

---

# Advanced Computing and Communications Research under NSF Support

S. Kamal Abdali

National Science Foundation, Arlington, VA 22230, USA

**Abstract.** This paper discusses the research initiatives and programs supported by the National Science Foundation to promote high-end computing and large-scale networking. This work mainly falls under the US interagency activity called High Performance Computing and Communications (HPCC). The paper describes the Federal government context of HPCC, and the HPCC programs and their main accomplishments. Finally, it describes the recommendations of a recent high-level advisory committee on information technology, as these are likely to have a major impact on the future of government initiatives in high-end computing and networking.

## 1 Introduction

A previous paper [1] described the activities of the National Science Foundation (NSF) in the U.S. High Performance Computing and Communications (HPCC) program until 1996. The purpose of the present paper is to update that description to cover the developments since then. While some management changes have taken place during this period, and there is some redirection of its thrusts, the HPCC program continues to flourish, to say the least. The main new activities at the NSF are Partnerships for Advanced Computing Infrastructures (PACIs), the Next Generation Internet (NGI), and the Knowledge and Distributed Intelligence (KDI) initiative, and there are renewed programs for Science and Technology Centers and Digital Libraries. New initiatives that may replace the program or change its direction substantially are also expected to result from the recommendations of the Presidential Information Technology Advisory Committee (PITAC). The paper is mainly concerned with these new issues. But to make it self-contained, the entire HPCC context is briefly described also.

## 2 The HPCC program

The US High Performance Computing and Communication (HPCC) program was launched in 1991. It operated as a congressionally mandated initiative from October 1991 through September 1996, following the enactment of the

High Performance Computing Act of 1991. Since October 1996, it has continued as a program under the leadership of the Computing, Information, and Communications (CIC) Subcommittee of the Committee on Technology (CT) which is itself overseen by the National Science and Technology Council, a US Cabinet-level organization. Instrumental in the establishment of the program was a series of national-level studies of scientific and technological trends in computing and networking [2–5]. These studies concluded and persuasively argued that a federal-level initiative in high-performance computing was needed to ensure the preeminence of American science and technology. Solving the challenging scientific and engineering problems that were already on the horizon required significantly more computational power than was available. Another factor was the progress made abroad, especially the Japanese advances in semiconductor chip manufacture and supercomputer design, and the Western European advances in supercomputing applications in science and engineering. It was also clear that the advances in information technology would have a far reaching impact beyond science and technology, and would affect society in general in profound, unprecedented ways. The HPCC program was thus established to stimulate, accelerate, and harness these advances for coping with scientific and engineering challenges, solving societal and environmental problems, meeting national security needs, and in improving the nations economic productivity and competitiveness.

As late as 1996, the goals of the HPCC initiative were stated separately (e.g., in [10]) from the CIC mission descriptions. Now that HPCC has become a CIC research and development (R&D) program, its goals are subsumed in the CIC goals, which are formally stated as follows ([13]):

- Assure continued US 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

### 3 HPCC Participants and Components

The HPCC program at present involves 12 Federal agencies, each with its specific responsibilities. In alphabetical order, the participating agencies are: Agency for Health Care Policy and Research (AHCPR), Defense Advanced Research Projects Agency (DARPA), Department of Energy (DOE), Department of education (ED), Environmental Protection Agency (EPA), Na-

tional Aeronautics and Space Administration (NASA), National Institute of Health (NIH), National Institute of Standards and Technology (NIST), National Oceanic and Atmospheric Administration (NOAA), National Security Agency (NSA), National Science Foundation (NSF), and Department of Veteran Affairs (VA). The activities sponsored by these agencies have broad participation by universities as well as the industry. The program activities of the participating organizations are coordinated by the National Coordination Office for Computing, Information, and Communications (NCO), which also serves as the liaison to the US Congress, state and local governments, foreign governments, universities, industry, and the public. The NCO disseminates information about HPCC program activities and accomplishments in the form of announcements, technical reports, and the annual reports that are popularly known as “blue books” [6–13]. The NCO also maintains the web site <http://www.ccic.gov> to provide up-to-date, online documentation about the HPCC program, as well as links to the HPCC-related web pages of all participating organizations.

The program currently has five components: 1) High End Computing and Computation, 2) Large Scale Networking, 3) High Confidence Systems, 4) Human Centered Systems, and 5) Education, Training, and Human Resources. Together, these components are meant to foster, among other things, scientific research, technological development, industrial and commercial applications, growth in education and human resources, and enhanced public access to information. In addition to these components, there is a Federal Information Services and Applications Council to oversee the application of CIC-developed technologies for federal information systems, and to disseminate information about HPCC research to other Federal agencies not formally participating in the program.

The goals of the HPCC components are as follows (see the “Blue Book” 99 [13] for an official description):

1. **High End Computing and Computation:** To assure US leadership in computing through investment in leading-edge hardware, software, and algorithmic innovations. Some representative research directions are: computing devices and storage technologies for high-end computing systems; advanced computing architectures; advanced software systems, algorithms, and software for modeling and simulation. This component also supports investigation of ideas such as optical, quantum, and bio-molecular computing that are quite speculative at present, but may lead to feasible computing technologies in the future, and may radically change the nature of computing.
2. **Large Scale Networking:** To assure US leadership in high-performance communications. This component seeks to improve the state-of-the-art in communications by investing in research on networking components, systems, services, and management. The supported research directions include: advanced technologies that enable wireless, optical, mobile, and

wireline communications; large-scale network engineering; system software and program development environments for network-centric computing; and software technology for distributed applications, such as electronic commerce, digital libraries, and health care delivery.

3. **High Confidence Systems:** To develop technologies that provide users with high levels of security, protection of privacy and data, reliability, and restorability of information services. The supported research directions include: system reliability issues, such as network management under overload, component failure, and intrusion; survival of threatened systems by adaptation and reconfiguration; technologies for security and privacy assurance, such as access control, authentication, and encryption.
4. **Human Centered Systems:** To make computing and networking more accessible and useful in the workplace, school, and home. The technologies enabling this include: knowledge repositories and servers; laboratories that provide access to information repositories and that facilitate sharing knowledge and control of instruments at remote labs; systems that allow multi-modal human- system interactions; and virtual reality environments and their applications in science, industry, health care, and education.
5. **Education, Training, and Human Resources:** To support HPCC research that enables modern education and training technologies. All levels and modes of education are targeted, including elementary, secondary, vocational, technical, undergraduate, graduate, and career-enhancing education. The education and training also includes the production of researchers in HPCC technologies and applications, and a skilled workforce able to cope with the demands of the information age. The supported research directions include information-based learning tools, technologies that support lifelong and distance learning for people in remote locations, and curriculum development.

## 4 HPCC at NSF

As mentioned above, NSF is one of the 12 Federal agencies participating in the HPCC program. The total HPCC budget and the NSF share in it since the inception of the program are shown in Table 1. Thus, during this period, NSF's share has ranged approximately between one-fourth and one-third of the total Federal HPCC spending. The HPCC amount has remained approximately 10% of the NSF's own total budget during the same period.

**Table 1.** HPCC Investment: Total budget and NSF's share (in \$M)

| Fiscal Year       | 1992 | 1993 | 1994 | 1995 | 1996 | 1997 | 1998 | 1999 |
|-------------------|------|------|------|------|------|------|------|------|
| Total HPCC budget | 655  | 803  | 938  | 1039 | 1043 | 1009 | 1070 | 830  |
| NSF's HPCC share  | 201  | 262  | 267  | 297  | 291  | 280  | 284  | 297  |

The NSF objectives for its HPCC effort are:

- Enable U.S. to uphold a position of world leadership in the science and engineering of computing, information and communications.
- Promote understanding of the principles and uses of advanced computing, communications, and information systems in service to science and engineering, to education, and to society.
- Contribute to universal, transparent, and affordable participation in an information-based society.

Thus NSF's HPCC-related work spans across all of the five HPCC program components.

HPCC research penetrates to varying depth nearly all the scientific and engineering disciplines at NSF. But most of this research is concentrated in the NSF's Directorate of Computer and Information Science and Engineering (CISE). This directorate is organized into 5 *divisions* each of which is, in turn, divided into 2–8 *programs*. The work of the CISE divisions can be, respectively, characterized as: fundamental computation and communications research; information, knowledge, intelligent systems, and robotics research; experimental systems research and integrative activities; advanced computational infrastructure research; and advanced networking infrastructure research. While the phrase "high performance" may not be explicitly present in the description of many programs, the actual research they undertake is very much focused on HPCC. Indeed, the CISE budget is almost entirely attributed to HPCC. Representative ongoing research topics include: scalable parallel architectures; component technologies for HPCC; simulation, analysis, design and test tools needed for HPCC circuit and system design; parallel software systems and tools, such as compilers, debuggers, performance monitors, program development environments; heterogeneous computing environments; distributed operating systems, tools for building distributed applications; network management, authentication, security, and reliability; intelligent manufacturing; intelligent learning systems; problem solving environments; algorithms and software for computational science and engineering; integration of research and learning technologies; very large data and knowledge bases; visualization of very large data sets.

## 5 Large HPCC Projects

The HPCC program has led to several innovations in NSF's mechanisms for supporting research and human resources development. The traditional manner of funding individual researchers or small research teams continues to be applied for HPCC work too. But to meet special HPCC needs, NSF has initiated a number of totally new programs, such as supercomputing centers, partnerships for advanced computational infrastructures, science and

technology centers, and various “challenges”. Also launched were special initiatives such as digital libraries, knowledge and distributed intelligence, and the next generation internet. These projects are much larger than the traditional ones in the scope of research, number of participating investigators, research duration, and award size.

### 5.1 Science and Technology Centers (STCs)

The purpose, structure, and HPCC contributions of STCs were described in [1]. So here we mainly state the developments that have taken place since.

STCs are intended to stimulate “integrative conduct of research, education, and knowledge transfer.” They provide an environment for interaction among researchers in various disciplines and across institutional boundaries. They also provide the structure to identify important complex scientific problems beyond disciplinary and institutional limits and scales, and the critical mass and funding stability and duration needed for their successful solution. They carry out fundamental research, facilitate research applications, promote technology transfer through industrial affiliations, disseminate knowledge via visitorships, conferences and workshops, educate and train people for scientific professions, and introduce minorities and underrepresented groups to science and technology through outreach activities.

STCs are large research projects each of which involves typically 50+ principal investigators from 10+ academic institutions, and also has links to the industry. The participants work together on interdisciplinary research unified by a single theme, such as parallel computing or computer graphics. The projects are awarded initially for 5 years, are renewable for another 5 years, and are finally given an extra year for orderly phaseout. There is no further renewal, so a center has to shutdown definitely in at most 11 years. Of course, the investigators are free to regroup and compete again in the program in the future if it continues.

As a result of the competitions that took place in 1989 and 1991, 25 STCs were established by NSF. All of them have entered their final year now. The following four of those STCs were supported by the HPCC program: The Center for Research in Parallel Computation (CRPC) at Rice University; The Center for Computer Graphics and Scientific Visualization at the University of Utah; The Center for Discrete Mathematics and Theoretical Computer Science (DIMACS) at Rutgers University; and The Center for Cognitive Science at the University of Pennsylvania. These STCs have contributed numerous theoretical results, algorithms, mathematical and computer science techniques, libraries, software tools, languages, and environments. They have also made significant advances in various scientific and engineering application areas. Their output has been impressive in quality, quantity, and impact.

In 1995, NSF undertook a thorough evaluation of the STC program. For one study [14], Abt Associates, a private business and policy consulting firm was commissioned to collect various kind of information about the STCs, and

the National Academy of Science was asked to examine that data and evaluate the program. Another study [15] was conducted by the National Academy of Public Administration. Both studies concluded that the STC program represented excellent return on federal research dollar investment, and recommended that the program be continued further. The studies also endorsed most of the past guidelines regarding the funding level, award duration, emphasis on education and knowledge transfer (additionally to research), review and evaluation criteria, and management structure.

Based on these findings, NSF has decided to continue the STC program. A new round of proposal solicitations took place in 1998. The submitted proposal have been evaluated, and the awards are expected to be announced soon (as of March 1999).

## 5.2 Partnerships for Advanced Computational Infrastructures (PACIs)

The precursor to PACIs was a program called *Supercomputing Centers (SCs)* that was established by NSF in 1985 even before the start of the HPCC initiative. But the SC program greatly contributed to the momentum behind HPCC, and, since its launch, became a significant part of the initiative. For a 10-year duration, the program funded four SCs: Cornell Theory Center, Cornell University; National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign; Pittsburgh Supercomputer Center, University of Pittsburgh; and San Diego Supercomputer Center, University of California–San Diego. Several of their accomplishments and HPCC contributions have been reported in [1].

A Task Force to evaluate the effectiveness of the SC program was commissioned by NSF in 1995. This resulted in a document which is popularly known as the “Hayes Report [16]. The study considered the alternatives of renewing the SCs or having a new competition, and recommended the latter. For a more effective national computing infrastructure development, it also recommended funding fewer but larger alliances of research and experimental facilities and national and regional high- performance computing centers. Based on these findings, NSF instituted the PACI program in 1996, as the successor to the SC program. The aim of the PACIs is to help maintain US world leadership in computational science and engineering by providing access nationwide to advanced computational resources, promoting early use of experimental and emerging HPCC technologies, creating HPCC software systems and tools, and training a high quality, HPCC-capable workforce.

After holding a competition, NSF made two PACI awards in 1997. These are the National Computational Science Alliance (Alliance) led by the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign, and the National Partnership for Advanced Computational Infrastructure (NPACI) led by the San Diego Supercomputer Center at the University of California at San Diego. Each consists of more

than 60 partner institutions, including academic and government research labs, national, state-level and local computing centers, and business and industrial organizations. The leading sites, which maintain a variety of high-performance computer systems, and the partners which maintain smaller configurations of similar systems, jointly constitute a metacomputing environment connected via high-speed networks. The partners contribute to the infrastructure by developing in-house, using, and testing the necessary software, tools, environments, applications, algorithms, and libraries, thereby contributing to the further growth of a “national grid” of networked high-performance computers.

The initial mission of the SCs was to satisfy the supercomputing needs of US computational scientists and engineers. The major role of the PACIs continues to be to provide supercomputing access to the research community in all branches of science and engineering. But their expanded mission puts a heavy emphasis on education and training at all levels.

### 5.3 Next Generation Internet (NGI)

The NGI initiative, a multi-agency Federal R&D program that began in October 1997, is the main focus of LSN. It represents consolidation and refinement of ideas behind the vision of a National Information Infrastructure. This infrastructure is a subject of various studies, most importantly [17,18]. The NGI initiative supports foundational work to lead to much more powerful and versatile networks than the present-day Internet. To advance this work, the initiative fosters partnerships among universities, industry and the government. The participating federal government agencies include: DARPA, DOE, NASA, NIH, NIST and NSF. The NGI goals are:

1. Promote research, development, and experimentation in networking technologies.
2. Deploy testbeds for systems scale testing of technologies and services.
3. Develop “revolutionary” applications that utilize the advancements in network technologies and exercise the testbeds.

The aim of the advancement stipulated in Goal 1 is to dramatically improve the performance of networks in reliability, security, quality of service/differentiation of service, and network management. Two testbeds are planned for Goal 2. The first testbed is required to connect at least 100 sites and deliver speeds that are at least 100 times faster end-to-end than the present-day Internet. The second testbed is required to connect about 10 sites with end-to-end performance speed faster than the present Internet by at least a factor of 1000. The “revolutionary” applications called for in Goal 3 are to range over enabling applications technologies as well as disciplinary applications. Suggested examples of the former include collaboration technologies, digital libraries, distributed computing, virtual reality, and remote



operation and simulation. Suggested application areas for the latter include basic science, education, health care, manufacturing, electronic commerce, and government information services.

The NGI work in progress was showcased in the *Supercomputing 98* conference in a special session called *Netamorphosis*. The “Netamorphosis” demonstrations consisted of 17 significant NGI applications, ranging over visualization, scene analysis, simulation, manufacturing, remote operation, etc. For example, a demonstration entitled “Real-Time Functional MRI: Watching the Brain in Action” showed how one could remotely view brain activity while a patient was performing cognitive or sensory-motor tasks. The system could process functional MRI data in real-time, though the data acquisition, main computations, and visualization all took place at different sites connected by advanced networks. Another demonstration entitled “Distributed Image Spreadsheet: Earth Data from Satellite to Desktop” showed how scientists could analyze, process, and visualize massive amounts of geologic, atmospheric, or oceanographic data transmitted to their workstations from Earth Observing System satellites.

#### 5.4 Digital Libraries Initiative (DLI)

The original DLI, now referred to as DLI Phase 1, started as a joint venture of NSF, DARPA, and NASA. Now the initiative is in Phase 2, and includes as sponsors those agencies as well as the National Library of Medicine, the Library of Congress, and the National Endowment for the Humanities.

The initiative seeks to advance the technologies needed to offer information essentially about anything, to anyone, located anywhere around the nation and the world. A digital library is intended to be a very large-scale storehouse of knowledge in multimedia form that is accessible over the net. The construction and operation of digital libraries requires developing technologies for acquiring information, organizing this information in distributed multimedia knowledge bases, extracting information based on requested criteria, and delivering it in the form appropriate for the user. Thus, the DLI promotes research on information collection, analysis, archiving, search, filtering, retrieval, semantic conversion, and communication.

The Phase 1 is supporting 6 large consortia consisting of academic and industrial partners. Their main project themes and their lead institutions are: geographic information systems, maps and pictures, content-base retrieval (University of California–Santa Barbara); intelligent internet search, semantic retrieval, scientific journal publishing alternatives (University of Illinois); media integration and access, new models of “documents,” natural language processing (University of California–Berkeley); digital video libraries, speech, image and natural language technology integration (Carnegie Mellon University); intelligent agent architecture, resource federation, AI service market economies, educational impact (University of Michigan); uniform access, dis-

tributed object architectures, interface for distributed information retrieval (Stanford University).

The Phase 1 of the initiative was mainly concerned with learning, prototyping, and experimenting in the small. The Phase 2 expects to put this experience into actually building larger, operational, and usable systems and testbeds. There is emphasis on larger contents and collections, interoperability and technology integration, and expansion of domains and user communities for digital libraries. The supported activities are expected to range through the full spectrum of fundamental research, content and collections development, domain applications, testbeds, operational environments, and applications for developing educational resources and preserving the national cultural heritage.

### 5.5 Knowledge and Distributed Intelligence (KDI)

KDI is a new initiative that NSF established in 1998. The HPCC research has traditionally been concentrated in the NSF's Computer and Information Science and Engineering directorate. The KDI initiative stems from the realization that the advances in computing, communications, and information technologies provide unprecedented possibilities for accelerating progress in all spheres of human thought and action. KDI stresses knowledge as opposed to information, but realizes, of course, that intelligent gathering of information is a prerequisite to creating knowledge. Thus, a goal of KDI is to improve the human ability to discover, collect, represent, store, apply, and transmit information. This is to lead to improvements in the ways to create knowledge and in the actual acquisition of new knowledge. The KDI research is classified into three components:

1. Knowledge Networking (KN)
2. Learning and Intelligent Systems (LIS)
3. New Computational Challenges (NCC)

The KN component aims at building an open and context-rich environment for online interactions among individuals as well as groups. For such an environment to arise, advances have to be made in the techniques for collecting and organizing information and discovering knowledge from it. The KN-enabled vast scale of information acquisition and the power to uncover knowledge buried in collected data has grave implications for privacy and other human interest matters. Hence, KN is also concerned with research on social, societal, ethical, and other aspects of networked information.

The focus of the LIS component of KDI is to better understand the process of learning itself, as it occurs in humans, animals, and artificial systems. This understanding is to be used for improving our own learning skills, developing better teaching methods, and creating intelligent artifacts.

The NCC component is in the spirit of NSF's "Challenges programs, such as Grand Challenges, National Challenges, and Multidisciplinary Challenges.

In [1], these programs were described, and their impact and some of their accomplishments were stated. The NCC component continues to seek solutions of very complex scientific and engineering problems, ones that are computationally expensive, data intensive, and require multidisciplinary team approaches. The Challenges research and the advance in high-performance computing and communications system have a mutually benefiting push-pull relationship; the former stress tests the latter, and the latter helps the former grow in scale and scope. NCC research aims to improve our ability to model and simulate complex systems such as the oceans or the brain. In adopting the Challenges research, the KDI initiative sees it as another knowledge creation activity.

In 1998, NSF made 40 awards for KDI research for a total funding of \$51.5M. The awards span a broad range of topics, vast scopes of research, and investigators representing diverse disciplines and institutions. The 1999 KDI competition is in process.

## 6 HPCC Evaluation

General directions as well as clear objectives were defined for the HPCC program from the very beginning. Thus, some evaluation is built into the program. Some objectives naturally lead to quantifiable measures of progress, such as computation speeds in teraflops, communication bandwidth in gigabits, network extent in number of connected nodes, etc. On the other hand, there are qualitative aspects of progress, such as scientific breakthroughs, innovative industrial practices, societal penetration of knowledge and technology, quality of work force trained, etc.

The evaluation of the STC and SC programs has already been mentioned. Other parts of the NSF HPCC program have also produced impressive results. For the effectiveness of the HPCC program as whole, a number of evaluation studies have been done. The “Branscomb Report [19], is devoted to studying the means for making the program more productive. A thorough assessment of the effectiveness of the program is undertaken in the “Brooks-Sutherland Report“ [20]. The purpose of a more recent recent study [21] is to suggest the most important future HPCC applications, specially the ones with highest national, societal, and economic impact.

There is consensus that the HPCC program has been successful on most fronts. Not only the year by year milestones for quantifiable progress have been met, but the activities undertaken by the program have led to several significant, unanticipated beneficial developments. The launch of new important HPCC-inspired initiatives witnesses the programs strong momentum. But as the next section shows, there is a perception that the HPCC program is underfunded and the progress resulting from it is going to decelerate unless newer and larger investments are added to it.

## 7 Presidents Information Technology Advisory Committee (PITAC)

PITAC was established in February 1997 to provide advice to the Administration on all areas of computing, communications, and information technology. This committee at present consists of 26 research leaders representing academia and the industry. It issued an interim report in August 1998 and a final one in February 1999 [22], after a series of meetings and broad consultations with the research community. This report examines the impact of R&D in Information Technology (IT) on US business and science, and makes a number of recommendations for further work.

The PITAC report observes that the past IT R&D through HPCC and other programs is a significant factor in the nations world leadership position in science, industry, business, and the general well-being of the citizenry. IT advances are responsible for a third of the US economic growth since 1992, and have created millions of high-paying new jobs. The computational approach to science in conjunction with the HPCC algorithms, software, and infrastructure have helped the US scientists make new discoveries. The competitiveness of US economy is owed much to the efficiencies resulting from IT in engineering design, manufacturing, business, and commerce.

If IT is the engine that is driving the economy, then obviously it needs to be kept running by further investment. The PITAC report argues that the IT industry is spending the bulk of its own resources, financial and human, on near-term development of new products for an exploding market. The IT industry can contribute only a small fraction of the long-term R&D investment needed. Moreover, the industry does not see any immediate benefits of the scientific and social components of IT, and therefore has no interest in pursuing them. After estimating the total US R&D expenditure on IT, and the Federal and industrial shares of it, the PITAC conclusion is that the Federal support of the Information Technology (IT) R&D is grossly inadequate. Moreover, it is focused too much on near-term and applied research.

PITAC has recommended increments of about \$1.3 billion per year for the next 5 years. PITAC has also identified the following four high priority areas as main targets of increased investment.

**Software:** Software production methodologies have to be dramatically improved, by fundamental research, to deliver robust, usable, manageable, cost-effective software.

**Scalable Information Infrastructure:** With the ever increasing size, complexity, and sheer use of networks, research is needed on how to build networks that can be easily extended yet remain reliable, secure, and easy to use.

**High-End Computing:** Scientific research and engineering design are becoming more and more computational. The increasing complexity of problems demand ever faster computing and communications. Thus, sustained

research is needed on high performance architectures, networks, devices, and systems.

**Socioeconomic Impact:** Research is needed to exploit the IT advances to serve the society and to spread its benefits to all citizens. The accompanying social, societal, ethical, and legal issues have to be studied, and ways have to be sought for mitigating any potential negative impact.

Based on the PITAC recommendations, a new Federal interagency initiative called Information Technology for the Twenty-first Century (IT<sup>2</sup>) is being developed, as a possible successor to the HPCC program.

## 8 Conclusion

Scientific and engineering work is becoming more computational, because, increasingly, computation is replacing physical experimentation and the construction and testing of prototypes. (Indeed, the US Accelerated Strategic Computing Initiative plans to depend totally on computational simulation in its weapons research program for those weapons whose physical testing is banned by international treaties.) Several recent scientific discoveries have been possible because of computation. The HPCC program has played a key role in the rise of computational science and engineering.

In [1], it was observed that collaboration and team work emerged as an important modality of HPCC research. In particular, the HPCC programs have emphasized 1) multi-disciplinary, multi-investigator, multi-institution teams, 2) partnerships among academia, business, and industry, and 3) cooperative, interagency sponsorship of research. In recent years, the collaboration has increased in intensity and scale. The transition from SCs to PACIs is a good example.

The previous Challenge projects tended to be computation-intensive. In a number of NCC projects, the data-intensive aspect dominates the computation-intensive one. Because of this situation, data mining has emerged as a key solution strategy for many Challenge-scale problems.

In practice, the HPCC program has so far been focused on applications and infrastructure development. Partly this is because most of the participating agencies in the HPCC program have special missions, and have rightly emphasized the fulfillment of their missions rather than basic research. The development of high performance computing infrastructure has also served some critical research needs. But there is need now to bolster fundamental research in order to stimulate further progress towards the original HPCC goals. The PITAC report urges this.

## References

1. Abdali S.K.: High Performance Computing Research at NSF, In G. Cooperman, G. Michler and H. Vinck (Eds.), *Proc. Workshop on High Performance*

- Computation and Gigabit Local Area Networks*, Lect. Notes in Control and Information Sci. # 226, Springer-Verlag Berlin, 1997.
2. *A National Computing Initiative: The Agenda for Leadership*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1987.
3. *Toward a National Research Network*, National Academy Press, Washington, D.C., 1988.
4. *Supercomputers: Directions in Technology and Applications*, National Academy Press, Washington, D.C., 1989.
5. *Keeping the U.S. Computer Industry Competitive: Defining the Agenda*, National Academy Press, Washington, D.C., 1990.
6. *Grand Challenges: High Performance Computing and Communications* ("FY 1992 Blue Book"), Federal Coordinating Council for Science, Engineering, and Technology, c/o National Science Foundation, Washington, D.C., 1991.
7. *Grand Challenges 1993: High Performance Computing and Communications* ("FY 1993 Blue Book"), Federal Coordinating Council for Science, Engineering, and Technology, c/o National Science Foundation, Washington, D.C., 1992.
8. *High Performance Computing and Communications: Toward a National Information Infrastructure* ("FY 1994 Blue Book"), Office of Science and Technology Policy, Washington, D.C., 1993.
9. *High Performance Computing and Communications: Technology for a National Information Infrastructure* ("FY 1995 Blue Book"), National Science and Technology Council, Washington, D.C., 1994.
10. *High Performance Computing and Communications: Foundation for America's Information Future* ("FY 1996 Blue Book"), National Science and Technology Council, Washington, D.C., 1995.
11. *High Performance Computing and Communications: Advancing the Frontiers of Information Technology* ("FY 1997 Blue Book"), National Science and Technology Council, Washington, D.C., 1996.
12. *Technologies for the 21st Century* ("FY 1998 Blue Book"), National Science and Technology Council, Washington, D.C., 1997.
13. *Networked Computing for the 21st Century* ("FY 1999 Blue Book"), National Science and Technology Council, Arlington, VA, 1998.
14. *National Science Foundation's Science and Technology Centers: Building an Interdisciplinary Research Program*, National Academy of Public Administration, Washington, D.C., 1995.
15. *An Assessment of the National Science Foundation's Science and Technology Centers Program*, National Research Council, National Academy Press, Washington, D.C., 1996.
16. *Report of the Task Force on the Future of the NSF Supercomputing Centers* ("Hayes report"), Pub. NSF 96-46, National Science Foundation, Arlington, VA.
17. *The Unpredictable Certainty: Information Infrastructure through 2000*, National Research Council, National Academy Press, Washington, D.C., 1996.
18. *More Than Screen Deep: Toward Every-Citizen Interfaces to the Nation's Information Infrastructure* ("Biermann Report"), National Research Council, National Academy Press, Washington, D.C., 1997.
19. *From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing* ("Branscomb Report"), Pub. NSB 93-205, National Science Foundation, Washington, D.C., August 1993.

20. *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure* ("Brooks-Sutherland Report"), National Research Council, National Academy Press, Washington, D.C., 1995.
21. *Computing and Communications in the Extreme: Research for Crisis Management and Other Applications*, National Research Council, National Academy Press, Washington, D.C., 1996.
22. *Information Technology Research: Investing in Our Future*, President's Information Technology Advisory Committee Report to the President, National Coordination Office, Arlington, VA, 1999.