

High Performance Computing and Communication Research at the National Science Foundation

S. Kamal Abdali
National Science Foundation
kabdali@nsf.gov

ABSTRACT: This paper describes the High Performance Computation and Communications (HPCC) of the National Science Foundation (NSF). The rationale and organization of the US level HPCC program are outlined to present the context. Then NSF's HPCC-related activities are given in more detail.

1 Background

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. From October 1996, it continues as a program under the leadership of the Computing, Information, and Communications (CIC) Subcommittee of the US National Science and Technology Council's Committee on Computing, Information, and Communications (CCIC). The program followed a series of national-level studies of scientific and technological trends in computing and networking[1, 2, 3, 4]. It was clear that the advances in information technology would affect society in profound, unprecedented ways. The program was thus established to stimulate, accelerate, and harness these advances for coping with societal and environmental challenges, meeting national security needs, and in increasing economic productivity and competitiveness. Formally, the goals of the HPCC program are to:

- Extend US technological leadership in high performance computing and computer communications
- Provide wide dissemination and application of the technologies to spread the pace of innovation and to improve the national economic competitiveness, national security, education, health care, and the global environment
- Provide key enabling technologies for the National Information Infrastructure (NII) and demonstrate select NII applications

2 Program Participants and Components

The HPCC program at present involves 12 Federal agencies, each with its specific responsibilities. The participating agencies are: Defense Advanced Research Projects Agency (DARPA), National Science Foundation (NSF), National Aeronautics and Space Administration (NASA), Department of Energy, National Institute of Health (NIH), National Security Agency (NSA), National Institute of Standards and Technology (NIST), Department of Veteran Affairs, Department of Education, National Oceanic and Atmospheric Administration (NOAA), Environmental Protection Agency (EPA), and Agency for Health Care Policy and Research (AHCPR). 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 High Performance Computing and Communication (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" [5, 6, 7, 8, 9]. It also maintains the website <http://www.hpcc.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. Specifically, the goals of these components are the following (see Blue Book 97 [9] for an official description):

1. **High End Computing and Computation:** To assure US leadership in computing through investment in leading-edge hardware and software innovations. Some representative research directions are: computing devices and storage technologies for high-end computing systems, advanced software systems, algorithms and software for modeling and simulation. This component also supports investigation of “ultra-scale computing” ideas such as quantum and DNA-based 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 network components and services. The supported research directions include: technologies that enable wireless, optical, mobile, and wireline communications; large-scale network engineering, management, and services; 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; technologies for security and privacy assurance, such as access control, authentication, 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; collaboratories 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 research that enables modern education and training technologies. The education and training is targeted to produce 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.

The original HPCC program components, in force from 1991 to 1996, were: 1) High Performance Computing Systems, 2) National Research and Education Network, 3) Advanced Software Technology and Algorithms, 4) Information Infrastructure Technology, and 5) Applications, and Basic Research and Human Resources. The new program component structure reflects a refocusing of the HPCC activities in view of the experience and progress of the last five years.

3 HPCC at NSF

As stated above, NSF is one of the 12 agencies participating in the HPCC program. In the total HPCC budget request of \$1143M in FY 96 for all agencies, NSF’s share is \$314M. This represents nearly 10% of NSF’s annual requested budget for FY 96. NSF’s HPCC-related work spans across all of the five HPCC program components. The objectives of NSF’s HPCC effort include 1) fundamental research in high-end computing, 2) technological advances in hardware and software that are prerequisite for HPCC applications, 3) development of national HPCC facilities and services so as to make HPCC accessible to scientific and industrial researchers, educators, and the citizenry, 4) creating partnerships among universities, research labs, and industry to develop advanced computational infrastructure for HPCC, and 5) training of a scientific workforce conversant in HPCC.

HPCC research penetrates to some extent nearly all the scientific and engineering disciplines at NSF, and most of the work undertaken by NSF’s Directorate of Computer and Information Science and Engineering is related 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.

4 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 fulfill HPCC needs, NSF has initiated a number of totally new projects, such as 1) supercomputing centers and partnerships for advanced computational infrastructures, 2) science and technology centers, 3) various "challenges", and 4) the digital libraries initiative. These projects are much larger than the traditional ones in scope of research, number of participating investigators, research duration, and award size.

4.1 Supercomputing Centers (SCs) and Partnerships for Advanced Computational Infrastructure (PACIs)

NSF SCs were actually started in 1985, before the congressional formulation of the HPCC initiative. But they greatly contributed to the momentum behind HPCC, and, since its launch, have served to advance its cause.

Currently NSF supports 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. They are working cooperatively in a loose "metacenter" alliance, producing the impression of a single supercomputing environment to the users. Together the SCs provide a wide spectrum of high performance architectures, such as traditional (vector) supercomputers, massively parallel processors, and networks of high performance workstations. The SCs cater to the supercomputing needs of US computational scientists and engineers. But their customers range from neophytes to demanding experts. The rapid pace of parallel architecture advance does not allow vendors much time to devote to software. Thus the SCs obtain essentially bare machines, and build operating software to make the machines usable by the research community.

The SCs go much beyond offering a service facility. They are aggressively engaged in software development, user education, and research collaboration. They have been very effective in spreading the parallel computation culture. They have introduced many innovative, influential software tools. For example, the web browser Mosaic has greatly popularized web use, has had revolutionary impact on information dissemination and acquisition, and has given rise to an entirely new web-based software industry. Another significant product is Cave Automatic Virtual Environment (CAVE). CAVEs are multi-user room-size virtual reality environments with 3-D projections and sound surrounding the users. CAVEs are being used in highly interactive applications such as drug design and medical imaging. There is also research on linking multiple remote CAVEs into a collaboratory. The SCs have contributed significantly to scientific visualization, math libraries, parallelization tools, and communication software. They have also led many standardization efforts, such as the development of the Hierarchical Data Format (HDF) as a file format standard, or the "national file system."

In 1996, the SC program is being replaced by the Partnerships for Advanced Computational Infrastructure (PACI) program. This program represents an extension and refinement of the metacenter alliance of the SCs. The program aims to create a nationwide high performance computational facility with participation by universities, research labs, state and local governments, and the private sector. The facility will 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.

4.2 Science and Technology Centers (STCs)

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. STCs 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.

STCs 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.

As a result of competitions that took place in 1989 and 1991, 25 STCs were established by NSF. The following four of those STCs are 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.

For example, DIMACS devoted a “Special Year” in FY 95 to molecular biology. The series of research seminars and workshops during that year has contributed much to the acceleration of interest in molecular computing, DNA sequencing, and protein structure studies. Also at DIMACS, a concerted effort directed to the traveling salesperson problem led to a breakthrough on this topic.

Similarly, the work at CRPC has been instrumental in the development of parallel languages, compilers, libraries, systems, and tools. CRPC has provided key leadership in industry-wide standardization of High Performance Fortran (HPF), and the prototype compiler Fortran D produced at CRPC is serving as a model for industrial HPF compiler development. Other well-known languages and systems which CRPC pioneered or significantly contributed to include ParaScope, PVM, MPI, HPC++, and ADIFOR. CRPC is also a force behind the National High Performance Software Exchange, various “infoservers,” solvers for problems posed over the net, and several other innovative knowledge sharing schemes. In addition, there is notable research produced at CRPC on parallel algorithms for physical simulation and optimization.

4.3 “Challenge” Programs

The term “Grand Challenges” was introduced very early during the HPCC initiative. It is meant to characterize problems that are generally recognized as very difficult based on current standards of science and technology. They are computationally intensive, and are bound to “stress test” high performance computing hardware and software. They also require multidisciplinary approaches. Moreover, their solutions promise high payoffs in terms of scientific advances and crucial societal and industrial benefits. Later, the term “National Challenges” was also introduced. “Grand Challenges” is now used for computationally intensive, usually scientific, problems, while “National Challenges” is used for informationally intensive, usually engineering, problems. The distinction is often blurred.

A “Grand Challenge” problem list published in the Blue Book 93 [5] includes the following: magnetic recording technology, rational drug design, high speed civil transports, ocean modeling, ozone depletion, digital anatomy, and design of protein structures. The list is intended to be only representative, as many more such problems exist, and are actually under investigation by various HPCC research teams.

NSF has funded about 30 high-visibility group effort projects under the “Challenge” program. First, there were two rounds of “Grand Challenges” competition in FY 1990 and 1992 that mainly emphasized scientific problems. Then, there was a competition called “National Challenges” in FY 1993 that emphasized engineering problems such as manufacturing, civil infrastructure and health care delivery. Finally, there was a competition called “Multidisciplinary Challenges” in FY 1994 that encompassed physical sciences, engineering, computer science, and problem solving environments.

Unleashing large masses of scientific talent in collaborative attack on problems has indeed brought results. The challenge projects, specially those existing for over 3 years, have been demonstrably very productive. The projects have yielded, for example, highly accurate simulations in biomolecular dynamics, pollutant chemistry, materials science, radiation, and cosmology. Advanced computational techniques have been developed for modeling ocean and atmosphere phenomena, oil reservoirs, combustion, etc. For example, the weather forecasting research has led to the development of a hurricane modeling system that has become operational in 1995. As another example, computational advances in fluid dynamics have provided new understanding of the mechanisms controlling the solar heliosphere.

Expected challenge outputs also include technological artifacts, such as robots that assist surgery or advanced microscopes that incorporate pattern recognition. Information systems are under development for health care delivery, civil infrastructures, education, manufacturing, etc. The problem solving environment research is developing some systems targeted to broad problem classes and some that are quite generic.

4.4 Digital Libraries

The Digital Libraries Initiative is a joint venture of NSF, ARPA, and NASA. Its purpose is to advance the technologies needed to offer information essentially about anything, anywhere, to anyone. 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 Digital Libraries Initiative promotes research on information collection, analysis, archiving, search, filtering, retrieval, semantic conversion, and communication.

The initiative is supporting 6 large consortia consisting of academic and industrial partners. Their main project themes and their lead institutions are: scalable, intelligent, distributed library focusing on the California environment (University of California–Berkeley); multimedia testbed of earth and space science data (University of Michigan); “Alexandria Project,” maps and pictures (University of California–Santa Barbara); integrated virtual library of networked information resources and collections (Stanford University); internet browsing library, starting with contents of scientific and engineering periodicals (University of Illinois); information digital video library, starting with videos from public and educational TV channels (Carnegie Mellon University).

5 Evaluation and Impact

General directions as well as clear objectives were defined for the HPCC program from the very beginning. Thus, evaluation is built into the program. Some objectives naturally lead to quantifiable measures of progress, such as computation speeds in gigaflops, 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.

There is much experience with the SCs and STCs as most of these have existed for 5–10 years. They have also been subjected to much formal evaluation, and have been found successful. All these projects have external advisory committees which review their operation, progress and planned future directions periodically. NSF also monitors their performance via periodic progress reports and site visits. In addition to evaluating them on a per project basis, NSF has conducted rigorous evaluation of the SC and STC programs themselves using various means. Results of two recent assessments by blue ribbon panels, with recommendations for the future, are published in [12] for SCs and [13] for STCs. Other parts of the NSF HPCC program have also produced impressive results.

Collaboration is emerging as an important theme of HPCC. Most HPCC programs have emphasized 1) multi-disciplinary, multi-investigator, multi-institutional research teams, 2) partnerships between academia and industry, and 3) cooperative, interagency sponsorship of research. So much success is attributable to the multidisciplinary approach that this mode of research has de facto become a requirement in “challenge-scale” projects.

The HPCC program's effectiveness is being monitored by various kinds of studies and review panels. There is consensus that the 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 had noticeable impact and have led to several significant, unanticipated beneficial developments.

6 Critique and Conclusion

The initiative has been less successful in the area of software systems and tools. No unifying model of practical parallel computing has emerged that encompasses all the dominant current architectures. The variety of architectures has therefore caused fragmentation of the software development efforts. Lack of suitable software has retarded the spread of parallel computation. High performance computing has been very effective in some niche civilian markets. For example, large data bases, commercial transaction processing, and financial computing have been found well-suited for the new high performance machines. But in spite of the successful exploitation of high performance computing in such specialized applications, parallel computation has not yet become the paradigm for the masses. As a result, the parallel computing industry is far from being a commercial success.

In scientific computing and computer science, the exploitation of HPCC is very uneven. This difference is generally accounted for by the nature of the subject applications. Some problems lend themselves naturally to HPCC because of the regularity in their underlying data structures. But many problems with inherently irregular, heterogeneous data remain daunting to parallelization efforts.

In several numerical and optimization areas, there is considerable progress in adapting computations to high performance machines. But in fields such as symbolic and algebraic computation, automated deduction, and computational geometry, practical parallelization efforts are still in initial stages, though there is a lot of theoretical research. With the emergence of standards such as MPI that have been implemented on most architectures, it now seems possible to have "quick and dirty" parallel versions of important applications in the above mentioned areas. In theoretical computer science, the common PRAM model is not sufficiently realistic for analysis of parallel algorithms, and there is currently no agreement on a general model that is sufficiently realistic. Farther out in the future, there are "ultra-scale" paradigms, such as molecular computing, that could radically change the nature of computing. At present these are highly speculative ideas, as yet undeserving to be called technologies, and aggressive research programs are needed to define and exploit their potential.

References

- [1] *Supercomputers: Directions in Technology and Applications*, National Academy Press, Washington, D.C., 1989.
- [2] *Toward a National Research Network*, National Academy Press, Washington, D.C., 1988.
- [3] *Keeping the U.S. Computer Industry Competitive: Defining the Agenda*, National Academy Press, Washington, D.C., 1990.
- [4] *Scaling Up: A Research Agenda for Software Engineering*, National Academy Press, Washington, D.C., 1989.
- [5] *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.
- [6] *High Performance Computing and Communications: Toward a National Information Infrastructure* ("FY 1994 Blue Book"), National Science and Technology Council, Washington, D.C., 1993.
- [7] *High Performance Computing and Communications: Technology for a National Information Infrastructure* ("FY 1995 Blue Book"), National Science and Technology Council, Washington, D.C., 1994.

- [8] *High Performance Computing and Communications: Foundation for America's Information Future* ("FY 1996 Blue Book"), National Science and Technology Council, Washington, D.C., 1995.
- [9] *High Performance Computing and Communications: Advancing the Frontiers of Information Technology* ("FY 1997 Blue Book"), National Science and Technology Council, Washington, D.C., 1996.
- [10] *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.
- [11] *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.
- [12] *Report of the Task Force on the Future off the NSF Supercomputing Centers* ("Hayes report"), Pub. NSF 96-46, National Science Foundation, Arlington, VA.
- [13] *An Assessment of the National Science Foundation's Science and Technology Centers Program*, National Research Council, National Academy Press, Washington, D.C., 1996.