Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure

Computer Science and Telecommunications Board

Executive Summary

Information technology drives many of today's innovations and offers still greater potential for further innovation in the next decade. It is also the basis for a domestic industry of about $500 billion,[1] an industry that is critical to our nation's international competitiveness. Our domestic information technology industry is thriving now, based to a large extent on an extraordinary 50-year track record of public research funded by the federal government, creating the ideas and people that have let industry flourish. This record shows that for a dozen major innovations, 10 to 15 years have passed between research and commercial application (see Figure ES.1). Despite many efforts, commercialization has seldom been achieved more quickly.

Publicly funded research in information technology will continue to create important new technologies and industries, some of them unimagined today, and the process will continue to take 10 to 15 years. Without such research there will still be innovation, but the quantity and range of new ideas for U.S. industry to draw from will be greatly diminished. Public research, which creates new opportunities for private industry to use, should not be confused with industrial policy, which chooses firms or industries to support. Industry, with its focus mostly on the near term, cannot take the place of government in supporting the research that will lead to the next decade's advances.

The High Performance Computing and Communications Initiative (HPCCI) is the main vehicle for public research in information technology today and the subject of this report. By the early 1980s, several federal agencies had developed independent programs to advance many of the objectives of what was to become the HPCCI. The program received added impetus and more formal status when Congress passed the High Performance Computing Act of 1991 (Public Law 102-194) authorizing a 5-year program in high-performance computing and communications. The initiative began with a focus on high-speed parallel computing and networking and is now evolving to meet the needs of the nation for widespread use on a large scale as well as for high speed in computation and communications. To advance the nation's information infrastructure there is much that needs to be discovered or invented, because a useful "information highway" is much more than wires to every house.

As a prelude to examining the current status of the HPCCI, this report first describes the rationale for the initiative as an engine of U.S. leadership in information technology and outlines the contributions of ongoing publicly funded research to past and current progress in developing computing and communications technologies (Chapter 1). It then describes and evaluates the HPCCI's goals, accomplishments, management, and planning (Chapter 2). Finally, it makes recommendations aimed at ensuring continuing U.S. leadership in information technology through wise evolution and use of the HPCCI as an important lever (Chapter 3). Appendixes A through F of the report provide additional details on and documentation for points made in the main text.

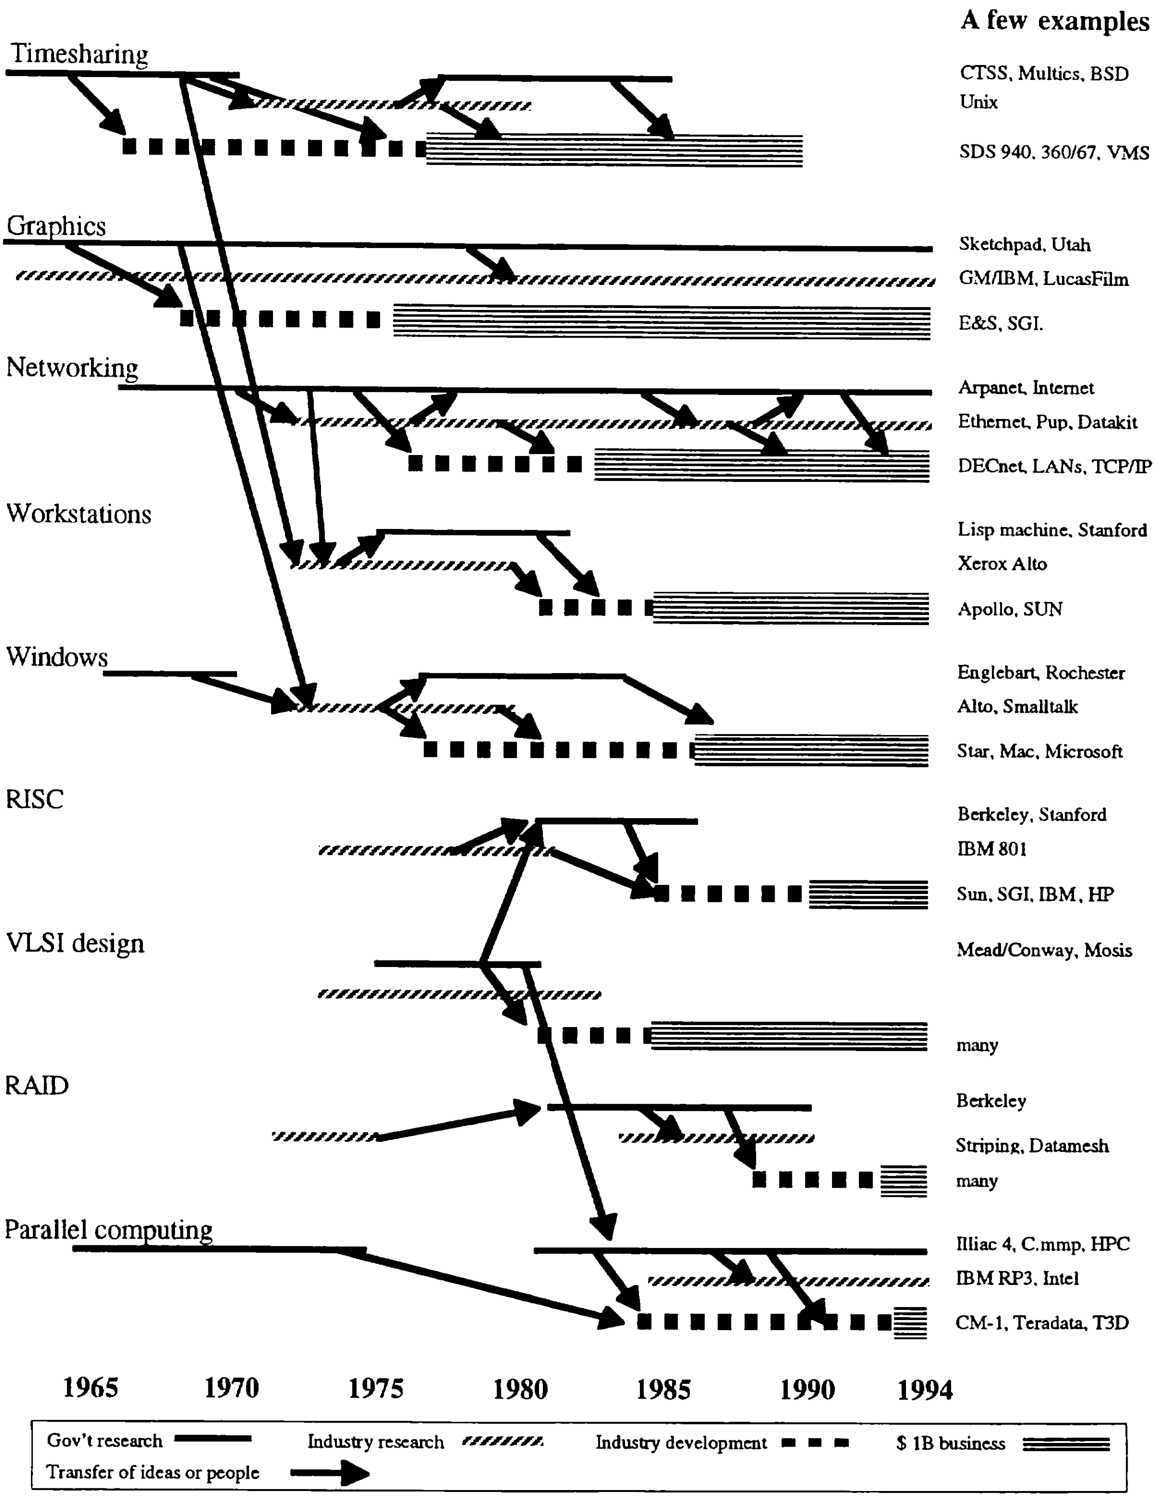

FIGURE ES.1 Government-sponsored computing research and development stimulates creation of innovative ideas and industries. Dates apply to horizontal bars, but not to arrows showing transfer of ideas and people.

INFORMATION TECHNOLOGY—FUNDAMENTAL FOR SOCIETY AND THE ECONOMY NOW AND TOMORROW

Computers, the devices that process information, affect our lives both directly and indirectly. Today, more than 70 million microcomputers are installed in the United States, and between one-fifth and one-third of U.S. households have one.[2] Entertainment, education, communications, medicine, government, and finance are using computers in more ways to enhance our lives directly through the provision of such services as distributed learning and remote banking. Computers are also used to make essential products and activities cheaper and better: airplanes, molded plastics, automobiles, medical imaging, and oil exploration are only a few of many examples. A broader benefit is the $500 billion industry's creation of jobs, taxes, profits, and exports.

Clearly, the uses and applications of information technology will continue to grow. In fact, the information revolution has only just begun. Computers will become increasingly valuable to industries and to citizens as their power is tapped to recognize and simulate speech, generate realistic images, provide accurate models of the physical world, build huge automated libraries, control robots, and help with a myriad of other tasks. To do these things well will require both computing and communications systems many times more powerful than we have today. Ongoing advances in knowledge will constitute the foundation for building the systems and developing the applications that will continue to advance our quality of life and ensure strong U.S. leadership in information technology. Strong leadership in information technology in turn supports other sectors including industry, health, education, and defense by serving their needs for equipment, software, and know-how.

The Basis for Continuing Strength— A Successful Government-Industry Partnership

Federal investment in information technology research has played a key role in the U.S. capability to maintain its international lead in information technology. Starting in World War II publicly funded research has helped to stock the nation's storehouse of trained people and innovative ideas. But our lead is fragile. Leadership can shift in a few product generations, and because a generation in the computing and communications industry is at most 2 years, our lead could disappear in less than a decade.

Since the early 1960s the U.S. government has invested broadly in computing research, creating new ideas and trained people. The result has been the development of important new technologies for time-sharing, networking, computer graphics, human-machine interfaces, and parallel computing, as well as major contributions to the design of very large scale integrated circuits, fast computers and disk systems, and workstations (see Figure ES.1; see also Chapter 1, Box 1.2 for details). Each of these is now a multibillion-dollar business. From these successes we can learn some important lessons:

- *Research has kept paying off over a long period.*

- *The payoff from research takes time.* As Figure ES.1 shows, at least 10 years, more often 15, elapse between initial research on a major new idea and commercial success. This is still true in spite of today's shorter product cycles.

- *Unexpected results are often the most important.* Electronic mail and the "windows" interface are only two examples; Box 1.2 in Chapter 1 outlines more.

• *Research stimulates communication and interaction.* Ideas flow back and forth between research programs and development efforts and between academia and industry.

• *Research trains people,* who start companies or form a pool of trained personnel that existing companies can draw on to enter new markets quickly.

• *Doing research involves taking risks.* Not all public research programs have succeeded or led to clear outcomes even after many years. But the record of accomplishments suggests that government investment in computing and communications research has been very productive.

Government Support of Research Is Crucial

The information technology industry improves its products faster than most others: for the last 40 years a dollar has bought hardware with twice as much computation, storage, and communication every 18 to 24 months, offering a 100-fold gain every decade (Patterson and Hennessy, 1994, p. 21). This rate will continue at least for the next decade (see Chapter 1, Figure 1.1). Better hardware in turn makes it feasible to create software for new applications: electronic and mechanical design, climate mapping, digital libraries, desktop publishing, video editing, and telemedicine are just a few examples. Such applications are often brought to market by new companies such as Microsoft and Sun Microsystems, both of which produce revenues of more than $4 billion per year (Computer Select, 1994) and neither of which existed 15 years ago.

The information technology industry is characterized by great importance to the economy and society, rapid and continuing change, a 10- to 15-year cycle from major idea to commercial success, and successive waves of new companies. In this environment a broad program of publicly funded research is essential for two reasons:

• First, industrial efforts cannot replace government investment in basic computing and communications research. Few companies will invest for a payoff that is 10 years away, and even a company that does make a discovery may postpone using it. The vitality of the information technology industry depends heavily on new companies, but new companies cannot easily afford to do research; furthermore, industry in general is doing less research now than in the recent past (Geppert, 1994; Corcoran, 1994). But because today's sales are based on yesterday's research, investment in innovation must go forward so that the nation's information industry can continue to thrive.

• Second, it is hard to predict which new ideas and approaches will succeed. The exact course of exploratory research cannot be planned in advance, and its progress cannot be measured precisely in the short term. The purpose of publicly funded research is to advance knowledge and create new opportunities that industry can exploit in the medium and long term, not to determine how the market should develop.

THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS INITIATIVE

Goals and Emphases

The HPCCI is the current manifestation of the continuing government research program in information technology, an investment that has been ongoing for more than 50 years. Although it emphasizes research in high-performance computing and communications, the HPCCI now has in its budget nearly all of the federal funding for computing research of any kind. The wisdom of this arrangement is doubtful.

The HPCCI was initiated to serve several broad goals (NCO, 1993):

- Extend U.S. leadership in high-performance computing and networking;
- Disseminate new technologies to serve the economy, national security, education, health care, and the environment; and
- Spur gains in U.S. productivity and industrial competitiveness.

The original plans to achieve these goals called for creating dramatically faster computers and networks, stretching their limits with Grand Challenge problems in scientific computing, setting up supercomputer centers with the machines and experts needed to attack these challenges, and training people to build and exploit the new technology. More recently the focus has been shifting toward broader uses of computing and communications.

High Performance

"High performance"—which involves bringing more powerful computing and communications technology to bear on a problem—has enabled advances on several fronts. High-performance systems, for example, deliver answers sooner for complex problems that need large amounts of computing. Timely and accurate forecasting of weather, mapping of oil reservoirs, and imaging of tumors are among the benefits encompassed by the goals listed above. But "high performance," which is broader than supercomputing, is a moving target because of the steady and rapid gains in the performance/cost ratio. Yesterday's supercomputer is today's personal computer; today's leading-edge communications technology will be among tomorrow's mainstream capabilities.

Information technology evolves as new and valuable applications are found for hardware that gets steadily more powerful and cheaper. To benefit, users need affordable hardware, but they also need the software that implements the new applications. Yet learning how to build software takes many years of experimentation. If this process starts only when the hardware has already become cheap, the benefits to users will be delayed by years. Research needs to treat today's expensive equipment as a time machine, learning how it will be used when it is cheap and widely available, as it surely will be tomorrow. Knowing how to use computers for new tasks sooner can help many industries to become more competitive.

To date, the HPCCI's focus has been mainly on speed, but speed is not the only measure of high performance. Both speed, measured today in billions of operations per second or billions of bits per second, and scale, measured by the number of millions of users served, are important research issues. However, for the nation's information infrastructure, scale now seems more difficult to achieve. Information technology can be thought of as a tent, with the height of the center pole as speed and the breadth of the base as scale. Widening the tent to allow more work on scale without decreasing the work on speed requires more cloth; with the same resources, widening the tent would sacrifice research on speed for research on scale. This report recommends ways to reallocate funds within the HPCCI so as to accommodate greater emphasis on scale.

Accomplishments to Date

The HPCCI has focused mainly on parallel computing, networking, and development of human resources. Building on progress in research begun before the HPCCI, work and accomplishments to date reveal two key trends: better computing and computational infrastructure and increasing researcher-developer-user synergy.

Despite the difficulty of measuring impact at this early stage, it is the committee's judgment that the HPCCI has been generally successful so far. That assessment is necessarily qualitative and experiential now. Because the HPCCI is only 3 years old, results that can be measured in dollars should not be expected before the next decade.

The HPCCI has contributed substantially to the development, deployment, and understanding of computing and communications facilities and capabilities as infrastructure. It has helped transform understanding of how to share resources and information, generating proofs of concept and understanding that are of value not only to the scientific research community but also to the economy and society at large.

In parallel computing the fundamental challenge is not building the machines, but learning how to program them. Pioneering users and their software developers must be motivated by machines that are good enough to reward success with significant speedups.[3] For this reason, a great deal of money and effort have had to be spent to obtain parallel machines with the potential to run much faster than existing supercomputers. From the base built by the HPCCI, much has been learned about parallel computing.

The HPCCI has fostered productive interactions among the researchers and developers involved in creating high-performance computing and communications technology and researchers who use the technology. Building on the varying perspectives of the three groups, complex problems are being solved in unique ways. In particular, the HPCCI has funded cross-disciplinary teams associated with the Grand Challenge projects to solve complex computational problems and produce essential new software for the new parallel systems.

More specifically, the HPCCI has:

- Increased the nation's stock of expertise by educating new students and attracting new researchers;

- Made parallel computing widely accepted as the practical route to achieving high-performance computing;

- Demonstrated the feasibility of and initiated deployment of parallel databases;

- Driven progress on Grand Challenge problems in disciplines such as cosmology, molecular biology, chemistry, and materials science. Parallel computation has enhanced the ability to attack problems of great complexity in science and engineering;

- Developed new modes of analyzing and visualizing complex data sets in the earth sciences, medicine, molecular biology, and engineering, including creating virtual reality technologies. Many supercomputer graphic techniques of the 1980s are now available on desktop graphics workstations;

- Through the gigabit network testbeds associated with the National Research and Education Network component, demonstrated the intimate link between computing and communications systems;

- Built advanced networks that are the backbone of the Internet and the prototypes for its further evolution into the basis for a broader information infrastructure;

- Deployed a high-speed backbone that has kept up with the yearly doubling of the size of the Internet, and organized the impending transition of this backbone away from government funding; and

- Created the Mosaic browser for the World Wide Web, the first new major application in many years that promises to greatly increase access to the resources available on the Internet. This was an entirely unexpected result.

Evolution

A large-scale, integrated information infrastructure designed to serve the entire nation is becoming a high priority for government and industry as well as a source of challenges for research. Complex systems with millions of users pose many problems: performance, management, security, interoperability, compatible evolution of components, mobility, and reliability are only a few. Today's technology can solve these problems for systems with a few thousand users at most; to do so for millions or hundreds of millions of users is far beyond the current state of the art.[4] Providing users with high-bandwidth connections is itself a major problem, but it is only the beginning. There is a wide gap between enabling a connection and providing a rich array of useful and dependable services.

Because the HPCCI has become the rubric under which virtually all of the nation's research in information technology is conducted, it is not surprising that its focus has been changing in response to past successes, new opportunities, and evolving societal needs. The recently added Information Infrastructure Technology and Applications (IITA) program, broadly construed, addresses many of the problems just mentioned; it is already the largest component of the HPCCI,[5] and its continued evolution should be encouraged.

But with the policy focus—in the government, the press, and in most of the agencies—centered on information infrastructure,[6] high-performance computing seems to have been downplayed. The committee emphasizes the importance of retaining the HPCCI's momentum at just the time when its potential to support improvement in the nation's information infrastructure is most needed.

Organization

Several federal agencies participate in the HPCCI, most notably the National Science Foundation (NSF), the Advanced Research Projects Agency (ARPA), the National Aeronautics and Space Administration (NASA), and the Department of Energy (DOE) (see Appendix A for a full list). Because of its successes the HPCCI has become a model for multiagency collaboration and for the "virtual agency" concept advanced through the *National Performance Review* (Gore, 1993). Each participating agency retains responsibility for its own part of the program, but the agencies work together in joint funding of projects, such as the federal portions of the Internet; joint reviews of grants and contracts, such as the NSF-ARPA-NASA digital library initiative; joint testbeds; and consortia, such as the consortium for Advanced Modeling of Regional Air Quality that joins six federal agencies with several state and local governments.

The HPCCI supports a diverse set of contractors at universities, companies, and national laboratories throughout the country. It provides project funding in varying amounts through contracts, grants, and cooperative agreements awarded according to diverse methods. This diversity is healthy because it allows many views to compete, resulting in a broad research program that ensures a continuing flow of advances in information technology.

Some have argued for a more centrally managed program, with thorough planning, precise milestones, and presumably no wasted effort. Tighter management would cost more in bureaucracy and turf wars, but the essential question is whether it would produce better or worse results for the money spent. The committee believes that because of the long time scale of research, diversity is essential for success. No one person or organization is either smart or lucky enough to plan the best program, no single approach is best, and success often comes in unanticipated ways. Because it is a national research program and because of the many different but interdependent underlying technologies, the HPCCI is necessarily and properly far more diverse than a focused effort such as the Apollo moon landing program or a commercial product development program.

In contrast to central management, coordination enhances the benefits of diversity by helping to prevent unintended duplication, redundancy, and missed opportunities. The HPCCI's National Coordination Office (NCO) serves this purpose, aiding interagency cooperation and acting as liaison for the initiative to the Congress, other levels of government, universities, industry, and the public. Its efforts are reflected in its impressive FY 1994 and FY 1995 "Blue Books" describing the program's activities and spending.[7] Strengthening the NCO and appointing an advisory committee, as recommended in the committee's interim report (CSTB, 1994c), would facilitate regular infusions of ideas and advice from industry and academia and enable better communication of the HPCCI's goals and accomplishments to its many constituents. This committee should consist of a group of recognized experts that is balanced between academia and industry and balanced with respect to application areas and the core technologies underlying the HPCCI.

Budget

Because it grew from earlier programs, a significant portion of the HPCCI budget is not new money. The budget grew from a base of $490 million in preexisting funding in FY 1992 to the $1.1 billion requested for FY 1995.[8] Each year agencies have added to the base by moving budgets for existing programs into the HPCCI and by reprogramming existing funds to support the HPCCI. Congress has also added funding each year to start new activities or expand old ones.

The result is that much of the $1.1 billion requested for FY 1995 is money that was already being spent on computing and communications in FY 1992. The request has three elements: (1) funds for activities that predate the HPCCI and were in the FY 1992 base budget, (2) funds for activities that have since been designated as part of the HPCCI, and (3) new funds for new activities or for growth. Although dissecting the budget in this way would shed light on the program, the committee was unable to do so because each participating agency treats the numbers differently.

It appears that the FY 1995 request breaks down roughly as one-third for applications, one-third to advance the essential underlying computing and communications technologies, one-quarter for computing and communications infrastructure, and small amounts for education and electronics (see Appendix C).

THE FUTURE OF THE HPCCI: RECOMMENDATIONS

The committee believes that strong public support for a broadly based research program in information technology is vital to maintaining U.S. leadership in information technology. Incorporating this view of the importance and success of the government's investment in research, the 13 recommendations that follow address five areas: general research program, high-performance computing, networking and information infrastructure, the supercomputer centers and the Grand Challenge projects, and program coordination and management. Within each area the recommendations are presented in priority order.

General Recommendations

1. Continue to support research in information technology. Ensure that the major funding agencies, especially the National Science Foundation and the Advanced Research Projects Agency, have strong programs for computing and communications research that are independent of any special initiatives.

The government investment in computing research has yielded significant returns. Ongoing investment, *at least as high as the current dollar level*, is critical both to U.S. leadership and to ongoing innovation in information technology. Today the HPCCI supports nearly all of this research, an arrangement that is both misleading and dangerous: misleading because much important computing research addresses areas other than high performance (even though it may legitimately fit under the new IITA component of the HPCCI), and dangerous because reduced funding for the HPCCI could cripple all of computing research. The "war on cancer" did not support all of biomedical research, and neither should the HPCCI or any future initiative on national infrastructure subsume all of computing research.

2. Continue the HPCCI, maintaining today's increased emphasis on the research challenges posed by the nation's evolving information infrastructure. The new Information Infrastructure Technology and Applications program of the HPCCI focuses on information infrastructure topics, which are also addressed in the initiative's other four components. The committee supports this continued evolution, which will lead to tangible returns on existing and future investments in basic hardware, networking, and software technologies.

High-Performance Computing

3. Continue funding a strong experimental research program in software and algorithms for parallel computing machines. Today a crucial obstacle to widespread use of parallel computing is the lack of advanced software and algorithms. Emphasis should be given to research on developing and building usable applications-oriented software systems for parallel computers. **Avoid funding the transfer ("porting") of existing commercial applications to new parallel computing machines unless there is a specific research need.**

4. Stop direct HPCCI funding for development of commercial hardware by computer vendors and for "industrial stimulus" purchases of hardware. Maintain HPCCI support for precompetitive research in computer architecture; this work should be done in universities or in university-industry collaborations and should be driven by the needs of system and application software. HPCCI funding for stimulus purchase of large-scale machines has been reduced, as has the funding of hardware development by vendors. The committee supports these changes, which should continue except when a mission need demands the development of nonstandard hardware.

Public research is best done in universities. Not only are academic organizations free to think about longer-term issues, but they also stimulate technology transfer through publication and placement of graduates. The national experience supports Vannevar Bush's basic tenet: publicly funded research carried out in universities produces a diversity of excellent ideas, trained people, research results, and technologies that can be commercially exploited (OSRD and Bush, 1945).

5. Treat development of a teraflop computer as a research direction rather than a destination. The goal of developing teraflop capability has served a valuable purpose in stimulating work on large-scale parallelism, but further investment in raw scalability is inappropriate except as a focus for precompetitive, academic research. Industrial development of parallel computers will balance the low cost of individual, mass-produced computing devices against the higher cost of communicating between them in a variety of interesting ways. In the near future a teraflop parallel machine will be built when some agencies' mission requirements correspond to a sufficiently economical commercial offering. Continued progress will surely lead to machines even larger than a teraflop.

Networking and Information Infrastructure

New ideas are needed to meet the new challenges underlying development of the nation's information infrastructure. The HPCCI can contribute most by focusing on the underlying research issues. This shift has already begun, and it should continue.

This evolution of the research agenda, which would support improvement of the nation's information infrastructure, is partly under way: in the *FY 1995 Implementation Plan* (NCO, 1994, p. 15), over one-quarter of the NSF and ARPA HPCCI funding is focused on the IITA component, and activities in other components have also evolved consistent with these concerns. The committee supports this increased emphasis.

6. Increase the HPCCI focus on communications and networking research, especially on the challenges inherent in scale and physical distribution. An integrated information infrastructure that fully serves the nation's needs cannot spring full-grown from what we already know. Much research is needed on difficult problems related to size, evolution, introduction of new systems, reliability, and interoperability. Much more is involved than simply deploying large numbers of boxes and wires. For example, both hardware and software systems must work efficiently to handle scheduling; bandwidth optimization for transmission of a range of data formats, including real-time audio and video data; protocol and format conversion; security; and many other requirements.

7. Develop a research program to address the research challenges underlying our ability to build very large, reliable, high-performance, distributed information systems based on the existing HPCCI foundation. Although a comprehensive vision of the research needed for advancing the nation's information infrastructure has not yet been developed, three key areas for research are scalability, physical distribution, and interoperative applications.

8. Ensure that research programs focusing on the National Challenges contribute to the development of information infrastructure technologies as well as to the development of new applications and paradigms. This dual emphasis contrasts with the narrower focus on scientific results that has driven work on the Grand Challenges.

Supercomputer Centers and Grand Challenge Program

The NSF supercomputer centers have played a major role in establishing parallel computing as a full partner with the prior paradigms of scalar and vector computing by providing access to state-of-the-art computing facilities. NSF should continue to take a broad view of the centers' mission of providing access to high-performance computing and communications resources, including participating in research needed to improve software for parallel machines and to advance the nation's information infrastructure.

The committee recognizes that advanced computation is an important tool for scientists and and engineers and that support for adequate computer access must be a part of the NSF research program in all disciplines. The committee did not consider the appropriate overall funding level for the centers. However, the committee believes that NSF should move to a model similar to that used by NASA and DOE for funding general access to computing. The committee prefers NASA's and DOE's approach to funding supercomputer centers, where HPCCI funds are used only to support the exploration and use of new computing architectures, while non-HPCCI funds are used to support general access.

9. The mission of the National Science Foundation supercomputer centers remains important, but the NSF should continue to evaluate new directions, alternative funding mechanisms, new administrative structures, and the overall program level of the centers. NSF could continue funding of the centers at the current level or alter that level, but it should continue using HPCCI funds to support applications that contribute to the evolution of the underlying computing and communications technologies, while support for general access by application scientists to maturing architectures should come increasingly from non-HPCCI funds.

10. The Grand Challenge program is an innovative approach to creating interdisciplinary and multi-institutional scientific research teams; however, continued use of HPCCI funds is appropriate only when the research contributes significantly to the development of new high-performance computing and communications hardware or software. Grand Challenge projects funded under the HPCCI should be evaluated on the basis of their contributions both to high-performance computing and communications technologies and to the application area. Completion of the Grand Challenge projects will provide valuable insights and demonstrate the capabilities of new high-performance architectures in some important applications. It will also foster better collaboration between computer scientists and computational scientists. The committee notes that a large share of HPCCI funding for the Grand Challenges currently comes from the scientific disciplines involved. However, the overall funding seems to come entirely from HPCCI-labeled funds. For the same reasons outlined in Recommendation 9, the committee sees this commingled support as unhealthy in the long run and urges a transition to greater reliance on scientific disciplinary funding using non-HPCCI funds.

Coordination and Program Management

11. Strengthen the HPCCI National Coordination Office while retaining the cooperative structure of the HPCCI and increasing the opportunity for external input. Immediately appoint the congressionally mandated advisory committee intended to provide broad-based, active input to the HPCCI, or provide an effective alternative. Appoint an individual to be a full-time coordinator, program spokesperson, and advocate for the HPCCI.

In making this recommendation, the committee strongly endorses the role of the current NCO as supporting the mission agencies rather than directing them. The committee believes it vital that the separate agencies retain direction of their HPCCI funds. The value of interagency cooperation outweighs any benefits that might be gained through more centralized management.

Diverse management systems for research should be welcomed, and micromanagement should be avoided. In the past, choosing good program officers and giving them freedom to operate independently have yielded good value, and the committee believes it will continue to do so.

Furthermore, independence will encourage diversity in the research program, thus increasing opportunities for unexpected discoveries, encouraging a broader attack on problems, and ensuring fewer missed opportunities.

12. Place projects in the HPCCI only if they match well to its objectives. Federal research funding agencies should promptly document the extent to which HPCCI funding is supporting important long-term research areas whose future funding should be independent of the *future* of the HPCCI.

A number of preexisting agency programs have entered the HPCCI, with two effects: the HPCCI's budget appears to grow faster than the real growth of investment in high-performance computing and communications research, and important programs such as basic research in computing within NSF and ARPA may be in jeopardy should the HPCCI end.

13. Base mission agency computer procurements on mission needs only, and encourage making equipment procurement decisions at the lowest practical management level. This recommendation applies equally to government agencies and to government contractors. It has generally been best for an agency to specify the results it wants and to leave the choice of specific equipment to the contractor or local laboratory management.

NOTES

1. See U.S. DOC (1994); the Department of Commerce utilizes data from the U.S. Bureau of the Census series, the *Annual Survey of Manufactures*. It places the value of shipments for the information technology industry at $421 billion for 1993. This number omits revenue from equipment rentals, fees for after-sale service, and mark-ups in the product distribution channel. It also excludes office equipment in total. It includes computers, storage devices, terminals and peripherals; packaged software; computer program manufacturing, data processing, information services, facilities management, and other services; and telecommunications equipment and services.

See also CBEMA (1994); CBEMA values the worldwide 1993 revenue of the U.S. information technology industry at $602 billion. In addition to including office equipment, it shows larger revenues for information technology hardware and telecommunications equipment than does the Department of Commerce.

2. Microcomputers (personal computers) are defined as computers with a list price of $1,000 to $14,999; see CBEMA (1994), pp. 60-61. Forrester Research Inc. (1994, pp. 2-3) estimates the share of households with PCs at about 20 percent, based on its survey of households and Bureau of Census data. Forrester's model accounts for retirements of older PCs and for households with multiple PCs. This is a lower estimate than the Software Publishing Association's widely cited 30 percent share. By definition, the microcomputer statistics exclude small computers and other general-purpose and specialized devices that also make use of microprocessors and would be counted in a more comprehensive measurement of information technology.

3. Earlier experience with three isolated computers, "Illiac 4" (built at the University of Illinois) and "C.mmp" and "Cm*" (both built at Carnegie Mellon University), bears out this point.

4. Of course, systems specialized for a single application or for homogenous technology, such as telephony, serve millions of users, but what is now envisioned is more complex and heterogenous, involving integration of multiple services and systems.

5. The other four programs of the HPCCI are Advanced Software Technology and Algorithms, Basic Research and Human Resources, High-Performance Computing Systems, and the National Research and Education Network.

6. Notably, references to the computing portion of the HPCCI have been overshadowed recently by the ubiquity of speeches and documents devoted to the notion of a national information infrastructure (NII). The NII has also been featured in the titles of the 1994 and 1995 Blue Books.

7. Each year beginning in 1991 the director of the Office of Science and Technology Policy submits a report on the HPCCI to accompany the president's budget. The FY 1992, FY 1993, and FY 1994 books were produced by the now-defunct Federal Coordinating Council for Science, Engineering, and Technology; the FY 1995 report was produced by the NCO (acting for the Committee on Information and Communications). The report describes prior accomplishments and the future funding and activities for the coming fiscal year. These reports have collectively become known as "Blue Books" after the color of their cover.

8. NCO (1994), p. 15. Note that figures represent the President's requested budget authority for FY 1995. Actual appropriated levels were not available at press time. Because the HPCCI is synthesized as a cross-cutting multiagency initiative, there is no separate and identifiable "HPCCI appropriation."

COMMITTEE TO STUDY HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS: STATUS OF A MAJOR INITIATIVE

FREDERICK P. BROOKS, JR., University of North Carolina at Chapel Hill, *Co-chair*
IVAN E. SUTHERLAND, Sun Microsystems Laboratories, *Co-chair*
ERICH BLOCH, Council on Competitiveness
DEBORAH ESTRIN, University of Southern California/Information Sciences Institute
JOHN HENNESSY, Stanford University
BUTLER W. LAMPSON, Digital Equipment Corporation
EDWARD D. LAZOWSKA, University of Washington
WILLIAM A. LESTER, JR., University of California at Berkeley
JANE PRESTON, Telemedical Interactive Consultative Services Inc.
W. DAVID SINCOSKIE, Bell Communications Research Inc.
LARRY SMARR, National Center for Supercomputing Applications/University of Illinois at Urbana-Champaign
JOSEPH F. TRAUB, Columbia University

Staff

MARJORY S. BLUMENTHAL, Director
JAMES E. MALLORY, Staff Officer
LESLIE M. WADE, Project Assistant

COMPUTER SCIENCE AND TELECOMMUNICATIONS BOARD