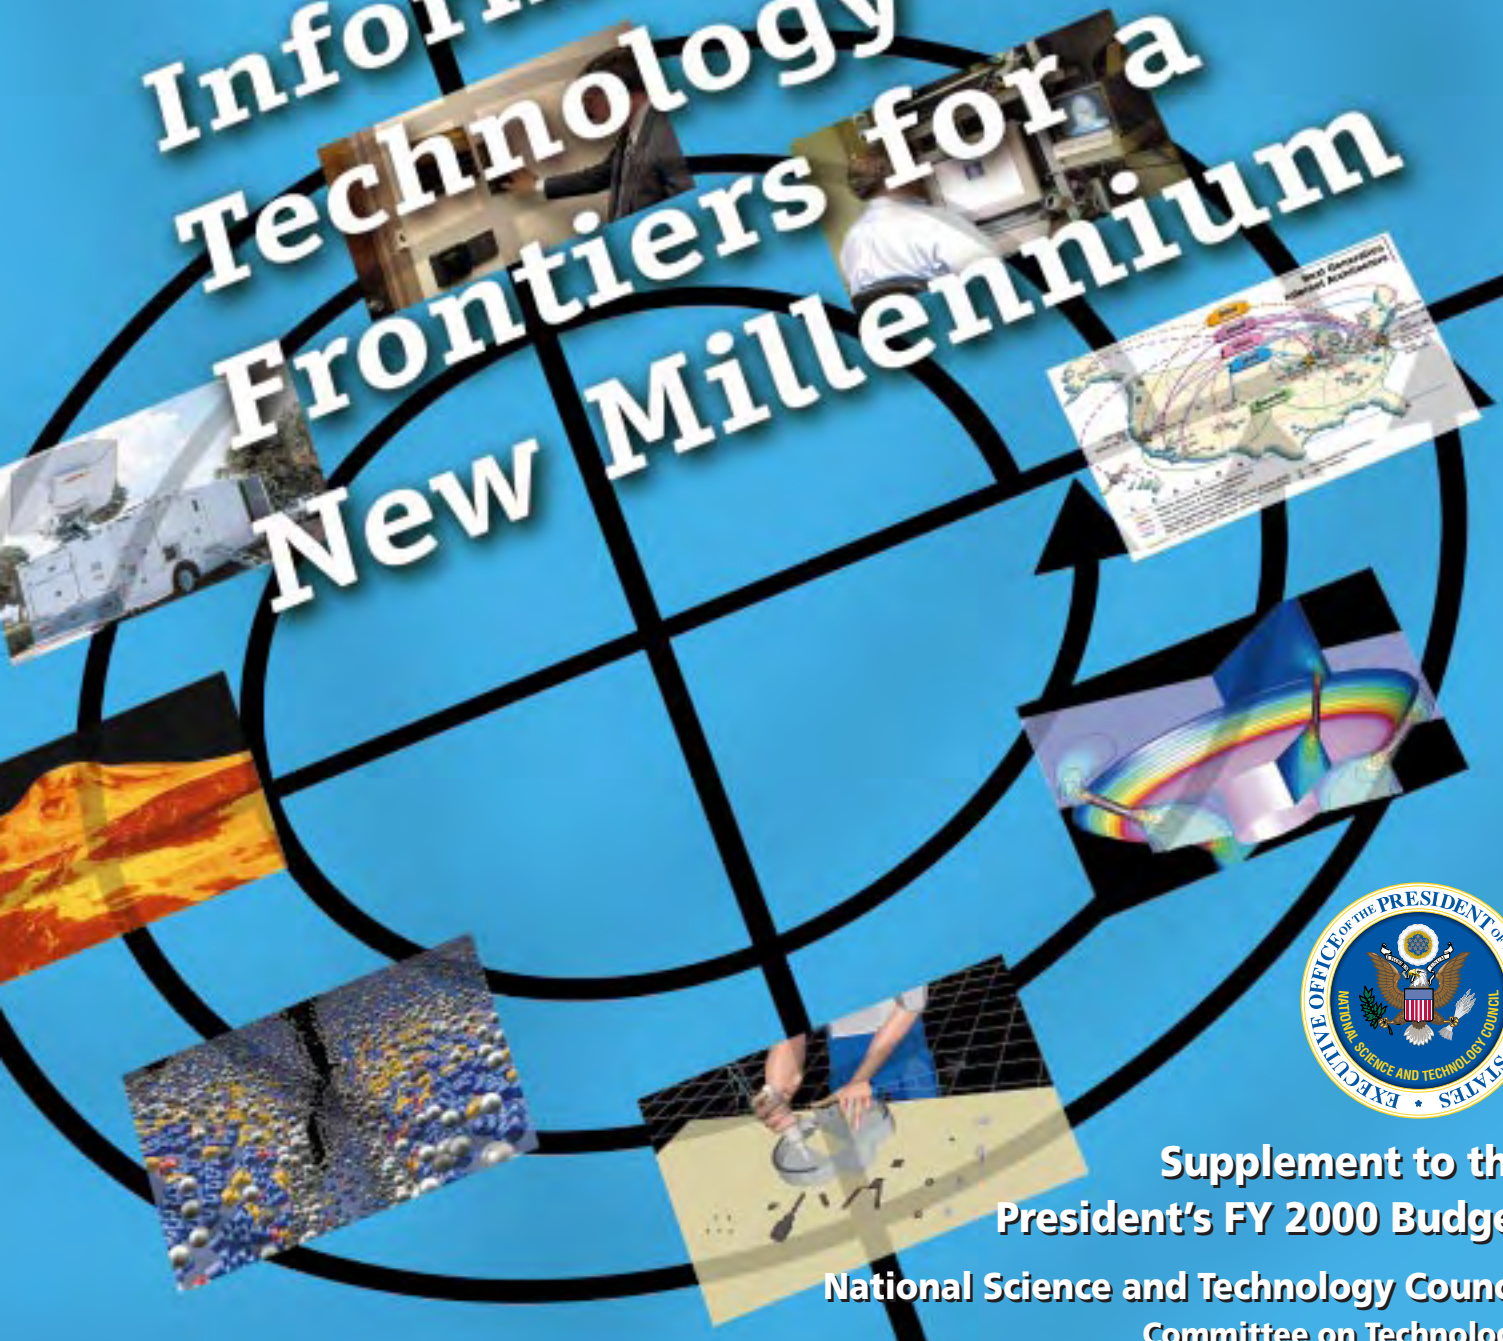

Supplement to the President's FY 2000 Budget

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, please 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 and pictured against the backdrop of a symbolic "technology spiral," whose components are research and development, partnerships, privatization, and commercialization.

HIGH PERFORMANCE COMPUTING
AND COMMUNICATIONS

INFORMATION TECHNOLOGY FRONTIERS FOR A NEW MILLENNIUM

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

Committee on Technology

National Science and Technology Council

April 1999

THE WHITE HOUSE

WASHINGTON

April 8, 1999

MEMBERS OF CONGRESS:

I am pleased to forward with this letter "High Performance Computing and Communications: Information Technology Frontiers for a New Millennium," a report prepared by the Subcommittee on Computing, Information, and Communications Research and Development of the National Science and Technology Council's Committee on Technology. This report supplements the President's FY 2000 Budget and highlights the High Performance Computing and Communications (HPCC) programs' FY 1999 accomplishments and FY 2000 plans.

Today's HPCC programs provide the research and development (R&D) foundations to help advance software, infrastructure, and applications that allow us to develop a world-class education system, build a stronger national defense, create access to affordable high quality health care, and improve the quality of life for all Americans. Working cooperatively with academia and industry, the multi-agency HPCC progams descrived in this report support important national missions and foster innovation and discovery through development of advanced computing infrastructures.

In the FY 2000 budget, the President proposes to expand and extend the HPCC programs with a $366 million initiative known as IT2, or Information Technology for the Twenty First Century. Developed in response to the 1999 report by the Congressionally-chartered President's Information Technology Advisory Committee (PITAC), the initiative will provide increased support of long-term, high-risk, fundamental research in computing; develop new advanced computational infrastructure for science and engineering; and fund research into the social, economic, and workforce implications of information technologies. The initiative is a vital first step in responding to the PITAC's call for increased funding for long-term, high-risk R&D that will enable our nation to remain at the forefront of the Information Revolution.

As we embark on the new information-driven millennium, the Administration looks forward to working with Congress to strengthen investments in fundamental information technology research to solidify further the technological base that lies at the heart of America's scientific and economic leadership.

Neal Lane
Assistant to the President
for Science and Technology

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

Delores M. Etter
Vice Chair
Deputy Undersecretary of Defense (Science 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. Ruzena Bajcsy
Assistant Director, Directorate for Computing and Information Science and Engineering
National Science Foundation

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

William Craft
Director, Multilateral Trade Affairs
Department of State

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

Elgie Holstein
Associate Director, Natural Resources, Energy, and Science
Office of Management and Budget

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

Mark McClellan
Deputy Assistant Secretary for Economic Policy
Department of Treasury

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

Ying Shih
Acting Director, Investment Program Office
Central Intelligence Agency

Samuel L. Venneri
Chief Technologist
National Aeronautics and Space Administration

James W. Vollman
Director for Labor Market Information
Department of Labor

Eugene Wong
Assistant Director, Directorate for Engineering
National Science Foundation

Subcommittee on Computing, Information, and Communications R&D

NCO

Chair
Kay Howell

NSF

Representative
Ruzena Bajcsy
Alternate
George Strawn

DARPA

Representative
David Tennenhouse
Alternate
Ronald L. Larsen

DOE

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

NIH

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

NASA

Representative
Lee B. Holcomb
Alternate
Betsy Edwards

NSA

Representative
George R. Cotter
Alternate
Norman S. Glick

NOAA

Representative
Thomas N. Pyke, Jr.
Alternate
William T. Turnbull

NIST

Representative
William Mehuron
Alternate
R. J. (Jerry) Linn

AHCPR

Representative
J. Michael Fitzmaurice

EPA

Representative
Joan H. Novak
Alternate
Robin L. Dennis

VA

Representative
Daniel L. Maloney

ED

Representative
Linda G. Roberts
Alternate
Alexis T. Poliakoff

OMB

Eric L. Macris
Elizabeth M. Robinson

OSTP

Henry A. Kelly
Lori A. Perine

High End Computing and Computation (HECC) Working Group

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

Large Scale Networking (LSN) Working Group

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

High Performance Networking Applications Team (HPNAT)

Co-Chairs
William T. Turnbull, NOAA
William F. Decker, NSF

Internet Security Team (IST)

Co-Chairs
Grant M. Wagner, NSA
Vacant

Joint Engineering Team (JET)

Co-Chairs
Javad Boroumand, NSF
Phil Dykstra, DoD

Networking Research Team (NRT)

Co-Chairs
Mari W. Maeda, DARPA
Vacant

High Confidence Systems (HCS) Working Group

Co-Chairs
Helen Gill, DARPA
Paul Pittelli, NSA (Acting)

Human Centered Systems (HuCS) Working Group

Chair
Michael Lesk, NSF

Education, Training, and Human Resources (ETHR) Working Group

Chair
John Cherniavsky, NSF

Federal Information Services and Applications Council (FISAC)

Co-Chairs
G. Martin Wagner, GSA
Vacant

Digital Government Liaison
Lawrence E. Brandt, NSF

Federal Statistics Team

Chair
Cathryn S. Dippo, BLS

Federal Web Consortium Liaison
Valerie Gregg, NSF

Information Technology for Crises Management Team

Chair
Anngienetta R. Johnson, NASA

Next Generation Internet Team

Chair
William T. Turnbull, NOAA

Universal Access Team

Co-Chairs
Susan B. Turnbull, GSA
Gary W. Strong, NSF/DARPA

Table of Contents

Executive Summary

As we near the dawn of a new millennium, advances made possible by computing, information, and communications research and development (R&D) — once barely noticeable to the general public — are reaching dramatically past the scientists, engineers, and educators who fostered them, forever transforming every aspect of our work and our lives in ways that are exciting but unpredictable. Seemingly overnight, the Industrial Age has become the Information Age, and a shining future beckons on the horizon, a future where neither distance, disability, nor economic difficulties will be barriers to a brighter tomorrow. All citizens, wherever they may live or work, will have equal access to the developing knowledge and learning tools that will assure U.S. competitiveness and economic well-being far into the next century.

Much of this promising future is a direct result of decades of investments in information technology research by the Federal government, working cooperatively with academia and the private sector. The high performance computing technologies so vital in bringing the Cold War to a successful conclusion are now transforming our social and economic lives through a vast network of computers and servers that conduct much of our business and manage the Internet, providing us with unprecedented and inexpensive access to critical information. Indeed, the Internet itself grew out of the military's need to move information among different computing systems. Private sector demand for these technologies grew at an unexpectedly rapid pace when Web browsers, initially developed with Federal funding, took the world by storm in 1994.

Federally-supported long term, fundamental research in high performance computing and communications has been a critical driving force behind today's impressive economic and social transformations. But in recent years, Federal investments in vital high performance computing and communications research have not kept pace with the rapid development of computing, information, and communications technologies. To assure the continued health, prosperity, and economic competitiveness of future generations, the Federal government must again take the lead in this area, significantly increasing its investments in the same type of R&D that catapulted the U.S. into the leadership of the Information Age.

The Information Technology for the Twenty-first Century (IT^2) initiative

A highlight of the President's FY 2000 R&D budget is the proposed Information Technology for the Twenty-first Century (IT^2) initiative. This research initiative proposes $366 million in increased investments to help advance the knowledge base in fundamental information science and to train the next generation of researchers who will sustain the Information Revolution well into the 21st Century. Building on the Government's previous accomplishments and existing investments in High Performance Computing and Communications (HPCC), including the Next Generation Internet (NGI), and the Department of Energy's (DOE's) Accelerated Strategic Computing Initiative (ASCI), the initiative will extend some existing research and provide opportunities to address new, complementary research topics in three key areas:

- Long term information technology research that will lead to fundamental advances in computing and communications
- Advanced computing infrastructure as a tool to facilitate scientific and engineering discoveries of national interest
- Research on the economic and social implications of the Information Revolution, and efforts to help train additional IT workers at our universities

The IT^2 research agenda responds directly to the findings and recommendations of the President's Information Technology Advisory Committee (PITAC), which concluded in a report released in February 1999 that the Federal Government is underinvesting in long term IT research relative to its importance to the Nation. The initiative is garnering widespread support from industry and academia. If approved by Congress, the coordination and implementation of the IT^2 initiative will be integrated with the HPCC R&D programs.

The HPCC R&D programs

The Federal government coordinates multiagency research in computing and communications through the HPCC R&D programs (formerly known as the Computing, Information, and Communications [CIC] R&D programs). HPCC-coordinated activities are organized into five Program Component Areas (PCAs):

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

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

HECC: High End Computing and Computation

HECC R&D provides the foundation for U.S. leadership in high end computing, promoting its use in government, academia, and industry. HECC researchers are developing computation-intensive algorithms and software for modeling and simulating complex physical, chemical, and biological systems; information-intensive science and engineering applications; and advanced concepts in quantum, biological, and optical computing.

The HECC Working Group (HECCWG) coordinates Federal R&D dedicated to maintaining and expanding U.S. leadership in high performance computing and computation, which includes algorithms, architecture, components, software, and high end mission applications. The HECCWG also promotes cooperation in high end computing and computation R&D among Government laboratories, academia, and industry.

LSN: Large Scale Networking

LSN R&D provides the leadership in networking technologies, services, and performance to meet Federal agency mission needs and to develop technologies that enable the future growth of the Internet. Key LSN R&D areas include technologies for highly capable very high speed networks and applications that require such technologies.

LSN activities are coordinated by a Working Group and four Teams, each of which includes non-Federal participants:

- The Joint Engineering Team (JET) coordinates connectivity among the Federal agency networks (FedNets) — DOE's Energy Sciences network (ESnet), the National Aeronautics and Space Administration's (NASA's) Research and Education Network (NREN), the National Science Foundation's (NSF's) very high performance Backbone Network Services (vBNS), and the Department of Defense's (DoD's) Defense Research and Engineering Network (DREN) — with the Abilene network (a university/industry partnership), with international networks at the Chicago-based Science, Technology, and Research Transit Access Point (STAR TAP), and with geographically disadvantaged states such as Alaska and Hawaii.
- The Networking Research Team (NRT) coordinates agency networking research programs, shares networking research information among Federal agencies, and supports NGI networking R&D activities. It provides outreach to end users to promote dissemination of networking research information and to promote coordination among end users and applications developers.
- The High Performance Networking Applications Team (HPNAT) coordinates Federal R&D in high performance networking applications in science and engineering, weather and the environment, biomedicine, and healthcare.
- The Internet Security Team (IST) facilitates testing and experimentation with emerging advanced security technologies and serves as a focal point and clearinghouse for application and engineering requirements for security systems.

NGI: Next Generation Internet Initiative

The 1991 High Performance Computing Act was amended with bipartisan support in FY 1998 by the Next Generation Internet Research Act, which authorizes the NGI Initiative. The NGI Initiative — coordinated under the LSN Working Group — is:

- Conducting R&D and experimentation in networking technologies in order to add functionality and improve performance. This includes hybrid networks (including satellites and terrestrial components), Internet security, Internet protocol (IP) over wave division multiplexing (WDM), mobile networks, multicast, network management and modeling (for reliability and robustness), optical add-drop multiplexers, quality of service (QoS), and test and evaluation.
- Developing two testbeds for system-scale testing and for developing and demonstrating advanced applications. The 100x testbed includes the FedNets and connects about 130 universities and Federal facilities with end-to-end performance that will be 100 times faster than the Internet of 1997; 25 more connections are expected in FY 2000. The Defense Advanced Research Projects Agency's (DARPA's) SuperNet, the 1,000x testbed, connects about 20 sites nationwide at speeds that will be 1,000 times faster than 1997's Internet.
- Developing and demonstrating revolutionary applications — in basic science, crisis management, education, the environment, Federal information systems, healthcare, and manufacturing — and their enabling technologies — collaboration technologies, digital libraries, distributed computing, privacy and security, and remote operation and simulation.

SC98 demonstrations

More than a dozen NGI and NGI-related applications were demonstrated at several research exhibits at the SC98 supercomputing conference held in Orlando, Florida, including:

- Realtime functional Magnetic Resonance Imaging (fMRI): Watching the brain in action
- Distributed Image Spreadsheet: Earth data from satellite to desktop
- Remote access multidisciplinary microscopy (RAMM): Viewing the changing threads of life with 4-D telemicroscopy
- Broadcast news navigator
- Collaborative remote robotic arc welding
- Exploring the Earth system on the "second Web"
- GeoWorlds: Integrated digital libraries and geographic information systems for disaster relief operations
- Testing and measuring Internet security technologies
- Testing and measuring Internet Protocol (IP) QoS.
- iGrid: Managing and measuring geographically distributed networking, computing, storage, and display resources accessible from the STAR TAP

HCS: High Confidence Systems

HCS R&D is developing technologies for achieving predictably high levels of computing and communications system availability, reliability, safety, security, and survivability. Such systems must withstand internal and external threats and natural disasters. These technologies are needed as we increasingly rely on our information infrastructure to support our financial, healthcare, manufacturing, power (electricity, natural gas, nuclear power), and transportation infrastructures.

HCS activities include National Security Agency (NSA) work in active network defense, secure network management, network security engineering, cryptography, and secure communications; DARPA work in information survivability; NSF's computing-communications research; the National Institute of Standards and Technology (NIST)-NSA National Information Assurance Partnership (NIAP); and National Institutes of Health (NIH) research in protecting patient records and in collaborative technologies for telemedicine.

The HCS Working Group is preparing an agenda for new research in theoretical foundations, tools and techniques, engineering and experimentation, and pilots and demonstrations, which has contributed to the IT^2 initiative.

HuCS: Human Centered Systems

HuCS R&D focuses on improving the interactions among humans, computing systems, and information resources. Federal investments in HuCS R&D benefit scientists, physicians, engineers, educators, students, the workforce, and the general public in dozens of disciplines, including biomedicine, defense, manufacturing, education, library sciences, law enforcement, weather forecasting, and crisis response. Areas of HuCS research are:

- Knowledge repositories and information agents, including the multiagency Digital Libraries Initiative Phase Two; the National Library of Medicine's (NLM's) Visible Human and Unified Medical Language System (UMLS) projects; and DARPA's Text, Radio, Video, and Speech program supporting information extraction for battlefield awareness

- Multimodal human-computer interfaces, including NSF's Speech, Text, Image, and MULtimedia Advanced Technology Effort (STIMULATE) and other speech recognition technologies developed by NIST, DARPA, and NSA
- Multilingual technologies, including Spanish language interfaces developed by NSF, DARPA, and other U.S. and international organizations
- Universal access, including the Web Accessibility Initiative to assure Internet accessibility to all people, sponsored by NSF, the Department of Education (ED), several U.S. corporations, and the European Commission; and ED's National Institute on Disability and Rehabilitation Research
- Visualization, virtual reality, and robotic tools, including "Virtual Los Angeles," NSF's realtime visualization system for large scale urban environments, NSF's robotic surgery, NASA's "software scalpel" and computerized breast cancer diagnostic tools, and NOAA's Virtual Worlds
- Collaboratories, including DOE's Materials Microcharacterization Collaboratory (MMC) and NIST's Systems Integration for Manufacturing Applications

ETHR: Education, Training, and Human Resources

ETHR R&D supports education and training in computing, communications, and information technologies from K-12 through postgraduate training and life-long learning. ETHR facilitates the development of software learning tools, modeling of education and learning, research on cognitive processes, and demonstrations of innovative technologies and applications. Activities include NSF's studies of learning and intelligence in systems, DOE's Computational Science Graduate Fellowships, NIH/NLM's biomedical informatics training grants, and NASA's Learning Technologies Project for using its vast data collections.

FISAC: Federal Information Services and Applications Council

The FISAC helps apply HPCC technologies to Federal information systems and services to support agency missions by providing two-way communications channels with agencies that are either wholly or in part outside of the HPCC budget crosscut. The HPCC R&D programs use FISAC to disseminate information about their research agendas, priorities, and results, while the FISAC provides feedback and identifies research needed by Federal applications. FISAC works primarily through its FedStats, Information Technologies for Crises Management, NGI Applications, and Universal Access Teams, and its liaison with NSF's Digital Government program and the Federal Web Consortium. Crises management and Federal statistics research agendas are being developed through NSF-sponsored workshops held by the National Research Council's Computer Science and Telecommunications Board. The FedStats Web site (*www.fedstats.gov*), a gateway to the 70-plus Federal agencies that compile and disseminate Federal statistics, is funded by the Interagency Council on Statistical Policy chaired by the Office of Management and Budget (OMB). The Universal Access Team has merged with the General Services Administration's (GSA's) Accessible Technology working group.

President's Information Technology Advisory Committee (PITAC)

The PITAC was established by Executive Order in February 1997, as authorized by the HPC Act of 1991. Renewed for a second two year term in February 1999, the Committee is comprised of 26 academic and industry leaders who are charged to provide the President with an independent assessment of the Federal government's role in HPCC, information technology, and Next Generation

Internet R&D. The PITAC released its Report to the President in February 1999. It recommended creating a strategic research initiative in software, scalable information infrastructure, high end computing, and the socioeconomic impact of IT. The report also addressed management and modes of funding for Federal IT research. The PITAC's recommendations are the primary driver of efforts to enhance the Federal IT research portfolio, primarily through the IT^2 initiative, which will help to refocus and strengthen existing Federal investments in the HPCC R&D programs, including the NGI initiative.

HPCC R&D budget and coordination

The proposed FY 2000 multiagency HPCC R&D budget is $1,462 million, representing an 11 percent increase over the estimated $1,314 million in FY 1999. The HPCC R&D programs are coordinated by the Subcommittee on CIC R&D. This Subcommittee reports to the Committee on Technology of the President's National Science and Technology Council (NSTC). The Subcommittee works through its five PCA Working Groups and their Teams.

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

The NCO is charged by the White House Office of Science and Technology Policy (OSTP) to ensure coordination of multiagency Federal information technology R&D by supporting the Subcommittee on CIC R&D, the PITAC, and activities related to the development of the IT^2 initiative. It facilitates the preparation of multiagency planning, budget, and assessment documents. The NCO Director, who reports to the Director of OSTP, chairs the Subcommittee on CIC R&D.

The NCO provides a central point of contact about HPCC and PITAC activities for the Congress; Federal, state, and local organizations; academia; industry; professional societies; foreign organizations; and others to exchange technical and programmatic information.

Each year the NCO responds to thousands of information requests with Web, print, and video material that includes Congressional testimony; HPCC, NGI, and NCO publications; and PITAC publications and meeting material.

Purpose of this report

This report highlights ongoing and proposed HPCC efforts, focusing on representative FY 1999 accomplishments, key FY 2000 R&D areas, and the budget crosscut.

Information now on the Web

Copies of HPCC, NGI, CIC, PITAC, and IT^2 publications, links to participating agency and related Web sites, and this report can be found at:

http://www.ccic.gov

and

http://www.ngi.gov

IT²: A Bold Investment in America's Future

With the Information Technology for the Twenty-first Century (IT^2) initiative, the Federal Government is planning a recommitment to fundamental research in information technology. Building on the Government's previous accomplishments and current investments in the High Performance Computing and Communications (HPCC) Program, including the Next Generation Internet (NGI) Initiative, the proposed IT^2 initiative responds directly to the findings and recommendations of the President's Information Technology Advisory Committee (PITAC) (described beginning on page 101.

PITAC recommendations

In its February, 1999, report to the President, the PITAC — which includes leaders from academia and industry — concluded that the Government is underinvesting in the kind of long term information technology research that is critical to the Nation's future. The PITAC recommended that the Government undertake a bold strategic initiative supporting long term research in fundamental computing, information, and communications areas. The Committee also recommended increased Federal investments in powerful high end computing systems to support advanced research. Finally, the PITAC recommended conducting research to identify, understand, anticipate, and address the socioeconomic effects that the rapid development of information technology will have on U.S. citizens in the coming millennium.

Information Technology (IT) is critical to the Nation's future

Leading-edge information technology (IT), due to its enormous and profound socioeconomic benefits, has, seemingly in a few short years, become critical to our Nation's continued well-being and prosperity. IT is also changing the way we live, work, learn, and communicate with each other. For example, advances in IT can improve the way we educate our children, allow people with disabilities to lead more independent lives, and improve the quality of healthcare for rural Americans through technologies such as telemedicine.

IT advances in supercomputers, simulations, and networks are creating a new window into the natural world, making high end computational experimentation a vital tool for pathbreaking scientific discoveries. U.S. leadership in information technology is also essential for our national security. Our military strategy now relies on information superiority to gain crucial advantages over our adversaries and to keep our troops out of harm's way.

Benefits of past Government investments

The technologies resulting from past Federally sponsored research (for example, the Internet, the first graphical Web browser, and advanced microprocessors) have helped strengthen American leadership in the IT industry. IT now accounts for one third of U.S. economic growth and employs 7.4 million Americans at wages that are more than 60 percent higher than the private-sector average. All sectors of the U.S. economy are using information technology to compete and win in global markets, and business-to-business electronic commerce in the U.S. alone is projected to grow to $1.3 trillion by 2003.

The Nation needs significant new investments in IT research to help ensure future economic growth and to address important and persistent national needs such as defense, education, the environment, healthcare, and transportation.

IT² investments

The IT² initiative will support:

- Long term information technology research leading to fundamental breakthroughs in computing and communications, in the same way that Government investment begun in the 1960s led to today's Internet
- Advanced computing for science, engineering, and the Nation, including supercomputers, software, networks, and the research teams needed to support and use them. Advanced computing will support applications such as reducing the time required to develop life-saving drugs; designing cleaner, more efficient engines; more accurately predicting hurricanes and tornadoes as well as long term climate change; and accelerating scientific discovery.
- Research on the economic and social implications of the Information Revolution and efforts to help train additional information technology workers at our universities

Proposed FY 2000 IT² budget and participating agencies

The proposed fiscal year (FY) 2000 budget for IT² is $366 million. The proposed participating agencies are:

- Department of Defense (DoD) (including the Defense Advanced Research Projects Agency [DARPA])
- Department of Energy (DOE)
- National Aeronautics and Space Administration (NASA)
- National Institutes of Health (NIH)
- National Oceanic and Atmospheric Administration (NOAA)
- National Science Foundation (NSF)

The IT² Draft Implementation plan and more details about the IT² initiative can be found at:

http://www.ccic.gov/

High End Computing and Computation

HECC research and development (R&D) provides the foundation for U.S. leadership in high end computing, promoting its use in government, academia, and industry. HECC researchers are developing computation-intensive algorithms and software for modeling and simulating complex physical, chemical, and biological systems; information-intensive science and engineering applications; and advanced concepts in quantum, biological, and optical computing.

The HECC Working Group (HECCWG) coordinates Federal R&D dedicated to maintaining and expanding U.S. leadership in high performance computing and computation, which includes systems architecture, hardware, foundational and algorithmic research, software, and high end mission-oriented applications. The HECCWG also promotes Federal cooperation in high end computing and computation R&D among Government laboratories, academia, and industry.

This section describes the FY 1999 accomplishments and FY 2000 plans in HECC R&D.

High end architectures

Hybrid Technology Multithreaded (HTMT) architecture

Researchers funded by DARPA, NASA, and the National Security Agency (NSA) are evaluating the feasibility of constructing a computing system capable of a sustained rate of 10^{15} floating point operations per second (one petaflop). Preliminary evaluations of conventional architectures and mainstream technologies indicate that such a system will necessitate a radical departure from current research. HTMT architecture would blend modified semiconductor technology with leading-edge hybrid technologies including superconducting technology, optical interconnects, high speed very large scale integration (VLSI) semiconductors, and magnetic storage technology configured to satisfy the architecture requirements. Fundamental drivers are multi-gigahertz speeds, exceptional bandwidths, and very large and cost-efficient memory size. If results from FY 1999 R&D are encouraging, researchers are expected to build major sections of a prototype beginning in FY 2000.

Quantum computing

Current trends in miniaturizing computer components suggest that they will reach the quantum scale by 2020. Today's computers count in binary digits - zeroes and ones. A theoretical quantum computer (QC) counts in "qubits," which are superpositions of zeroes and ones that might be represented, for example, by the direction of a spinning electron. Quantum computers will be able to perform non-classical logic operations and use the phenomenology of quantum parallelism to enable the solution of problems that are intractable with even the most powerful conventional parallel supercomputers that can be envisioned today.

Supported by NSA, more than nine universities and companies are conducting QC research. Projected accomplishments include:

Artist's 3-D depiction of a prototype computing system employing HTMT architecture. DARPA, NASA, and NSA researchers are evaluating the feasibility of constructing a hybrid HTMT system that would be capable of a sustained rate of 10^{15} floating point operations per second (one petaflop).

- Developing the quantum analogues of classical information and communication theory to help in future QC design and evaluation
- Applying the "hidden subgroup" problem to the design of new quantum algorithms
- Studying the computational power achievable in models positioned between classical and quantum computing
- Verifying a new pulsed-laser technique to implement quantum logic gates between photons
- Creating entanglements of three photons
- Achieving coherent quantum state manipulations of a few ions
- Uncovering clear evidence of quantum superposition, entanglement, and multi-qubit operations
- Characterizing the resulting quantum coherence

Future QC research includes demonstrating 1-qubit operations and performing experiments to achieve 2-qubit operations using the optical-lattice method of trapping atoms and simulating the dynamics of a set of qubits in finer detail than has previously been achieved.

The initial five-year phase of NSA's quantum computing project ends in FY 1999. Full-scale development of a practical quantum computer may require 20 or more years of research.

Quantum teleportation

DOE's Oak Ridge National Laboratory (ORNL) has established a state-of-the-art laboratory for quantum teleportation, one research area in the field of quantum computing and communications. Quantum theory indicates that one can instantaneously teleport a quantum state over arbitrary distances, but one can use it only after a subsequent "classical" communication is carried out. To attempt to understand what is actually teleported, and what can and cannot be achieved supraluminally, ORNL is conducting experiments in which a signal transmitter and the transmitted signal are quantum mechanically entangled. High power femtosecond lasers are used for the experiments.

Commodity clusters: high performance computing on desktop systems

The high performance virtual machine (HPVM) software project, supported by NSF, allows off-the-shelf desktop computers to be used as high performance clusters. Researchers at the NSF-funded National Computational Science Alliance (NCSA) and National Partnership for Advanced Computational Infrastructure (NPACI) have built a Windows NT supercluster consisting of hundreds of workstations running on Intel Pentium II processors, and have run the complex astrophysical hydrodynamics code, ZEUS-MP, at the rate of several gigaflops. HPVM allows users to easily port their Message Passing Interface (MPI) codes to the NT cluster environment, providing a low-cost alternative to conventional supercomputing.

Evaluation of high end architectures

ORNL is evaluating new high end architectures to assess their applicability for DOE and other Government applications. One of these architectures is the SRC-6, which combines commodity microprocessor technology with a new high performance Memory Algorithm Processor (MAP) subsystem, a programmable hardware component that allows programmers new flexibility to customize the hardware for specific applications. ORNL is evaluating the MAP prototype subsystem and plans to take delivery of an SRC-6.

NPACI, along with DOE's National Energy Research Supercomputer Center (NERSC), Cal Tech, and Boeing Computer Services are also evaluating the Tera Multi-Threaded Architecture (MTA), located at the San Diego Supercomputer Center (SDSC). This system seeks to achieve high degrees of parallelism by splitting a program into hundreds of individual execution threads and having the system hardware execute whichever thread is ready to progress.

High end hardware components

Automating component design

NSF supports basic research in Electronic Design Automation (EDA) and in applicable VLSI design technologies, such as systems-on-a-chip, embedded systems, and multi-technology optical and micro-electromechanical design methods. Research areas include scientific methods, intellectual processes, abstractions, search paradigms, and information models used in VLSI design, covering all phases of the EDA design cycle for integrated circuits and systems from conception through manufacturing and testing. Other areas of research include:

- Theoretical foundations models (algorithms, tools, analysis, synthesis, simulation, validation, and verification)

Pictured is the diamond Multi-Chip Module (MCM) aerosol spray cooled Cray J90 supercomputer mounted vertically in the black "clamshell enclosure" on the left. One redundant aerosol spray pump/heat exchanger is visible as a silver cylinder on the lower right. The pump controller-fail safe electronics are visible in the upper right. Mounted on the far left wall of the integration platform are the internal power supplies for the pump control electronics.

- ❏ System design methods (systems-on-a-chip, multi-chip, and multi-technology systems)
- ❏ Manufacturing (fault models and algorithms for diagnosis and test in digital, analog, and mixed signal designs)
- ❏ Design and system prototyping methods, tools, and environments, including information infrastructure

Diamond-based Multi-Chip Modules (MCMs)

NSA's MARQUISE/SOLITAIRE program repackages a high performance computer using diamond-based MCMs and thin film spray cooling. The SOLITAIRE prototype will demonstrate the maturity of high density packaging technologies incorporating die-last MCMs, diamond substrates, and spray cooling. Environmental testing of the integrated system will demonstrate its reliability.

Micro spray cooling

NSA's spray cooled power converter research program is developing a high density, high efficiency power converter using spray cooling and planar laminated transformers. Current research focuses on applying isothermal spray cooling to overcome the fundamental coefficient of thermal expansion mismatch encountered during hybrid integration of silicon or gallium arsenide (GaAs) devices with diamond substrates.

Superconductive crossbar switch

Research suggests that superconducting supercomputers can deliver very high performance with very low power requirements. NSA supports research in superconducting electronics to provide high performance computing alternatives to current silicon and GaAs technologies, which have speed and power limitations. NSA is constructing a 128 x 128 superconductive crossbar switch that operates at 2.5 gigabits per second (Gbps) per port for use in supercomputing and network applications. It is self-routing, is scalable in size from 32 x 32 to more than 1024 x 1024, and has a latency of less than 4 nanoseconds.

Although the crossbar electronics operate at 4 degrees Kelvin, its input and output ports are at room temperature, and the cryogenic elements are cooled by a

refrigerator, allowing use in a standard room temperature environment. Extended to higher speed and size, this switch is a candidate element for use in HTMT architectures.

Optoelectronic research

NSA supports research on optical techniques that can theoretically support data rates of several hundred Gbps for future communications and computing systems by providing logical functions that expedite the data routing process. Transmission at high data rates demands techniques that can restore data signals so that low bit error rates can be maintained without compromising signal quality. NSA is also conducting research on a high performance spectrometer on a chip and a semiconductor optical amplifier.

Smart memories

NSA supports research to produce a flexible computing architecture more power-efficient than the current evolutionary path of reduced instruction set chip (RISC) architectures while remaining programmable in a high level language. This architecture will be able to reconfigure itself to optimize for currently executing computations in order to accommodate a wide class of coarse- to fine-grained algorithms.

High Performance Storage System (HPSS)

A consortium of DOE national laboratories, in collaboration with International Business Machines (IBM), has developed the High Performance Storage System (HPSS). With more than two years of operation and deployment at about twenty sites, HPSS has become a standard for storage systems in the high performance computing community. HPSS 4.1 contains new features such as support for the Distributed File System (DFS), file families, Message Passing Interface-Input/Output (MPI-IO), scalable accounting, and performance improvements. In FY 1999, Sun Microsystems and Storage Technology Corporation joined the HPSS community, where research continues on large database storage and transfer.

Advanced scalable computing for weather forecasting

Building on HPCC advances, NOAA will lease a new IBM high performance scalable computer that will significantly improve weather, flood, and climate forecasting for the country. Expanded NOAA forecast models are computationally demanding and require high reliability. This new high performance computing system will use a highly parallel computer architecture to run models with improved physics and greater resolution, producing forecasts with better accuracy and allowing NOAA to operate more sophisticated models of the atmosphere and oceans.

Foundational and algorithmic research

Theory of computing

NSF supports fundamental research in computing theory in three areas:

- Application-specific theory supports developing models and techniques for solving problems that arise in areas of science and engineering such as molecular biology, communications networks, and computational linguistics. Of particular interest are theoretical developments having potential impact on experimental or applied areas of computer science research, and strategies that mix theory with experimentation.
- Core theory covers computational complexity, cryptography, interactive computation, computational learning, parallel and distributed computation, computation on random data, online computation, and reasoning about knowledge.

- Fundamental algorithms include developing combinatorial, approximation, parallel, online, numerical, geometric, and graph algorithms that span application domains.

Numeric, symbolic, and geometric computation

NSF supports fundamental research in computation and graphics, including computer algebra, numerical computation and modeling of physical processes, mathematical optimization, computational geometry, imaging, deductive methods for reasoning in computational logics, and automated deduction. These areas combine mathematical analysis with advanced algorithms. NSF also supports integrating research results into problem-solving environments to support computational science and engineering. Supported research addresses generic as well as scientific and engineering discipline-specific computation, including innovative applications of advanced computational and graphic techniques in scientific and engineering systems such as manufacturing and design, proof support systems, prototyping, and design verification.

Whisker Weaving (WW)

Research in mathematical and computational sciences at DOE's Sandia National Laboratories (SNL) is currently focused on WW — an algorithm for generating 3-D unstructured hexahedral meshes — and mesh optimization. Most hexahedral meshing algorithms work on only a small class of geometries, restrict the surface mesh, and require extensive user input. WW technology is designed to eliminate these shortcomings and could significantly reduce the time needed to generate hexahedral meshes for complex geometries. As time to mesh is currently the most serious impediment to developing the automated design capability of DOE's Accelerated Strategic Computing Initiative (ASCI) and other major programs, the potential impact of both WW and mesh optimization is substantial.

NSF research in parallel and distributed computing

NSF's FY 1999 HPCC research includes parallel computation models, parallel algorithms and software for scientific computing, dynamic compilation and optimizing parallel compilation, distributed operating systems, superscalar architectures, high performance memory systems, and signal processing and communications systems research, especially in wireless information technology and networks.

NSF FY 2000 research plans include a wireless center, hardware/software codesign, and high performance scientific and commercial applications.

Global optimization

The aim of SNL's global optimization project is to develop new, innovative methods representing better algorithmic trade-offs between the ability to find a global optimum and the ability to generate near-optimal solutions quickly. The algorithmic development will focus on domain-independent methods applicable to a wide range of applications. SNL will develop a global optimization method on applications with continuous and combinatorial search domains to provide a practical evaluation of the analyses. This project has recently begun to focus on practical folding methods, including methods for off-lattice models.

Large scale scientific and engineering design optimization

A result of years of progress in modeling and simulating large structural mechanics, thermal analysis, and fluid dynamics problems is that these codes can now be used to design systems and develop predictive computational models for these systems. This in turn has created a need for new optimization algorithms to overcome difficulties such as noisy and expensive function evaluations and lack of derivative information. To address this need, SNL has developed an object-oriented library, OPT++, for nonlinear optimization. The library includes a wide variety

of methods available for both unconstrained and bound constrained optimization, large scale constrained optimization algorithms for computational chemistry problems, and new methods for rapid, robust parallel optimization.

Data mining

Data mining is the process of automatically extracting large amounts of logically coherent data from large and sometimes confusing databases. Data mining is becoming a vital process to large corporations that often store many years of data that has been compiled but never analyzed. There are many different techniques used for data mining, and often the technique that is chosen depends on the type of data and information that is to be extracted from the database. Some of the most popular techniques are association, sequence-based analysis, clustering, classification, estimation, fuzzy logic, genetic algorithms, fractal-based transforms, and neural networks.

Among the best of NCSA's data mining tools are machine learning algorithms, so-called because they are primarily adaptive systems that can learn on the job. These algorithms rely on an iterative process through which the programs learn from examples. Machine learning algorithms can be used for tactical and strategic decision making.

The ability to analyze very large data collections — billions of objects — is needed to support research in such areas as astronomical digital sky surveys and high energy physics. The challenge is to identify a relevant subset of data that is represented by a large number of attributes. Scalable statistical algorithms for clustering data are being developed within NPACI at the University of California-Davis.

Software

Next Generation Software (NGS) program

NSF's NGS program fosters research to develop performance engineering technologies for the design, management, and control of computing systems; and to create a new system software architecture to provide runtime support for complex applications executing on complex computing platforms, such as computational grids and future petaflop platforms.

NGS views the computing system as having several layers: the application layer, grid platform architecture, the system software, processing nodes, and interconnect layers. NGS research will lead to software systems that understand the interrelation among these layers as it affects the behavior of a computing system, and will guide the design, management, operation, and control of the system.

Operating systems and compilers

NSF research in operating systems and distributed systems includes developing mechanisms and applications programming interfaces (APIs) for uniform access and management of resources in local area networks (LANs) and wide area networks (WANs); middleware infrastructure for building scalable services; resource management for new applications and quality of service requirements; and security and electronic commerce. Research in compilers and runtime systems includes dynamic compilation, models of storage consistency and storage-hierarchy performance, and compiler support for programming on the World Wide Web.

Programming methods and languages

NSA supports research in computational methods and languages for massively parallel, distributed heterogeneous computing platforms and special-purpose processors. Research includes developing the compiler AC — that targets the Silicon Graphics/Cray Research (SGI/Cray) T3D and T3E architectures, among others — and implementing the "futures" model of message-passing programming.

Software engineering and languages

NSF supports fundamental research to develop quality software-based systems, including domain-specific languages for specification and design; constructive approaches to software design and evolution; issues of software modularity and composition; enhancement of confidence and quality; automating stages of software development; distributed and network environment issues, including distributed development and software security; and formal foundations for all aspects of software engineering and programming languages.

Java numerics

The rapid and widespread adoption of the Java language and environment for network-based computing has created a demand for reliable and reusable numerical software components to support the development of scientific software. The National Institute of Standards and Technology (NIST) works with the Java Grande Forum to enable the use of Java in high performance numerical computing.

Aztec Iterative Solver

SNL is conducting R&D on Aztec, a library of state-of-the-art iterative methods for solving linear systems. Its goal is to perform well on parallel computers while being easy for application engineers to use. Aztec's simplicity comes from employing a global distributed matrix that allows a user to specify pieces (different rows for different processors) of an application matrix exactly as would be done in the serial setting (that is, using a global numbering scheme). Speed comes from using standard distributed memory techniques, including having a transformation function compute local addresses, ghost variables, and message information, increasing the speed of calculations and communication of data dependencies. Aztec takes advantage of advanced partitioning techniques and uses efficient dense matrix algorithms when solving block sparse matrices.

In applications, Aztec performs a critical task requiring a significant portion of the total simulation time. Aztec has facilitated the work of application engineers by providing them with leading-edge iterative methods and by freeing them from cumbersome programming tasks associated with parallel iterative methods. In addition to internal use, including DOE's Accelerated Strategic Computing Initiative (ASCI) applications, there are more than 100 licensed Aztec external users worldwide.

Distributed memory software for ocean modeling

NOAA's new Modular Ocean Model, version 3 (MOM 3) improves model physics and for the first time can be run on distributed-memory supercomputers. Tests using message passing have been done on the SGI/Cray T3E-900 and T90 supercomputers at NOAA's Geophysical Fluid Dynamics Laboratory (GFDL). The baroclinic and tracer portions of the model achieve near linear scaling. Continuing areas of research include better scaling of the barotropic equation and developing parallel techniques for outputting data efficiently during model execution. The latter issue, being addressed jointly with the NERSC, will also lead to high resolution modeling of the southern ocean.

New fill-reducing ordering algorithms

DARPA-funded researchers at ORNL are working with Xerox Palo Alto Research Center (PARC) and the University of Tennessee to develop a family of modular parallel sparse matrix solvers based on preconditioning iterative methods by combining direct sparse solvers with inherent iterative parallelism and scalability. This will enable defense and industry engineers to use parallel machines to solve large problems effectively, without writing complicated domain-specific parallel linear system solvers.

Scalable Unstructured Mesh Algorithms and Applications (SUMAA3d)

SUMAA3d is a suite of software tools for unstructured mesh computations on distributed-memory computers that emphasizes scalability on widely varying architectures ranging from tightly coupled massively parallel processors (MPPs) to networks of workstations, techniques to achieve accurate solutions efficiently, and usability and interoperability. Applications include modeling piezoelectric crystals, high temperature superconductors, and cracked pressure vessels. The SUMAA3d project is a collaboration among researchers at DOE's Argonne National Laboratory (ANL), the Pennsylvania State University, the University of British Columbia, and Virginia Tech.

New SUMAA3d algorithms and software include:

- ❏ Automatic mesh generation tools for unstructured meshes with mixed element types
- ❏ Mesh optimization procedures to improve mesh quality
- ❏ Adaptive mesh refinement/de-refinement using edge bisection techniques to model rapidly changing solutions and multiscale geometries accurately
- ❏ Dynamic mesh partitioning/repartitioning after mesh modification

Recent efforts have focused on interoperability with the Advanced Large-scale Integrated Computational Environment (ALICE) (page 21) for solving large scale partial differential equation (PDE) based applications. Current efforts focus on interfacing SUMAA3d with solver and visualization components and using the knowledge gained to develop design specifications for a general mesh component. Efficient interfaces between SUMAA3d and both the Portable Extensible Toolkit for Scientific computation (PETSc) (page 18) and BlockSolve95 have been completed.

Dynamic load balancing and data migration for adaptive numerical methods

Adaptive numerical methods that automatically refine or coarsen meshes and/or vary the order of accuracy of a method offer greater reliability, robustness, and computational efficiency than traditional approaches for solving partial differential equations. On parallel computing systems, however, adaptive algorithms introduce complications because they decompose and redistribute dynamic rather than static data structures. SNL is investing in hierarchical adaptive mesh refinement of unstructured grids for the application codes MPSalsa, ALEGRA, and SIERRA. Each of these projects uses quadtree/octree data structures to represent the refined meshes. That is, during refinement, elements are split into smaller elements that are stored as "children" of the initial "parent" elements in a hierarchical data structure. Quadtrees and octrees are widely used in both mesh generation and adaptive refinement on serial computers. Using local load-balancing techniques, quadtrees for 2-D adaptive, structured-mesh methods with local time stepping were successfully balanced. Issues that must be resolved for more general 3-D unstructured grids, steady-state problems, and method-of-lines time-stepping include:

- ❏ Designing efficient dynamic load-balancing methods that account for data's current location to reduce migration costs
- ❏ Developing load-balancing heuristics for hierarchical meshes
- ❏ Efficiently distributing octree data structures on parallel computers
- ❏ Incorporating accurate models and measurements of processor work loads and migration costs into load-balancing algorithms

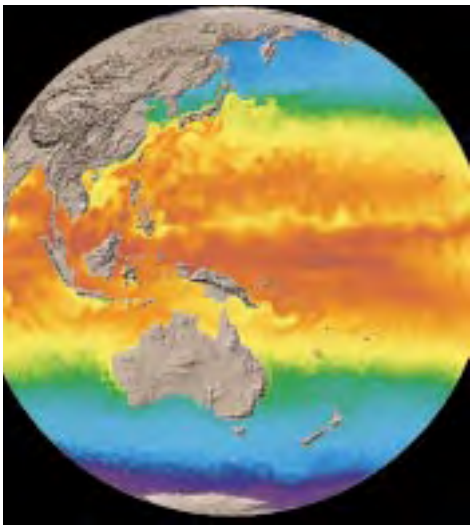

Los Alamos National Laboratory's (LANL's) POPTEX tool provides interactive visualization capabilities for the DOE Global Climate Modeling Grand Challenge. Historically, global climate data has been viewed using video technology. While useful for viewing the progress of the simulation, however, videos cannot be modified without creating a new video. LANL's goal was to provide the benefit of video visualizations (putting the results of the simulation into motion) while adding capabilities that enable dynamic, flexible, and interactive exploration of scientific data. The figure shows a POPTEX result, the temperature of the ocean with low values at the blue end of the spectrum and high values at the red end.

To address these and related issues, a dynamic load balancing tool called Zoltan is being developed for use in a range of applications. Adaptive PDE solvers are the most significant target, but the tool will also be used to support particle calculations, contact detection in crash simulations, and other dynamic and adaptive codes. In FY 2000, additional load-balancing algorithms, support for heterogeneous computers, tools for data migration, and other enhancements will be added.

PETSc, the Portable, Extensible Toolkit for Scientific computation

PETSc is a suite of interoperable software components for the large scale simulation of physical phenomena modeled by partial differential equations. PETSc includes an expanding suite of data structures and numerical kernels for use in application codes written in C, C++, and Fortran. PETSc is relatively easy for beginners to use, yet allows advanced users detailed control over solutions.

PETSc 2.0 uses an MPI for all message-passing communication, enabling portability from networks of workstations through large scale parallel processors. The PETSc programming model is also the most appropriate for non-uniform memory access (NUMA) shared-memory machines, which require the same attention to memory hierarchies as distributed-memory machines.

ANL-led multisite computational science projects built around the PETSc software include research on multi-model multi-domain computational methods in aerodynamics and acoustics (an NSF Multidisciplinary Challenges Program) conducted jointly with Old Dominion University, Courant Institute, University of Colorado at Boulder, University of Notre Dame, and Boeing Computer Services.

Scalable tools for weather modeling

NOAA's Forecast Systems Laboratory (FSL) has published advanced components of its Scalable Modeling System (SMS) of tools for developing and porting geophysical fluid dynamics models on parallel computing architectures. SMS is available as experimental public domain software.

Parallel input/output (I/O) for ocean and weather data

Improving single-processor performance is more difficult for the full-featured, 3-D fluid dynamics applications that are typical of the simulation and prediction models of climate and weather research. Better compilers and tools and more on-site technical expertise are needed to address this system limitation for NOAA's data-intensive applications — a challenge as important as good scalability in achieving full potential in these systems. Scientists at NOAA's GFDL are col-

laborating with DOE scientists at the Lawrence Berkeley National Laboratory (LBNL) to develop a parallel I/O implementation of netCDF, a common data format widely used in the oceanographic and meteorological research communities.

Information based computing

A growing need within scientific disciplines is support for information discovery and data handling. Supercomputer data management technologies that automate discovery and retrieval are being deployed in the NPACI Data Intensive Computing Environment. The infrastructure is composed of archival storage systems (HPSS), data movement systems (SDSC Storage Resource Broker), meta-data catalogs (SDSC MCAT), digital libraries for data discovery services, and information mediators for data presentation. The resulting infrastructure supports the publication of scientific data sets, making it feasible to integrate scientific data collections with traditional text and image digital libraries.

APOALA

The Environmental Protection Agency (EPA) -supported APOALA project at Penn State is developing the ability to integrate a temporal geographic information/visualization environment in order to analyze complex, large scale environmental processes. The prototype will integrate a new database model for space-time data, a visual space-time query-building language, parallelism for efficient retrieval and manipulation of data, visualization that facilitates traditional and exploratory analysis of time sequences of 3-D data, and data exploration and knowledge discovery capabilities that are closely coupled with space-time visualization methods for Earth data.

The Globus project

Supported by DARPA, DOE, and NSF, the Globus project is developing fundamental technologies needed to build computational grids and execution environments that enable applications to integrate tens or hundreds of geographically-distributed computational and information resources, instruments, and displays. Globus has evolved out of the I-Soft software environment developed for the I-WAY networking experiment demonstrated at SC95 in San Jose, California, and is currently running on machines on every continent in the world except Antarctica.

The Globus team, from the Information Sciences Institute at the University of Southern California (USC) and ANL, won the High Performance Computing (HPC) Challenge Award at SC98 for their demonstration of wide area applications. The team showcased the Globus metacomputing toolkit and the associated Globus Ubiquitous Supercomputing Testbed Organization (GUSTO) testbed, the first large scale realization of a high performance distributed computing infrastructure, by performing three unique computations:

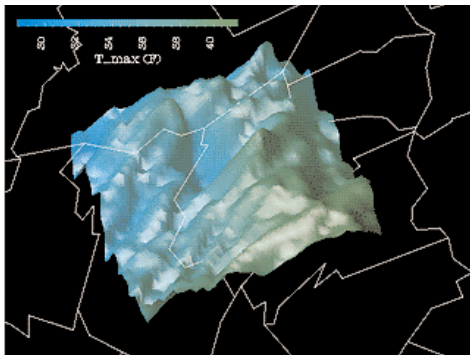

This 3-D shaded relief representation of a portion of Pennsylvania uses color to show maximum daily temperature. Displaying multiple data sets at once, and interactively changing the display, helps users quickly and intuitively explore their data, either to formulate or confirm hypotheses.

- Collaborative online analysis of data from a microtomographic beamline at ANL's Advanced Photon Source (APS)
- Record-setting distributed interactive simulation using multiple supercomputers
- High-throughput computing for crystallographic phase problems

ACTS Toolkit and Globus: Controlling remote instruments

The coupling of a scientific instrument with remote supercomputers and visualization devices can transform instrument capabilities by allowing quasi-realtime (taking minutes rather than hours) and remote analysis, imaging, and control. DOE's Advanced Computational Testing and Simulation (ACTS) Toolkit combines with the Globus toolkit to achieve such a coupling. Brilliant Xrays, now available at facilities such as ANL's APS, make it feasible to record high resolution 3-D tomographic data on time scales of less than a second per image. This technology is used to image and simulate biological and archaeological data, providing high performance parallel implementations of reconstruction algorithms for microtomographic datasets by means of filtered back projection techniques. Globus software moves data efficiently among detector, secondary storage, supercomputer, and remote users. An interactive analysis system operates on the ImmersaDesk and Cave Automatic Virtual Environment (CAVE) virtual reality systems to integrate data as it is acquired from the APS and reconstruction system, allowing users to check intermediate results. Together, these components reduce the time required for reconstruction and analysis from hours to minutes.

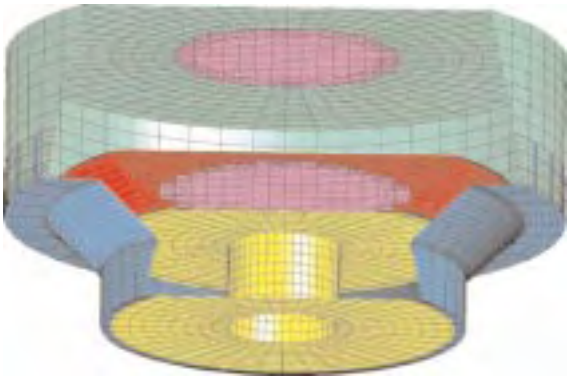

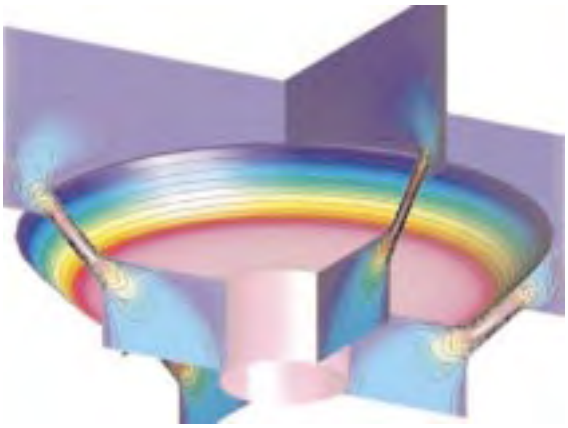

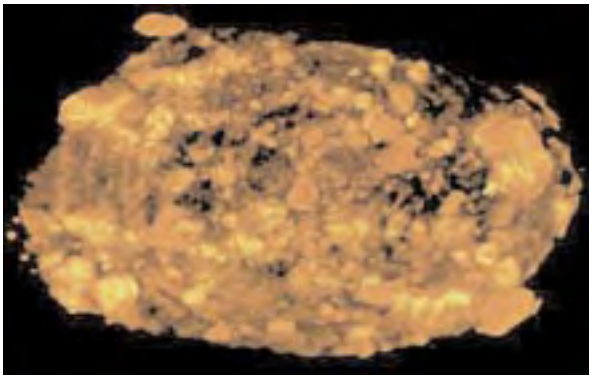

Volume rendered view (lower left) of the 3-D structure of a meteorite determined by computed microtomography. This image was produced as part of the DOE Grand Challenge project "Supercomputer Solution of Massive Crystallographic and Microtomographic Structural Problems," which is using components of the Globus toolkit, including the Nexus communication library, to enable quasi-realtime reconstruction of microtomography data from the Advanced Photon Source. Examples of imaging diesel valves using DOE's ACTS toolkit appear above. The image at the right shows pressure and velocity around a moving valve in a diesel engine. The flow here was found as part of a computational fluid dynamics (CFD) effort to simulate the flow within the complex 3-D geometry of a diesel engine. The computation was carried out using the Overture Framework and the PADRE library for parallel data distribution.

Advanced Large-scale Integrated Computational Environment (ALICE)

The goal of the ALICE project at ANL is to eliminate barriers in using independently developed software to construct high performance numerical applications, laying the groundwork for widespread exploitation of teraflop-scale computational resources and the resulting new scientific insights. ALICE addresses the data management and interoperability problems when combining multiple software packages for mesh manipulation, numerical solution of partial differential equations, optimization, sensitivity analysis, and visualization.

ALICE research focuses on low-overhead integration of extensible software for scientific problem solving and component-based toolkits that encapsulate expert knowledge in numerical algorithms and parallel computing. ALICE development will ensure the relevance and practicality of large scale scientific applications design, supporting both new and legacy applications, and enabling scientists to reuse legacy kernels and program more comfortably in, for example, traditional Fortran.

ALICE Memory "Snooper" (AMS)

The ALICE Memory "Snooper" (AMS) is an application programming interface (API) that aids in writing computational steering, monitoring, and debugging tools. The AMS is a client/server, multithreaded API that also supports parallel applications using MPI.

Distributed problem-solving system

At ORNL, NetSolve is transforming disparate computers and software libraries into a unified, easy-to-access computational service. By aggregating the hardware and software resources of any number of computers loosely connected across a network, NetSolve can tap their combined power through a familiar client interface that hides the underlying complexity of the system, making supercomputing transparently available to a broad range of users on ordinary network platforms.

Collaborative User Migration, User Library for Visualization and Steering (CUMULVS)

Computational steering could revolutionize computational experiments by allowing scientists to interactively explore (steer) a simulation in time and/or space. Instead of the typical simulation mode — manually setting input parameters, computing results, storing data to disk, visualizing the results in a separate visualization package, then starting again from the beginning — computational steering allows the scientist or engineer to interactively manipulate algorithmic and model parameters beyond their initial values. The CUMULVS collaborative computational steering tool provides remote visualization, steering, and heterogeneous checkpoint/restart to large distributed simulations. CUMULVS is being used at NSF's NPACI and at DOE and DoD sites across the country.

ALICE Differencing Engine (ADE) visualization toolkit

ANL researchers are developing the ADE visualization toolkit so that a scientist or engineer can visualize pairwise differences in up to ten data sets containing vector and scalar data in an immersive, interactive, virtual reality environment. Standard techniques such as vector field glyphs, animated streamlines and flow fields, and scalar field color remappings provide insight into individual data sets. Multiple data sets can be displayed simultaneously along an interactively defined cutting plane. The user may cull away portions of the data set to investigate regions of maximum difference, magnify the differences, and remap the coloring to correspond to alternative scalar fields.

The toolkit has been used to design a new aluminum furnace in a joint project between DOE and Air Products and Chemicals, Inc. Air fuel, an air/oxygen mix, and pure oxygen fuel data sets were displayed and analyzed in order to understand the effect of different fuel types on furnace efficiency.

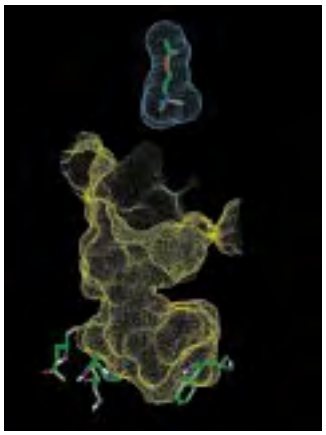

NPACI researchers are unlocking the mystery surrounding one of the fastest, most efficient enzymes in the human body. Some of the largest acetylcholinesterase (AChE) enzyme simulations to date have been conducted by researchers at the University of Houston and UCSD. An open "side door" in one of the AChE subunits was recorded at 152 picoseconds of one such molecular dynamics simulation. A frame from this animation shows the "breathing" motions of the gorge or channel that leads from the region outside the AChE enzyme to the active site. These fluctuations in the width of the channel allow the substrate acetylcholine (ACh) to move from the outside into the active site. They also contribute to the selectivity of the enzyme by slowing the entrance of substrates that are larger than ACh. Full animation can be seen at http://chemcca10.ucsd.edu/ache_animated.html

By incorporating the communication mechanisms of the ALICE Memory Snooper, the ADE toolkit will soon be able to dynamically retrieve data from multiple timesteps of an ongoing simulation, allowing an application scientist to monitor the progress of the application's solution and investigate the differences between timesteps.

Computer simulation demonstrates "breathing" enzyme action

Using powerful supercomputers, NSF- and NIH-supported researchers have unlocked the mystery surrounding one of the fastest, most efficient enzymes in the human body. Acetylcholinesterase (AChE) works to instantly stop transmissions from one nerve cell to the next nerve or muscle cell by catalyzing the chemical reaction that breaks up the neurotransmitter acetylcholine (ACh), serving as the off-switch for the transmission. The active site where the reaction takes place was found deep inside a groove on AChE — a groove too narrow to admit ACh quickly. Using large scale computer simulations run at the NSF-supported NPACI, researchers have shown that AChE does its job by "breathing." That is, this flexing motion causes AChE to inhale and exhale ACh molecules almost as fast as if the groove were always open.

The supercomputer simulation showed that this motion keeps larger molecules out of the space intended for ACh. This finding shows that enzymes can use movements to select particular substrates in crowded environments like the inside of a living organism.

Parallel volume-rendering system for scientific visualization (ParVox)

The demand for parallel supercomputing in interactive scientific visualization is increasing as the ability of the machines to produce large output datasets has dramatically increased. NASA's distributed volume rendering system, ParVox, will provide a solution for distributive visualization of large time-varying datasets on a scientist's desktop even when using low-speed networks and low-end workstations.

NASA is currently working on the functional pipelining of ParVox. It will be separated into three modules, an input module that reads data from files or from the network, a rendering module that performs the rendering and classification functions, and a compression/output module that compresses images and outputs them. The input module will support data-input on demand, allowing out-of-core rendering for very large datasets. The major milestone for FY 1999 is to add rendering functions for unstructured grid 4-D datasets.

ParAgent

Developed at Iowa State University with EPA support, ParAgent is an interactive tool for automatic parallelization of specific classes of programs that follows a knowledge-based approach, using program characteristics to guide the automatic

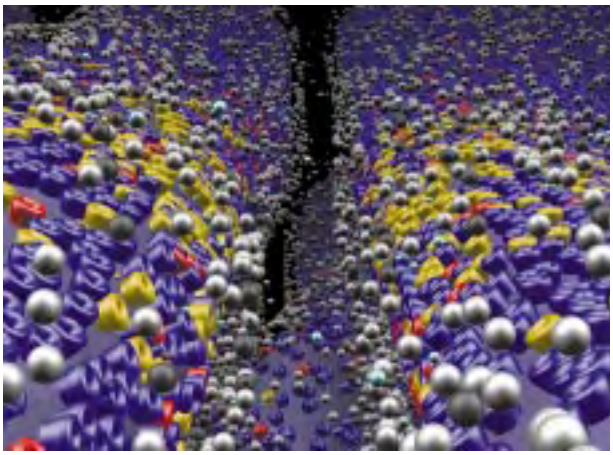

Current generation where nerve meets muscle. An MCell simulation of ionic current generation at a neuromuscular junction in a rat diaphragm muscle shows the neurotransmitter ACh (cyan spheres) 300 microseconds after release of 6,200 ACh molecules. The highly convoluted membrane of the junction is covered with acetylcholine receptors (the cup-shaped objects on the membrane). After release, ACh molecules diffuse away from their release point and bind to receptors and AChE. The image shows bound AChE (black spheres), unbound AChE (gray spheres), and receptors—unbound (purple), singly bound (red), doubly bound closed (green), and doubly bound open (yellow). Doubly bound open receptors conduct an ionic current that initiates a cascade of events leading to muscle fiber contraction. AChE breaks down ACh and prevents prolongation of the current. MCell was developed by NPACI-supported researchers at the Salk Institute for Biological Studies and at Cornell University. In the NPACI Neuroscience thrust area, they are collaborating with NetSolve researchers and developers at the University of Tennessee and UCSD.

parallelization process. ParAgent includes tools for Web-based performance monitoring, performance analysis, and visualization. The current version is designed for programs based on the Finite Difference Method (FDM), and has been implemented for parallel computing on 64 node cluster of 200 MHz Pentium II workstations at NASA's Ames Laboratory.

VisAnalysis Systems Technology/Space-Time Toolkit (VAST/STT)

The EPA-supported VAST team at the University of Alabama in Huntsville (UAH) is developing Java-based software for visual integration of multi-source data. This VAST/STT accepts data in native space-time domains and allows "on-the-fly" data transformations into a user-selected interactive display domain that can integrate 3-D model grids, 2-D maps, Geographic Information System (GIS) vector data, vertical and horizontal profiles, dynamic particles, episodic events, and satellite and aircraft sensor data. In particular, the STT can be used to explore relationships among dynamic spatial data with temporal resolutions from microseconds to centuries.

The STT is being tested by interrelating two months of 1995 data — obtained from aircraft and satellite-based sensors, meteorological and environmental model output, vertical wind profilers, point source measurements, and Doppler radar volumes — for the Nashville Southern Oxidant Study (SOS).

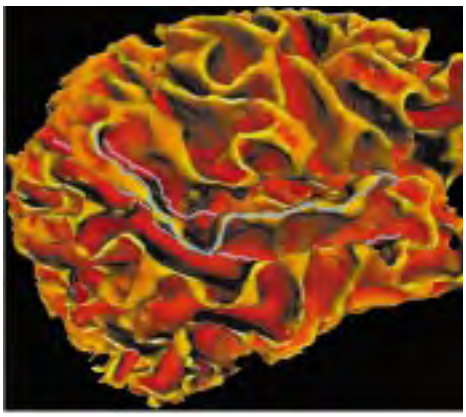

NPACI researchers have developed software to extract the grey-white matter cortical surface from high resolution, 3-D human brain data obtained from UCLA. Dynamic programming methods were used to define the maximal contour of the grey-white surface from a quarter-resolution version of the UCLA data. Processing of the full-resolution data will require NPACI infrastructure.

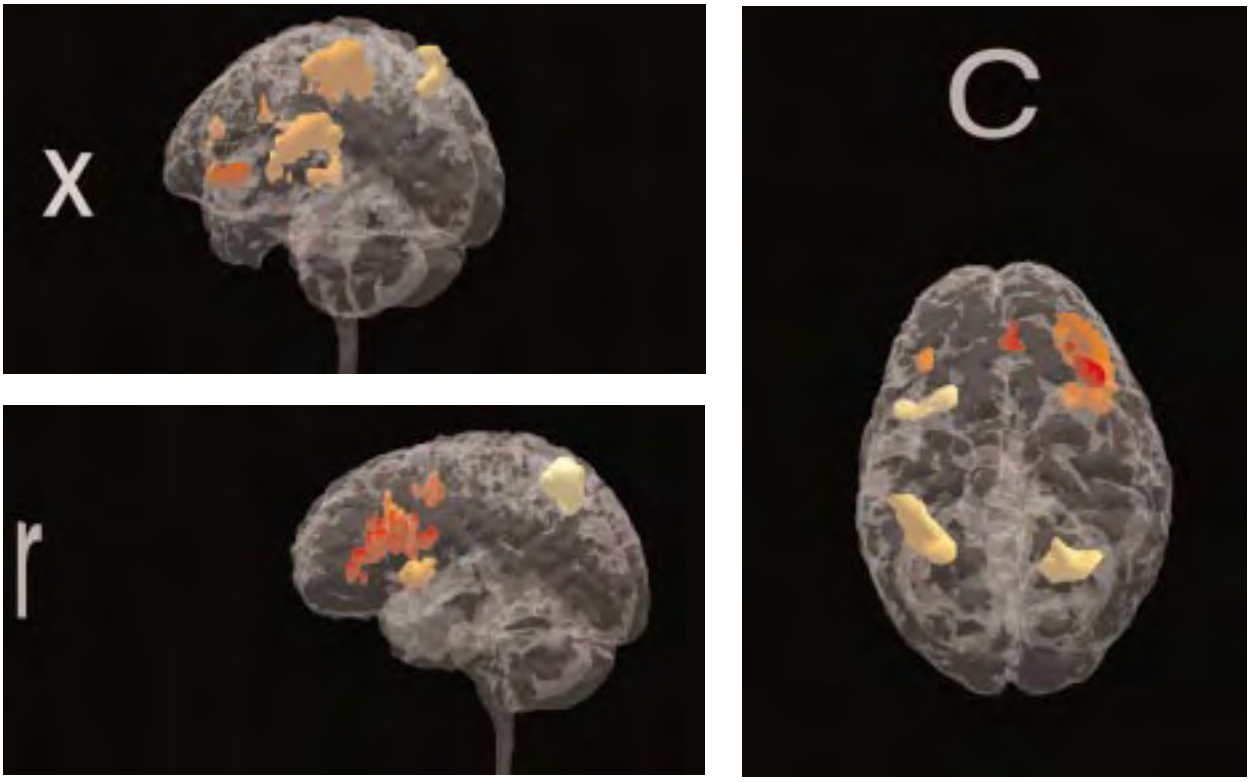

By combining the high temporal resolution of electroencephalography (EEG) and the high spatial resolution of functional magnetic resonance imaging (fMRI), psychologists are able to study the series of events that underlie the formation of memories in the brain. Here we see the cortical regions that are activated as a subject remembers a sequence of letters presented at 3-second intervals. Researchers at the Pittsburgh Supercomputing Center are focusing on improving the rapid, online analysis and visualization capability of fMRI. As microprocessor performance increases, realtime fMRI technology that was pioneered on supercomputers may soon be incorporated into dedicated hardware bundled with MRI scanners. This will allow immediate assessment of scan quality and facilitate the use of fMRI in psychiatry, drug evaluation, and neurosurgical planning.

Advanced data visualization and exploration

Today's limited capabilities for extracting and assimilating the information buried in data means that much scientific data is either never examined or examined only superficially. SNL research focuses on developing advanced capabilities for visualizing and exploring data in order to improve productivity and/or create new technologies in particular fields.

High End Applications

Code that keeps blood flowing (NekTar)

New NekTar software for modeling blood and other complex fluid flows may lead to changes in accepted surgical practices that will extend the life expectancy of those suffering from arterial diseases like atherosclerosis. These codes can produce animated simulations of flows that once could only be captured accurately in "still" formats like magnetic resonance imaging (MRI). Research at Brown University and simulations at NCSA will allow scientists to more accurately model many types of fluid flows and to do computational steering — that is, to zero in on specific areas of a calculation while the computation is running. This is especially useful in studying the complex human circulatory system which, with its miles of arteries, veins, and vessels, not only carries oxygen but also transports wastes and aids in heating and cooling the body.

Making and testing neuron models

Computational models describe the behavior of neurons and the communications among neurons in networks and are essential for investigating the functions of brain regions. Models assist in studying how interactions between regions underlie the functions of the entire nervous system, including behavior and thought. Models can be used to perform virtual experiments that are too difficult or impossible to conduct using biological tissue or living subjects. Computationally intensive simulations can incorporate greater numbers of neurons to model increasingly complex and realistic properties, both electrical and chemical. NSF-funded NPACI neuroscience researchers using high performance computing and infrastructure resources are conducting this work.

Road maps for understanding the human brain

To understand how the brain allows humans to perform multiple tasks (such as simultaneously observing traffic, listening to the radio, and stopping at a traffic signal) and to investigate changes in the brain related to pathological conditions and disease, NSF-funded NPACI neuroscientists are creating tools to produce road maps of the structures and connections in the brain. The researchers are producing software to analyze data collected from MRI scans, digitized images of cryosectioned brains — brains frozen and cut into thin slices — and high powered microscopes. The software can re-create a 3-D model of the brain and generate a road map for comparing a particular brain with others.

Catalytic antibodies

When viruses, bacteria, or other harmful invaders enter the bloodstream, the immune system fends them off by generating antibodies — proteins shaped to latch onto the intruder (analogous to a lock fitting a key) so that its ability to cause harm is immobilized. What if it were possible to harness this remarkable ability to manufacture proteins matched to the 3-D features of these microscopic intruders? For the pharmaceutical industry, these "catalytic antibodies" offer promise for rational drug design — the creation of molecules with 3-D features sculpted to interact with other molecules.

Taking the pulse of a red giant

NSF-funded University of Minnesota and NCSA astrophysicists have generated a 3-D simulation of a red giant with such detail that they could watch it pulsate. In a region that encompassed nearly the entire star and was the size of the orbital radius of Jupiter around the Sun, super hot gases were in lava-like turmoil. This global convective pattern flowed asymmetrically, with gas flowing outward from

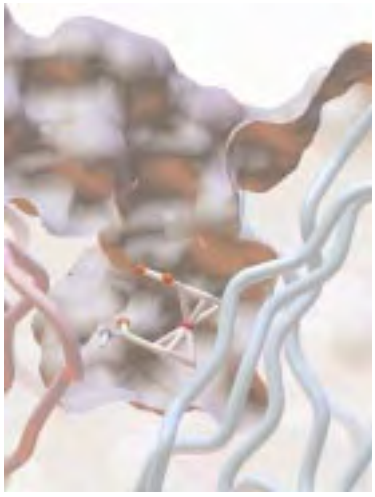

Close-up view of antibody 13G5, with the cyclopentadienyl ferrocene molecule deeply buried in the antibody's binding-site. Scripps Research Institute scientists determined the structure of this antibody, opening the way for computational studies of its catalytic effect by UCLA scientists using the SGI/Cray Origin2000 at NCSA.

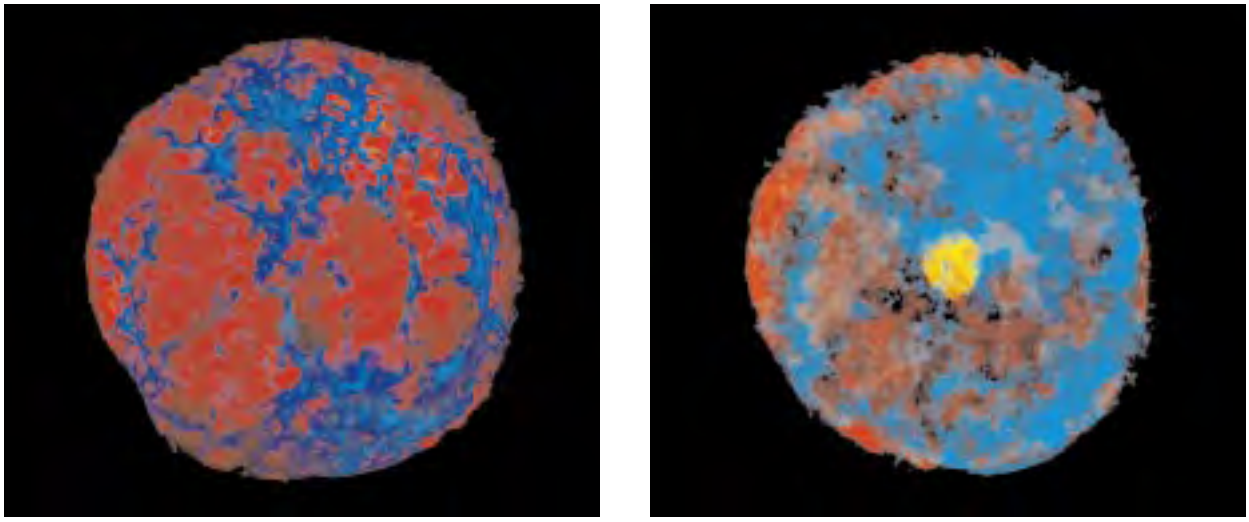

Model of a red giant star. At left, the relatively diffuse envelope has contracted to about its minimum size. At right, the same stage rendered transparently so that its hot stellar core is visible. Relatively warm temperatures are red and yellow, while relatively cool temperatures are blue and aqua. Researchers from the University of Minnesota generated the largest simulation to date of highly convective red giant stars — 134 million computational cells — by adapting their numerical approach to the distributed shared-memory capabilities of the SGI Cray Origin2000 at NCSA.

the center of the generally hotter side of the star to the cool side, emitting heat along the way. Once on the cool side, the gas sank, forming a funnel that reheated upon passing the hot, stellar core. This may help explain differences in illumination within these large, pulsating stars — upon whose "standard candles" astronomers rely for mapping distances in the universe. It will also help scientists know what to expect from our own Sun that eventually will also become a red giant.

Assisted Model Building with Energy Refinement (AMBER)

The NSF- and NIH-supported AMBER computational chemistry project has developed software to study protein folding and the relative free energies of binding of two ligands to a given host (or two hosts to a given ligand), to investigate the sequence-dependent stability of proteins and nucleic acids, and to find the relative solvation-free energies of different molecules in various liquids.

Researchers recently used AMBER to make the longest molecular dynamics simulation of a molecule, a 36-residue protein called the villin headpiece subdomain. The simulation — carried out over 100 days at the Pittsburgh Supercomputing Center on an SGI/Cray T3D and a T3E at SGI/Cray Research — followed the molecule for a full microsecond, which was possible due to improvement of the code's efficiency for parallel computing.

NASA Grand Challenges

SGI is collaborating with NASA and Grand Challenge investigator teams to implement highly scalable applications on the Cray T3E series of scalable systems, using advanced systems and operating software to implement a wide range of models based on standard languages and libraries to achieve new levels of performance. By achieving high performance levels, teams will be able to perform more accurate modeling.

Currently, the testbed at NASA's Goddard Space Flight Center (GSFC) includes 1,024 processors supporting the Grand Challenge teams as well as NASA-directed programs. Applications are implemented using SGI-developed scalable

operating systems, compilers, and development tools to achieve scaling and performance. The challenge is to achieve effective distribution of the task and associated interprocessor communication. Long latencies for communication among processors impede the ability of science models to scale in performance. NASA's Earth and Space Sciences (ESS) Round-2 teams, in concert with SGI analysts, are implementing Earth and space science models on the Cray T3E. Most 50 gigaflops milestone efforts have been completed and several 100 gigaflops milestones have been conducted. Work on the remaining 50 and 100 gigaflops milestones will continue through FY 1999.

National Center for Environmental Prediction (NCEP)

Numerical weather models guide weather forecasters, incorporating observations of temperature, wind, precipitation, pressure, and other meteorological information from sources on the ground, in the air, and in space. These observations are processed by powerful computers that generate predictions for forecasters, allowing them to anticipate weather conditions from hours to weeks in advance — or, as in the case of the recent El Niño, many months in advance.

NOAA's NCEP has worked with the Naval Research Laboratory (NRL) and NASA/GSFC to develop techniques to implement weather prediction models on parallel supercomputers. NCEP provided benchmarks for its most recent procurement that were independent of the architecture of the target machine, leading to the first-ever selection of a scalable architecture computer to run the U. S. weather forecasting suite of models. At the time of the announcement, Department of Commerce Secretary William M. Daley said, "Accurate weather forecasting is one of the great scientific achievements of the 20th century. We have reached unprecedented levels of accuracy in recent years as a result of much research, modernization and improvements such as supercomputers, radars, satellites and other technologies....We eagerly await the next generation of computational power because we know we can do even better in the future."

Predicting weather that affects aircraft

With support from the Federal Aviation Agency (FAA), NOAA's FSL has improved the Rapid Update Cycle (RUC-2) model to help meteorologists better predict weather that affects aircraft including icing conditions, winds, clouds, and clear air turbulence. While the initial version ran only once every three hours, RUC-2 can run hourly, has higher resolution, captures greater detail, and produces more accurate information.

In one ISE demonstration, exhibit staff and researchers at three NASA field centers will become virtual astronauts in a simulated spacewalk rehearsal of a future space station assembly flight. Subjects will wear virtual-reality headgear displaying mock-ups of space station parts derived from computer-aided design files. The "astronauts" will "maneuver" around the components while computers at each site exchange positional information over the Internet, continuously updating the scene from each user's point of view.

Hundreds of thousands of asteroids cross Earth's orbit around the Sun. NASA-supported researchers modeled Castalia as a peanut-shaped solid rock 1.6 kilometers long and 0.8 kilometers wide, shown here in a cutaway view being hit by a house-sized rock traveling at 5 kilometers per second. The scattering white dots are fragments from the smaller rock. Fractures appear throughout the asteroid, with the greatest damage shown in red. Lasting merely a second, the collision approximates the force of an early atomic bomb. Nuclear weapons have been proposed for breaking up, or at least diverting, asteroids headed towards Earth. These simulations show that such an impact will fracture a solid asteroid, but, later, gravity will reassemble the pieces.

Hurricane prediction

The 1998 season included several major hurricanes that caused considerable property damage to the southeastern United States and an enormous death toll in Central America. The GFDL Hurricane Prediction System was used extensively by NOAA/NCEP as part of its modeling suite to predict the track of these storms.

Realtime weather forecasting

Advances in time-urgent weather forecasting were demonstrated at the January American Meteorological Society (AMS) meeting in Dallas, TX. The input data were the detailed data stream from NOAA's new Doppler radar system. Each day, realtime, high resolution numerical weather forecasts were generated remotely using an entire 128-processor SGI/Cray Origin2000 supercomputer running almost non-stop at the NSF-funded NCSA.

NASA Intelligent Synthesis Environment demonstration

NASA's proposed Intelligent Synthesis Environment (ISE) program focuses on how future missions will be developed from initial concept through operations and conclusion. ISE will conduct research and develop tools and processes to allow scientists, technologists, and engineers with diverse expertise and interests to function as a team. By functioning as a networked collaboration among geographically dispersed and professionally diverse personnel involved in defining, designing, executing, and operating NASA's missions, development time will be shortened and life-cycle costs reduced. When fully developed, collaboration will take place in a full-sensory, immersive, virtual environment in which humans and computers interact through 3-D sight, sound, and touch in a computationally rich mission life-cycle simulation. Methods for trading cost, risk, and performance over total mission life cycles will be developed, enabling NASA to fully understand missions prior to committing to development.

Efficient electromigration simulations using approximate Green's functions

In some applications of the boundary integral method, most notably when there is symmetry present, Green's function can lessen the computational effort, since it exactly satisfies the imposed boundary conditions on part of the surface. This piece of the boundary can then be eliminated from the calculations, reducing the time required to construct and solve the matrix system. While this is a powerful method, such exact Green's functions, which must also be simple enough to work with, are rare.

ORNL researchers have demonstrated that the benefits of this technique can be extended by means of "approximate Green's functions." They have employed this new algorithm in modeling the electromigration problem in materials science, in which a microelectronics wire has been deformed by a void. Here, an approximate Green's function can be constructed to reduce the computational domain to

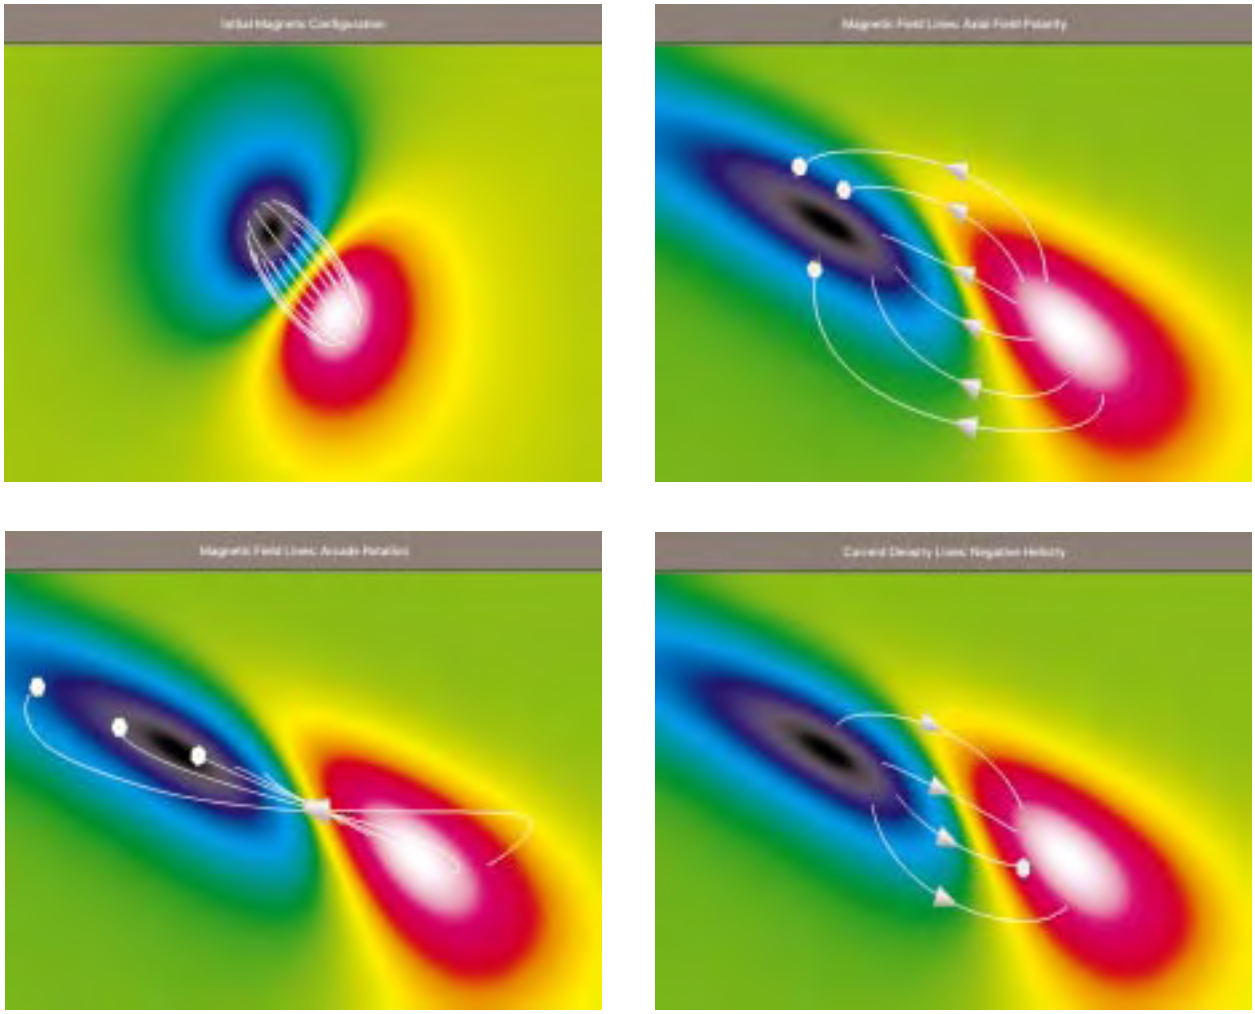

Sheared Active Region Fields: Pictured above is the evolution of an initially potential (current-free) dipole magnetic field typical of northern-hemisphere active regions on the Sun. The field was subjected to differential rotation and supergranular diffusion, modeling the effects of the Sun's surface convection, which stretch and spread the flux threading the solar surface. The objective is to demonstrate the development of properties of the magnetic field that show distinct preferences by hemisphere, independent of the sunspot cycle. This calculation used FCTMHD3-D, NASA's parallel, finite-volume 3-D code.

Upper left: Seen from overhead, white lines of magnetic force connect the regions of positive (red) and negative (blue) polarity in the initial state. There is no twist in the field. The force lines cross the magnetic neutral line at the surface — the yellow-green boundary between the red and blue polarities — at a 90 degree angle.

Upper right: After the shear and diffusion at the base have been applied for a time, the lines of force cross the magnetic neutral line at the surface at an acute angle. A component of the field now points to the right along the neutral line, when observed from the positive polarity (red) side.

Lower left: A sequence of lines of force drawn at increasing heights above the surface shows an apparent counterclockwise rotation of the loops as seen from above.

Lower right: Several lines of electric current are drawn here, showing that the current flows antiparallel to the lines of magnetic force - that is, the current points from the negative polarity (blue) side to the positive polarity (red) side. This imparts a left-handed twist, or negative helicity, to the magnetic field.

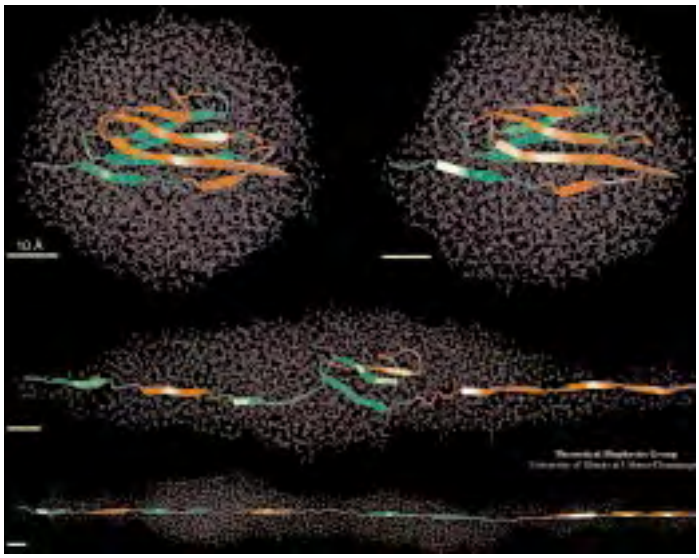

Molecular dynamics: Unfolding of Ig and Fn-III protein domains by steered molecular dynamics simulation. The architecture of immunoglobulin-like (Ig) and fibronectin type III-like (Fn-III) domains constitutes possibly the most prevalent structural motif of proteins that have a role in cell-cell signaling, cell-cell aggregation, and embryogenesis, as well as in mechanically coordinating and strengthening cells and tissues. The proteins are implicated in the etiology of many diseases, ranging from heart insufficiency to cancer. Recently, atomic force microscopists have investigated the mechanical properties of connected Ig and Fn-III domains, by simulating experiments to explain the specific mechanical properties of these proteins using steered molecular dynamics, a new simulation technology developed at the NIH Resource for Macromolecular Modeling and Bioinformatics (Beckman Institute, University of Illinois), that permits researchers to mechanically manipulate models of proteins. The technology combines its high end graphics tool with its parallel simulation software.

simply the void surface. Despite using an approximation to the true Green's function, comparisons with finite element results show that there is no loss of accuracy in the analysis. Moreover, by eliminating the corner formed by the intersection of the void and wire boundaries, the results in this area are much more accurate.

Electroforming simulation is another application where this technique will prove useful. A typical geometry consists of a tank and electrodes. With this new method, all tank surfaces can be eliminated from the calculation, saving time to discretize the model and in the analysis.

Biology WorkBench

The NSF-supported Biology WorkBench at NCSA is a revolutionary Web-based tool for biologists. The WorkBench allows biologists to search many popular protein and nucleic acid sequence databases. Database searching is integrated with access to a wide variety of analysis and modeling tools, all within a point and click interface that eliminates file format compatibility problems. This product was first released in 1996 and a new version 3.0 was released recently with many enhanced functionalities.

Neuroimaging Analysis Center (NAC)

NAC, supported by NIH's National Center for Research Resources (NCRR), develops image processing and analysis techniques for basic and clinical neurosciences. Core NAC research includes algorithms and techniques for segmenting brain structures, registration methods for associating image data to patient anatomy or one set of image data to another, visualization tools for displaying anatomical structures and quantitative information, software and hardware infrastructures for high performance computing, and digital anatomy atlases for interactive and algorithmic computational tools. NAC emphasizes the dissemination of concepts and techniques.

NAC technologies are employed in research on Alzheimer's disease and the aging brain, morphometric measures in schizophrenia and schizotypal disorder, quantitative analysis of multiple sclerosis, and interactive image-based planning and guidance in neurosurgery.

Laboratory of Neuro Imaging Resource (LONIR)

There is a rapidly growing need for brain models comprehensive enough to represent brain structure and function as they change across time, in large populations, in different disease states, across imaging modalities, across age and gender, and across species. The NIH-supported LONIR develops novel strategies to investigate brain structure and function in full multidimensional complexity. LONIR provides international networks of collaborators with a diverse array of tools to create, analyze, visualize, and interact with models of the brain. A major focus of these collaborations is to develop 4-D brain models that track and analyze complex patterns of dynamically changing brain structure in development and disease, expanding investigations of brain structure-function relationships to four dimensions.

Modeling research focuses on new strategies for surface and volume parameterization that provide an advanced analysis of surface and volumetric brain models, tracking their change across time. Additional research focuses on anatomical fundamentals, analyzing anatomical and cytoarchitectural attributes across multiple scales and across time; and visualization and animation, for the dissemination of brain models that visualize complex variations in brain structure and function across time. Ongoing national and international collaborations are analyzing normal and aberrant growth processes, brain development, tumor growth, Alzheimer's disease, and related degenerative disease processes, schizophrenia, and brain structure in normal and diseased twins. The range and sophistication of these strategies matches the broad scope of studies that focus on mapping and modeling the dynamically changing brain.

Virtual Cell

The NIH-supported National Resource for Cell Analysis and Modeling (NRCAM) is developing methods to model cell physiological processes in the context of the actual 3-D structure of individual cells. Approaches in computational cell biology are coupled with high resolution light microscopy to facilitate the interplay between experimental manipulation and computational simulation of specific cellular processes.

NRCAM is developing the Virtual Cell, a general computational framework for modeling cell biological processes. This new technology associates biochemical and electrophysiological data describing individual reactions with experimental microscopic image data describing their subcellular locations. Individual processes are integrated within a physical and computational infrastructure that will accommodate any molecular mechanism. Current development of the Virtual Cell

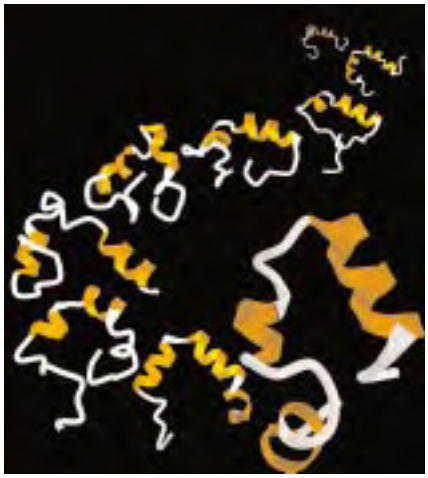

The protein-folding problem — how a protein's amino-acid sequence relates to its folded shape — is the most important unsolved problem facing computational biochemistry. NSF-supported researchers from the University of California, San Francisco, using PSC's T3D and resources at Cray Research, performed a microsecond simulation of the folding of a small protein, the villin headpiece. This image shows a snapshot along the computed folding pathway. The computation is two orders of magnitude longer than the longest previous calculation of its kind and has identified the initial collapse of the protein to a "molten globular state" from which it searches for the ultimate fold.

will expand the generalized mathematical descriptions to include additional cell biological mechanisms, increasing accessibility to biologists studying different biological processes, and integrating the interface with a database of images and reaction mechanisms. Current applications of Virtual Cell include studies of calcium dynamics in neuroblastoma cells and cardiomyocytes and studies of intracellular RNA trafficking in oligodendrocytes. Additional collaborative research projects include modeling diffusional processes in mitochondria, examining the role of $[Ca^{2+}]_i$ in triggering Ca^{2+} sparks in the heart, and analyzing structural changes in the endoplasmic reticulum during egg activation.

Optical information processing

One of the major barriers to commercial application of optical technology to information processing is the high cost of system development and manufacture. This problem has been solved in other industries through the use of computer-aided design (CAD) and the integration of system design with manufacturing. The development of better system level metrology is needed to allow more computer-based methods to be used in the design and manufacture of optical information processing systems.

NIST is designing an optical pattern recognition system that will test an input image (at video rates) against a large reference set — 1,000 human faces, for example — with images of 640 by 480 pixels or larger. NIST has constructed realtime video system testbed versions of the two subsystems needed in a high speed optical pattern recognition application. NIST will evaluate the subsystems and components for future design work.

Micromagnetic modeling

NIST is developing computational tools for accurate and efficient micromagnetic calculations. Such calculations are essential to achieving higher densities and faster read-write times in magnetic disk drives. Given an earlier NIST study that showed how unreliable software for micromagnetic calculations is, NIST is now developing a reference code that will be thoroughly tested, compared to results from experimental measurements, and made publicly available.

Modeling the high speed machining process

A combined experimental and analytical program to study the dynamics of high speed machining is currently underway at NIST. Its goal is to provide accurate measurements and obtain a better scientific understanding of fundamental metal-cutting processes while evaluating the predictive capabilities of finite-element software for high-speed machining operations. NIST has derived a mathematical model of the basic process of orthogonal machining, and finite-element software for its simulation has been developed and implemented on a NIST supercomputer.

Developing pollutant control strategies

EPA-funded researchers at North Carolina State University have been exploring how computer-based decision support systems (DSSs) can improve strategy design. In prototype DSSs focused on air quality management, alternative management strategies such as command-and-control and emissions trading programs are designed and tested with state-of-the-art air quality models. The prototype DSSs also have optimization components that can be used to identify cost-effective benchmark strategies, characterize the tradeoffs among design goals, and incorporate consideration of uncertainty in the design process. Techniques used for these analyses, such as genetic algorithms, are highly computationally intensive. Their use is made practical through high performance computing on a heterogeneous, distributed network of computers.

Large Scale Networking

LSN R&D provides the technological leadership in high performance network communications that will develop the networking technologies, services, and performance needed for the future growth of the Internet and network requirements of Federal government agencies. Early Federal networking R&D investments were instrumental in building the technological foundation of today's global Internet. 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 changed the way people use computer networks. This paved the way for our Nation's current leadership in the multi-billion dollar information technology industry.

LSN R&D in conjunction with academia, industry, and government is transitioning leading-edge networking technologies and capabilities to the private sector where their use is transforming the way we live and work. Key LSN research areas include advanced network components and technologies for engineering and managing large scale networks of the future. The LSN programs will:

- ❏ Increase the effectiveness of Federally-funded network technology research
- ❏ Increase the effectiveness of Federal research networks
- ❏ Enable network-intensive applications that advance Federal goals
- ❏ Facilitate interagency collaborations in LSN R&D
- ❏ Provide mechanisms for cooperation in LSN R&D among Federal agencies, Government laboratories, academia, and industry

Since FY 1998, the Next Generation Internet (NGI) initiative has been a major LSN focus. NGI builds on the LSN programs to provide the R&D and advanced networking testbeds in new technologies and applications to rapidly expand the capabilities of the Internet. This section describes the FY 1999 accomplishments and FY 2000 plans in LSN R&D and the NGI initiative.

LSN Support Teams

The LSN Working Group (LSNWG) coordinates Federal LSN R&D efforts. Four teams report to the LSNWG, helping implement different aspects of Federal R&D. They are:

Joint Engineering Team (JET)

The JET coordinates the network architecture, connectivity, exchange points, and cooperation among Federal agency networks (FedNets, page 38) and other high performance research networks. The JET provides close coordination of connectivity and services among vendors, academia, and industry to improve end-to-end user performance and avoid duplication of resources and efforts in providing high performance networking services. In addition, the JET cooperates with the

academic community's Gigabits per second points of presence (Gigapops), the Abilene Network (a consortium among Qwest, Cisco, Nortel, and Indiana University), and Internet2 (I2). Currently, the JET is helping to implement NGI testbeds to provide high performance end-user-to-end-user services. The JET also supports cooperation among the agencies and with Abilene to provide improved, lower cost services to geographically challenging areas such as Alaska and Hawaii. It provides cooperation on critical connectivity requirements such as for demonstrations at Supercomputing (SC) conferences and for the 1999 Global Information Network (GOIN) Workshop.

Networking Research Team (NRT)

The NRT coordinates agency networking research programs, shares networking research information among Federal agencies, and supports NGI Goal 1 (page 43) activities. It provides outreach to end users to promote dissemination of networking research information and to promote coordination among end users and applications developers. The NRT is active in developing agency workshops on middleware.

High Performance Networking Applications Team (HPNAT)

The HPNAT coordinates Federal R&D to maintain and extend U.S. technological leadership in high performance networking applications through research that employs advanced networking technologies, services, and performance to support leading-edge applications. Advances in these areas will lead to new and more capable network applications to support Federal agency missions, helping build the foundation for the continued evolution of a national information infrastructure.

The HPNAT provides mechanisms for cooperation in large scale networking applications development among Federal agencies, Government laboratories, academia, and industry, and organizes information dissemination activities including technology demonstrations, workshops, and seminars. The HPNAT supports NGI Goal 3 (page 43) by helping to organize NGI demonstrations at conferences such as SC98, held in November, 1998, in Orlando, Florida.

Internet Security Team (IST)

The IST facilitates testing and experimentation with emerging advanced security technologies and serves as a focal point and clearinghouse for application and engineering requirements for security systems. It provides the LSNWG with feedback and direction for NGI research in network security by serving as a forum for the exchange of security requirements/needs and available and emerging security technologies. The IST encourages development and use of Internet security testbeds by working closely with LSN agencies and the JET to help implement these testbeds and publicize testbed activities to national and international security research communities.

LSN R&D

LSN R&D programs address the mission requirements of the participating agencies and the development of enabling technologies and applications to expand the capabilities of the global-scale Internet. This section describes some of the Federal agency activities.

Advanced Networking Infrastructure and Research (ANIR) program

NSF's Advanced Networking Infrastructure and Research (ANIR) program consolidates and integrates the NSFNET program and associated research to advance fundamental network research and the networking infrastructure for the science and engineering community. ANIR emphasizes the development and deployment of high performance networking for cutting-edge research in all disciplines, testing and development of prototype networks, and fundamental

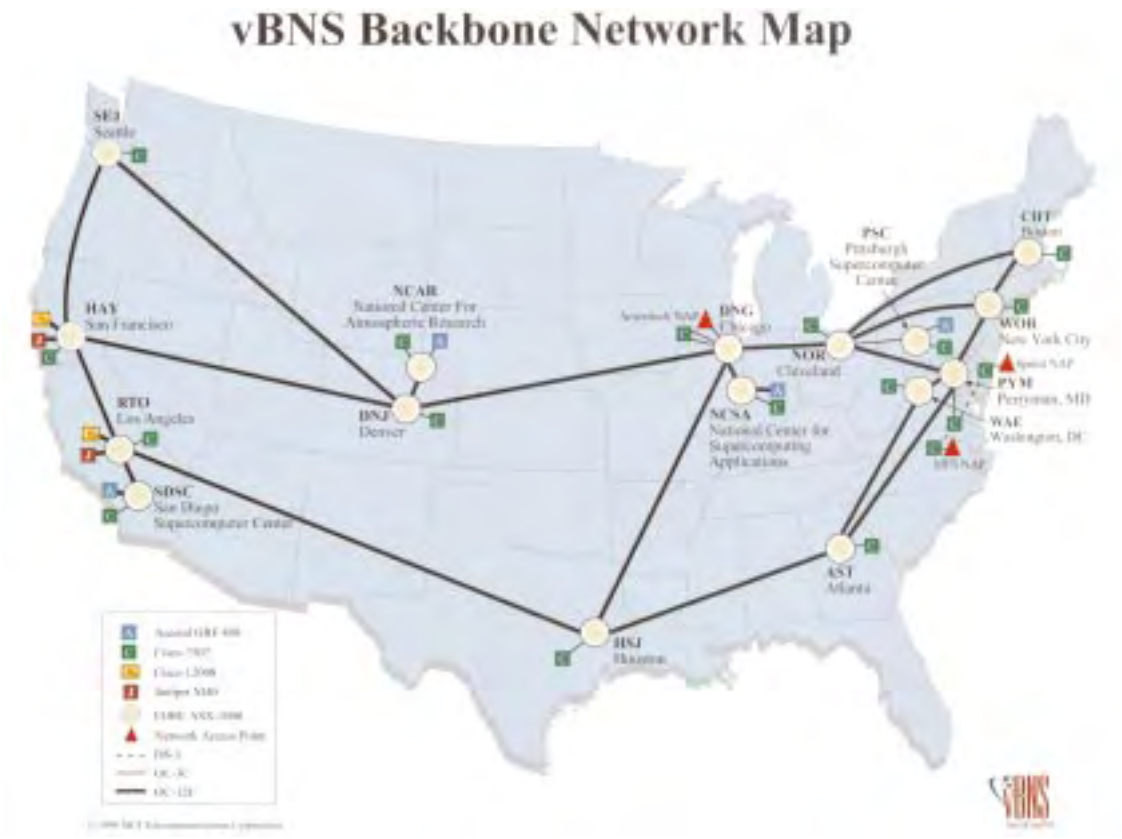

The very high performance Backbone Network Service (vBNS) backbone network service that provides high performance connectivity among NSF research sites and Internet2 gigapops.

research for practical future high capability networks. Under this program, the vBNS links NSF-supported high performance computing centers and almost 100 research institutions engaged in research needing next generation networking capabilities. Links are being established on a competitive basis to research institutions with scientific applications demanding the highest performance networks. The activity supports collaborative development of national and international networks with other agencies and countries, the NGI initiative, and the university-based I2 program.

ANIR supports research to develop network access and control protocols, network management tools and techniques, wireless networks, mobile computing, optical systems, software to support distributed computing, software to support resource discovery and access to networked resources, and I/O devices and subsystems.

Networking technology development

NSF's Internet Technologies program focuses on the fundamental science and technology needed to facilitate the efficient, high speed transfer of information through networks and distributed systems. It supports development of complex network monitoring, problem detection, and resolution mechanisms; development of automated and advanced network tools, network-based middleware, and networked applications tools; and creation of usable and widely deployable networking applications that promote collaborative research and information sharing.

Research areas include agent-based networks, high speed networks, multicast, multimedia applications, multiple access protocols, network architectures, network design, network management, network security, network systems, object-oriented frameworks for networks, optical networks, performance evaluation, protocols, quality of service, resource management, traffic control, and wireless and mobile networks. The program encourages collaboration with other disciplines of computer science and engineering such as communications, control theory and devices, databases, distributed systems, operating systems, signal processing, and software. It also supports NSF's computing-communications research program, including access to high performance networks and computing systems for teams of university researchers; computer system interfaces to communications networks and other high speed peripherals; and interconnection structures among processors, memories, and I/O channels.

Active Networks

DARPA is developing a new network architecture based on programmable infrastructure. Through large scale testbeds, the Active Networks program is advancing active networking and network management, innovative network infrastructure services, and high value end-user services. This research is coordinated with network technology and service deployments made by DoD, NASA, and other Federal agencies.

Global Mobile information systems

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 by developing nomadic technologies and techniques at the applications, networking, and wireless link/node levels.

Extensible networking

DARPA's Extensible Networking program supports the underlying networks and network services needed to accommodate the large scale changes necessitated by ever-increasing — and increasingly diverse — network traffic. Advances in networking will enable revolutionary applications and a vast increase in the geographic scope and heterogeneity of access to the information infrastructure, ensuring that the capacity of the core network and its services can be efficiently and robustly scaled to accommodate accelerated growth. To support extensible networking, DARPA initiated in FY 1999 programs in Gigabit capacity wireless networking, internetworking with low earth orbiting (LEO) satellites, and deeply networked systems.

Very High Speed Networking

NSA's Very High Speed Networking program will provide a high performance network infrastructure characterized by multi-gigabit per second trunking speeds and the ability to support sustained data flows of at least hundreds of megabits per second today, and ultimately, multi-gigabits per second.

In FY 1999, NSA increased network efficiency by reducing the networking protocol layers and moving the control back to the endpoints of the network. In cooperation with NRL and the Defense Intelligence Agency (DIA), NSA used Asynchronous Transfer Mode (ATM) over a wavelength, without intervening Synchronous Optical Network Transmission (SONET) terminals by transmitting an HDTV 720 progressive digitized signal at 1.5 Gbps through the equivalent of 400 kilometers and eight ATM switches. NSA is also operating an all-optical transparent Internet on the Advanced Technology Demonstration network (ATDnet) in the Washington, D.C., area, employing two optical networking technologies:

- Prototype wavelength routers from Lucent Technologies provided under the DARPA-funded MONET Consortium project
- NSA's optical crossbar network from Optical Networks, Inc.

NSA will demonstrate end-to-end communication with no intervening electro-optical conversion, employing "just in time signaling" with "optical burst switching," to provide both packet service and circuit service on the same infrastructure. This project provides quality services on a single network and addresses the latency issue for future networks. NSA will also demonstrate optical multicasting, employing the "drop and continue" feature of the wavelength routers acting as a public network and the natural multicasting capability of an optical crossbar switch acting as a private all-optical network. Participating ATDnet sites include NSA (at its new Laboratory for Telecommunications Science), NASA's GSFC, NRL, and DIA.

Based on the results of experiments in FY 1999, NSA will study new approaches to congestion control in FY 2000 and address multi-domain network management. While the Internet permits read-only Simple Network Management Protocol access across network boundaries for access to select data, it is usually insufficient to debug a connection end-to-end. NSA will examine peer relationships between network management centers that exchange information in a controlled way, enabling end-to-end monitoring and fault isolation.

Networks for biomedical research

NIH's National Cancer Institute (NCI) is implementing evolving networking technologies and high speed interfaces to the computational infrastructure of the NCI Frederick Biomedical Supercomputing Center (FBSC) to improve access for members of the biomedical research community. This will provide advances in data communications technologies including local area networking, wide area networking for multimedia data transmission, dedicated specialized high speed interfaces for local computer to computer connections (for example, high speed crossbar switches, High Performance Parallel Interface [HiPPI], and fiber channels), and the use of evolving data communications standards such as ATM and SONET. In FY 1999, this program expanded the use of visual and voice interaction systems among biomedical computing researchers and projects.

NOAA's advanced ATM-based network

As a result of early experimentation and testing supported by NOAA's HPCC program, NOAA is installing an advanced ATM-based network in its Boulder laboratories. The network has 2,400 nodes participating in 80-90 Virtual Local Area Networks (VLANs). This design provides redundancy and the flexibility to allocate bandwidth for individual needs, a requirement for an infrastructure that must support realtime surface observations, satellite feeds, and the needs of a massively parallel processing supercomputer.

Adaptive wireless technologies for hazard response

NOAA is exploring adaptive wireless technologies for use in hazardous spill responses. NOAA has developed a rapidly deployable wireless local area network (LAN) for use at the on-scene headquarters that has been successfully transitioned to operations and is now in routine use. In FY 1999, research expanded the reach of the network to the mobile field personnel evaluating a spill and provided ubiquitous high speed communications from the mobile headquarters back to the central facility.

NOAA's Coastal Services Center in Charleston, South Carolina, teamed with the Florida Marine Research Institute to prototype and evaluate an emergency response system involving leading-edge computing, communications, and Geographic Information System (GIS) technologies. The tested system included a wireless LAN with a range of three miles, Global Positioning System (GPS) integrated with wearable computers, and video and voice over IP, all integrated with a GIS system running at the central facility and connected via a Very Small Aperture Terminal (VSAT). This system demonstrated both the potential of these technologies and the limits of currently available bandwidth. Pictured above left is NOAA's rapidly deployable wireless LAN based in its on-scene mobile headquarters. It has been successfully transitioned to operations and is now in routine use. In FY 1999, research expanded the reach of the network to the mobile field personnel evaluating a spill (above, center and right) and provided ubiquitous high speed communications from the mobile headquarters back to the central facility.

FedNets

FedNets — Federal agency networks — include Federal agency mission networks and agency high performance research networks. FedNets coordinate closely to support participating agency mission and R&D requirements. The FedNets are:

- ❏ vBNS: NSF's very high performance Backbone Network Service
- ❏ DREN: Defense Research and Engineering Network
- ❏ NREN: NASA Research and Education Network
- ❏ NISN: NASA Integrated Services Network
- ❏ ESnet: Energy Sciences Network

Science, Technology, and Research Transit Access Point (STAR TAP) and International Grid (iGrid)

NSF has established the Science, Technology, and Research Transit Access Point (STAR TAP) at the Ameritech Network Access Point (NAP) in Chicago to interconnect the vBNS with international advanced networks that support high performance applications and develop new networking technologies. STAR TAP is managed by the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago, Argonne National Laboratory, and the Ameritech NAP — a Next Generation Internet Exchange Point (NGIX) that connects to the FedNets and to Abilene, enabling international collaborations with other Federal agencies, universities, and industry.

STAR TAP facilitates the long term interconnection and interoperability of advanced international networking in support of applications, performance measuring, and technology evaluations, anchoring the international vBNS connections program and allowing collaboration with the NGI initiative and the I2 communi-

ty. More than 15 country networks will interconnected at STAR TAP by the end of 1999. These include networks from the Asian-Pacific Advanced Network consortium (APAN), Canada (CA*Net — Canada's high performance network), the European Laboratory for Particle Physics (CERN), France (Renater), Israel, The Netherlands (SURFnet), the Nordic countries (NORDUnet), Singapore (SingaREN), Taiwan (TANet), a U.S./Asia-Pacific consortium (TransPAC), and a U.S.-Russia consortium (MirNET).

One of STAR TAP's principal contributions is designing and enabling an integrated approach to the management, performance measuring, scheduling, and consumption of geographically-distributed network, computing, storage, and display resources — a collection of resources called the International Grid (iGrid). To showcase the iGrid, EVL collaborated with Indiana University on a major research demonstration booth at SC98. The booth highlighted case studies from the U.S. and Australia, Canada, Germany, Japan, The Netherlands, Russia, Singapore, Switzerland, and Taiwan.

iGrid applications instrumental in the development of high speed international networks and services include:

- ❏ Parallel applications running on globally-distributed computing elements
- ❏ Remote access to instruments such as electron microscopes and accelerators
- ❏ Tele-immersion and shared workspaces
- ❏ Collaborative medical diagnosis systems
- ❏ New approaches to high quality digital video over commercial networks
- ❏ APIs and toolkits for authorization, authentication, resource allocation, and communications to support the development of distributed applications
- ❏ Multimedia conferencing tools that require quality of service and differential handling of individual streams in a bundle (for example, audio quality is usually preferred over video, and reliable protocols are needed for text conferencing and whiteboards)

High Performance Network Service Providers (HPNSPs)

NSF has designated a category of commercial sector High Performance Network Service Providers (HPNSPs) that provide advanced network services over broadband networks to university and Federal agency sites and provide the high performance services needed by the NGI. The Abilene network is the first HPNSP. Vendors coordinate closely with the HPNSPs, through the JET and other LSN teams, to provide connectivity and advanced services to high performance network users.

Energy Sciences Network (ESnet)

DOE's ESnet is a service-oriented production network that supports 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 adopts and integrates leading-edge technologies to support DOE's mission applications and will continue to implement and enhance advanced interconnection and peering with NGI networks and other Federal research networks, as well as university networks and aggregation points (such as Gigapops) to support collaborations among DOE mission programs and university programs.

DOE conducts networking research, advanced network deployment, and advanced application support for more than 20,000 users of dozens of DOE experimental facilities and high performance computing resources. DOE's core network and network security research programs include high speed services to applications, routing and congestion control, differentiated services to applications, manageable security infrastructure and architecture, integration of services across autonomous systems and networks, network performance measurement and management, and infrastructure to support both mission science and networking R&D.

Connectivity to Alaska and Hawaii

The JET has promoted cooperation among the Federal and university research communities for access to Hawaii and Alaska by establishing a consortium of the Federal agencies needing network connectivity to these states. The agencies have cooperatively developed and maintained high bandwidth/high performance network service to sites and facilities within both states to address agency mission requirements, significantly improving the network performance of their science and research sites.

Distributed data access improvements

NASA, NIH, NOAA, and other agencies project huge increases in demand for data via networks over the next few years due to the launch of Earth observation satellites and projected exponential increases in demand for weather, environmental, and health data. These agencies coordinate within the JET to develop and implement improved high performance networking architecture, bandwidth, user applications, and network management.

Over the past four years, multiple NOAA sites have experienced order of magnitude increases in demand for weather and climate data by highly distributed end users. This increased demand is projected to continue over the foreseeable future. In particular, demand at its Silver Spring Metro Center facility in Maryland is expected to grow from today's 150 gigabytes per day (GBpd) to more than 600 GBpd in FY 2000. To prepare for this demand, NOAA has cooperated with the other LSN agencies to implement improved high bandwidth connectivity, networking, and dissemination from its facilities. Advanced communications technologies such as ATM, as well as the latest Internet software technologies, are also being implemented at NOAA's Boulder, Colorado, facility for improved environmental data dissemination.

Experimental Program to Stimulate Competitive Research (EPSCoR)

NSF's Experimental Program to Stimulate Competitive Research (EPSCoR) fosters improvements in high performance networking and network user services at the state and institutional level and helps states increase their R&D competitiveness by identifying and developing the science and technology resources at each state's major research universities.

Applications

A major objective of LSN R&D is to develop advanced applications and user services to address critical networking needs of the research community, academic and industrial users, and society in general. Each agency focuses on applications addressing its mission requirements and the development of enabling network services — such as Quality of Service (QoS), security, multicasting, and hybrid communications — needed to support them.

High Performance Applications for Science and Engineering (HPASE)

NSF supports an extensive program of research on applications that pursue fundamental knowledge in science and engineering. These applications include High Performance Applications for Science and Engineering (HPASE) that will

advance computational capabilities to enable new discoveries in science and engineering. Thus, they require access to the highest performance computing systems available, interconnected by high speed networks.

In FY 1999, HPASE continued development of Earth system models and high spatial resolution meso-scale forecast models in a collaborative effort among the university community, the National Center for Atmospheric Research (NCAR), and other Federal laboratories. By the end of FY 1999 a set of simulations covering the period from 1860 to 2300 will be available to the research community, providing examples of the changes to the Earth's climate under various anthropogenic influences. A continuous process of model improvements will be guided in part by the results from these major simulations. The model will have an increased capability to incorporate additional observation sets (for example, the NOAA Next Generation Weather Radar [NEXRAD]) into the initial conditions to improve data assimilation. Additional R&D will create a high efficiency space weather environment model.

In FY 2000 under the HPASE program, NCAR will obtain a major new supercomputing system to develop new computational techniques that allow large codes to run efficiently on the new machine. This effort will be closely coupled to ongoing activities in model development and improvement. Researchers will run new scenarios and ensembles with the community climate model to explore natural and anthropogenic effects on the climate system. These simulations will help reduce the uncertainty of current model results and explore the model's sensitivity to natural variations, such as fluctuations in solar radiance.

Weather and the environment

NOAA supports a program of research in advanced networking applications to make its vast environmental and weather data resources easily, quickly, and completely available to a wide range of end users. In pursuing this goal, NOAA exploits existing and developing access and information technologies including the NGI, distributed computing, the Web, collaboration, mirroring, multicast, and digital libraries. NOAA's advanced networking applications research areas include collaborative data visualization technologies over advanced networks, multicasting to disseminate weather model data output to Federal and university users, testing and evaluating security functions for a range of users including remote and mobile users, and prototyping use of technologies for improved efficiency and robustness.

China Clipper

The DOE China Clipper program will accelerate the use of applications that require substantial computing resources, involve high rate and high volume data flows, involve human interactions, and require coordination of distributed resources.

Telemedicine for healthcare applications

LSN R&D at NIH's National Library of Medicine (NLM) supports testbed networks to link hospitals, clinics, doctor's offices, medical schools, medical libraries, and universities to enable healthcare providers and researchers to share medical data and imagery and access medical literature. NLM also supports the development of collaborative network technologies to allow healthcare providers in remote locations to provide realtime treatment to patients, including technologies for visualizing human anatomy and analyzing images from X-rays, computerized axial tomography (CAT) scans, positron emission tomography (PET) scans, and other diagnostic tools, as well as database technologies for storing, accessing, and transmitting patients' medical records while protecting the accuracy and privacy of those records.

Additional NLM R&D focuses on evaluating telemedicine techniques. In FY 2000, NLM will continue funding projects promoting the application of LSN technologies to healthcare, telemedicine, digital libraries, and methods of protecting the privacy of electronic health data.

Computer-based patient records

The objective of the Agency for Health Care Policy and Research's (AHCPR's) computer-based patient record program is to improve the uniformity, accuracy, and retrievability of data about patient care in the community and promote its use for improved clinical decisions. It requires the integration of the networked information systems of healthcare providers and institutions in remote and urban areas.

Integrated Academic Information Management Systems (IAIMS) grants

The goal of NLM's IAIMS program is to develop, test, and implement generic information flow management systems within university health science centers and major teaching hospitals to increase research productivity, improve access to patient data for technology assessment and health outcomes research, and provide more efficient patient care and healthcare resource use. The 120-plus Academic Medical Centers are comprised of health profession schools and their associated teaching and research hospitals, clinics, and laboratories. These centers need up-to-the-minute information on patient care, research, education, and administration, regularly drawing upon databases of bibliographic and factual information, molecular databases, patient records, and laboratory and clinical data.

Research funded by IAIMS grants is expected to benefit all health delivery organizations, including community hospitals and outpatient services.

Visible Human (VH) project

The large size of NLM's VH image set and other medical images challenges storage and network transmission technologies. Since the full set of Visible Human images would require a capacity of more than 100 CD-ROMs — an impractical distribution option — NLM is investigating advanced compression and networking techniques to minimize storage capacity and improve transmission speed over the Internet.

Bioinformatics

Bioinformatics is an essential component of genome research, protein engineering, and drug design through its use of analytical and predictive methods to identify key molecular patterns associated with health and disease. NLM's National Center for Biotechnology Information (NCBI) focuses on automated systems to store and analyze the vast and growing volume of molecular biology, biochemistry, and genetics data. Within a distributed database architecture, NCBI collects sequence data from researchers worldwide and incorporates them into GenBank, NIH's DNA sequence data bank — a key data resource of the Human Genome Project — to produce an integrated database system consisting of GenBank, the genetic scientific literature in Medline, taxonomy, and 3-D molecular structures. These databases are accessed daily over the Internet from more than 90,000 sites and account for more than 4 million hits per day. Basic research on efficient data analysis techniques and large scale genome analysis conducted within NCBI's Computational Biology Branch has been a key factor in gene discovery.

Agency workshops

Many LSN agencies held workshops in FY 1999 to promote progress and collaboration on advanced networking technology R&D. For these and other workshops, please see the special section on Workshops beginning on page 93.

The Federal NGI initiative, together with the country's other networking R&D investments, is creating 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. Tightly coupled with the networking research and R&D infrastructure support funded under the LSN budget, the NGI is helping to build partnerships among academia, industry, and government that will keep the U.S. at the cutting edge of information and communications technologies and stimulate the introduction of new multimedia applications in our schools, businesses, and homes.

The NGI initiative has three goals:

NGI Goals

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

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

Goal 2. To develop two NGI testbeds for system-scale testing of advanced technologies and services and for developing and testing advanced applications:

- A 100x testbed that will connect at least 100 NGI sites with end-to-end performance at least 100 times faster than the Internet of 1997
- A 1000x testbed that will connect about 20 sites with end-to-end performance at least 1,000 times faster than 1997's Internet

Goal 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 1999 are DARPA, NSF, DOE, NASA, NIH (NLM and NCRR), and NIST. The NGI initiative is managed by the participating agencies and coordinated by the LSN Working Group, whose teams in turn coordinate closely with experts from academia, industry, and Federal laboratories.

NGI Goal 1 accomplishments and plans

Goal 1 activities focus on research, development, and testbed deployment and demonstration of technologies to permit the effective, robust, and secure management, provisioning, and end-to-end delivery of differentiated classes of service.

Quality of Service (QoS)

NASA has used dedicated circuits to support critical network applications such as mission control. This expensive approach results in overprovisioned networks that are idle most of the time. NASA QoS technologies offer the promise of more cost-efficient sharing of network resources, enabling select applications to receive preferential treatment when networks are congested. NASA is developing a QoS testbed between its Ames Research Center in California and its Goddard Space Flight Center in Maryland to test various QoS mechanisms and to determine how to achieve QoS in both low- and high-bandwidth environments.

NASA participates in the I2 QBone, an end-to-end QoS testbed to accelerate the development of interdomain quality of service. NASA's primary contributions are: 1) incorporating the NGIX-West into the I2 QBone infrastructure, providing a means of testing delivery of differentiated services over a network exchange; and 2) deploying applications on the QBone that have stringent QoS requirements and evaluating end-to-end performance of these applications to determine the efficacy of various QoS approaches. NASA is also sponsoring university research to produce end-to-end QoS by developing an interface to translate application QoS requirements into requirements for using network resources, with a focus on multimedia applications. A distributed visual tracking application, wherein a remote user tracks an object in a video stream, is being used to validate the interface.

NIST QoS research focuses on the emerging Multi-Protocol Label Switching (MPLS) technology, a candidate for supporting QoS on an Internet scale. NIST developed NIST Switch, a public domain prototype for research in QoS routing and signaling protocols that serves as a reference implementation to foster commercial implementations. NIST has also developed the Distributed Internet Protocol and PERformance (DIPPER) test system to enable easy testing of distributed, topology-sensitive QoS routing and signaling protocols by allowing the user to write protocol test scripts that are downloaded and executed at multiple points in a network's topology.

In FY 1999 more than 300 organizations in the Internet research and product development community adopted, developed, and released NIST QoS testing tools, including:

- ❏ NIST Net — a general-purpose tool to emulate performance dynamics in IP networks, allowing controlled, reproducible experiments with QoS sensitive applications and protocols
- ❏ ISPI — an interactive measurement tool for experiments in realtime transport and resource reservation protocols

Bandwidth Broker

NASA is sponsoring research to develop techniques for allocating bandwidth and services across network domains. A "Bandwidth Broker" based on a QoS routing technique provides detailed information about the local network and less pre-

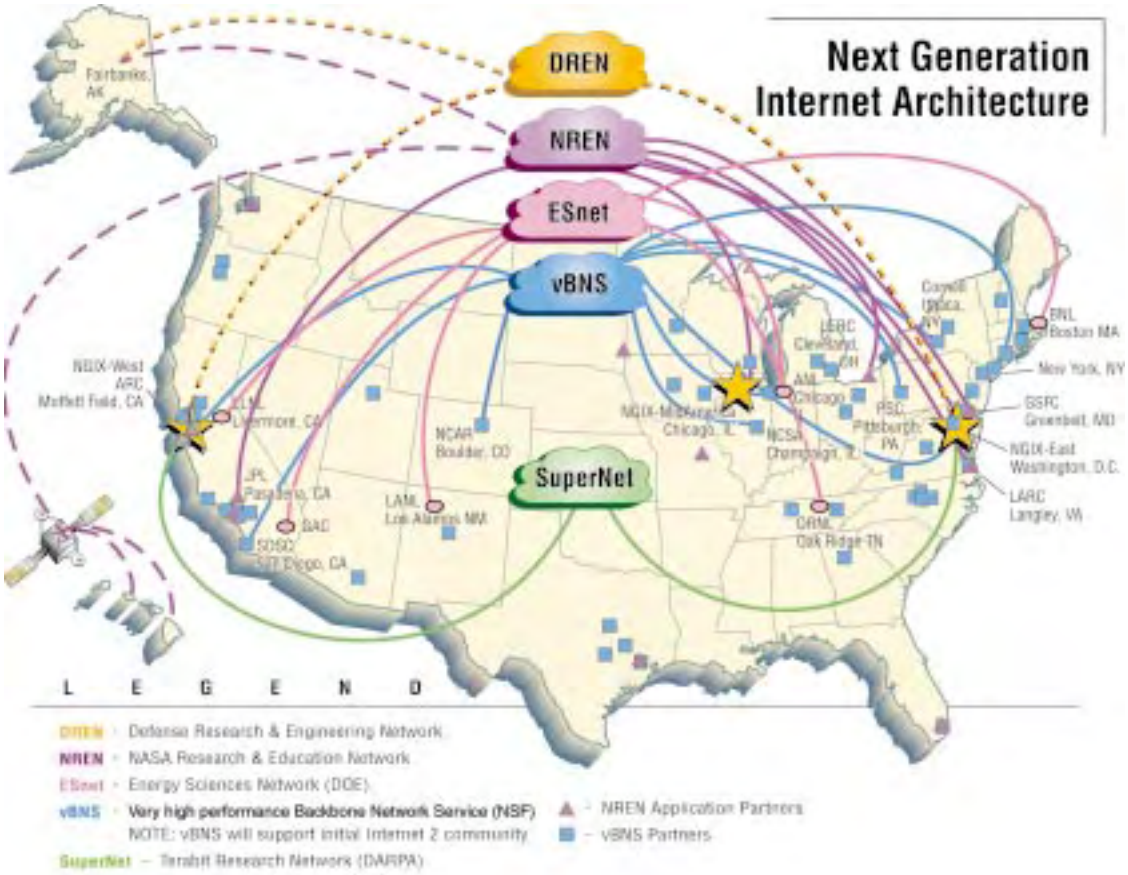

cise data about remote networks. Such a broker can reach coarse-grained preliminary bilateral agreements within a short time and refine such agreements when more information about bandwidth, services, and prices becomes available.

DARPA NGI research

In FY 1999, DARPA contributed R&D in multi-gigabit broadband access technologies, assured service mechanisms, integrated network management, and QoS to the NGI.

Multidomain multicast

Since NASA applications increasingly involve collaboration among groups of researchers from multiple scientific disciplines at distributed and remote locations, efficient multicast has become essential. In addition, since multiple network-domain boundaries must be crossed to provide end-user-to-end-user high performance networking, the integration of services across autonomous networks is critical. A portion of NASA's NREN R&D is focused on these challenges.

Hybrid networks

NASA's R&D in hybrid networking focuses on satellite and terrestrial components to provide seamless high performance networking for end-to-end hybrid connectivity to highly distributed sites and QoS across hybrid networks.

Internet security and mobile networks

NIST's Cerberus reference implementation of the Internet Engineering Task Force (IETF) Internet Protocol security (IPsec) protocols was extended with the addition of PlutoPlus, NIST's prototype Internet Key Exchange (IKE) protocol. The integrated IPsec/IKE reference implementation provides a platform for

research into Internet security systems integration. In addition, the integrated Cerberus and PlutoPlus were added to the IPsec Web-based Interoperability Tester (IPsec WIT), allowing users to execute, remotely over the Internet, more than 400 interoperability tests against an IPsec/IKE implementation that is under development.

In FY 1999, NIST began research in Mobile Ad-Hoc Networks (MANET) in partnership with DARPA. This project will produce techniques and metrics for evaluating the MANET protocols proposed to IETF and similar protocols proposed to DARPA as part of their Global Mobile program. NIST also began work on metrology for service-rich, agile Dense Wave Division Multiplexing (DWDM) based access and metropolitan networks to promote technology development, interoperability, and standards and to accelerate deployment of Wave Division Multiplexing (WDM) in NGI and commercial networks.

DOE's networking research

In FY 1999 DOE invited proposals for research grants in three areas:

- Research in basic networking, including technologies that allow very high speed interfaces to connect devices to networks, protocols and techniques for coordinating multiple heterogeneous network-attached devices, software to allow applications to adapt to changing network conditions, and network performance characterization
- University networking technology testbeds that focus on developing and testing techniques and technologies to allow deployment of advanced network services across independently administered, interconnected networks
- Network testbeds and partnerships among applications developers and network researchers that focus on integrating advanced applications with leading-edge network research to test wide area data intensive collaborative computing technologies

NGI Goal 2.1 accomplishments and plans

NGI Goal 2.1 is to develop the 100x testbed that will connect at least 100 NGI sites with end-to-end performance at least 100 times faster than the Internet of 1997. The JET coordinates engineering implementations and architecture of FedNets and the Goal 2.1 testbed. The NREN, vBNS, and DREN, respectively the NASA, NSF, and DoD FedNets, are interconnected to provide the fabric for the 100x testbed. They cooperate closely with ESnet and NISN (DOE and NASA production networks) to provide end-to-end user connectivity, performance, and services.

NREN

NASA conducts network research through its NREN program. NREN, which interconnects NASA's Grand Challenge centers, is a high performance network testbed employing advanced telecommunications technologies such as ATM over SONET services, with a backbone capacity of 155 megabits per second (Mbps). Over the next several years, increases in bandwidth to 622 Mbps and beyond are anticipated.

Next Generation Internet Exchange Points (NGIXs)

The FedNets and Abilene interconnect at the NGIX-West, maintained by NASA's Ames Research Center in California, and at the NGIX-Midwest at the Ameritech NAP in Chicago. Most FedNets plan to interconnect at NGIX-East, expected to be located in the Washington, D.C., area. Network connectivity does

not necessarily imply peering (that is, exchange of traffic). Peering between any two networks must be established by agreement based on each network's policies.

Cooperation among the FedNets and with Abilene, I2, and the Gigapop operators provides a coordinated advanced network architecture and engineering of network services to assure high performance end-to-end connectivity and services among users, preventing duplication of costly resources to end users, fostering innovative, cooperative approaches to remote users, and significantly improving Federal agency capabilities.

Connections Program

By the beginning of FY 1999, 131 NSF Connections Program awards resulted in:

- 70 institutions connected to the NSF vBNS
- 19 additional planned connections to the vBNS
- 42 pending connections to the vBNS, Abilene, or another HPNSP

By the end of FY 1999, NSF will have made 150 awards for high performance connections to the vBNS and other high performance networking testbeds. In FY 2000, additional NSF connections awards and connections to NASA, NIH, and NOAA sites are expected to result in more than 150 institutions connected to the 100x testbed, which significantly exceeds NGI Goal 2.1. Since its cooperative vBNS agreement with MCI concludes in March of 2000, NSF will initiate a plan for the post-vBNS cooperative agreement era.

NSF coordination with universities and research institutions

The University Corporation for Advanced Internet Development (UCAID) fosters university development of high performance networking. It promoted the development of Gigapops for university and research institutions in local or metropolitan areas. In FY 1999, the Abilene network, in coordination with UCAID, began providing high performance backbone network service among the Gigapops. Abilene and any other qualifying high performance network service providers can help support NSF high performance networking research.

NGI Goal 2.2 accomplishments and plans

Goal 2.2 technologies address ultrahigh speed switching and transmission technologies and the demonstration of end-to-end connectivity at speeds over one Gbps. DARPA leads this multiagency effort, with participation by NSF, DOE, NASA, and other Federal agencies.

SuperNet

DARPA is deploying SuperNet, which is expected to provide 10 to 100 Gbps speeds in FY 1999-FY 2000. Other NGI networks and programs coordinate with SuperNet to provide end user connectivity, applications, and user feedback. In FY 2000, DARPA will continue developing technologies such as optical Add-Drop Multiplexers (ADMs), improved I/O devices, and innovative network architecture and management software, and will encourage use of the network to demonstrate innovative high performance technologies and applications.

SuperNet will provide a testbed using the resources of some existing research testbeds and developing additional component testbeds, including NTON, HSCC, ONRAMP, BOSSNET, and MONET/ATDNET. These testbeds will be implemented with gigabit access technologies, IP over WDM, and innovative network engineering management and modeling. SuperNet has demonstrated QoS negotiation over wide area ATM networks, a baseline QoS architecture for the full network, operating system kernel adaptation tools, and mechanisms to achieve a factor of

NASA's NREN strategy is to conduct leading-edge systems engineering and applications engineering R&D and develop and demonstrate advanced applications, as illustrated in the graphic to the right.

three to five reduction in communications overhead. Other SuperNet milestones include simulating WDM transmission in WANs and developing cascaded transparent optical ADMs and amplifiers.

Test and evaluation

NIST supports the U.S. information technology industry by testing and evaluating new NGI networking and infrastructure technologies at all stages of development and by fostering their rapid commercialization and deployment. Advances in measurement and testing technologies will enable rapid evaluation of research designs and prototypes and will facilitate the transfer of new technologies to the NGI Goal 2 testbeds and to the communications industry. The goal is for test and instrumentation technologies to become part of the protocol design and specification process and be integrated into the implementation and deployment of the network itself.

NGI Goal 3 accomplishments and plans

NGI Goal 3 is to develop and demonstrate revolutionary applications, some of which are described in the following sections. Successful applications and technologies are being transitioned to the Federal agency operational networks and to the public domain for commercialization.

SC98 demonstrations

The HPNAT supported eleven demonstrations of NGI applications at SC98, held in November, 1998, in Orlando, Florida, including:

NSF: Science and engineering applications

DARPA: Innovative data search and language processing

NIH: Biomedical applications including telemedicine and data access

NIST: Remote manufacturing, QoS, and Internet security

NASA: Data visualization and aeronautics and space applications

A detailed description of these SC98 demonstrations begins on page 51.

Aeronautics and space applications

NREN collaborates closely with NASA programs including Earth Sciences, NASA centers, and other Federal agencies such as the FAA to identify networking research needs. Through its systems engineering and applications engineering (for example, hybrid systems/high bandwidth/enhanced network services for the space station), NREN helps develop and demonstrate revolutionary applications.

In FY 2000, NASA plans to demonstrate a 500 times end-to-end performance improvement of several Grand Challenge mission applications over FY 1996 performance measurements across the NREN testbed (a 622 Mbps and greater wide area network). Revolutionary NREN applications include:

- Establishing a virtual institute model for scientists collaborating with each other and interacting with their data and models in real time. NREN is working with the new Astrobiology Institute, a partnership among NASA and academic and research organizations, to conduct interdisciplinary research in astrobiology. The Institute's focus on life in the universe brings together geographically dispersed laboratories and teams of astronomers, chemists, physicists, geologists, biologists, and exobiologists to operate as a virtual institute linked by a high performance Internet.
- Prototyping a geographically distributed heterogeneous information and computational capability to be known as the Computational Aerosciences Information Sciences Power Grid.
- Demonstrating realtime and post-mortem access by distributed engineering and operations teams to high-definition cameras and sensors to support space launches and missions. NREN seeks to expand the capabilities of the Kennedy Space Center's Checkout and Launch Control System (CLCS) by adapting QoS and multicast technology to enable WAN realtime interactive access to multiple video streams.
- Proving ubiquity of information flow by leveraging LEO constellations of networking satellites connecting virtually every site and mobile platform on Earth, and Geosynchronous Earth Orbiting (GEO) networking satellites that enable low-cost multicast distribution of data services to teams, aircraft, spacecraft, and communities worldwide.
- Prototyping remote visualization and manipulation of extreme environments and hypothetical worlds in virtual reality settings. Currently, NREN supports the ESS "Turbulent Convection and Dynamos in Stars," which will model turbulent flow in stars at three sites, each with a different data set; and the virtual Distributed Online Clinic (vDOC), which will display time correlated 3-D anatomical images at five sites in real time for remote manipulation, diagnosis, and treatment. Delay constraints are crucial to provide the performance guarantees required by the teleconferencing and video distribution portions of the application. These projects involve high bandwidth, realtime reliable multicast and are expected to provide insight into techniques for supporting multiple applications running concurrently over an enhanced network services environment.

Environmental data processing

NOAA plans to expand its NGI network connections to enable advanced data aggregation, dissemination, collaboration, and computing applications. NOAA is also exploiting World Wide Web (WWW) software technologies to implement advanced visualization of NOAA environmental information on the Web.

Medical Connections program

NLM is demonstrating the use of the NGI for health and medicine related applications. The program spans telepresence, tele-immersion, teletrauma, telemammography, internetworking, and nomadic computing, and is intended to improve the cost, quality, usability, efficacy, and security of healthcare, health edu-

cation, and health research systems. It is designed to lead to new applications based on the ability to control, feel, and manipulate devices at a distance, and is expected to enable the transfer of massive amounts of data accurately, securely, and almost instantaneously, by providing virtually error-free service, security and medical data privacy, network management, and infrastructure for collaboratories.

Applications linking major scientific facilities

DOE's NGI research program is focused on discovering, understanding, developing, testing, and validating the networking technologies needed to enable wide area, data intensive, and collaborative computing that is not currently possible. This program will integrate scientists working on fundamental research in applied mathematics, computer science, and networking with scientists working on DOE applications to develop new ways to link scientists with DOE's major scientific user facilities and computational centers. Such research is needed to enable effective use of petabyte/year facilities such as the Relativistic Heavy Ion Collider, to provide remote visualization of terabyte to petabye data sets from computational simulation, to develop advanced collaboratories, and to enable effective remote access to tomorrow's advanced scientific computers. These applications involve extremely large data sets and require that scientists be able to interact with the data in (nearly) real time.

NGI logo

The NGI R&D initiative links Federal agencies with academic and industry partners — private sector entities that receive NGI agency funding or who contribute cooperatively to the NGI program. For example, the vBNS component of the 100x testbed is funded cooperatively by NSF and MCI, the 100x testbed cooperates with university Gigapop operators to provide high performance network connectivity and end-to-end user services, and Cisco participates in NREN's QoS testbed. To foster awareness of the NGI initiative and to promote such cooperative programs, LSN plans to offer the use of the NGI logo by approved NGI partners.

NGI Demonstrations at SC98

On November 9-12, 1998, six agency members of the LSN Working Group — DARPA, NASA, NIH, NIST, NOAA, and NSF — held eleven NGI demonstrations at SC98 in Orlando, Florida. This marked the first time a group of NGI demonstrations was held outside of Washington, D.C. The NCO hosted the demonstrations in its research exhibit in an effort to foster industry awareness about the Federal government's investments in networking technologies. Six additional NGI demos were held at agency booths in the research area, including NASA, NCSA, the National Laboratory for Applied Networking Research (NLANR), NPACI, NSF's Integrative Graduate Education and Research Training (IGERT), and the Pittsburgh Supercomputing Center.

SC98 demonstrations included:

Realtime Functional MRI (fMRI): Watching the Brain in Action

The "brain in action" allows remote viewing of brain activity that occurs while a patient is executing cognitive or sensory-motor tasks. Neurosurgeons, neurologists, psychiatrists, and brain scientists can use this technology to investigate brain function and diagnose and treat brain diseases. With this technology, neurosurgeons will be able to plan the surgical removal of a tumor by pinpointing — and avoiding — those unique cognitive and sensory-motor abilities located near a tumor site.

NGI advances will improve interactive realtime capability, patient confidentiality, and reliable data delivery between the fMRI acquisition site, the processing site, and the visualization site.

Participants

Pittsburgh Supercomputing Center
Carnegie Mellon University
University of Pittsburgh

Sponsors

NIH: National Center for Research Resources, National Institute on Drug Abuse, and National Institute of Mental Health
NSF

Pictured left is the entrance to the NCO-supported exhibits at the Orlando Convention Center at SC98. NGI technologies were demonstrated here and at other agency booths throughout the exhibit area.

NASA- and NOAA-supported researchers demonstrate the Distributed Image Spreadsheet (DISS) at SC98.

Distributed Image Spreadsheet: Earth Data from Satellite to Desktop

With the Distributed Image Spreadsheet (DISS), scientists visualize, manipulate, and analyze massive geologic, atmospheric, and oceanographic data sets transmitted to their desktops from NASA's Earth Observing System satellites. Government agencies, universities, corporations, and weather services use DISS data for atmospheric, oceanographic, biospheric, and land use studies.

NGI advances will improve I/O performance of multimedia digital libraries, the use of distributed file systems, and visualization technologies.

Partners

NASA: Goddard Space Flight Center and Ames Research Center
NOAA: Hurricane Research Division
University of Missouri-Columbia

Hurricane Forecasting

While it is not a funded NGI agency, NOAA will derive benefits from NGI R&D. NOAA/GFDL's SC98 exhibit highlighted its Hurricane Prediction System (HPS) — a comprehensive computer forecast system that predicts the behavior of hurricanes, such as the recent Georges and Bonnie, up to three days in advance. Predictions include not only storm motion, but also intensity, precipitation, and 3-D fields, such as wind, temperature, and humidity. The system incorporates information about atmospheric conditions ranging from large scale to hurricane scale. Since 1995, NOAA's NCEP has used the HPS for operational hurricane prediction at the National Weather Service. Since then, the GFDL model has, on average, predicted storm tracks more accurately than any other model and has helped save lives, national resources, and millions of dollars in evacuation costs.

NGI technologies will lead to improved collection of observational data and more effective dissemination of storm warnings to the public.

Participants

NOAA: GFDL, NCEP and National Hurricane Center (NHC)
DoD: U.S. Navy Fleet Numerical Meteorology and Oceanography Center
University of Rhode Island

Interactive Video Dialogues

Voice-controlled multimedia scenarios engage users and virtual characters in face-to-face, realistic, dramatic dialogue for education and training using the Conversim™ interface software developed by Interactive Drama, Inc. Interactive Video Dialogues applications include combat casualty triage training, knowledge systems, and language training. For example, military linguists located anywhere in the world can sustain language proficiency by routinely conversing in virtual dialogue with native speakers.

NGI will enhance network capacity to accommodate large quantities of full-motion, broadcast-quality video in a speech recognition environment.

Sponsors
DARPA
NIH

Participant
Defense Language Institute

Remote Access Multidimensional Microscopy (RAMM): Viewing the Changing Threads of Life with 4-D Telemicroscopy

This online microscope system non-invasively digitizes 3-D images of living organisms as their cellular and subcellular structures evolve. Examples include a worm embryo undergoing division and a fruit fly's wings taking form. The system can remotely acquire, view, and analyze 4-D data sets (3-D space plus time). The RAMM project allows medical research groups, scientists, and educators to access the system remotely and collaboratively for observation and experimentation, without having to invest in technology or equipment.

NGI will improve realtime imaging essential for remotely manipulating the microscope without affecting an organism's development and will increase bandwidth to transmit 100+ MB data sets.

Sponsor
NIH: National Center for Research Resources

Participants
Digital Microscopy Group - Integrated Microscopy Resource (IMR)
University of Wisconsin Madison

Broadcast News Navigator

A search and retrieval application, Broadcast News Navigator indexes, summarizes, and displays recorded broadcast news stories. Government agencies and broadcasters identify and view news stories on issues, people, organizations, and locations of interest.

The objective is a breakthrough in language understanding technologies to increase one's ability to find information in speech, text, or printed source; analyze it; and produce reports based on it. This will result in high performance, task-independent, language-independent algorithms that are trained from example data to perform information extraction of events from news. Those results will have the

Researchers from the University of Wisconsin-Madison demonstrated the use of Remote Access Multidimensional Microscopy (RAMM) technology to remotely view 4-D datasets.

Supported by DARPA and NSA, the Broadcast News Navigator indexes, summarizes, and displays recorded broadcast news stories, an effort aimed at more efficiently extracting specific content from the massive information-overload faced today by both the Government and consumers.

following impact: 1) technology that successfully extracts the content in context, 2) more portable software because of the combination of domain-independent algorithms and statistical learning techniques, and 3) far more affordable language understanding software, since the learning algorithms will greatly reduce the cost of applying the Navigator to a given task.

NGI will provide access to user groups with advanced networking capabilities who can evaluate and help improve application performance.

Sponsor — *DARPA*

Participants — *The MITRE Corporation*
DARPA
NSA

Collaborative Remote Robotic Arc Welding

Collaborative robotic arc welding brings engineers and equipment at a welding cell site together with weld engineers at a remote laboratory to collaborate on R&D to improve industrial welding practices. Networking technologies enable video connectivity between the two sites, along with remote monitoring, sharing, and analysis of weld quality information. The collaborative arc welding program develops and demonstrates methods for improving practices of the American welding community, which includes the automobile, heavy equipment, and shipbuilding sectors.

NGI will provide high performance networking with high speeds, low latencies, guaranteed quality of service, and secure connections to enable full-motion video between sites and remote instrument/cell control.

Sponsors and participants — *NIST: part of the National Advanced Manufacturing Testbed Program jointly conducted by NIST's Manufacturing Engineering Laboratory (MEL) and Materials Science and Engineering Laboratory (MSEL)*

Exploring the Earth System on the "Second Web"

3-D virtual worlds on the Web teleport viewers into high resolution, stereo/3-D animations of tropical storms, forest fires, clear air turbulence, cyclones, and El Niño. Researchers and educators create and share Earth system data to study the patterns and behaviors behind naturally occurring, and sometimes dangerous, phenomena. New Web technologies make it possible to share detailed, intricate Earth system observations and simulations in stereo/3-D with other researchers, educators, and the public.

NGI will provide high bandwidth wide area networks coupled with 3-D Web technologies and realtime data compression, allowing multiple remote users to share and explore science in virtual 3-D worlds.

Sponsors

NSF
U.S. Forest Service
DOE
Silicon Graphics, Inc.

Participant

NCAR

GeoWorlds: Integrated Digital Libraries and geographic information systems for disaster relief operations

Creating synergy between two technologies — digital libraries and geographic information technologies — GeoWorlds retrieves, organizes, and displays everything that is known about a region in a rich display format, allowing teams of users in disparate locations to collaboratively assess disaster situations and develop appropriate responses. GeoWorlds supports humanitarian assistance and disaster relief by helping response teams assess the impact of disasters, identify assets and partners that can contribute to a response, and evaluate geographic constraints on response plans. GeoWorlds can also be applied to business planning, local government land use, law enforcement, and intelligence analysis.

For maximum effectiveness, GeoWorlds needs extremely high bandwidth and controlled QoS to move massive amounts of map, image, and document information. NGI technologies will improve access to realtime geographic information system data from remote sources and support multiple collaborative disaster relief sessions.

Participants

University of Southern California Information Sciences Institute
University of Southern California, Department of Geography
University of California Santa Barbara Alexandria Digital Library Project
University of Illinois at Urbana-Champaign Digital Library Initiative
University of Arizona Artificial Intelligence Laboratory
University of California at Berkeley Digital Library Project
University of Illinois NCSA

Testing and measuring Internet security

More than 200 researchers and product developers use NIST tools for R&D on new security protocols and products. For example, NIST Cerberus/Pluto++ is a prototype reference implementation of IP security and IKE protocols. IPSec WIT, a Web-based interoperability test system, lets researchers and developers conduct

NIST conducts NGI research on the integration of security protocols/key management/certificate systems, the configuration and management of security policies, and the incorporation of security into QoS and other new NGI network services.

NCO staffers assisting SC98 conference attendees at the NCO booth on the floor of the Orlando Convention Center.

interoperability tests of security protocols anytime and anywhere without relocating systems or software.

NGI supports research on the integration of security protocols/key management/certificate systems, the configuration and management of security policies, and incorporation of security into QoS architecture and other new NGI network services.

Sponsors and participants: *NIST*
NSA

Testing and measuring IP QoS

More than 300 Internet researchers and product developers use NIST's IP QoS testing tools and prototypes when developing QoS-sensitive products. NIST Net emulates arbitrary performance characteristics of complex IP networks, enabling controlled and reproducible QoS sensitivity experiments for application/protocol R&D. NIST Switch is an experimental, prototype platform for research in MPLS, QoS routing, and QoS signaling protocol mechanisms. DIPPER, a Distributed Internet Protocol and PERformance test system, allows a researcher to test IP QoS signaling and forwarding mechanisms between multiple remote locations.

NGI will support research on scalable architectures and mechanisms to support QoS routing, signaling, and management, as well as on advanced test, measurement, and analysis tools for new NGI protocol capabilities and networks.

Sponsor and participant: *NIST*

High Confidence Systems

HCS R&D focuses on the critical information technologies necessary to achieve predictably high levels of system availability, reliability, safety, security, and survivability. A high confidence system is one in which the consequences of its behavior are well understood and predictable. It must withstand internal and external threats and must deal with naturally occurring hazards as well as malicious attacks from a sophisticated and well-funded adversary. Systems that employ HCS technologies will be resistant to component failure and malicious manipulation and will respond to damage or perceived threat by adaptation or reconfiguration.

Security-critical, safety-critical, and life-critical systems are needed in chemical production, electric power generation, financial services, healthcare, manufacturing, oil and gas production, and transportation, as well as in emergency services, law enforcement, and national defense. Systems for power generation and distribution, banking, medical implants, automated surgical assistants, and transportation also need reliable computing and telecommunication technologies.

This section describes representative HCS FY 1999 accomplishments and FY 2000 goals.

NSA HCS research program

NSA supports and participates in the Information Security (INFOSEC) Research Council to coordinate its research program with DARPA, DOE, NIST, and the DoD service laboratories. The goal of NSA's HCS research program is to ensure that information assurance (IA) solutions keep pace with leading-edge information technology to provide essential security services. The IA program includes the technology areas of active network defense, secure network management, and network security engineering, and areas of enabling research in cryptography and secure communications technology.

Active network defense

Active network defense provides and coordinates research and advanced technology development to support the DoD's Defensive Information Operations. Recent accomplishments include completion of a study for a DoD Minimum Essential Information Infrastructure (MEII) in response to a Defense Science Board recommendation, and the establishment of the Pacific Institute of Computer Security at the University of California at Davis and SDSC to conduct research and develop tools to support intrusion analysis and computer forensics. Future research efforts will develop new tools and techniques for analyzing various types of attacks, their sources and objectives, and technologies to support manual and automatic response.

NSA has applied the PARENTAGE visualization tool, originally developed for SIGnals INTelligence (SIGINT) applications, to network attack analysis. NSA will investigate additional SIGINT technologies for network defense applications. Future R&D in visual analysis of network attacks will focus on prototypes that dis-

Face recognition can provide instant hands-free access control to computer workstations. Pictured above left and right are two views of an experimental 3-D face recognition system being tested at NSA. The system uses infrared imaging techniques. Pictured left is an NSA researcher preparing to test the face recognition system using a computer-controlled precision positioning platform.

play the massive sets of multivariate data associated with very large scale systems. NSA has also begun R&D to determine appropriate automated network responses under different intrusion scenarios. Research in mobile agents will investigate applying this technology to network attack detection and response.

Secure network management

Secure network management R&D supports a security management infrastructure by developing secure protocols for information sharing, network control, and monitoring events within information systems. NSA's development of the Internet Security Association and Key Management Protocol (ISAKMP) standard through the IETF provides a flexible capability to secure network connections in order to meet the needs of national security users. Future research will help produce security enhanced Internet protocol specifications, reference implementations, and support at international standards bodies.

NSA is developing a reference implementation of an IPSEC. Along with routing security mechanisms and group key management services, ongoing R&D will focus on proofs-of-concept for key management, fractional keying, and non-cryptographic techniques for multicast communications. Additional research into securing the key management infrastructure will include American National Standards Institute (ANSI) X.509 certificate directory services; management, revocation, and secure binding to security attributes; key generation and recovery; cross-certification from multiple security domains; and trusted time-stamping.

Network security engineering

Network security engineering is concerned with providing information security in a networked environment characterized by globally distributed systems and services coupled with dynamic and pervasive information sharing and collaboration. This R&D addresses issues critical to secure hardware, software, and networked systems. Boundary definition addresses the problem of identifying and protecting network borders in order to establish points for monitoring, controlling, and defending against cyber attack. Boundary protection is currently managed primarily by high performance ATM firewalls that filter communications based upon addressing data, and NSA uses the commercialized results of research in this area to protect its internal networks. New research will develop high assurance, high performance boundary protection devices that add the ability to filter on the data itself or on specific protocols, with a goal of higher efficiency and effectiveness and much higher data rates than currently possible. In order to develop appropriate IA solutions, NSA is also undertaking a program to assess the security implications of advanced ATM network switching technologies, such as IP switching.

NSA research addresses security issues in the use of object technology. Researchers are identifying security problems of greatest concern to distributed object-based computing and developing solutions for adoption by the Object Management Group's (OMG) Common Object Request Broker Architecture (CORBA).

In operating system security R&D, NSA created the security-related components incorporated into the Fluke operating system developed by the University of Utah. Future research will focus on prototyping additional security components such as a security policy server, a security services negotiation server, and a cryptographic subsystem that includes operating system servers for cryptography and authentication.

Network security engineering research also includes identification and authentication (I&A) technologies applied to people, files, security policies, and hardware. NSA transfers robust and reliable biometric and smartcard I&A technologies to industry. The FY 2000 program highlights research on basic biometrics identification technologies, fingerprint recognition, face and speaker recognition, faster algorithms, and evaluation of currently installed prototype systems. Proof-of-concept models will be available to demonstrate a network computer capable of high assurance switching among different classification levels of information. NSA assurance research includes work on verifying key generation and nuclear command and control designs, reducing risks in executing untrusted Java code downloaded from remote Internet sites, and high assurance modeling of security and trust policies. Finally, NSA research aimed at providing protection from hands-on physical tampering will employ industry partnerships to leverage the commercial market to develop a protective coating for semiconductor chip wafers and to develop a unique, easily produced, secure coating for Government use.

Cryptography

As the Nation's primary resource for cryptography, NSA provides the Federal government's cryptographic algorithms. Backed by the highest level of cryptomathematics expertise, NSA designs new algorithms for the unique requirements of the military, DoD, and the intelligence community. NSA's multi-year public key cryptography research effort will produce designs for efficient public key algorithms and protocols, faster and more efficient arithmetic techniques, elliptic curve software, proactive authentication techniques, related technical support, and public key cryptography standards support.

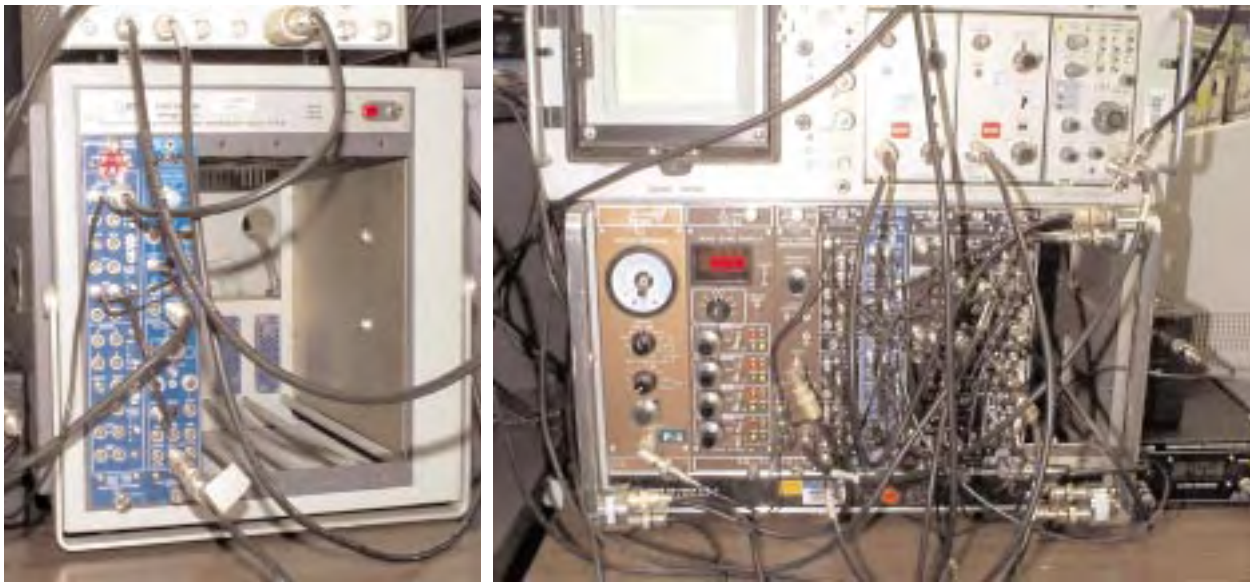

The unit pictured at left is part of an NSA research effort to perfect the use of single photons to carry encrypted messages across fiber optics. This state-of-the-art technique defeats attempts to intercept communications. Pictured right is a closeup of an experimental quantum cryptography system. The next research phase will attempt to shrink this lab bench prototype into a transportable unit. NSA's quantum cryptography experiments have demonstrated the feasibility of this new technology. With continued research, this technique may soon prove to be a practical and powerful method to protect critical information.

Quantum computing is a new and powerful threat to traditional cryptography. NSA researchers have begun to counter this threat by employing key exchange techniques that use quantum physics as a basic protection mode. Recent research has demonstrated the feasibility of key exchanges over 47 kilometers of fiber optic cable. NSA will also devise new classes of cryptoalgorithms that are not susceptible to attacks by quantum computing techniques.

NSA leads government-wide efforts to develop standardized interfaces to integrate cryptography into widely used applications software. Security services such as authentication and encryption can be more easily incorporated into commercial products using a cryptographic API (CAPI). This allows users to easily select the level of protection they need, ranging from commercially developed software algorithms from companies such as RSA Data Security, Inc., to DoD-developed hardware technology such as Fortezza. NSA will develop high assurance reference implementations of CAPIs and standards for broader classes of security service, such as those being studied in the Common Data Security Architecture (CDSA). Researchers will provide a comprehensive suite of security capabilities to protect data ranging from simple files and email to complex multimedia communications.

Secure communications technology

Information transport and its associated infrastructure must demonstrate high assurance capabilities in times of crisis and attack. NSA research encompasses the following enabling technologies: speech coding, wireless communications, high speed cryptography, and optical networking.

Coding research will develop low bit rate algorithms required for digitizing, encrypting, and transmitting tactical voice communications. Wireless research will investigate and counteract the vulnerabilities of the wireless services, use the results to influence standards, provide select demonstrations of critical wireless technologies, and perform the testing, evaluation, and verification needed to ensure the solutions work effectively. NSA has developed technologies to allow the use of

STU-III secure voice services over the European GSM cellular communications system, thereby extending the life of DoD's secure voice technology. Two new technology developments include a terminal for demonstrating wireless multimedia communications for the military and a collaborative effort with the Army Communications Electronics Command to overcome some of the denial-of-service vulnerabilities of tactical cellular communications.

Research in high speed secure communications techniques includes higher performance microelectronics, advanced packaging, and highly efficient cryptographic algorithms. Researchers are consulting with customers to address their security problems. In cryptography, NSA will develop a proof-of-concept 10 Gbps ATM encryptor. Optical communications researchers will develop proof-of-concept optical logic technology and switching devices for cryptographic applications, with a longer range goal of incorporating this technology into a completely photonic key generator.

Information survivability

DARPA is developing technologies to protect DoD's mission-critical systems against attack upon or through the supporting information infrastructure. This will lead to stronger protection, higher performance, and more cost-effective security solutions scalable to several thousand sites, satisfying defense requirements for secure and survivable systems.

Information survivability focuses on early prototypes of hardware and software technologies to protect large scale heterogeneous systems used over a wide performance range in diverse threat environments. DARPA is developing survivability technologies to mitigate national and defense computing infrastructure vulnerabilities that could be exploited by an information warfare enemy. Intrusion detection systems will allow attacks on the defense infrastructure to be detected, damage to be assessed, and an appropriate response to be taken while allowing crisis-mode operation of critical infrastructure components.

DARPA is developing high confidence networking technologies — including security mechanisms, value-added security services, and robust networking protocols designed to facilitate continuous operations in hostile environments — that will be integrated into the network infrastructure.

High confidence computing systems that provide modular security services and mechanisms, provide high reliability for distributed computations, and allow

NSA's Technology Demonstration Center provides a showcase for many of the agency's most innovative developments. Access to the facility is controlled by a customized fingerprint recognition system, pictured at the left.

geographically-separated parts of an organization to interact as if they shared a common security perimeter are also under development. This also includes secure and fault-tolerant operating systems, firewalls, and system management tools. Assurance and integration tools will aid the development of high assurance and trusted systems and the ability to reason about their security properties.

High performance networking environments

NASA is developing technologies to help achieve high confidence in system safety, including new techniques and applications for network security and reliability in high performance networking environments and effective network management to implement administrative policies including security, QoS, and routing in complex, high performance networks.

NSF computing-communications research

NSF's Computing-Communications Research (C-CR) program supports research on fault-tolerant and redundant hardware structures and high confidence systems.

Java security

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. Researchers examining the Java language and both the HotJava and Netscape browsers that support it have discovered a number of flaws that compromise security. These include implementation errors, unintended interactions between browser features, differences between the Java language and bytecode semantics, and weaknesses in the design of the language and the bytecode format. Research also examines the underlying tension between the openness desired by Web application writers and the security needs of their users and explores ways both might be accommodated.

National Information Assurance Partnership (NIAP)

NIST and NSA are partners in the National Information Assurance Partnership (NIAP), a program to enhance the quality of information security products and increase consumer confidence in those products that have been objectively evaluated. To help businesses and consumers choose commercial off-the-shelf (COTS) computer security products ranging from firewalls to database management systems, the two agencies are developing a program to ensure that these information technology products meet international standards. The program will be centered on the Common Criteria, an international standard for computer security products.

Common Criteria-based evaluations will take place in accredited private-sector laboratories, and NIAP will validate the results. NIAP will issue Common Criteria certificates, which will be recognized by other signatory countries in the Common Criteria Mutual Recognition Agreement.

The goals of the program include:

- ❏ Operating a Common Criteria-based evaluation system
- ❏ Providing for security evaluations in private-sector laboratories
- ❏ Ensuring that these evaluations meet consistent standards, resulting in increased confidence in the products
- ❏ Increasing availability of evaluated products
- ❏ Creating a climate conducive to the export of the products

- Developing high quality, cost effective, public domain test methods and tests

Role Based Access Control (RBAC)

In the increasingly complex IT environment, careful and correct specification of access control rules for access to online documents, capabilities, or systems is both critical and increasingly difficult, since traditional methods focus on individual users, files, or other system objects. In the real world, access is managed based on the role or roles that a user assumes in the course of work. NIST has pioneered a new access control model, Role Based Access Control (RBAC), that better meets the needs of user organizations. NIST is implementing this model in a number of environments, including a Web-based application.

NIST software technologies and standards

NIST works in several areas to create software development and analysis tools, testing technologies, and standards:

- *Software quality.* NIST is developing models, methods, and tools for tracing software processes to variables and resources, helping industry to improve the quality of software development and maintenance. Topics include formal methods, semantic correctness, performance assessment, and benchmarking.
- *Software analysis.* NIST researchers are developing tools for static and dynamic software analysis, focusing on measuring conformance to specifications and diagnosing causes of deviations from specifications. Initial R&D will be conducted on static analysis tools for program slicing, generating testing paths, and on object classes to detect pre- and post-condition violations in Web applets. NIST is developing experimental software designs and standard reference software with known errors for measuring the effectiveness of software development and testing methods.
- *Software assurance.* NIST provides technologies to produce high integrity, affordable software for productive use. NIST will provide guidance to establish fundamental life cycle processes to develop and maintain quality software and advanced development, evaluation, and measurement technologies to address specific assurance problems.
- *Conformance testing.* NIST is developing performance testing scenarios, testing procedures, and test suites to help industry, the user community, and testing laboratories with conformance standards testing. NIST works with other standards organizations to capture and incorporate conformance criteria early in the test cycle.
- *Software standards.* NIST makes technical contributions to standards-making bodies, representing the interests of the Federal user community, and serves as liaison with standards committees. NIST developed and maintains the online retrieval system for Federal Information Processing Standards (FIPS).

Telemedicine and secure patient records

In FY 1999, NLM continued to support research in technologies for storing and transmitting patients' medical records while protecting the accuracy and privacy of those records. Current projects promote the application of HCS technologies to healthcare, including telemedicine collaboration technologies to allow healthcare providers to deliver realtime treatment to patients in remote locations.

HCS national research agenda

In addition to the base research program already underway, the HCS agencies are developing an HCS national research agenda for new research in technologies to provide assured construction of high confidence systems. This agenda focuses on the critical information technologies needed to address challenges such as increased reliance on software and on a commodity technology base, increased scale and complexity, stress due to system performance demands, demand for interconnectivity, rush to market, and threat.

The goals of the agenda are to:

- ❏ Provide a sound theoretical and technological basis for assured construction of safe, secure systems
- ❏ Provide software, hardware, and system engineering tools that incorporate ubiquitous, application-based, domain-based, and risk-based assurance
- ❏ Reduce the effort, time, and cost penalty of system assurance activities
- ❏ Provide a technology base of public domain, advanced-prototype implementations of high confidence technologies to enable rapid adoption
- ❏ Provide measures or other evidence justifying confidence

To accomplish these goals, the agenda proposes research in HCS foundations to provide supporting theory and a scientific basis for achieving high confidence in safety-critical, security-critical, and other high-consequence systems. It proposes research in HCS design, tool, and language technologies for building assurance into systems; includes an engineering and experimentation component to provide reference implementations of components and system classes that typically require high confidence; illustrates HCS technology implementation and aids adoption; and provides for evaluation. The agenda includes pilot projects and demonstrations that will apply the HCS technologies to user domains, addressing agency systems of mission scale and importance.

The preparation of the HCS research agenda contributed to the design of the IT^2 initiative, of which it is a part.

The High Confidence Systems Working Group will dedicate the HCS National Research Agenda to the memory of Andy Arenth of NSA, who served as the Working Group's Co-Chair. Andy skillfully energized and guided the group in establishing a unified vision of the end goal. His contributions and leadership are gratefully acknowledged and will be sorely missed.

Human Centered Systems

HuCS R&D focuses on improving technologies that enable humans, computing systems, and information resources to work together more effectively and transparently. Federal investments in HuCS research lead to increased accessibility and usability of computing systems and communications networks. HuCS researchers endeavor to:

- Improve the ability to create and use worldwide information systems with enormous amounts of diverse content
- Increase the effectiveness and comfort with which people can exploit computers
- Extend computer availability to everyone in our society
- Create environments where sensors and mobile equipment adapt to humans, resulting in more rapid and effective completion of work tasks

Scientists, engineers, physicians, educators, students, the workforce, and the general public are all potential beneficiaries of HuCS technologies.

FY 2000 focus areas include:

- Knowledge repositories, information agents, and digital libraries that enable people to better collect, manage, analyze, summarize, and present data and images from complex databases
- Increased research data collection and annotation for training systems and for experimental evaluation
- Multimodal human-computer interfaces — including speech recognition tools, audio interfaces, and optic and haptic (touch-sensitive) devices — that help provide universal access to information technologies and enable people, regardless of their physical ability, education, or culture, to communicate with computers by sight, sound, touch, and gesture
- Visualization, virtual reality, and 3-D imagery tools that enhance human perception and understanding in simulation and problem-solving environments
- Collaboratories that facilitate knowledge sharing, group decision-making, control of remote instruments, and data sharing across large distributed systems
- Multilingual technologies that enhance foreign document and speech-to-speech translation, interpretation, and understanding, including collaboration across national boundaries

Knowledge and Distributed Intelligence (KDI)

Recent advances in computing power and Internet connectivity are reshaping relationships among people and organizations and transforming the processes of discovery, learning, and communication. These advances are creating unprecedented opportunities to provide rapid and efficient access to enormous amounts of knowledge and information, to study vastly more complex systems, and to increase our understanding of learning and intelligence in these systems.

To help exploit these opportunities, NSF's Foundation-wide Knowledge and Distributed Intelligence (KDI) program seeks to accelerate technological innovation and the adoption of technologies into society. Anticipated benefits of KDI research include:

- ❏ Deep, far-reaching scientific discovery
- ❏ Increases in scientific productivity and in the timeliness and quality of the results
- ❏ Increased ability to handle problems of greater complexity, scale, and structure
- ❏ The creation of new scientific and engineering communities to exploit novel discoveries
- ❏ Enhancements in science and engineering education through development of richer learning tools, technologies, and environments, and more universal access to these resources and tools
- ❏ Enhanced understanding of the processes and results of learning
- ❏ A more complete understanding of the fundamental processes of distributed intelligence in natural and artificial systems
- ❏ An understanding of the legal, ethical, and societal implications of the increased capability to gather and access information
- ❏ Enhanced ability to communicate and transfer new understanding and technological innovations to society
- ❏ Advances in statistical data reduction, data visualization, data mining, and data organization for retrieval to use vast stores of data
- ❏ Improved methods to express, compute with, and evaluate different types of uncertainties in real-world data

Knowledge Networking (KN)

In FY 1999, KDI focused on Knowledge Networking (KN), Learning and Intelligent Systems (LIS), and New Computational Challenges (NCC). HuCS R&D was primarily conducted in KN. NSF's LIS efforts are described in the ETHR section of this book, while NCC efforts are described in the HECC section.

The goals of NSF's KN efforts are to (1) understand the fundamental processes through which knowledge is created, communicated, validated, and valued in distributed information systems, and (2) improve the technical, social, educational, and economic performance of knowledge generation and use, collaborative computation, and remote interaction.

KN supports multidisciplinary research to develop and employ the next generation of communication networks, associated information repositories, collaborative technologies, and knowledge management techniques to gather, create, dis-

Human-computer interface researchers at Rutgers University have developed unique force feedback transducers to enable computer control and monitoring of graded exercises for extremities recovering from traumatic injury. The transducers, embedded in tight-fitting gloves for hands and boots for feet, apply programmed forces and movement to the injured extremities for prescribed durations. Medical specialists supervise rehabilitation remotely over computer networks, and patients have access to a library of computerized rehabilitation exercises. The photograph on the left illustrates the Rutgers University force feedback system. The figures on the right show a patient undergoing remote rehabilitation. This research is supported under the NSF STIMULATE initiative and the Aid to Persons with Disabilities program.

tribute, use, and evaluate knowledge in new and more secure ways. These activities include research on the human, behavioral, social, and ethical dimensions of knowledge networking.

STIMULATE

The Speech, Text, Image, and MULtimedia Advanced Technology Effort (STIMULATE) is a multiagency collaboration among researchers funded by NSF, DARPA, and NSA that seeks to improve human-computer interaction. Researchers focus on creating new, innovative computer interfaces using multiple languages and modalities, such as gestures, facial expressions, handwriting, images, and video. Multimodal devices developed through STIMULATE will help people with disabilities, language challenges, and reading deficiencies gain access to computing and communications technologies.

Digital Libraries (DL) Initiative, Phase Two

The original Digital Libraries Initiative (DLI), begun in 1995, was a joint four year collaborative program involving NSF, DARPA, and NASA. Its goal was to advance methods to collect, store, organize, and use widely distributed knowledge resources that contain diverse types of information and content stored in a variety of electronic forms.

Taking a walk downtown in "Virtual Los Angeles." The goal of this NSF-supported project at UCLA is to help solve civil engineering and urban design- and urban planning problems.

Currently in Phase Two, DL is now a multiagency effort of NSF, DARPA, NASA, NLM, the Library of Congress, and the National Endowment for the Humanities conducted in partnership with the National Archives and Records Administration, the Smithsonian Institution, and the Institute of Museum and Library Sciences. DL Phase Two will develop the next generation of digital libraries in such areas as education, engineering and design, Earth and space sciences, biosciences, geography, economics, and the arts and humanities; advance the use and usability of globally-distributed networked information resources; and encourage new and existing communities to focus on innovative applications areas. DL Phase Two will also address the digital libraries life cycle from information creation, access, and use, to archiving and preservation, and will study the long term social, behavioral, and economic effects of digital libraries in human activities such as research, education, commerce, defense, health services, and recreation.

A special feature highlighting the Digital Libraries Phase Two Initiative begins on page 80.

Virtual Los Angeles: A realtime visualization system for large scale urban environments

NSF-funded researchers on UCLA's Urban Simulation Team have developed "Virtual Los Angeles" — a long term urban simulation project that combines custom simulation software and realtime database technologies with efficient modeling methods. The simulation is derived from aerial photographs of Los Angeles with street level imagery and 3-D geometry, city engineering maps, internally generated vegetation libraries, and numerous site visits. Accurate to the level of graffiti on walls and signs in windows, Virtual Los Angeles is a realistic 3-D visual model of an extremely dense urban environment, containing data covering more than 4,000 square miles. It allows users to fly or walk through the virtual city, complete with details like trees, street signs, and other visual landmarks.

The Virtual Los Angeles database is a long term project that can be used to help address a multitude of real-world civil engineering, urban design, and urban planning problems, including the development of a "Virtual World" data server displaying a realtime virtual model of the entire Los Angeles Basin. Researchers are working on applications including a 3-D in-car navigation system, a helicopter

flight training system, and other transit and traffic management tools. The Los Angeles Fire Department and other emergency response teams will use these tools as an aid to emergency response — for example, to learn whether adjoining buildings are higher or lower than available rescue equipment can reach.

Virtual reality technology for simulating medical procedures

In medical procedures, the ability to visualize large amounts of data is often crucial. In clinical care, the ability to "see" areas deep inside the body—especially the brain—can determine whether a surgical procedure or other therapeutic intervention can be successfully performed with minimum patient trauma. These fundamental requirements are increasingly being aided by virtual reality technology that provides realistic visualization of data in real time. NIH/NCRR researchers use virtual reality technologies to simulate operations and other medical procedures, and to interface with high resolution instruments like atomic force microscopes. They are also building and evaluating surgical collaboratories that allow remote users to gain high speed access to scientific instruments for basic research, molecular visualization, and surgical and other therapeutic interventions such as radiation treatment planning.

In FY 1999, NCRR conducted research in augmented reality, a form of virtual reality, to support image-guided neurosurgery, initiating eight collaborative demonstrations at resource centers covering a wide range of technology R&D. Building on these efforts in FY 2000, NCRR will conduct additional R&D focusing on new uses for emerging virtual reality technologies and environments.

Robotic surgery

NSF-supported research at Johns Hopkins University focuses on Computer-Integrated Surgery (CIS), a rapidly expanding field where advances in computing and engineering technologies are helping to overcome the limitations of traditional surgery. By extending human surgeons' ability to plan and implement surgical interventions more accurately and less invasively, CIS systems address a vital national need to reduce surgical costs and improve both the clinical outcomes and the efficiency of healthcare delivery.

According to the November 1998 *Scientific American*, "coming changes in hospital technology may within a few years make television dramas...less exciting but should help operating rooms become healthier places for patients. The tense crowd gathered around the operating table is likely to be replaced by a silent robot cutting and suturing through an incision less than an inch long. The surgeon doing the operation will…be…seated at a console peering deep into the body through a virtual reality headset that combines the actual image of the patient with up-to-date diagnostic images showing the precise location of the problem. He may even have rehearsed the operation in cyberspace before the incision was made. The manipulators that control the surgical tools will be capable of the tiniest, most precise motions — and they will compensate for any tremors in the surgeon's hand."

To foster advances in robotic surgery, NSF is establishing a new Engineering Research Center in Computer-Integrated Surgical Systems and Technology. Headed by Johns Hopkins University, this collaborative project involves Johns Hopkins School of Medicine, Johns Hopkins Applied Physics Laboratory, MIT's Project on Image Guided Surgery, the Surgical Planning Group at Harvard University's Brigham and Women's Hospital, and Carnegie Mellon University and its affiliate, Shadyside Hospital. The Center draws upon experts in computer science and robotics; electrical, mechanical, and biomedical engineering; and physicians specializing in fields such as radiology, neurosurgery, urology, orthopedics, and ophthalmology.

NSF-funded HuCS research projects at Johns Hopkins include robotic steady hand assistance for neurosurgery and microsurgery; robotic camera and instrument holding and monitoring for laparoscopic surgical procedures; treatment planning for and automation of pattern placement therapy for minimally invasive cancer treatment; image processing techniques for the next generation RoboDoc® system for hip replacement surgery; automation of percutaneous needle insertion for kidney surgery; and modeling and instrument tracking for orthopedic and craniofacial procedures. MIT research includes using medical scans to model a patient's internal anatomy. One Carnegie Mellon project is the HipNav computer-based surgical assistant that helps more accurately plan and place the socket portion of a hip implant. The goals of this research are to develop a family of systems that extend the human ability to operate on very small structures inside the body without large incisions, assist in routine surgical tasks as "robotic residents" operating under a surgeon's supervision, help surgeons rehearse and optimize surgical plans based on preoperative patient-specific data, and permit surgeons to assist in procedures conducted in remote areas.

NASA's "software scalpel" and virtual reality tools

Researchers at the NASA Ames Center for Bioinformatics are developing a "software scalpel" that—combined with clear, accurate, 3-D images of the human head—is helping doctors practice complex facial reconstructive surgery and visualize its outcome more accurately. Using the new Virtual Surgery Cutting Tool software, a physician wearing 3-D glasses can see an image of a patient's head from all angles on a computer monitor or on the surface of a large "immersive virtual reality workbench." This computerized reconstruction can also be applied to mastectomy patients who require breast reconstruction and children who need reconstructive surgery to correct head and facial deformities.

In the future, virtual reality will allow surgeons to rehearse many complex procedures before operating, and will serve as a powerful teaching tool for medical students. A digital library of computerized "virtual patients" will be created so that physicians can share information about infrequently performed medical procedures.

NASA's computerized breast cancer diagnostic tool

NASA researchers are developing and testing a smart probe for use in breast cancer detection and analysis. A spin-off from a computerized robotic brain surgery "assistant," this probe is designed to see a lump, determine by its features if it is cancerous, and predict how the disease may progress, aiding doctors in expediting diagnoses and helping them to suggest individualized treatment. The predictive capability uses trainable neural net software to look for telltale characteristics or patterns.

Battlefield awareness

DARPA battlefield awareness R&D develops and demonstrates technologies to extract information from imagery, understand news broadcasts, and integrate information from heterogeneous data sources.

As part of DARPA's Text, Radio, Video, and Speech (TRVS) program, researchers are developing, demonstrating, and evaluating automated news information processing techniques that will:

- ❏ Determine the structure of and automatically index data from broadcast information sources with different speakers, topics, and distractions with 85 percent reliability
- ❏ Extract facts from diverse media with 90 percent accuracy and transcribe spoken language with less than a 20 percent word error rate

- Transcribe audio portions of radio and TV news broadcasts in multiple languages with a goal of handling two languages, Mandarin and Spanish, with less than a 35 percent error rate in FY 1999

Researchers are developing a plug-and-play framework to extract content from diverse TRVS inputs to provide for effective interaction with the extracted information. DARPA's long term goal is to intercept spoken language communications and automatically process the data to derive information from enemy sources. By processing the data through analysis algorithms, the military can extract key facts about the enemy's plans and infer the enemy intent.

Mobile Autonomous Robot Software (MARS)

DARPA is funding R&D in mobile autonomous robot software (MARS) that enables mobile robots to perform tasks such as the safe, reliable, realtime operation of unmanned military systems in dynamic, realistic, and unstructured operational environments. These tasks may include platform mobility, navigation, obstacle avoidance, payload operation, and human-robot interaction without dependence on a human operator. MARS techniques will aid DoD in domains such as reconnaissance, surveillance and target acquisition, countermine and explosive ordnance disposal, force protection and physical security, and logistics support operations.

Robots employing this software will exhibit:

- Reprogrammable, goal-directed behaviors that exploit symbolic, human-supplied information
- Reflexive, sensor-mediated behaviors that are stable in complex, uncertain, and dynamic environments
- A high degree of adaptability, including a facility for learning
- The ability to operate safely and reliably in close proximity to humans

MARS will develop reusable software for autonomous mobile systems that encompasses:

- A software framework that lets robots synthesize deliberative (symbol mediated) and reactive (sensor mediated) control
- Software composition methods that incorporate both hand-coded programming and automated learning-driven coding so robots can function better in unpredictable environments

Information visualization

Information visualization is the process of transforming inherently spatial data and information into a visual form, allowing users to understand the information better. As part of its HuCS R&D, NIST is examining the utility and feasibility of 3-D visualization and computer graphic techniques to access, manipulate, and exchange complex information.

In FY 1999, NIST is developing evaluation methods and test corpora for measuring the usability and scalability of interfaces for visualizing a statistical text retrieval system that accesses large collections of data. NIST researchers are also developing a visualization system that monitors user navigation paths through a Web site, helping usability engineers to determine the effectiveness of the site's user interface. NIST plans to investigate the use of large screen visualization of docu-

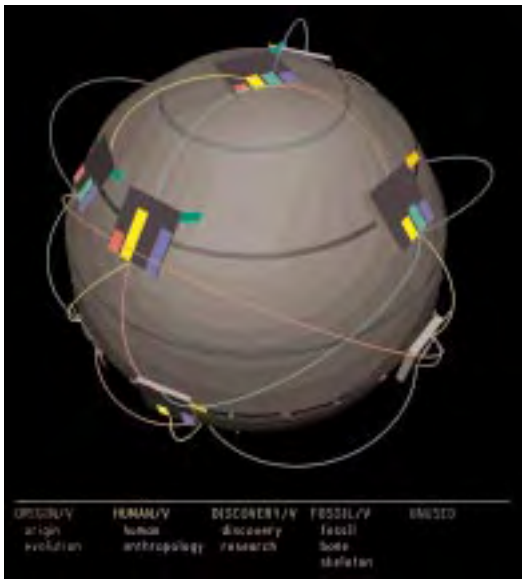

NIST is examining the usefulness and feasibility of employing 3-D visualization and computer graphic techniques to access, manipulate, and exchange complex information. In the global 3-D model to the left, clusters of conceptually similar documents, and the relationships among them, are portrayed.

ments and related information to support effective human-computer interaction in an electronic "smart" meeting room.

Visualization and virtual reality for collaboration and manufacturing

NIST is collaborating with industrial partners in applying 3-D visualization to accelerate the manufacturing process. NIST researchers are analyzing the usability and performance capability of a number of virtual environment and Web visualization and collaboration tools, including the Virtual Reality Modeling Language (VRML) to visualize distributed collaborative interactions in high performance "smart" spaces, facilitating design and engineering analyses for remote participants. These also include commercial off-the-shelf software to create visualizations of the factory floor assembly line and to support parts design and assembly.

Virtual Worlds

NOAA's Pacific Marine Environmental Laboratory (PMEL) and Old Dominion University are continuing their partnership to provide scientists with general purpose Java tools, enabling them to combine graphical objects created from commercial and customized routines, such as GIS and Matlab, to produce their own integrated, customized Virtual Worlds for interactive exploration on the Web.

Virtual Worlds constructed by scientists with this Java toolkit will include graphical objects (e.g., contour slices, vectors, bathymetry, topography), 3-D domains, specification of pre-defined viewpoints, proximity sensor triggers, user touch sensor triggers, and audio and video capabilities as appropriate. PMEL and Old Dominion will also explore realtime generation of Virtual Worlds from archived data.

NOAA collaborations using NCSA's Habanero

NOAA is using advanced distributed network collaboration tools, such as NCSA's Habanero, to enhance research collaborations needing access to NOAA's data holdings. Habanero is a collaborative framework and set of applications that can be used to create and work in shared applications from remote locations over the Internet. It enables developers of groupware applications to build powerful collaborative software in a reduced amount of time by allowing them to create or convert existing applications into collaborative applications. Useful Habanero tools

include a whiteboard, a collaborative text editor, a Java graph, and various chat utilities.

Tools are evolving to provide on-the-fly production of integrated browser images from profile datasets stored at two or more geographically-distributed data archives. This allows two or more scientists at different locations to view these images and have annotations simultaneously appear on all collaborating scientists' desktops. All collaborators will be able to interactively post messages and change the data being displayed.

NOAAServer

NOAA has implemented a prototype application that contains robust collaborative features such as text and line annotations, highlighting of individual profiles, display of data values, and support for multiple users, and plans to expand the prototype to include user feedback and access to multiple NOAAServer data servers. This project leverages the collaborative tools developed at NCSA and the ongoing development of graphics that integrate data from multiple archives at different locations under the NOAAServer project.

Speech recognition technology

NIST collaborates with DARPA and NSA on speech technologies by developing benchmark tests to track technology development and measures of progress. In its work with DARPA, NIST researchers focus on continuous speech recognition and on understanding the large and difficult vocabulary used in radio and TV news broadcasts. For NSA, NIST is implementing test protocols for transcribing conversational telephone-line speech and speaker recognition technologies. These tests involve scientists in the United Kingdom, France, and Germany.

In FY 1999, NIST is working on a testbed to develop and test spoken-language based information retrieval technologies. Research focuses on user interactions and on experiments with dialogue management protocols. FY 1999 and FY 2000 speech technology research projects focus on automatic speech recognition and transcription, spoken document retrieval, topic detection and tracking, and new information extraction measures and tests.

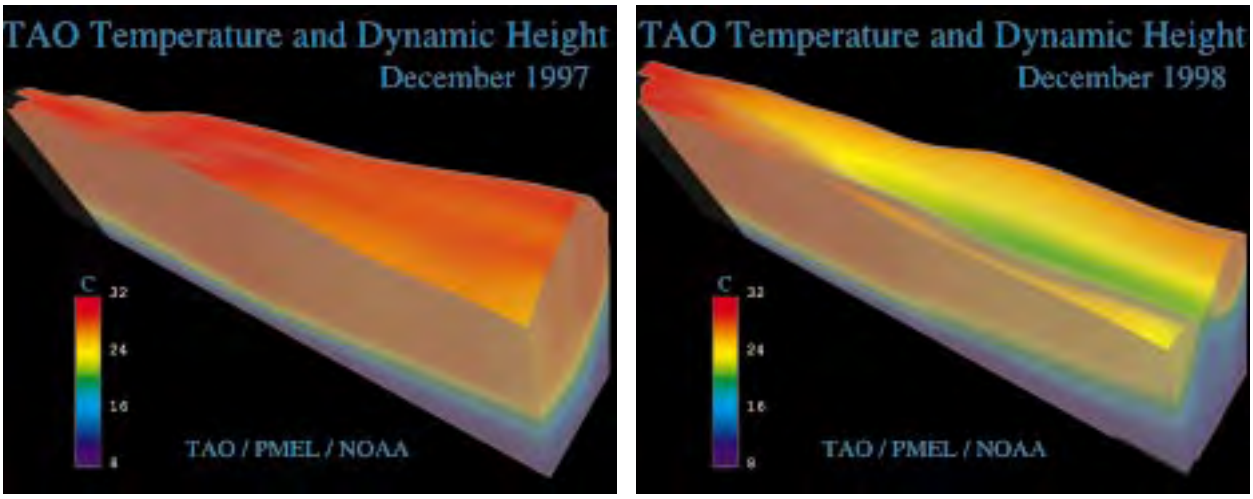

NOAA/PMEL and Old Dominion University have developed three dimensional representations of NOAA El Niño data in VRML and a prototype, customized, complex Virtual World containing these 3-D graphical objects. Users can interact with the graphical objects and with the virtual world via the Web. Depicted above is the Equatorial Pacific sea level colored with ocean temperatures from NOAA's network of El Niño buoys in the tropical Pacific Ocean. These 3-D VRML renderings show El Niño (December 1997) and La Niña (December 1998), along the equator from 8N to 8S, viewed from the Andes mountains in Peru.

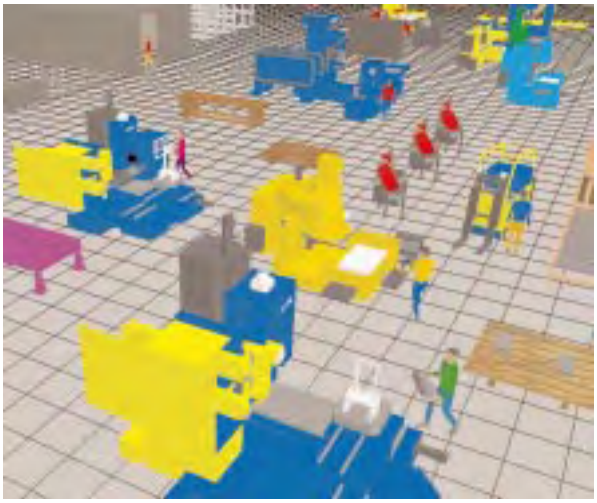

Used to model manufacturing assembly operations, simulation tools will help improve the productivity of engineers who plan manufacturing operations. The graphic on the left demonstrates an ergonomic modeling software system where a virtual human is assembling the base of a miter saw using assembly instructions specified in a NIST-developed computer language. Ergonomic simulation environments will be used to validate assembly processes, perform ergonomic analyses, and train workers on actual assembly operations.

The Virtual Machine Shop is being used to develop, test, and evaluate interfaces to manufacturing simulation environments. The graphic on the right shows a NIST-developed machine shop simulation where raw materials are prepared, cutting tools are set up and loaded into machine tool magazines, parts are machined, and materials are handled. NIST is developing interfaces between manufacturing simulation software and software that supports engineering and manufacturing activities.

Spanish language interfaces

To help support Spanish-speaking individuals in the emerging information society, researchers at the Oregon Graduate Institute's Center for Spoken Language Understanding (CSLU) and researchers at the Universidad de las Americas-Puebla—jointly funded by NSF, DARPA, Office of Naval Research (ONR), Fonix, and Intel—collaborated on a Spanish version of a spoken language toolkit for human computer interaction.

The CSLU Speech Toolkit is a comprehensive software environment for research, development, and education of spoken language systems that integrates speech recognition, speech synthesis, facial animation, and speaker recognition, and features authoring and analysis tools enabling quick and easy development of desktop- and telephone-based speech applications. The software is available free of charge for research and education at not-for-profit institutions.

The Spanish language CSLU Speech Toolkit, which incorporates Spanish speech recognition and synthesis, enables faculty, staff, and students to develop Spanish-spoken language systems and to teach courses in Mexico, accelerating work in speech technology infrastructure, research, system development, and education.

Systems Integration for Manufacturing Applications (SIMA)

At NIST, SIMA researchers are developing technologies to improve the interfaces between engineering and manufacturing software systems and between people and the applications themselves. Through the development of authoritative, Internet-accessible repositories of scientific and engineering knowledge and the development of computerized collaboration and remote operation tools, NIST supports interoperability among advanced manufacturing systems applications.

To assist businesses in plant and factory design, NIST R&D is enhancing commercial modeling and simulation software systems with programmable human ergonomic models, new human modeling programming languages, and the integration of physics-based models of machine and object behavior. These technologies will enable businesses to build "simulated factories" with prototyped plant layouts, material and component routings, and human factors prior to committing significant investments in plant redesigns or new factories. The results of these efforts are captured in consensus standards efforts conducted through standards organizations like the International Organization for Standardization.

NIST is continuing its work on a manufacturing collaboratory to provide manufacturers, distributors, and researchers with the structured methods, guidelines, and standards needed to implement collaboratory technologies in business applications. This has already helped research in robotic arc welding, which requires asynchronous and synchronous collaboration for local and remote support of diverse data formats such as video, still images, audio, text, database records, and image annotations. One product of this work will be a collaborative tool prototype with annotation capabilities that synchronize various time-indexed data streams, such as audio, video, and log data.

MMC Collaboratory

Linking instrumentation at six DOE user facilities nationwide, DOE's Materials Microcharacterization Collaboratory (MMC) enables scientists and researchers to access state-of-the-art tools remotely to study the structure and composition of metals, ceramics, and alloys at the atomic and molecular level. Scientists at other DOE and university labs can operate two electron microscopes in the National Center for Electron Microscopy at Lawrence Berkeley National Laboratory in real time on the Internet. For example, an aluminum-lead alloy can be remotely heated to lead's melting point, allowing observation of phase transformation using a Kratos EM-1500 transmission electron microscope — which uses the highest accelerating voltage available in the U.S. — with four angstrom resolution. (One angstrom is one ten-billionth of a meter. Most atoms are approximately two angstroms apart.) This work is helping materials scientists across the country better understand such manufacturing processes as soldering, welding, and brazing. The uniqueness of these demonstrations lies in the ability to perform dynamic realtime experiments with the microscope over the Internet by bringing high performance computing close to the instrument site for online analysis, and by using such analysis for closed-loop control.

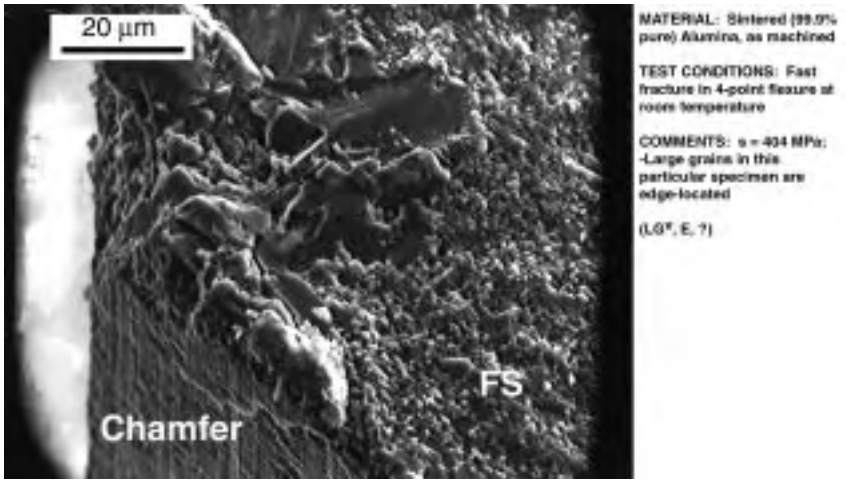

The NIST Ceramics WebBook provides scientists and engineers with Internet access to an extensive, authoritative information database. This image of a large-grain defect in ceramic material is located in the flaw catalog in the WebBook's Fractography section. The strength of many advanced ceramics reflects flaws present in the material and the material's intrinsic fracture toughness. Strength-test results must be interpreted in the context of these flaws whether this testing has been performed for quality control, materials development, or design purposes.

Biomedical technology collaboratories

NCRR has initiated eight collaboratory testbeds in NCRR-supported resource centers. These testbeds are designed to take full advantage of emerging Internet capabilities, such as increased bandwidth, security, access, and reliability, while testing technologies including instrument control from remote locations (NanoManipulator, electron and optical microscopes, and National Medical Resource [NMR] spectroscopes), collaborative software design, complete crystallographic experiments over the Internet, remote access to sophisticated visualization systems, and collaborative drug design and protein engineering.

Electronic Notebooks for collaborative environments

DOE is continuing its R&D on Electronic Notebooks used in collaborative environments. These notebooks — which require interoperability, integration of multiple technologies, and use of heterogeneous computing and communications systems — aid researchers in recording and sharing experiments that use remote instruments, serving as a medium for collaborative scientific inquiry and engineering design.

Security technology is also necessary to help detect possible tampering with data when notebooks are used as patent records for inventions or proof of conformance to procedures. Future R&D will focus on ease of use, privacy, security, and tools to coordinate access to multiple notebooks.

Unified Medical Language System (UMLS)

With the rapidly growing number of computerized database resources offering bibliographic, full text, and factual data via the Internet, it is often difficult for users to locate and process needed information due to unforgiving access protocols and retrieval languages that differ between databases. In biomedicine, the disparity in biomedical terminology used to describe related concepts in different machine readable files prevents practitioners and researchers from retrieving and integrating biomedical information from sources such as the biomedical literature, clinical records, medical data banks, and expert knowledge bases.

NLM is developing its Unified Medical Language System (UMLS), a long term project to compensate for dissimilar ways that related information is classified in different automated systems. Intelligent-agent-mediated gateways will provide users with a single point of access to information, freeing them from having to master multiple information sources.

NLM researchers are continuing to develop and deploy new capabilities for automatic source selection and for retrieving and sorting information from multiple databases available via NLM's Web site and its Internet Grateful Med (IGM), PubMed, and TOXNET retrieval services. NLM will focus on facilitating access for health professionals, researchers, and the public. FY 2000 plans include R&D in the use of UMLS knowledge sources and integrating multiple multimedia information sources into computer-based patient record systems.

Medline Plus

NLM has developed Medline Plus, an Internet site with information on common diseases and conditions, as well as links to Web sites, medical dictionaries, and reference tools used by medical librarians. Medline Plus developers are now packaging biomedical information and references so they are useful to people without medical degrees or backgrounds. Medline Plus contains information on approximately 50 health topics including AIDS, cancer, diabetes, eating disorders, Parkinson's Diseases, smoking, and tuberculosis. NLM is increasing Medline Plus coverage to some 400 health-related topics and is enlisting the aid of 207 public libraries to help make the site more user-friendly.

Researchers at Columbia University are generating a "gallery" of 3-D photographic quality color models of anatomical structures obtained from the Visible Human datasets. The 3-D visualizations are fully interactive and can be used and reused for undergraduate and graduate medical teaching, patient education, and in industry. The knowledge base design builds on work done at Columbia using NLM's UMLS. Pictured is a snapshot from a Quicktime movie that displays a rotational 3-D image of the bladder, prostate, and urethra.

The Visible Human Project

New computer-based technologies provide unprecedented opportunities to supplement traditional 2-D biological images with dynamic 3-D images that users can view, rotate, and reversibly dissect. As part of its Visible Human Project, NLM is building and evaluating digital image libraries of human anatomical structures by integrating advanced computing and communications technologies with medical imaging systems for computer tomography (CT) and magnetic resonance (MR) imaging. NLM is investigating both compression technologies for storing such massive datasets and communications technologies for transmitting them faster over the Internet.

NLM's FY 1999 activities in this ongoing project include segmenting and labeling male and female anatomical structures; converting the thorax sections from the Visible Human image file format into a national 3-D image file format; experimenting with interfacing the visualization subsystem with industrial image rendering software; initiating development of a platform-independent radiological atlas for epidemiological and general biomedical research; and developing tools to organize and statistically analyze information from text and X-ray image databases.

FY 2000 plans include beta testing the Visible Human anatomical database — including the 3-D file format — over the Internet, and completing the development and testing of an online, Java-based, platform-independent radiological atlas for biomedical research.

Universal access

In FY 1999, NSF began a multi-year research effort to help individuals find, manipulate, and use information in an efficient and comprehensive manner. The primary objective of this universal access research is to empower people with disabilities to participate in the emerging information society. The research also aims to benefit the entire Nation by advancing technologies that allow all people to become more productive through better and more effective use of multimodal information resources.

To achieve these goals, NSF plans to develop new models and architectures that emphasize interface usability and speed, define semantic structures for multimedia information to support cross-modal I/O, develop specific solutions to address the special needs of large disabled communities, and conduct experimental studies to evaluate the success of attempts to provide access. Research topics

include alternatives to the traditional desktop interface, access to auditory information in textual form for hearing impaired users, new I/O techniques for motor impaired users, and large scale tactile access to graphical information for visually impaired users.

NSF's High Capability Applications for the Individual will address society's need for universal, easy to use access to information resources, powerful methods of presenting information for ease of understanding, and customized "information space" for personal use. Examples include digital libraries and medical information servers. These and other applications will drive the enabling research of computing systems, human centered systems, and networking, communications, and the convergence of computing and communication, and may lead to fundamentally different ways of solving problems. The applications come from the physical and biological sciences, geosciences, social and behavioral sciences, and engineering.

Usability

As more companies and institutions rely on Web sites for electronic commerce, information dissemination, and networked applications, the ability to design usable Web sites is critical for increasing productivity and sales and lowering costs and user frustration. NIST is developing tools to automate and support the design and evaluation of Web sites. These tools help Web designers analyze sites for potential usability problems during the design phase and assist usability professionals in expediting their evaluations. In FY 1999, NIST expanded its automated analysis capabilities based on numerous guidelines and current Web-usability research, improved instrumentation tools, and conducted research on the effects of cultural icons on usability.

Fingerprint and mugshots standards

In the late 1960s, NIST partnered with the Federal Bureau of Investigation (FBI) to create a computerized approach for automatically comparing and matching fingerprint images. This work has expanded to include developing standards and specifications for capturing and storing fingerprint and mugshot images, as well as creating fingerprint and mugshot image databases for testing fingerprint processing techniques. NIST's R&D resulted in standards and specifications for the quality, format, and transmission of electronic images and related data, such as identifying marks.

NIST has coordinated the development of official ANSI data and image interoperability standards, and has developed data exchange standards for interoperable systems and evaluation methods for fingerprint systems. NIST maintains the ANSI standard for the data exchange of FBI fingerprint and mugshot data and publishes Standard Reference Material CDs containing over 50,000 fingerprint image pairs and mugshot evaluation data used throughout the industry. NIST has also developed public domain Pattern Level Classification Automation System (PCASYS) and has distributed over 300 copies of this system.

In 1999, NIST began to develop a system in which a smartcard containing an individual's fingerprint and digital signature is used in combination with the individual's realtime scanned fingerprint image to verify authenticity of the creator of the electronic record. The smartcard, together with a registry of legitimate digital signatures and the electronic signing of the transaction, will help ensure that neither the authenticity of the record nor the data were corrupted.

Regional Technology in Education Consortia (RTEC)

In FY 1999, ED is conducting R&D to provide the education community with resources and skills to improve teaching and learning. ED continues to support the Regional Technology in Education Consortia (RTEC), which provide state and local education agencies with technical assistance in developing comprehensive

educational technology plans, help integrate technology into the learning and teaching environments, and assist educators in acquiring, maintaining, and using new technologies.

Regional Educational Laboratories

ED's Regional Educational Laboratories assist educators and policy makers in implementing effective school reform strategies by testing new approaches to teaching and learning, providing training and technical assistance, and disseminating information. This is accomplished in part through educational technology services, including telecommunications for rural areas that lack modern technology, and videos and CD-ROMS that highlight exemplary teaching and learning methods.

National Institute on Disability and Rehabilitation Research (NIDRR)

ED's National Institute on Disability and Rehabilitation Research (NIDRR) investigates technologies to improve the lives of people with disabilities. In FY 1999, four of NIDRR's Rehabilitation Engineering Research Centers (RERCs) are focused on HuCS-related R&D: information technology access, communication enhancement, universal telecommunications access, and telerehabilitation.

NIDRR supports ABLEDATA, a database of assistive devices for individuals with disabilities, and the National Rehabilitation Information Center's REHABDATA, a rehabilitation bibliography database. In FY 2000, NIDRR proposes three initiatives in information technology and telecommunications accessibility: a demonstration center to bring together producers, researchers, and consumers; several industry consortia to help build accessibility into new products; and a technology transfer program to explore the use of existing technologies in new applications.

Web Accessibility Initiative (WAI)

The World Wide Web Consortium (W3C) is partnering with organizations around the world to implement the WAI to ensure that new Internet technologies are available to all individuals regardless of physical limitations. Sponsored by NSF, ED's NIDRR, IBM and Lotus Development Corporation, Microsoft, NCR, and the European Commission's TIDE Programme under Directorate Generale XIII, the WAI is coordinating efforts in five areas related to Web accessibility: technology, guidelines, tools, education and outreach, and R&D.

In the technology area, WAI is addressing issues such as HTML coding and style sheets, multimedia, graphics, mobile access, and other data formats and protocols. Researchers are also developing accessibility guidelines for browsers, authoring and content tools, and tools for evaluation, repair, and proxy conversions. The program office facilitates WAI technical activity and coordinates partnerships among stakeholders in Web accessibility including industry, disability organizations, governments, and research organizations.

Digital Libraries Initiative, Phase Two

To help advance the digital revolution, the Digital Libraries (DL) Phase Two Initiative is building upon the successes of previous Federally supported DL research. The Initiative provides leadership in research to develop next generation digital libraries, to advance the use and usability of globally distributed networked information resources, and to focus on innovative applications. The DL Phase Two Initiative is jointly supported by NSF, DARPA, NLM, the Library of Congress, NASA, and the National Endowment for the Humanities, in partnership with the National Archives and Records Administration, the Smithsonian Institution, and the Institute of Museum and Library Sciences.

DL researchers are faced with the continued challenge of applying increased computational capacity and network bandwidth to large amounts of distributed, complex data and transforming that data into coherent, usable, and accessible information. DL researchers will 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; and study the interactions between humans and digital libraries in social and organizational contexts. The Initiative encourages partnering arrangements to create 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 also addresses the digital libraries life cycle from information creation, access, and use, to archiving and preservation.

Focus areas

DL focus areas are:

- ❏ Human-centered research
- ❏ Content and collections-based research
- ❏ Systems-centered research

Research topics include intelligent user interfaces; collaboration technologies and tools; methods, algorithms, and software leading to wide-spectrum information discovery, search, retrieval, manipulation, and presentation capabilities; efficient data capture, representation, preservation, and archiving; metadata; interoperability of content and collections; technologies, methods, and processes to address social, economic, and legal issues associated with the creation and use of digital collections; intelligent agents; advanced multimedia information capture, representation, and digitization; and open, networked architectures for new information environments capable of supporting complex information access, analysis, and collaborative work. Current DL projects include:

Many of poet and engraver William Blake's best-known poems were self-published by the author as "illuminated books," limited editions engraved and hand-colored by Blake himself, perhaps with the assistance of his wife. Due to the rarity of these volumes, most readers have never seen Blake's poems as he actually wished to present them. This problem is now being addressed with digital library technologies. Digital reproductions of many of Blake's illuminated books now appear in The William Blake Archive, which, according to its creators, is a free Web site "conceived as an international public resource that would provide unified access to major works of visual and literary art that are highly disparate, widely dispersed, and often severely restricted. The Archive integrates editions, catalogues, databases, and scholarly tools into an archive with fully searchable texts and images." The image to the left was featured in the January, 1999, issue of "D-Lib Magazine: The Magazine of Digital Library Research," which is produced by the Corporation for National Research Initiatives and is sponsored by DARPA on behalf of the Digital Libraries Initiative.

William Blake's The Book of Thel*, copy O, plate 1 (detail). Lessing J. Rosenwald Collection, Library of Congress. Courtesy of the William Blake Archive* <http://www.iath.virginia.edu/blake/>. *Used with permission.*

Tracking footprints through an information space: leveraging the document selections of expert problem solvers

At the Oregon Graduate Institute of Science and Technology, NSF-supported researchers are helping problem-solvers find information in large, complex information spaces without the distraction of redundant or irrelevant information. Their research focuses on healthcare, where patients' medical records are typically large, complex, geographically distributed collections of documents created by numerous healthcare professionals for divergent purposes over an extended period. Their approach is to capture and trace the information used by experts as they solve problems and exploit this information to assist others. This research is conducted by a cross-disciplinary team comprised of a medical doctor who is focusing on the information-seeking behavior of physicians and computer scientists who are extracting and using regularly structured information.

Trusted Image Dissemination (TID)

Since modern computing has greatly facilitated the use of information in the form of images, filtering images in addition to text has become essential. To help extend security and privacy protection to multimedia databases, researchers at Stanford University are focusing on TID technologies to provide image-filtering capabilities that complement traditional means of checking the contents of multimedia documents. TID will be used to restrict or screen information contained in images that are part of electronic patient records to avoid violations of security or privacy. TID research efforts focus on:

- Further development of an existing wavelet-based algorithm for searching medical image databases and development of techniques to retrieve digital images and relevant textual information from multimedia medical databases
- Extracting textual information from retrieved images
- Developing techniques to edit digital medical images automatically and adapting and developing tools to manually edit the images to omit identifying information
- Defining rules for protecting the privacy of medical images and implementing them in a security mediator
- Developing a Web customer interface for the security mediator

Automatic reference librarians for the World Wide Web

At the University of Washington, researchers are developing more powerful automatic reference tools to help people retrieve high quality information efficiently from the Web. The central objective is to create software agents that possess "reference intelligence," which assumes only a limited understanding of complex topics but relies on a sophisticated understanding of how and where to find information. This reference intelligence would emulate human reference librarians, who are often not specialists in a topic of inquiry (for example, computational fluid dynamics) but are experts in identifying relevant resources on the topic (such as The International Journal of Fluid Dynamics).

Medical informatics

In FY 2000, NLM plans to fund DL research in medical informatics, including projects that benefit healthcare consumers. The UMLS Knowledge Sources and the Visible Human dataset will be made available by NLM for use and testing in DL Phase Two projects.

International digital libraries

To help avoid duplication of effort, prevent the development of fragmented digital systems, and encourage productive interchange of scientific knowledge and scholarly data around the world, NSF supports a program on International Digital Libraries Collaborative Research in FY 1999. This program will contribute to creating information systems that can operate in multiple languages, formats, media, and social and organizational contexts.

The program's goal is to enable users to access digital collections easily, regardless of location, language, or format, and to enable broad use in research, education, and commerce. A global information environment requires research on:

- Methods and standards to ensure long term interoperability among distributed and separately administered databases for advanced retrieval of many kinds of information; for worldwide self-organizing databases and data mining; and for organizing and preserving domain-specific content
- The development of linked, compatible databases with inherently regional information, such as databases of geographic, botanic, agricultural, demographic, and economic data
- Technology for intellectual property protection in a global marketplace

NSF will fund U.S. participation in multi-country, multi-team projects to help foster long term sustainable relationships between U.S. and non-U.S. researchers and research organizations. Specific research areas include multi-lingual information systems, cross-language retrieval systems, language translation, and language teaching software; multinational digital libraries including sound, data, image, multimedia, software, and other forms of content; interoperability and scalability technologies for extremely large worldwide collections; metadata (data about data) techniques and tools; geospatial, environmental, biological, historical, and other information systems in which location is highly relevant; preservation and archiving of digital scholarly information, including technologies and procedures for long term information asset management; social aspects of digital libraries and cross-cultural context studies; use of digital libraries in educational technology at all levels; economic and copyright issues, including authentication, payment, rights formalism, trust, and fair use; electronic publishing and scholarly communication technologies, including collaboratories, online repositories, and new methods for distributing scientific knowledge.

Education, Training, and Human Resources

ETHR R&D supports Federal activities in computer- and communications-related education and training to advance learning technologies at all levels including K-12, community college, technical school, trade school, university undergraduate and graduate, and lifelong learning. R&D in learning technologies is needed to enable citizens to use the Nation's information infrastructure and to provide universal access to the resources necessary for efficient and effective education and training. This will lead to more knowledgeable and productive citizens who will use cutting-edge information technologies to maintain U.S. competitiveness in today's highly aggressive market environment.

ETHR R&D encourages and facilitates interagency collaborations for Federal education and training R&D and evaluation of advanced technologies for high quality, affordable software learning tools; information-based models of educational systems and learning productivity; research on information technology applied to learning and cognitive processes; and demonstrations of innovative technologies and networking applications. The ETHR Working Group coordinates Federal cooperation with schools and other educational and training venues.

Experimental and integrative activities, education, and training

In the 21st century, there will be an ever-increasing need for a workforce trained in using advanced high performance computing and communications. NSF's ETHR R&D activities focus on increasing the pool of people with the knowledge, skills, and insights to lead research in science and technology needed to make high performance computing and information processing more useful and to pursue fundamental knowledge in all disciplines of science and engineering. Activities include new course and curriculum development, and collaborative research in high performance computing and communication and information processing in undergraduate education.

PACI outreach

The Education Outreach and Training activities of NSF's Partnerships for Advanced Computational Infrastructure (PACI) partners focus on enabling all citizens to use emerging computing technologies to advance their ability to understand and solve problems in education, science, business, government, and society.

Learning and Intelligent Systems (LIS)

NSF supports undergraduate and graduate research to provide the human resources needed for information technology research, education, and industry. Through its KDI initiative, NSF is investing in Learning and Intelligent Systems (LIS). The LIS goal is to stimulate research that will advance and integrate concepts of learning and intelligence emerging from theoretical and experimental work in education, cognitive science, computer science, neuroscience, engineering, social science, and physical science. LIS encompasses studies of learning and intelligence in systems, including (but not limited to) the nervous systems of humans or other animals; networks of computers performing complex computations; robotic devices that interact with their environments; social systems of human or non-

human species; and formal and informal learning situations. LIS also includes research that promotes the development and use of learning technologies across a broad range of fields.

DOE's Computational Science Graduate Fellowships

The Computational Science Graduate Fellowship Program encourages talented college students to undertake study and research in computational science accompanied by practical work experience at DOE research facilities. The program encourages students with outstanding academic records to continue their graduate studies in preparation for careers in computational science.

Biomedical informatics training grants

To help address the current shortage of biomedical professionals trained in using modern computing and telecommunications systems, NLM continues to expand its successful predoctoral and postdoctoral grants program for career training in biomedical informatics.

For example, the NLM-supported bioinformatics training program of the W. M. Keck Center for Computational Biology—a joint endeavor of Baylor College of Medicine, Rice University, and the University of Houston—will prepare researchers and clinicians to integrate high performance computing technologies into all levels of healthcare.

NLM supports medical informatics short courses at Stanford University that are also made available via distance learning technologies. A key aim of this course work is to provide hands-on experience with software programs in a training laboratory environment.

Medical informatics training at the Yale Center for Medical Informatics (YCMI) focuses on the creative use of computers in support of clinical medicine, biomedical research, and medical education. The goal is to provide fellows with experiences that will prepare them for careers based wholly or in part in medical informatics.

Learning Technologies Project (LTP)

The goal of NASA's Learning Technologies Project (LTP) is to promote the growth of a national information infrastructure using the vast amount of information NASA has acquired since its creation. The LTP projects will increase public access to scientific databases, develop new applications and pilot programs for using science data, and create new curriculum products and tools for K-12 and K-14 education via the Internet. LTP has five components:

- The Remote Sensing Public Access Center (RSPAC) is a focal point for all LTP projects, providing user-friendly Web sites, including the award-winning "Observatorium," that make up-to-date science information available to the public.
- Digital Library Technology (DLT) projects support the development of new technologies to facilitate public access to NASA data over the Internet. DLT focuses on developing tools, applications, and software systems that scale to accommodate evolving user requirements and order-of-magnitude increases in user access.
- K-14 Aeronautics Projects use the Internet as a tool for teaching science, mathematics, engineering, and aeronautics in grades K-14.
- K-12 Education Outreach Centers incorporate technology into curricula to enhance learning.

- LTP's Special Projects encourage the development of innovative applications of Earth and space science remote sensing data. By stimulating broad public use, over the Internet, of the databases maintained by NASA and other agencies, the program encourages schools, businesses, and citizens to access and use Earth and space science data with unprecedented temporal, spectral, and spatial resolution.

NASA's LTP Web sites provide a wealth of educational information to the general public. Here are four examples:

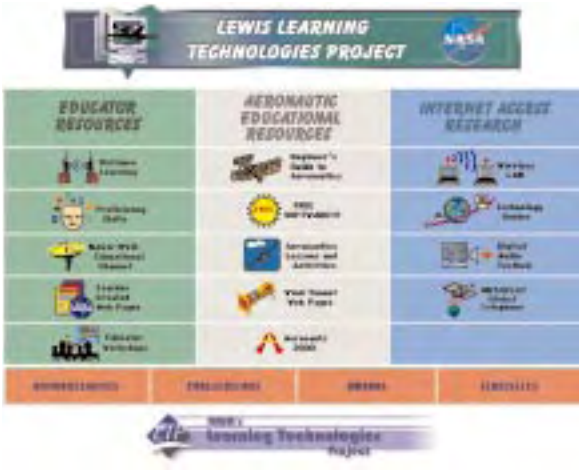

This Learning Technologies Web site, supported by NASA's Glenn Research Center at Lewis Field (formerly known as the Lewis Research Center) in Cleveland, Ohio, provides educational links to both students and teachers focusing on aeronautics and aerodynamics.

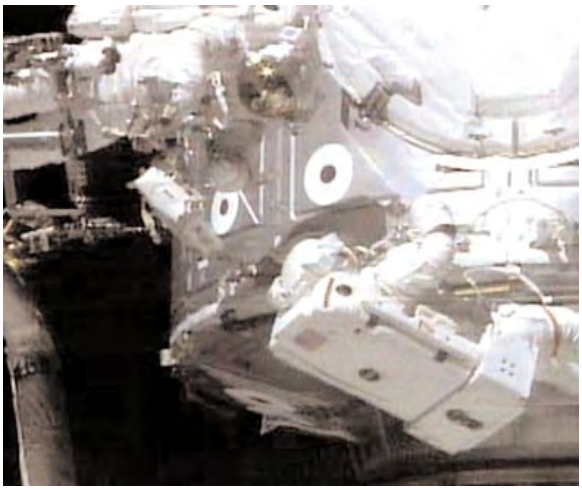

NASA's "Observatorium" site provides a wealth of scientific images and information, including this snapshot of astronauts Jerry Ross and Jim Newman working on the exterior of the Unity space station module during their first spacewalk on December 7, 1998.

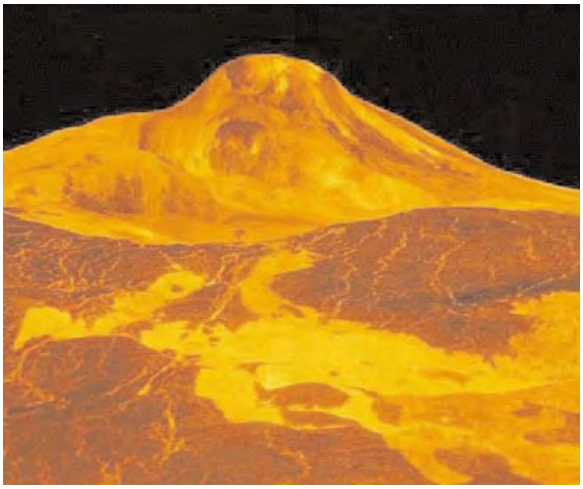

This 3-D look at the Venusian Volcano Maat Mons combines Magellan synthetic aperture radar data with radar altimetry. From the north side of the mountain, the lava flows extend hundreds of kilometers across the fractured plains in the foreground to the base of Maat Mons.

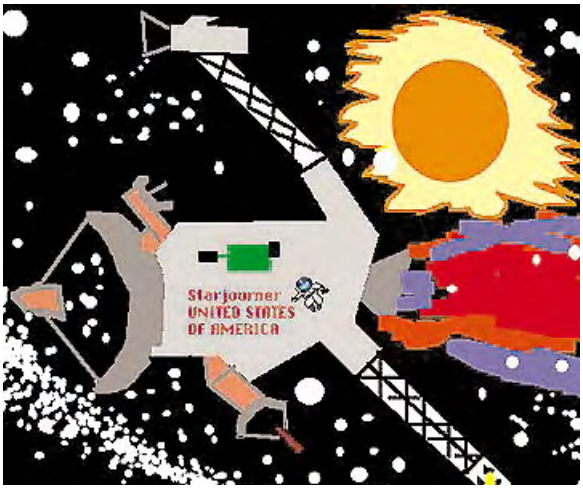

Information travels both ways at the "Observatorium" Web site. Here is a young Virginia artist's concept of a space probe called Starjourner, designed to travel to our nearest neighboring star. According to this promising engineer, "It would be launched with a powerful rocket ship. It would then be launched to the star with a nuclear rocket engine. It would arrive in fourteen years."

Federal Information Services and Applications Council

The Federal Information Services and Applications Council (FISAC) helps apply state-of-the-art information technology R&D to Federal information systems and services in support of agency missions and systems. FISAC provides a two-way channel of communications to agencies that are either wholly or in part not represented or participating in the HPCC R&D agenda. The Subcommittee on CIC R&D uses FISAC to disseminate information about its research agenda, priorities, and results to the broad Government community, while the FISAC provides feedback and identifies research needed to field the cutting-edge applications required by the Federal government. The FISAC:

- ❏ Provides a forum for information exchange among end users of HPCC R&D and the HPCC research community
- ❏ Identifies HPCC research needed to enable information technology applications critical to the Federal government
- ❏ Facilitates collaboration on information and communications applications research across the Government
- ❏ Identifies information and communications science and technology issues and requirements to be considered in developing Federal policies and programs
- ❏ Provides input and advice to the Subcommittee on CIC R&D on its R&D activities and priorities to ensure that applications' computing, information, and communications requirements are addressed

The FISAC accomplishes these objectives by:

- ❏ Promoting early application of HPCC R&D technologies to critical Federal government missions
- ❏ Supporting multiagency leadership to demonstrate and deploy advanced HPCC technologies that have the potential to be widely applicable to Federal agency missions
- ❏ Encouraging pilot projects to assess critical HPCC technologies needed by applications
- ❏ Supporting broad Administration goals to eliminate international barriers in information services and applications

FISAC membership is drawn from applications organizations throughout the Federal government. FISAC accomplishes its tasks through its Crises Management, Federal Statistics (FedStats), Next Generation Internet Applications, and Universal

Access teams, and liaison to NSF's Digital Government program and the Federal Web Consortium. Funding for FISAC activities comes from a combination of the HPCC R&D budget crosscut, other funds in HPCC R&D agencies, and funds from non-HPCC organizations.

Crises Management

FISAC's Information Technology for Crises Management (ITCM) Team collaborates with Federal, state, local, and international organizations and the private sector to promote collaborations that identify, develop, test, and implement HPCC technologies for crises management applications. The team consists of representatives from DoD, the Federal Emergency Management Agency (FEMA), NASA, NCO/CIC, NOAA, NSF, the United States Census Bureau, and the United States Geological Survey (USGS).

ITCM works to define HPCC research needs in emergency preparedness and response, mitigation, and recovery; guide the development of HPCC applications and tools to meet crises management demand; coordinate demonstrations of new HPCC crises management applications and tools; and disseminate ITCM applications and tools.

Digital Government

NSF's Digital Government program fosters joint research projects between the computer and information science research community and Federal agencies with significant information services missions. Research is focused on making government more accessible to the public without compromising privacy and security — either the individual's or the Government's. Participating agencies provide research testbeds.

Fifty project proposals were received in response to a 1998 digital government program announcement. While primarily for planning and workshop grants, they also included 15 research proposals. More than 20 Federal agencies participate in the grants that were awarded in areas such as Federal statistics, crises management, and environmental data integration. A round of FY 2000 proposals is expected in the summer of 1999.

The Digital Government program funded Crises Management, FedStats, Geospatial Information Systems, and state/local workshops conducted by the Computer Science and Telecommunications Board of the National Academy of Sciences in FY 1999. Further details about these and other workshops can be found beginning on page 93.

Federal Statistics (FedStats)

The FedStats R&D Team focuses on:

- ❏ Identifying IT research issues that need to be addressed to enable the FedStats and individual statistical agency Web sites to create a virtually-integrated national statistical information infrastructure
- ❏ Promoting partnerships with research institutions through NSF's Digital Government and similar programs

The FedStats Web site (*www.fedstats.gov*) provides a gateway to the 70-plus agencies of the Federal government that invest significantly in compiling statistics. The site is funded by the 14 member agencies of the Interagency Council on Statistical Policy (ICSP) chaired by the Office of Management and Budget. The ICSP co-sponsors the FedStats R&D Team with FISAC.

Next Generation Internet Applications

The Next Generation Internet Applications Team coordinates NGI applications with non-HPCC agencies. For example, NOAA, with the support of FEMA,

is developing a prototype complex realtime wind analysis system that will integrate realtime data from numerous sources including ships, buoys, and aircraft; allow expert human intervention for quality assurance; and provide capabilities for automated product development and tailored products. This work will then be made available to the National Hurricane Center and to FEMA's Natural Hazard Loss Estimation Methodology (HAZUS) system. Researchers are examining advanced distributed software technologies and integrating multiple data sources, automated quality control, and dynamic HTML to provide users with rapid wind analyses for forecasting and disaster recovery.

Universal Access

The Universal Access Team facilitates national research, development, and deployment of information technologies to enable broader participation by the public. A Universal Access workshop, developed by the Team and hosted by the Federal World Wide Web Consortium, accelerated the formation of partnerships for expanded initiatives in FY 1999-2000, including joint research with the Bureau of the Census, General Services Administration (GSA), Social Security Administration, and Stanford University to explore the design of universally accessible human computer interaction systems; and an R&D project for the Group of Seven (G7) economic leaders involving people with disabilities as co-developers of an accessible, collaborative, Web-conferencing environment. Participants included representatives from three countries.

Additional FY 1999 research is being conducted with the Department of Labor's Office of Worker's Compensation and Georgetown University Medical Center to establish national IT-ergonomic standards to reduce the Occupational Safety and Health Administration (OSHA) -estimated $100 billion annual IT-related business costs associated with worker injury. An Underwriter's Lab performance validation proposal is part of a Presidential Initiative on Universal Access in FY 2000. NSF has also announced a research program in Universal Access, for which the Team will promote proposal submissions.

The Universal Access Team merged with the Accessible Technology working group, established by the GSA Center for IT Accommodation in FY 1999. The combined group will also work with the Interoperability Committee of the CIO Council to integrate new electronic and information technology accessibility requirements mandated by Section 508 of the Rehabilitation Act to meet the goals of the Government-wide architecture activities of this committee.

Federal World Wide Web (WWW) Consortium

Founded in 1994, following release of NCSA's Mosaic — the first graphical Web browser — NSF's Federal WWW Consortium is a partnership between member agencies and NSF-sponsored academic and other research institutions whose goal is to accelerate the development and deployment of a Web-based digital Federal government. While not part of the FISAC, the Consortium provides advice to the FISAC on information technology issues of common interest. Collaboration tools, firewall research, universal access, and training workshops and seminars for Federal employees are past focus areas. Currently, the Consortium's member agencies use their resources to facilitate IT research with the PACI partners.

DOE's ASCI Program

DOE's Accelerated Strategic Computing Initiative (ASCI) began in FY 1996, but this is the first year that ASCI has been explicitly described in the President's HPCC Budget. This is in part a response to the PITAC's suggestion, but it also reflects ASCI's world leadership position in scientific and engineering computing.

ASCI's mission is closely tied to national security. On August 11, 1995, President Clinton announced the United States' intention to pursue a "zero yield" Comprehensive Test Ban Treaty and thus reduce nuclear danger. This decision ushered in a new era in the way the U. S. ensures confidence in the safety, performance, and reliability of its nuclear stockpile.

ASCI's role in the Nation's Stockpile Stewardship Program

ASCI is an essential element of the Department of Energy's Stockpile Stewardship Program (SSP), which was established to build on existing capabilities while developing new means of assessing the performance of nuclear weapon systems, predicting their safety and reliability, and certifying their functionality. The SSP not only must fulfill its responsibilities without nuclear testing, but also must deal with constraints on non-nuclear testing, the downsizing of production capability, and the cessation of new weapon designs to replace existing weapons. Further complicating matters, weapon components will exceed their design lifetimes, and manufacturing issues and environmental concerns will force changes in fabrication processes and materials of weapon components. The DOE selected computer simulation and modeling as the way to fulfill its responsibilities and established ASCI to provide high fidelity computer simulations of weapon systems that will enable scientists to make the necessary judgements to maintain the credibility of the nuclear deterrent.

The new capabilities must be developed within a fast-approaching deadline. Both the scientists who form the basis of our experience with nuclear weapon design and testing and the weapons themselves are aging. This leads to the year 2004 timeframe as the target date for having working ASCI computing systems and codes available so a smooth transition from "test based" certification and assessment can be made.

Computing and simulation with ASCI

ASCI is a mission-driven large-scale computing and simulation effort to deliver a set of advanced capabilities within a constrained time period. Because of this the ASCI effort focuses on creating and developing:

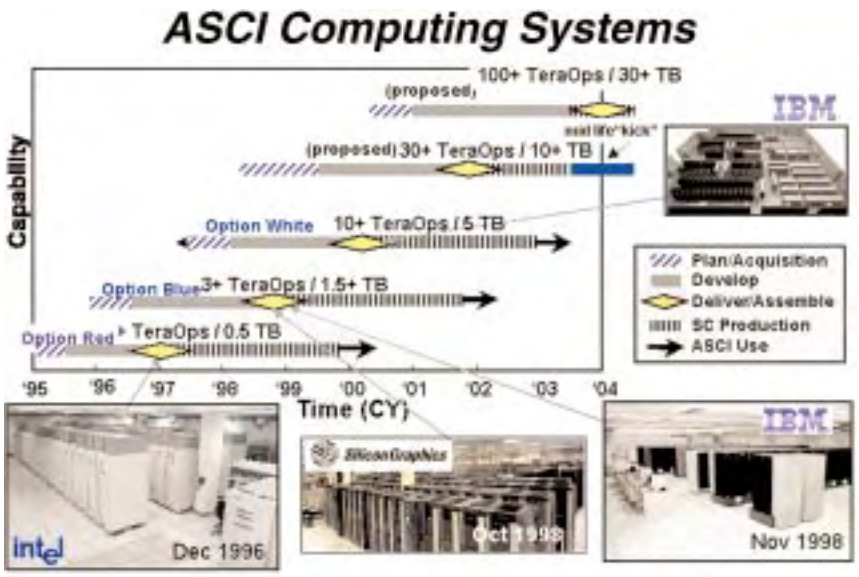

The year 2004 is the timeframe for having working ASCI computing systems and codes.

- *Advanced Applications Software.* ASCI is developing the high performance applications software needed to implement simulation capabilities for weapon systems.
- *The Highest End of Computing.* Computers that are more powerful are needed for simulation applications. Through partnerships between the national labs and the U.S. computer industry, ASCI is stimulating the more rapid development of high performance computers with speeds and memory capacities much greater than are anticipated from current development trends. Collaborations with industry, academia, and other Government agencies is spurring development of high performance access, communications, and presentation "supercorridors," making an information-rich connection between weapon scientists and simulations.
- *Problem-Solving Environments.* ASCI is developing a computational infrastructure to allow applications to be developed and execute efficiently on the ASCI computing systems and allow accessibility through the high performance, information-rich supercorridors.

Academic Strategic Alliance Program

The ASCI Academic Strategic Alliance Centers take the form of Centers of Excellence at five universities, focusing on the high performance, simulation-based, scientific and computational and computer science areas that strongly support the ASCI and Science Based Stockpile Stewardship objectives. The Centers were established with the intent of at least a five-year funding commitment, subject to a contract continuation review in the third year. At the end of five years, the program plan allows for renewal or re-bids for another five years.

Each Center focuses on one or more national-scale, multi-disciplinary applications for which the coupling and integration of computer-based simulations from multiple disciplines offers unprecedented opportunities for major advances and discoveries in basic and applied science areas that are important to ASCI, to the DOE SSP, and to the application areas. These applications are unclassified and relevant to nationally significant scientific, economic, and/or social priorities. A

primary goal of the Centers is to enable advances in mathematical modeling, numerical mathematics, computer systems, and computer and information science to help meet the high performance computing and problem solving environment needs of ASCI. The Centers are:

- Center for Integrated Turbulence Simulations (CITS), Stanford University
- Computational Facility for Simulating the Dynamic Response of Materials, California Institute of Technology
- Center for Astrophysical Thermonuclear Flashes, University of Chicago
- Center for Simulation of Accidental Fires and Explosions, University of Utah/Salt Lake City
- Center for Simulation of Advanced Rockets, University of Illinois at Urbana/Champaign

Brief descriptions of the Academic Strategic Alliance Centers and other information such as accomplishments, milestones, and personnel can be found in the ASCI Alliances FY 1999 Implementation Plan, which can be accessed on the Web at *http://www.llnl.gov/asci-alliances/* .

ASCI computing platforms

- The "ASCI Red" machine, built by the Intel Corporation and Sandia National Laboratories, links more than 9,000 Pentium Pro desktop microprocessors. In December 1996, this machine captured international headlines with its world record-breaking speed of one trillion operations per second — teraops — the first break into terascale computing.
- The "Blue Pacific" system, built by IBM and delivered to Livermore National Laboratory in the fall of 1998, incorporates 5,856 processors to achieve a peak speed of 3.9 teraops with 2.6 terabytes of memory.
- Los Alamos National Laboratory installed the "Blue Mountain" system in October, 1998. This SGI/Cray computer has 6,144 processors and over 1.5 terabytes of memory.

PathForward

The PathForward program is enabling U.S. computer companies to develop the technologies needed to produce the next generation ultra-scale computing systems for ASCI. PathForward draws on the capabilities, availability, expertise, and products currently being produced by leading computer companies, focusing on interconnect technologies, data storage technologies, systems software, and tools for large scale computing systems. These technologies, while critical to ASCI's platform needs, are areas in which private sector development would not otherwise

Los Alamos National Laboratory installed the "Blue Mountain" system in the fall of 1998.

take place, at least not in the time frame required by the SSP. At the same time, they are investments in which industry sees value for future products and markets — essential scaling and integration technologies that enable ultra-scale computing systems to be engineered and developed out of commodity computing building blocks.

HPCC Workshops

Federal agency workshops bring together researchers and technical experts from the university, commercial, and Federal laboratory communities to share information on the latest developments in technical and applications areas and foster, in real time, science discipline and interdisciplinary collaborations. The workshops provide a venue where technology users express their requirements to the technology developers, helping to keep technology development aligned with research needs. Workshops are also crucial to generate collaborative testbeds to prototype developing technologies, refine standards in an operational environment, and transition technology to the commercial sector.

As an example of Federal workshop dynamics, in FY 1999, with support from NSF and NASA, the Computer Science and Telecommunications Board of the National Academy of Sciences began a study of Federal information services, concentrating on two focus areas, Federal statistics and emergency management and response. A workshop on the latter topic occurred in December 1998, with a workshop on the former topic occurring in February 1999. The study will complete its report in late FY 1999.

Numerous Federal workshops have occurred over the past year, and many more are planned. The following is a representative sampling of recent or planned workshops:

Joint and multiagency supported workshops

The 2nd Conference on "Enabling Technologies for Peta(fl)ops Computing," sponsored by DARPA, DOE, NASA, NSA, NSF, and NIH, was held in Santa Barbara, California February 15-19, 1999. The California Institute of Technology and NASA's Jet Propulsion Laboratory jointly organized this conference, which followed a series of in-depth workshops and sponsored studies conducted to explore the factors that will determine the ultimate path to realizing petaflops capability and the means of effectively using it. The purpose of this second conference was to establish a community-wide consensus on the state and understanding of petaflops-scale computing approaches and determine directions for future research leading to practical petaflops performance systems. The four-day conference involved experts from a broad interdisciplinary community to provide in-depth coverage of a wide range of issues and foster detailed discussion across conventional discipline boundaries. More than 100 scientists and engineers from government, academic institutions, and computing industries participated in this successful conference.

The "National Workshop on Advanced Scientific Computation" was jointly sponsored by DOE and NSF and was held at the National Academy of Sciences in Washington, D.C., on July 30-31, 1998. The goal of the workshop was to examine promising high end scientific computation, simulation, and modeling, particularly for large scale and extraordinarily complex systems and processes. Additionally, the workshop was designed to address and outline the importance of next genera-

tion capabilities in science and engineering to academic, governmental, and industrial organizations.

NSF

Supported by NSF, the Computer Science and Telecommunications Board in collaboration with the Committee on National Statistics of the National Research Council held a "Workshop on Information Technology Research for Federal Statistics" on February 9-10, 1999, in Washington, D.C. Panel topics included "Information Technology Trends and Opportunities," "Study Design, Data Collection, and Data Processing," and "Creating Statistical Information Products." Another panel explored case histories with a goal of informing workshop participants about the processes used by the Federal statistical agencies to conduct their studies.

An NSF "Workshop on Middleware" was held at the Allen Conference Center on Northwestern University's Evanston, Illinois, campus on December 3-4, 1998. It was co-sponsored by Northwestern University-Evanston and Cisco. This workshop addressed developing an acceptable definition for networking middleware, identifying core components for a networking middleware infrastructure layer, and identifying needed research on components and their involvement in a networking testbed.

"Planning the Next Stage of NSF Advanced Networking Research and Development" was an NSF workshop co-sponsored by Educause in August, 1998. The goal of this workshop was to achieve a degree of consensus on the issues and opportunities that confront the NSF advanced networking program in the near term. It developed community input on the best general directions for follow-on programs and provided recommendations to NSF's ANIR on how to structure its post-vBNS program. More information can be found in the Educause report at:

http://www.educause.edu/netatedu/contents/reports/postvbnsrec981116.html

"Guidelines for NSF Networking Initiatives in the 21st Century" was an NSF workshop co-sponsored by the American Association for the Advancement of Science, on February 22-23, 1999. This workshop captured community input to produce a publicly available report summarizing visions of the future of NSF's ANIR programs following the present vBNS generation. A similar NSF-supported workshop on the "Post-vBNS Program for PACI" was held on February 17-19, 1999, in San Diego, California. Both workshops provided inputs from the general science and university research community, and will be followed by a similar workshop in the spring of 1999.

During the past year, NSF workshops also included "Research Priorities in Electronic Commerce," Austin, Texas, September 10-12, 1998; "Workshop on Distributed Information, Computation, and Process Management (DICPM) for Scientific and Engineering Environments," Herndon, Virginia, May 15-16, 1998; "1998 Information and Data Management Program Workshop," Washington, D.C., March 29-30, 1998; "NSF Workshop on Interfaces to Scientific Data Archives (ISDA'98)," Pasadena, California, March 24-27, 1998; "NSF Workshop on Information Retrieval Tools," Pittsburgh, Pennsylvania, March 20-21, 1998; "NSF Workshop on Managing and Mining Massive Data (M3D-98)," LaJolla, California, February 5-6, 1998; and "Digital Libraries Initiative All Project Meeting," Berkeley, California, January 5-6, 1998.

NASA

The NASA-sponsored "IEEE Frontiers '99 Conference," held February 21-25, 1999, in Annapolis, Maryland, provided a forum for exploring technical issues dri-

ving the outer boundaries of effective high performance computing. This conference provided a forum for the Federal information technology community in the Washington, D.C., area to share their latest findings in the area of high end computing, information, and communications. The spectrum of fields addressed by the Frontiers conferences includes applications and algorithms, system software and languages, component technologies, and system architectures.

As an introduction to ISE concepts, NASA hosted a series of demonstrations to help illustrate select aspects of the ISE vision on December 9-11, 1998, at NASA Headquarters. The focus was on geographically-distributed collaboration and advanced human-computer interfaces and interactions (that is, immersive, interactive interfaces and interactions) for advanced visualization of engineering processes—spanning mission conceptualization through operations—as well as visualization of science data that might be produced from a future mission. Though only a subset of the broader ISE vision, these demonstrations highlighted some of the emerging ISE technologies.

Examples of some of the demonstrations included: a collaborative, immersive demonstration of a Space Station assembly flight with four interactive, "virtual" Extra Vehicular Activity (EVA) crew members represented by an audience member at Headquarters (in a head-mounted display) and personnel from Langley Research Center in Virginia, Marshall Space Flight Center in Alabama, and Ames Research Center in California; a collaborative demonstration of a Space Station Intra Vehicular Activity (IVA) training scenario with individuals from Headquarters, Johnson Space Center in Texas, and Marshall Space Flight Center representing IVA crew members; a collaborative, interactive, 3-D visualization of a wind tunnel simulation; and interactive 3-D visualizations of a hurricane, a digital Earth, and a turbine engine.

At the "NASA Quality of Service Workshop," August 17-19, 1998, at NASA Ames Research Center, 18 technical papers reviewed the status of QoS technology and testbeds and provided recommendations for near-term extensions of QoS capabilities. A major result of the QoS workshop is the implementation of a multi-provider testbed to demonstrate DIFserve in the core network and RSVP for signaling at the edges of the network.

DARPA

A "Kickoff Workshop for the Principal Investigators of DARPA's Next Generation Internet Program (SuperNet) Awardees" was held on October 25-29, 1998, in Reston, Virginia. The focus of this workshop was on sharing new research ideas, discussing NGI SuperNet testbed deployment plans, and identifying synergies and collaborative opportunities across different projects. The workshop also provided a progress report for ongoing DARPA projects in its Broadband Information Technology Program. The meeting agenda, participants, and presentation topics are available at:

http://www.dyncorp-is.com/darpa/meetings/ngi98oct/agenda.html

Upcoming workshops

The purpose of the "Bridging the Gap from Networking Technologies to Applications" workshop co-sponsored by the NGI agencies, August 10-11, 1999, at Ames Research Center, California, will be to facilitate the transfer of next generation networking technologies to leading-edge revolutionary applications. Technologies of interest include QoS, scalable reliable multicast, audio/video/collaboration frameworks, and middleware. The workshop goal is to establish effective collaboration relationships and processes between the networking research and network applications communities so that applications are prepared to take

advantage of forthcoming advanced networking technologies implemented over high performance networking testbeds (NREN, NISN, ESnet, DREN, vBNS, Abilene, SuperNet, CAIRN). This workshop will bring together researchers and developers from industry, academia, and government to identify, define, and discuss future directions in collaboration and problem-solving technologies in support of scientific research. R&D in collaboration environments and integration frameworks have produced theories and technologies that provide basic levels of support and functionality. The workshop will identify where the current research and technology needs to evolve towards an architecture that can meet the needs of scientific work.

In the spring of 1999, DOE and the University of California-Berkeley will co-sponsor a conference on "Integrated Collaborative Problem Solving for Scientific Research."

Coordination of HPCC R&D Programs

President Clinton established the National Science and Technology Council (NSTC) by Executive Order on November 23, 1993. This Cabinet-level Council, chaired by the President, is the principal means within the Government for coordinating science, space, and technology efforts, as well as the diverse parts of the Federal research and development enterprise. Membership includes the Vice President, Assistant to the President for Science and Technology, cabinet secretaries and agency heads with significant science and technology responsibilities, and other White House officials.

The NSTC focuses on establishing 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. One of the NSTC's most important tasks is to prepare coordinated R&D strategies and budget recommendations to help science and technology achieve those goals. To accomplish this, the NSTC established five committees:

- Committee on Environment and Natural Resources
- Committee on International Science, Engineering, and Technology
- Committee on National Security
- Committee on Science
- Committee on Technology

A senior official (or officials) from a Federal agency or department chairs each committee, with a White House Office of Science and Technology Policy (OSTP), Senate-confirmed Associate Director serving as co-chair. In addition to the five standing committees, the NSTC establishes ad hoc working groups as needed to review and coordinate specific policies or programs that span the interests of the standing committees.

Committee on Technology (CT)

The NSTC's Committee on Technology (CT) provides overall technology policy, program, and budget guidance and direction for the Directors of OSTP and Office of Management and Budget (OMB). Comprised of senior-level representatives from the Federal government's technology R&D departments and agencies, the CT advises and assists the NSTC in increasing the overall productivity and effectiveness of Federal technology R&D. The Committee addresses significant national policy matters that cut across agency boundaries and provides a formal mechanism for interagency policy coordination and for the development of balanced and comprehensive technology R&D programs. The Committee promotes technology partnerships to leverage Federal R&D budgets more efficiently.

The CT oversees the activities of seven technology R&D Subcommittees:

- Interagency Working Group for the U.S. Innovation Partnership
- Interagency Working Group on Critical Infrastructure Protection R&D (co-managed with the Committee on National Security)
- Partnership for a New Generation of Vehicles
- Subcommittee on Building and Construction
- Subcommittee on Computing, Information, and Communications (CIC) R&D
- Subcommittee on Materials
- Subcommittee on Transportation R&D

Subcommittee on Computing, Information, and Communications (CIC) R&D

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

The Subcommittee has five Working Groups and several subgroups that meet regularly to coordinate activities, discuss new initiatives, and address specific program objectives:

- High End Computing and Computation (HECC) Working Group
- Large Scale Networking (LSN) Working Group
  - High Performance Networking Applications Team (HPNAT)
  - Information Security Team (IST)
  - Joint Engineering Team (JET)
  - Networking Research Team (NRT)
- High Confidence Systems (HCS) Working Group
- Human Centered Systems (HuCS) Working Group
- Education, Training, and Human Resources (ETHR) Working Group

The Federal Information Services and Applications Council (FISAC) also reports to the Subcommittee.

The FY 1999 accomplishments and FY 2000 plans for the Working Groups are highlighted in this book.

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

The NCO serves as a central point of contact for Federal HPCC R&D, and also provides guidance and support to the Subcommittee on CIC R&D. The Subcommittee is chaired by the NCO Director who reports to the Director of OSTP. The NCO helps the Subcommittee prepare multiagency planning, budget, and assessment documents. Together, the NCO and the agencies that participate in the HPCC R&D programs work to build and implement a balanced, comprehensive HPCC R&D agenda.

The NCO supports the President's Information Technology Advisory Committee (PITAC — described beginning on page 101), which is comprised of 26 industry and academic leaders. The PITAC is charged with providing an independent assessment of the Federal government's role in information technology R&D.

Outreach

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 HPCC needs, issues, and trends.

For example, during FY 1999, the NCO and the Subcommittee's LSN Working Group (LSNWG) participated in several activities at SC98, a national conference focused on supercomputing, scalable computing, networking, distributed comput-

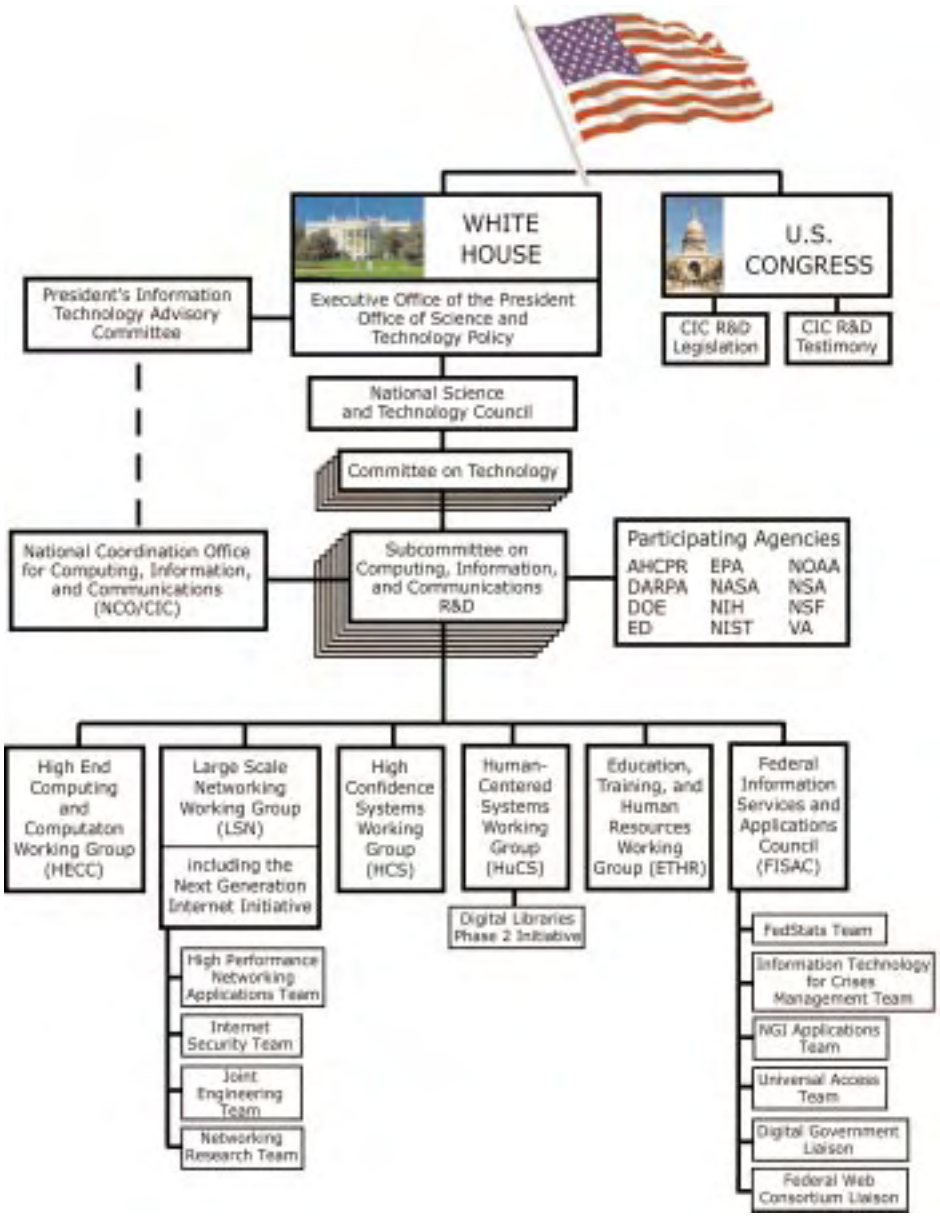

ing, data-intensive applications, and other emerging technologies that advance the frontiers of computing, communications, and computational science. The NCO and the LSNWG coordinated Next Generation Internet (NGI) applications to help educate researchers and the public about NGI and other Federal HPCC R&D investments. LSNWG representatives also held an NGI panel and participated in several panel discussions and tutorials. The NCO coordinated a PITAC Town Hall meeting where members discussed the findings and recommendations they were drafting for their Report to the President, described beginning on page 101.

Throughout each year, the NCO responds to thousands of information requests from Congressional offices, academia, industry, and the public. To fulfill these requests, the NCO provides print and video materials and maintains Web sites that include all NCO publications since 1994, information on the Subcommittee on CIC R&D and the organizations reporting to it, information on the NGI initiative, Congressional testimony, PITAC activities and reports, links to the servers of participating agencies, and other related information. The most up-to-date information on HPCC R&D programs is available at *http:/www.ccic.gov/* and *http://www.ngi.gov/* .

High End Computing Systems briefings

The NCO and the HECCWG organized High End Computing Systems briefings — held on August 17 and 18, 1998 — that were attended by more than 125 Federal agency scientists and engineers and select contractors. The event was hosted by the National Library of Medicine at its Lister Hill Center Auditorium in Bethesda, Maryland, and featured nine major U.S. high performance computing systems vendors: COMPAQ Computer Corporation, Hewlett-Packard Company, IBM Corporation, Intel Corporation, Sequent Computer Systems, Silicon Graphics, Inc., Sun Microsystems, Inc., SRC Computers, Inc., and Tera Computer Company. Similar briefings were held in February 1993 and October 1995. In the intervening years, there has been a wide range of technological developments, a major shakeout in the industry, and evolution of both the Federal government's R&D interests and needs, and industrial product lines and markets. Through these briefings, Government agencies updated their information about major trends and issues in high performance computing, which assists them in planning Federal high performance computing systems R&D and procurement activities.

Buy American Report

Congress requires information concerning non-U.S. high performance computing and communications funding activities. In FY 1999, DARPA was the only HPCC R&D agency that entered into grants, contracts, cooperative agreements, or cooperative research and development agreements for HPCC 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 HPCC R&D-related awards of $200,000 to University College London (UK) and $105,000 to the University of Warwick (UK). In FY 1999, no HPCC R&D procurement exceeds $1 million for unmanufactured articles, materials, or supplies mined or produced outside the U.S., or for manufactured articles, materials, or supplies other than those manufactured in the U.S. substantially all from articles, materials, or supplies mined, produced, or manufactured in the U.S.

President's Information Technology Advisory Committee (PITAC)

President Clinton established the PITAC in February 1997 to provide guidance and advice to the Administration on high performance computing, communications, information technology, and the Next Generation Internet (NGI), with particular emphasis on strengthening future IT R&D. The PITAC reports to the President through the President's Science Advisor.

PITAC FY 1998-FY 1999 highlights

On July 24, 1998, President Clinton amended Executive Order 13035, which established the Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet, renaming it the President's Information Technology Advisory Committee (PITAC) and enlarging its membership.

In August 1998, the PITAC issued an Interim Report to the President that set out a bold agenda for ensuring America's leadership in the Information Age by expanding Federal investments in long term R&D in technologies such as computers, networks, and software.

The Committee held a public meeting on November 4, 1998, and a Town Hall meeting at the SC98 conference in Orlando, Florida to review recommendations from the Interim Report and to solicit feedback from the general public.

PITAC report

On February 24, 1999, the PITAC issued its report, "Information Technology Research: Investing in Our Future," to President Clinton. The committee stated that the economic and strategic importance of information technology to our society help make a persuasive case for increasing Federal support for information technology R&D. The committee emphasized that a Federal role is critical due to industry's focus on the near term in today's competitive environment. Therefore, the Government is better suited to invest in solving problems of long term importance to society as a whole.

The PITAC concluded that Federal support for research in information technology is seriously inadequate at present and recommended that the Federal government create a strategic initiative and increase the total funding base for long term, basic information technology R&D by approximately $1.4 billion per year by FY 2004. The report stated that the following areas of the overall research agenda particularly need attention and should be a major part of such a long term R&D initiative:

Software

The PITAC stated that demand for software has grown far faster than our ability to produce it. Furthermore, the Nation needs software that is far more usable, reliable, and powerful than what is being produced today. The PITAC noted that the U.S. has become dangerously dependent on large software systems whose

behavior is not well understood and that often fail in unpredicted ways. Therefore, increases in research on software should be given a high priority, including software for managing large amounts of information, for making computers easier to use, for making software easier to create and maintain, and for improving the ways humans interact with computers.

Scalable information infrastructure

The PITAC noted that our Nation's dependence on the Internet is growing well beyond the intent of its original designers. Furthermore, our ability to extend its use has created enormous challenges. As the size, capability, and complexity of the Internet grows, it is imperative that U.S. researchers and scientists conduct the necessary research to learn how to build and use large, complex, highly-reliable, and secure systems.

High end computing

The PITAC observed that extremely fast computing systems that employ rapid calculation and rapid data movement are essential to provide accurate weather and climate forecasting, support advanced manufacturing design, design new pharmaceuticals, conduct scientific research in a variety of areas, and support critical national interests. To ensure that U.S. scientists continue to have access to computers of the highest possible power, the PITAC recommended that funding be focused on innovative architectures, hardware technologies, and software strategies that overcome the limitations of today's systems.

Socioeconomic impact

Information technology will significantly improve the flow of information to all people and institutions in the U.S., whose well-being depends on understanding the potential socioeconomic benefits and risks of ongoing advances in information technology. To realize the promise of new technologies, the PITAC thinks that the U.S. must invest in research to identify, understand, anticipate, and address these technologies and develop objectives and metrics to assess the ongoing transformations brought about by the integration of information technology into the lives of all citizens. The Nation must conduct research on the impact of the transformations against these objectives and develop appropriate policies to deal with the knowledge gained from this research.

Management and implementation of Federal information technology research

According to the PITAC, building a Federal IT program suited to the needs of the U.S. in the 21st century will require new management strategies, new modes of research support, and new implementation strategies, due to a combination of Federal budget requirements, the need for more long term cross-disciplinary team research, and the need to maintain a small, efficient, and coordinated research management process. It is essential that the Federal organizations responsible for managing and implementing a new IT program be positioned to review the entire Federal information technology research budget in order to restore the balance between fundamental and applied research, to encourage long term and high risk collaborative research projects, and to employ a systematic review by participating Federal agencies and the private sector.

Obtaining the report

The PITAC's report to the President is available at:

http://www.ccic.gov/ .

PITAC members at their February, 1999 meeting held at NSF in Arlington, Virginia.

Committee membership

Committee Co-chairs

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

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

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 WorldCom.

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

David M. 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 the AT&T-British Telecom Global Venture.

Robert H. Ewald is President and CEO of E-Stamp Corporation.

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

Committee members (continued)

Sherrilynne S. Fuller is Head, Division of Biomedical Informatics, School of Medicine, and Director, Health Sciences Libraries and Information Center 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.

Robert E. Kahn* is Chairman, CEO, and President of the Corporation for National Research Initiatives.

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 and Chief Technology Officer of AT&T.

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 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 an Advisor to CenterPoint Ventures.

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

**New member effective July 1998.*

Agency HPCC R&D Budgets by Program Component Area

FY 1999 Budget (Dollars in Millions)

Agency	HECC	LSN [a]	*Other PCAs* [b]	TOTAL
NSF	224.7	72.0	*4.3*	**301**
NASA	71.4	20.6	*0.6*	**93**
DARPA	48.0	82.2	*10.4*	**141**
DOE Office of Science	91.9	33.9		**126**
NIH	27.1	67.9	*8.0*	**103**
NSA	24.0	3.0		**27**
NIST	3.5	5.2	*4.3*	**13**[d]
NOAA	8.8	2.7		**12**[d]
AHCPR		3.1	*4.9*	**8**
EPA	4.2			**4**
SUBTOTAL	**503.6**	**290.6**	***32.5***	**828**[e]
DOE ASCI				484
TOTAL				**1,312**

FY 2000 Budget Request (Dollars in Millions)

Agency	HECC	LSN [a]	HCS	*Other PCAs* [c]	TOTAL
NSF	216.8	76.4	21.1		**314**
NASA	108.2	20.4	7.8		**136**
DARPA	38.4	75.5	16.0		**130**
DOE Office of Science	82.6	33.8			**116**
NIH	27.7	69.0	5.3	*2.0*	**104**
NSA	27.5	1.7	47.8		**77**
NIST	3.5	5.2	5.5		**14**
NOAA	10.3	2.7			**13**
AHCPR		5.2		*5.5*	**11**
EPA	4.2				**4**
SUBTOTAL	**519.2**	**289.9**	**103.5**	***7.5***	**919**[e]
DOE ASCI					543
TOTAL					**1,462**

[a] The NGI budget is part of the LSN budget.

[b] Several agencies included some HCS, HuCS, and ETHR budgets in the President's FY 1999 Budget.

[c] Several agencies included some HuCS and ETHR budgets in the President's FY 2000 Budget.

[d] The FY 1999 DOC total for NIST and NOAA has been updated since the President's Budget was released to be $25 million.

[e] Subtotals do not add to total due to roundoff.

HPCC R&D Summary

HPCC 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

HPCC 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 HPCC R&D Programs

Relevance/Contribution

The research must significantly contribute to the overall goals of the Federal High Performance Computing and Communications (HPCC) 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 HPCC 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 HPCC 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.

Glossary

2-D
Two dimensional.

3-D
Three dimensional.

4-D
Four dimensional.

Abilene network
A consortium among Qwest, Cisco, Nortel, and Indiana University.

AC
A compiler developed by NSA researchers that targets the Silicon Graphics/Cray Research (SGI/Cray) T3D and T3E architectures.

ACh
Acetylcholine.

AChE
Acetylcholinesterase.

ACTS
DOE's Advanced Computational Testing and Simulation Toolkit.

ADE
ALICE Differencing Engine Visualization Toolkit.

ADMs
Add-Drop Multiplexers.

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

ALICE
Advanced Large-scale Integrated Computational Environment.

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

AMBER
Assisted Model Building with Energy Refinement.

AMS
ALICE Memory Snooper.

AMS
American Meteorological Society.

ANIR
Advanced Networking Infrastructure and Research Division, part of NSF/CISE.

ANL
DOE's Argonne National Laboratory.

ANSI
American National Standards Institute.

APAN
Asian-Pacific Advanced Network.

API
Applications programming interface.

APOALA
An EPA-supported research project at Pennsylvania State University using an integrated approach for representation and analysis of space-time environmental data.

APS
Advanced Photon Source.

ASAP
ASCI's Academic Strategic Alliance Program.

ASCI
DOE's Accelerated Strategic Computing Initiative.

ASDE
ASCI Simulation Development Environment.

ATDNet
Advanced Technology Demonstration Network.

ATLAS
An ORNL project focused on automatically generating and optimizing numerical software for processors with deep memory hierarchies and pipelined functional units.

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

Aztec
SNL's library of state-of-the-art iterative methods for solving linear equations.

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
An acronym for binary digit.

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

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

BT
British Telecom.

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.

CA*Net
Canada's high performance network.

CAD
Computer-aided design.

CalTech
California Institute of Technology.

CAPI
Cryptographic API.

CAT
Computerized axial tomography.

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

C-CR
Computing-Communications Research Division, part of NSF/CISE.

CD
Compact disc.

CDSA
Common Data Security Architecture.

CERN
European Laboratory for Particle Physics.

CIC
Computing, Information, and Communications.

CIC R&D, Subcommittee on
See Subcommittee on CIC R&D.

CIO
Chief Information Officer.

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

CITS
Center for Integrated Turbulence Simulations.

CLCS
Kennedy Space Center's Checkout and Launch Control System.

CORBA
Common Object Request Broker Architecture.

CoS
Class of Service.

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

CSLU
Oregon Graduate Institute's Center for Spoken Language Understanding.

CSTB
Computer Science and Telecommunications Board, part of the National Academy of Science/National Research Council.

CT
Committee on Technology of the NSTC.

CT
Computer tomography.

CUMULVS
Collaborative User Migration, User Library for Visualization and Steering.

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

DFS
Distributed File System.

DHHS
Department of Health and Human Services.

DIA
Defense Intelligence Agency.

DICPM
Distributed Information, Computation, and Process Management.

DIPPER
NIST'S Distributed Internet Protocol and PERformance test system.

DISS
Distributed Image Spreadsheet.

DL
Digital Libraries.

DLI
Digital Libraries Initiative.

DLT
Digital Library Technology.

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

DoD
Department of Defense.

DOE
Department of Energy.

DREN
DoD's Defense Research and Engineering Network.

DSSs
Computer-based Decision Support Systems.

DWDM
Dense Wave Division Multiplexing.

ED
Department of Education.

EDA
Electronic Design Automation.

EPA
Environmental Protection Agency.

EPSCoR
NSF's Experimental Program to Stimulate Competitive Research.

ESNet
DOE's Energy Sciences Network.

ESS
NASA's Earth and Space Sciences.

ETHR
Education, Training, and Human Resources, one of the five HPCC Program Component Areas.

EVA
Extra Vehicular Activity.

EVL
Electronic Visualization Laboratory.

Exa-
A prefix denoting 10^{18}, or a million trillion; for example, exabytes.

FAA
Federal Aviation Administration.

FBI
Federal Bureau of Investigation.

FBSC
NIH's Frederick Biomedical Supercomputing Center.

FDM
Finite Differences Method.

FedNets
Federal agency networks.

FedStats
Federal Statistics.

FEMA
Federal Emergency Management Agency.

FIPS
Federal Information Processing Standards.

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.

fMRI
Functional Magnetic Resonance Imaging.

Fn-III
Fibronectin type III-like.

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

FY
Fiscal Year.

G7
Group of Seven Economic Leaders.

G, or Giga-
A prefix denoting 10^{9}, or a billion; for example, Gflops or gigaflops, gigabytes, gigabits.

GaAs
Gallium Arsenide.

GB
An acronym for Gigabyte.

Gb
An acronym for Gigabit.

GBpd
Gigabytes per day.

Gbps
Gigabits per second.

GEO
Geosynchronous Earth Orbiting, referring to a satellite.

GFDL
NOAA's Geophysical Fluid Dynamics Laboratory.

Gflops
Gigaflops, billions of floating point operations per second.

Gigapops
Gigabit per second points of presence.

GIS
Geographic Information System.

GOIN
Global Observation Information Network.

GPS
Global Positioning System.

GSA
General Services Administration.

GSFC
NASA Goddard Space Flight Center.

GSM
A European cellular communications system.

GUSTO
Globus Ubiquitous Supercomputing Testbed Organization.

HAZUS
Hazard loss estimation methodology.

HCS
High Confidence Systems, one of the five HPCC Program Component Areas.

HECC
High End Computing and Computation, one of the five HPCC Program Component Areas.

HECCWG
High End Computing and Computation Working Group.

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

HiPPI
High Performance Parallel Interface.

HPASE
NSF's High Performance Applications for Science and Engineering.

HPC
High performance computing.

HPCC
High Performance Computing and Communications. The Federal HPCC R&D programs were for several years known as the CIC R&D programs.

HPNAT
High Performance Networking Applications Team, which reports to the LSN Working Group.

HPNSP
High Performance Network Service Provider.

HPS
NOAA's Hurricane Prediction System.

HPSS
High Performance Storage System.

HPVM
High Performance Virtual Machine.

HSCC
A SuperNet testbed.

HTMT
Hybrid Technology Multithreaded Technology.

HuCS
Human Centered Systems, one of the five HPCC Program Component Areas.

IA
Information assurance.

I&A
Identification and authentication.

IAIMS
NLM's Integrated Academic Information Management System.

IBM
International Business Machines Corporation.

ICSP
Interagency Council on Statistical Policy.

IEEE
The Institute of Electrical and Electronics Engineers, Inc.

IETF
Internet Engineering Task Force.

Ig
Immunoglobulin-like.

IGERT
NSF's Integrative Graduate Education and Research Training.

IGM
NLM's Internet Grateful Med.

iGrid
International Grid.

IKE
Internet Key Exchange protocol.

IMR
Integrated Microscopy Resource.

INFOSEC
INFOrmation SECurity.

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

I2
Internet2.

I/O
Input/Output.

IP
Internet Protocol.

IPsec
IP Security Protocol.

IPsec WIT
IPsec Web-based Interoperability Tester.

ISAKMP
Internet Security Association and Key Management Protocol.

ISE
NASA's proposed Intelligent Synthesis Environment.

ISPI
NIST's Integrated Services Protocol Instrument, an interactive measurement tool for experiments in realtime transport and resource reservation protocols.

IST
Internet Security Team, which reports to the LSN Working Group.

IT
Information Technology.

IT^2
"Information Technology for the Twenty-First Century," a proposed Presidential initiative to increase the Government's investment in information technology R&D.

ITCM
FISAC's Information Technology for Crises Management Team.

ITL
NIST's Information Technology Laboratory.

IVA
Intra Vehicular Activity.

Java
An operating system-independent programming language. Initially designed for use on the Internet as a simpler, object-oriented alternative to C++, Java can be used to create complete applications that run on one computer or on a network of computers, and can be employed in coding small interactive application modules (applets) used in Web pages.

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

K, or Kilo-
A prefix denoting 10^3, or a thousand; for example, kilobits.

Kbps
Kilobits per second or thousands of bits per second.

KDI
NSF's Knowledge and Distributed Intelligence Program.

Km
Kilometers.

KN
NSF's Knowledge Networking program.

LAN
Local area network.

LANL
DOE's Los Alamos National Laboratory.

LBNL
DOE's Lawrence Berkeley National Laboratory.

LEO
Low Earth Orbiting, referring to satellite.

LIS
NSF's Learning and Intelligent Systems program.

LLNL
DOE's Lawrence Livermore National Laboratory.

LONIR
Laboratory of Neuro Imaging Resource.

LSN
Large Scale Networking, one of the five HPCC Program Component Areas.

LSNWG
Large Scale Networking Working Group.

LTP
NASA's Learning Technologies Project.

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

MANET
Mobile Ad-hoc Networks.

MAP
Memory Algorithm Processor.

MARS
Mobile Autonomous Robot Software.

MB
An acronym for Megabyte.

Mb
An acronym for Megabit.

Mbps
Megabits per second or millions of bits per second.

MCM
Multichip module.

MEII
Minimum Essential Information Infrastructure.

MEL
NIST's Manufacturing Engineering Laboratory.

Mflops
Megaflops, millions of floating point operations per second.

MirNET
U.S.-Russian network consortium.

MIT
Massachusetts Institute of Technology.

MMC
DOE's Materials Microcharacterization Collaboratory.

MOM 3
NOAA's Modular Ocean Model, Version 3.

MONET
Multiwavelength Optical Networking.

MPI
Message Passing Interface.

MPICH-G
Globus enabled version of MPI.

MPI-I/O
Message Passing Interface-Input/Output.

MPLS
Multi Protocol Label Switching.

MPP
Massively parallel processors.

MRI
Magnetic Resonance Imaging.

MSEL
NIST's Materials Science and Engineering Laboratory.

MTA
Multi-threaded architecture.

NAC
Neuroimaging Analysis Center supported by NIH's NCRR.

NAP
Network Access Point.

NASA
National Aeronautics and Space Administration.

NCAR
National Center for Atmospheric Research.

NCBI
NLM's National Center for Biotechnology Information.

NCC
NSF's New Computational Challenges program.

NCEP
NOAA's National Center for Environmental Protection.

NCI
National Cancer Institute, part of NIH.

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

NCRR
National Center for Research Resources, part of NIH.

NCSA
National Computational Science Alliance, Urbana-Champaign, Illinois, successor to the National Center for Supercomputing Applications. Part of NSF's PACI.

NERSC
DOE's National Energy Research Supercomputer Center.

netCDF
Common Data Format widely used in the oceanographic and meteorological research communities.

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

NEXRAD
NOAA's Next Generation Weather Radar.

NGI
Next Generation Internet, a Presidential initiative that is part of the HPCC R&D programs.

NGIX
Next Generation Internet Exchange Point.

NGS
NSF's Next Generation Software Program.

NHC
NOAA's National Hurricane Center.

NIAP
National Information Assurance Partnership.

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

NIH
National Institutes of Health, part of DHHS.

NISN
NASA's Integrated Services Network.

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.

NMR
National Medical Resource.

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

NORDUnet
A network spanning the Nordic countries.

NPACI
NSF's National Partnership for Advanced Computational Infrastructure. Part of NSF's PACI.

NRCAM
NIH-supported National Resource for Cell Analysis and Modeling.

NREN
NASA's Research and Education Network.

NRL
Naval Research Laboratory, part of DoD.

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

NSA
National Security Agency, part of DoD.

NSF
National Science Foundation.

NSFNET
An NSF computer network program, predecessor to vBNS.

NSTC
The Presidential National Science and Technology Council.

NTON
National Transparent Optical Network.

NUMA
Non-uniform Memory Access.

ODEs
Ordinary Differential Equations.

OMB
White House Office of Management and Budget.

OMG
Object Management Group.

ONR
Office of Naval Research.

ORNL
DOE's Oak Ridge National Laboratory.

OS
Operating system.

OSHA
Occupational Safety and Health Administration.

OSTP
White House Office of Science and Technology Policy.

PACI
NSF's Partnership for Advanced Computational Infrastructure.

ParAgent
An EPA-supported interactive tool for automatic parallelization of specific classes of programs.

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

PARC
Xerox Palo Alto Research Center.

ParVox
NASA's Parallel Volume Rendering System for Scientific Visualization.

PathForward
An ASCI program to develop the technologies needed to produce the next generation ultra-scale computing systems.

PCA
Program Component Area. The HPCC R&D programs are organized into five PCAs: 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). Each PCA spans an area in which multiple agencies have activities.

PCASYS
NIST's public domain Pattern-level Classification Automation SYStem.

PDEs
Partial Differential Equations.

Penn State
Pennsylvania State University.

PET
Positron emission tomography.

Peta-
A prefix denoting 10^{15}, or a thousand trillion; for example, petabits.

PETSc
The Portable Extensible Toolkit for Scientific computation.

pflops
Petaflops, 10^{15} flops.

PITAC
President's Information Technology Advisory Committee.

PKI
Public key infrastructure.

PMEL
NOAA's Pacific Marine Environmental Laboratory.

POPTEX
A software tool developed at LANL that provides interactive visualization capabilities.

QBone
I2's end-to-end QoS testbed to accelerate the development of interdomain quality of service.

QC
Quantum Computing.

QoS
Quality of Service.

R&D
Research and development.

RAMM
Remote Access Multidimensional Microscopy.

RBAC
Role Based Access Control.

RERC
NIDRR's Rehabilitation Engineering Research Centers.

RISC
Reduced instruction set chip, a type of microprocessor.

RSPAC
Remote Sensing Public Access Center.

RTEC
Regional Technology in Education Consortium.

RUC-2
Rapid Update Cycle.

SC
Acronym for supercomputing.

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 mass storage.

SDE
Simulation Development Environment.

SDSC
San Diego Supercomputer Center.

SGI
Silicon Graphics, Inc.

SIGINT
SIGnals INTelligence.

SIMA
NIST's Systems Integration for Manufacturing Applications.

SingaREN
Singapore network.

SMS
Scalable Modeling System.

SNL
Sandia National Laboratories.

SNMP
Simple Network Management Protocol.

SONET
Synchronous Optical Network Transmission.

SOS
Nashville Southern Oxidant Study.

S&T
Science and technology.

SSP
ASCI's Stockpile Stewardship Program.

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

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

STT
Space-Time Toolkit.

STU-III
Secure Telephone Unit, 3rd generation, the standard secure telephone for the U.S. Government.

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

SUMAA3d
Scalable Unstructured Mesh Algorithms and Applications.

SuperNet
A DARPA network expected to provide 10 to 100 Gbps speeds in FY 1999-FY 2000

SURFnet
Netherlands network.

T, or Tera-
A prefix denoting 10^{12} or a trillion; for example, terabits, teraflops.

TANnet
Taiwan network.

Tbps
Terabits per second.

TCP
Transmission Control Protocol.

TID
Trusted Image Dissemination.

TOXNET
An NLM toxic substances information retrieval service.

TransPAC
U.S./Asia-Pacific Consortium.

TRVS
DARPA's Text, Radio, Video, and Speech Program.

UAH
University of Alabama in Huntsville.

UCAID
University Corporation for Advanced Internet Development.

UCLA
University of California-Los Angeles.

UCSD
University of California-San Diego.

UMLS
NLM's Unified Medical Language System.

USC
University of Southern California.

USGS
United States Geological Survey.

U.S.
United States.

VA
Department of Veterans Affairs.

VAST
VisAnalysis Systems Technology.

vBNS
NSF's very high performance Backbone Network Services.

vDOC
Virtual Distributed Online Clinic.

VH
NLM's Visible Human Project.

Virginia Tech
Virginia Polytechnic Institute and State University.

VLANs
Virtual Local Area Networks.

VLSI
Very Large Scale Integration.

VRML
Virtual Reality Modeling Language.

VSAT
Very Small Aperture Terminal.

WAI
Web Access Initiative.

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.

WW
Whisker Weaving, an algorithm for generating 3-D unstructured hexahedral meshes.

WWW
World Wide Web.

W3C
World Wide Web Consortium.

YCMI
Yale Center for Medical Informatics.

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 C. Fenton
Systems Administrator

Vicki L. Harris
Administrative Assistant

Kristin Janger
Director's Assistant

Catherine W. McDonald
Office Manager/Technical Staff

Betty McDonough
Assistant Systems 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

Ann Rutherford
Administrative Officer

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

AHCPR

J. Michael Fitzmaurice, Ph.D.
Senior Science Advisor for Information Technology
Immediate Office of the Administrator
Agency for Health Care Policy and Research
2101 East Jefferson Street, Suite 600
Rockville, MD 20852
(301) 594-3938
FAX: (301) 594-2168

BLS

Cathryn S. Dippo, Ph.D.
Associate Commissioner, Office of Survey Methods Research
Bureau of Labor Statistics
2 Massachusetts Avenue N.E., Suite 4915
Washington, DC 20212
(202) 606-7372
FAX: (202) 606-7426

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

Helen Gill, Ph.D.
Program Manager, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-7461
FAX: (703) 696-2202

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

Gary M. Koob, Ph.D.
Program Manager, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-7463
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

José L. Muñoz, Ph.D.
Assistant Director, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-4468
FAX: (703) 696-2202

Stephen L. Squires, Ph.D.
Special Assistant for Information Technology
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2400
FAX: (703) 696-2209

Gary W. Strong, Ph.D.
Program Manager, Information Technology Office
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714
(703) 696-2259
FAX: (703) 696-0564

DoD

Phillip C. Dykstra
Chief Technologist, High Performance Computing Division
Army Research Laboratory
Attn: AMSRL-CI-HC
Aberdeen Proving Ground, MD 21005-5067
(410) 278-8608
FAX: (410) 278-9199

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, SC-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

Richard A. Carlson
Program Manager, Mathematical, Information, and Computational Sciences (MICS) Division, Office of Computational and Technology Research (OCTR)
Department of Energy
19901 Germantown Road
SC-31
Germantown, MD 20874
(301) 903-0073
FAX: (301) 903-7774

Frederick C. Johnson, Ph.D.
Mathematical, Information, and Computational Sciences (MICS) Division, Office of Computational and Technology Research (OCTR)
Department of Energy
OCTR/MICS, SC-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-3601
FAX: (301) 903-7774

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, SC-31
19901 Germantown Road
Germantown, MD 20874-1290
(301) 903-5152
FAX: (301) 903-7774

Norman H. Kreisman
Advisor, International Technology
Department of Energy, SC5
Mailstop 3H049-FORS
1000 Independence Avenue, S.W.
Washington, DC 20585
(202) 586-9746
FAX: (202) 586-7152

Paul C. Messina, Ph.D.
Senior Technical Advisor, Defense Programs, Office of Strategic Computing and Simulation
Department of Energy
Room 1J077
1000 Independence Avenue, S.W.
Washington, DC 20585
(202) 586-1101
FAX: (202) 586-2168

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, SC-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, SC-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

GSA

G. Martin Wagner
Associate Administrator for Governmentwide Policy
General Services Administration
1800 F Street, N.W. (M) Room 5240
Washington, DC 20405
(202) 501-8880
FAX: (202) 501-8898

Keith Thurston
Assistant to Deputy Associate Administrator, Office of Governmentwide Policy
General Services Administration
1800 F Street, N.W. (MK) Room 2239
Washington, DC 20405
(202) 501-3175
FAX: (202) 501-2482

Susan B. Turnbull
Director, Center for IT Accommodation
General Services Administration
1800 F Street, N.W. (MKC) Room 1234
Washington, DC 20405-0001
(202) 501-6214
FAX: (202) 501-6269

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

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

Anngienetta R. Johnson
Acting Director, Program Planning and Development Division, Office of Earth Science
National Aeronautics and Space Administration
Code YF
300 E Street, S.W.
Washington, DC 20546
(202) 358-4717
FAX: (202) 358-2769

Kul B. Bhasin
Acting Branch Chief,
Satellite Networks and Architectures Branch
National Aeronautics and Space Administration
NASA Glenn Research Center at Lewis Field
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

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

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

Eugene Tu, Ph.D.
Program Manager, High Performance Computing and
Communications Program
National Aeronautics and Space Administration
Mail Stop 258-3
Moffett Field, CA 94035-1000
(650) 604-4486
FAX: (650) 604-4036

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

Don R. Preuss
Chief Technology Officer, Center for Information
Technology (CIT)
National Institutes of Health
12 South Drive, MSC 5654
Building 12A, Room 3033
Bethesda, MD 20892-5654
(301) 496-5703
FAX: (301) 402-1754

Alan S. Graeff
Chief Information Officer, Center for Information
Technology (CIT)
National Institutes of Health
12 South Drive, MSC 5654
Building 12A, Room 3033
Bethesda, MD 20892-5654
(301) 496-5703
FAX: (301) 402-1754

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

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

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

William Mehuron, Ph.D.
Director, Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8900
Gaithersburg, MD 20899-8900
(301) 975-2900

R. J. (Jerry) Linn
Associate Director for Program Implementation, Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8901
Gaithersburg, MD 20899-8901
(301) 975-3624
FAX: (301) 216-2075

Kevin L. Mills, Ph.D.
Chief, Advanced Network Technologies Division , Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8920
Gaithersburg, MD 20899-8920
(301) 975-3618
FAX: (301) 590-0932

Ronald F. Boisvert, Ph.D.
Chief, Mathematical and Computational Sciences Division, Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8910
Gaithersburg, MD 20899-8910
(301) 975-3812
FAX: (301) 990-4127

James Fowler
Manager, Systems Integration for Manufacturing Applications Program (SIMA), Manufacturing Engineering Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8260
Gaithersburg, MD 20899-8260
(301) 975-3180
FAX: (301) 926-3842

Craig W. Hunt
Internetworking Technologies Group, Advanced Network Technologies Division, Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8920
Gaithersburg, MD 20899-8920
(301) 975-3600
FAX: (301) 590-0932

Kathy L. Lyons-Burke
Computer Scientist, Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8930
Gaithersburg, MD 20899-8930
(301) 975-4611
FAX: (301) 948-1233

Doug Montgomery
Manager, Internetworking Technologies Group, Advanced Network Technologies Division
National Institute of Standards and Technology
100 Bureau Drive, Stop 8920
Gaithersburg, MD 20899-8920
(301) 975-3630
FAX: (301) 590-0932

Steven R. Ray, Ph.D.
Chief, Manufacturing Systems Integration Division, Manufacturing Engineering Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8260
Gaithersburg, MD 20899-8260
(301) 975-3524
FAX: (301) 258-9749

Robert Rosenthal
Computer Security Division, Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8930
Gaithersburg, MD 20899-8930
(301) 975-3603
FAX: (301) 948-1233

David C. Stieren
Program Manager, National Advanced Manufacturing Testbed, Manufacturing Engineering Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8202
Gaithersburg, MD 20899-8202
(301) 975-3197
FAX: (301) 926-8730

Shukri A. Wakid, Ph.D.
Chief Information Officer
National Institute of Standards and Technology
100 Bureau Drive, Stop 8900
Gaithersburg, MD 20899-8900
(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

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

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 08542-0308
(609) 452-6502
FAX: (609) 987-5070

Louis Uccellini
Director, National Centers for Environmental Prediction
National Oceanic and Atmospheric Administration
5200 Auth Road, Room 101
Camp Springs, MD 20746
(301) 763-8000
FAX: (301) 763-8434

Bruce B. Ross, Ph.D.
Assistant Director, Geophysical Fluid Dynamics Laboratory
National Oceanic and Atmospheric Administration
Forrestal Campus, U.S. Route 1
P.O. Box 308
Princeton, NJ 08542-0308
(609) 452-6504
FAX: (609) 987-5070

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

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

Paul Pittelli
Deputy Chief,
Office of INFOSEC Research and Technology (R2)
National Security Agency
9800 Savage Road
Fort Meade, MD 20755-6000
(301) 688-0840
FAX: (301) 688-0255

William J. Semancik
Director, Laboratory for Telecommunications Sciences
National Security Agency
c/o U. S. Army Research Laboratory
Adelphi Laboratory Center
2800 Powder Mill Road, Building 601, Room 131
Adelphi, MD 20783-1197
(301) 688-1410
FAX: (301) 291-2591

Grant M. Wagner
Technical Director, INFOSEC Computer Science Research
National Security Agency
9800 Savage Road, Suite 6534
Fort Meade, MD 20755-6534
(301) 688-0844
FAX: (301) 688-0255

NSF

Ruzena Bajcsy, 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

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

Kamal Abdali, Ph.D.
Deputy Division Director, Computer-Communications Research Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1145
Arlington, VA 22230
(703) 306-1912
FAX: (703) 306-1947

Frank D. Anger, Ph.D.
Program Director, Software Engineering and Languages, Computer-Communications Research Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1145
Arlington, VA 22230
(703) 306-1911
FAX: (703) 306-1947

Robert Borchers, Ph.D.
Division Director, Advanced Computational Infrastructure and Research Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1122
Arlington, VA 22230
(703) 306-1970
FAX: (703) 306-0632

Javad Boroumand
Associate Program Director, Advanced Networking Infrastructure and Research Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1175
Arlington, VA 22230
(703) 306-1949
FAX: (703) 306-0621

Lawrence E. Brandt
Program Director for Digital Government, Division of Experimental and Integrative Activities, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1160
Arlington, VA 22230
(703) 306-1981
FAX: (703) 306-0589

Aubrey M. Bush, Ph.D.
Acting Division Director, Advanced Networking Infrastructure and Research Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1175
Arlington, VA 22230
(703) 306-1950
FAX: (703) 306-0621

John Cherniavsky, Ph.D.
Senior Advisor for Research, Directorate for Education and Human Resources
National Science Foundation
4201 Wilson Boulevard, Suite 875.29
Arlington, VA 22230
(703) 306-1690
FAX: (703) 306-0589

Frederica Darema, Ph.D.
Senior Science and Technology Advisor, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard
Arlington, VA 22230
(703) 306-1981

William F. Decker
Program Director, Advanced Network Infrastructure, Advanced Networking Infrastructure and Research Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1175
Arlington, VA 22230
(703) 306-1949
FAX: (703) 306-0621

Valerie Gregg
Program Manager, Federal World Wide Web Consortium, Advanced Computational Infrastructure and Research Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
(on detail from Bureau of the Census)
4201 Wilson Boulevard, Suite 1122s
Arlington, VA 22230
(703) 306-1980
FAX: (703) 306-0610

Stephen M. Griffin
Program Director, Digital Libraries Initiative, Information and Intelligent Systems Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1115
Arlington, VA 22230
(703) 306-1930
FAX: (703) 306-0599

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

Anne Richeson
Assistant Program Director, Advanced Network Infrastructure, Advanced Networking Infrastructure and Research Division, Directorate for Computer and Information Science and Engineering
National Science Foundation
4201 Wilson Boulevard, Suite 1175
Arlington, VA 22230
(703) 306-1949
FAX: (703) 306-0621

Gary W. Strong, Ph.D.
(on detail to DARPA)

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) 395-4652

Elizabeth M Robinson
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) 395-3810
FAX: (202) 395-1086

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 2000 Editorial Group

Executive Editor
Sally E. Howe
National Coordination Office

Editor
Terrence L. Ponick
National Coordination Office

Editorial Associate
Kristin Janger
National Coordination Office

NCO Writing Group
Sally E. Howe
Yolanda Comedy
Kristin Janger
Grant Miller
Krish Namboodiri
Terrence L. Ponick

Contributors
Kamal Abdali, NSF
Jon Bashor, LBNL
Larry Bergman, JPL
Frederica Darema, NSF
Bill Decker, NSF
Dick Dubois, NIH
Rob Eades, DOE/PNL
Betsy Edwards, NASA
Jim Fischer, NASA
Jim Fowler, NIST
Helen Gill, DARPA
Norman Glick, NSA
Marty Herman, NIST
Richard Hicks, DOE/ORNL
Rich Hirsh, NSF
Craig Hunt, NIST
Tom Kitchens, DOE
Jerry Linn, NIST
Jim McGraw, DOE/LLNL
John Melchi, Alliance/NCSA
Joan Novak, EPA
Rod Oldehoeft, DOE/LANL
Matthew Onstott, ED
Gail Pieper, DOE/ARNL
Anne Richeson, NSF
Betsy Riley, DOE/ORNL
Mary Anne Scott, DOE
Allen Sears, DARPA
Margaret Simmons, NPACI/SDSC
Paul Smith, DOE
Thomas Sterling, Caltech/JPL
Gary Strong, NSF
Bill Turnbull, NOAA
David Womble, DOE/SNL

Acknowledgments

Many individuals have contributed to this Supplement to the President's FY 2000 Budget, and we thank all of them for their efforts. Contributors to the FY 1999 acccomplishments highlighted in this report include the agency program managers who funded projects, researchers and their students, and industrial collaborators. Contributors to sections detailing FY 2000 HPCC R&D plans include the senior agency personnel who developed the research goals and program plans.

The Editorial Group thanks their colleagues at the NCO for their support in all aspects of this process, especially proofreading, preparation of the Web version, and distribution.

A special thanks goes to Norman Glick of NSA who has for years materially assisted in preparing the final document for publication.

Abstract

The Federal High Performance Computing and Communication (HPCC) programs invest 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 new millennium. This report summarizes the goals and objectives of the HPCC R&D programs by Program Component Area (PCA) and presents a condensed view of some of the programs' FY 1999 accomplishments and FY 2000 plans, including the proposed Information Technology for the Twenty-first Century (IT^2) Initiative.

The 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 President's Information Technology Advisory Committee (PITAC), the Federal Information Services Applications Council (FISAC) and DOE's Accelerated Strategic Computing Initiative (ASCI). The document includes a comprehensive list of Government personnel who are involved in the HPCC 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

This document is 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