Testimony Hearings on the President's Information Technology Advisory Committee, Interim Report to the President

Details

October 6, 1998

Testimony before the U.S. House of Representatives
Committee on Science, Subcommittee on Basic Research

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

Introduction

Chairman Pickering and distinguished members of the Basic Research Subcommittee, my name is Ken Kennedy. I am a professor of Computer Science at Rice University, where I direct the center for Research on Parallel Computation, a National Science Foundation Science and Technology center. I also serve as the academic co-chair of the President's Information Technology Advisory Committee. It is in the latter capacity that I appear before you today to discuss the interim report of the Advisory Committee and its implications for high-end computation. In April of this year, I received a letter from Chairman Sensenbrenner and Ranking Minority Member Brown of the House Committee on Science, which asked several questions about high end computing. On behalf of the committee, my co-chair Bill Joy and I responded in a letter dated August 17, 1998, a copy of which is attached to the testimony. As a part of my testimony today I will review that response. However, before turning to specifics, let me provide you with some background on our committee and a report on its recent activities.

The Advisory Committee

The Advisory Committee was appointed by the President and charged with the task of reviewing the research and development programs in information technology funded by the Federal government to determine whether these programs are meeting the needs of the Nation. A specific concern is whether research and development programs are helping to maintain United States leadership in advanced computing and communications technologies and their applications.

To address this question, we conducted a yearlong review of Federal information technology research and development programs. One focus of this review was to understand the balance in the program, both among different areas of investigation and among different categories of research. In carrying out the evaluation, our principal strategy was to have Federal funding program managers brief the committee on their activities and to question them on the directions and risk levels of the research being funded. Our preliminary findings are contained in the President's Information Technology Advisory Committee Interim Report to the President, dated August 6, 1998.

The Interim Report

The principal finding of the Interim Report was that over the past decade there has been a pronounced shift in Federal funding programs away from long-term high-risk projects toward short-term, applied research. This shift has happened for understandable reasons. The majority of funding for information technology research is allocated to mission agencies. During the past decade, the funding for information technology has grown at about the rate of inflation, while the size of the information technology endeavor, as measured by its impact on the economy and the number, size, and scope of the problems it can solve, has experienced explosive growth. Information technology has accounted for over a third of the growth of the Nation's gross domestic product since 1992 and some estimates indicate that it now represents as much as 46 percent of the economy today. To quote Federal Reserve Board Chairman Alan Greenspan's recent testimony before Congress: "...our nation has been experiencing a higher growth rate of productivity-output per hour worked-in recent years. The dramatic improvements in computing power and communication and information technology appear to have been a major force behind this beneficial trend."

But the impact of information technology on our society goes far beyond economic benefit. It is critical to the solution of problems in business, science, medicine, and education. It is also critical to the operations of most government agencies. When faced with rapidly expanding requirements for information technology research and provided with relatively flat budgets, mission agency mangers-understandably and correctly-give priority to the short-term needs of their mission. Given that information technology receives only about one in every seventy-five dollars of the total Federal research and development investment, it is not too surprising that a larger and larger fraction of it has been focused on short-term needs.

The Advisory Committee believes that unless this shift away from fundamental high-risk research is reversed, it will threaten the health and welfare of the Nation, along with its economic leadership, in the coming decades. In the past, sound federal R&D investment strategies, carried out over long periods of time, have laid the foundations for dramatic advances in many diverse areas including farm production, aeronautics, space, and health disciplines. Similarly, the funding for federal R&D in information

technology has been instrumental in creating the boom in computing and communication which is responsible for much of the current national economic growth as well as for major advances in health care, public safety, and other critical areas. This relationship between information technology research and economic prosperity a decade later has been clearly documented by the Brooks-Sutherland report, "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure" (NRC 1995). Thus, if it is left unchecked, the trend away from innovative, long-term research in information technology will interrupt the flow of ideas that are needed to fuel the information economy and solve critical national problems in the twenty-first century.

It is tempting to think that this problem could be addressed by simply reallocating existing information technology R&D funds. However, this would be shortsighted. The short-term problems to which funding has been redirected represent critical national needs. We cannot simply abandon work on secure computers and networks, supercomputer software, or the Year 2000 problem. Furthermore, the nature of the information technology industry makes it unlikely that corporations will fill the void. The primary engine for innovation in the US economy is the venture start-up, which cannot afford to engage in long-term research. In fact, most start-ups are funded to capitalize on ideas that have spun out of research programs in universities and government labs. Nor will larger information technology corporations solve the problem, because these companies live with profit margins that are so slim that they cannot afford to engage in much research that pays out on a schedule beyond 2 to 4 years.

To address these problems, the Advisory Committee has recommended doubling the annual information technology R&D budget to two billion dollars over a period of five years starting in FY 2000, with a substantial fraction of the proposed increase committed to high-risk long term research. We envision this work going to the kinds of projects supported by the Defense Advanced Research Projects Agency in the sixties, seventies and eighties. Our report summarizes many areas of computing and communications that could make dramatic strides forward with increased support. Three topics of particular importance are:

- **Software**: Methods for efficiently creating and maintaining high-quality software of all kinds and for ensuring the reliability of the complex software systems that now provide the infrastructure for much of our government and our economy.

- **Scalable Information Infrastructure**: Techniques for ensuring that the National Information Infrastructure-consisting of communications systems, the Internet, large data repositories, and other emerging systems-is reliable and secure, and can grow gracefully to accommodate the massive numbers of new users (perhaps billions) and applications expected over the coming two decades.

- **High End Computing and Communications**: Continued invention and innovation in the development of fast, powerful computing systems and the accompanying communication

systems needed to implement critical science, engineering, and business applications ranging from aircraft design to weather and climate modeling.

In addition, the report recommends that funding be directed toward understanding the sociological and economic impacts of innovations in information technology, and toward growing the workforce to meet the national need for information technology professionals. A significant component of the latter would go to programs for retraining existing professionals and programs to increase the number of people in the educational pipeline, particularly underrepresented minorities and women. Although there is evidence that the number of college students interested in careers in information technology is growing rapidly, we will not have the resources to properly educate all of them without substantive increases in the size of college and university faculties in computer science, applied mathematics, and computer engineering. This can only happen if more students choose to go to graduate school and pursue careers in teaching. The increased research funding we envision could help achieve this goal by making graduate schools and universities a more exciting option for undergraduate students from the U.S., who are increasingly choosing careers in industry over those in academia.

Not only is the overall level of funding for information research inadequate, but the current system for managing information R&D across all federal agencies is not ideally suited to the task. In the past, information technology has been considered largely the byproduct of research in other areas. More recently funding for information technology research has been combined with programs whose primary function has been to fund infrastructure and technology transfer. This must change if we are to make progress on fundamental problems in computing and communications. We recommend developing new approaches to management and funding of civilian information technology R&D, with the goal of providing adequate levels of multi-year funding to pursue longer-range research agendas. Both large multi-investigator projects, such as centers, and smaller efforts involving teams of a few researchers should be supported. While the new program should be responsive to application needs, it must have information technology as its fundamental objective. Moreover, the program needs to be managed in a way that allows greater flexibility and more risk-taking.

One compelling idea proposed in our report is to establish investigations of promising future technologies by geographically distributed teams of information technology researchers and other scientists and engineers. These "Expeditions to the 21st Century" would permit researchers to "live in the technological future" and report back to the Nation on what could be accomplished if high-risk technologies are developed to maturity. These efforts would include investigation of the sociological and economic impacts of these technologies. The report compares these virtual centers to the Lewis and Clark expedition, which opened up our nation to unanticipated expansion and economic growth. If these expeditionary centers are as successful as similarly bold projects of the past-such as Project MAC, Xerox PARC, and the Internet itself-they

could help extend the boom in information technology that Chairman Greenspan has identified as a major force behind the overall growth in the economy.

The US is approaching the 21st century with one of the most prosperous economies in history. Information technology presents enormous opportunities for growing the economy and improving health care, education, public safety, the environment, and many other areas of importance to the Nation. To capitalize on these opportunities, we should be increasing, federal investments in this critical area to levels commensurate with its importance to the Nation and to the government. We urge you to take steps now to ensure that the people of the United States continue to enjoy the fruits of the information revolution into the next millennium.

Implications for the High End

As I indicated earlier, high end computing is an important focus for increased investment by the Federal government. It is of special importance to government because it is needed to solve critical national problems, particularly in defense. Yet in spite of their importance, high-end computers do not have a large enough market to sustain even the medium-term research and development needed to continue the pace of dramatic progress in science, engineering, and applications critical to the Federal government. Because of this, computer companies have adopted the approach of using large numbers of commodity chips to build high-performance processors. This strategy, often referred to as "scalable parallel computation," was a major focus of the HPCC program and is the primary approach of most U.S. companies that now produce systems at the high end.

However, this strategy has a number of problems. Scalable parallel computers are much harder to use than previous generations of parallel machines. Although the HPCC program was extremely successful in bringing about dramatic increases in computational power, it did not have sufficient resources to develop all the software needed to make these machines generally applicable. This "usability gap" is compounded by evidence that many applications, including some critical to our national defense, cannot be made to work well on scalable parallel machines as currently implemented. In particular, scalable parallel machines do not do well on problems that require high rates of data movement between memory and processors.

These difficulties cannot be addressed by simply purchasing more computers for scientists. Such purchases must be balanced against investments in the long-term research that will overcome the limitations of today's high-end architectures. To ensure the continued pace of growth of the power of high-end computer systems and their effective use for science and engineering, the Advisory Committee Interim Report recommends five specific actions:

1. Fund research into innovative architectures and new computing technologies that overcome the limitations of today's high-end systems.

2. Increase support for R&D on software technologies that will improve the performance, range of applicability and usability of high-end systems.

3. Drive the high-end hardware and software computing research by establishing the goal of a thousandfold increase in the sustained performance of applications by the year 2010. This goal will force major paradigm shifts in both architecture and software.

4. Fund the acquisition of state-of-the-art high-end computing systems to support science and engineering research and ensure that these systems are networked and available to the civilian and government research communities.

5. Expand the Federal High End Computing and Computation program to include all the major elements of the government's investment in high-end computing.

Let me dwell on the final two for a moment. I suspect that some of the impetus for today's hearing came from the stories that circulated this spring about the power and usability of computer systems available to the civilian scientists. These systems were contrasted with those available in the DOE ASCI program for stockpile stewardship. The ASCI machines are by most estimates an order of magnitude more powerful than those available in the NSF centers.

The concerns about the dwindling relative power of computers available to civilian scientists resulted in proposals to replicate ASCI-class facilities for use in civilian science. I can say without hesitation that the Advisory Committee would support such an initiative. Our report states clearly that the Federal Government should purchase the high-end computer systems that scientists, both within and outside government, need to do their work. However, I can also state unequivocally that such investments will only be effective if accompanied by larger investments in software and architectures that will be needed to make these machines usable for the entire spectrum of high-end applications. It is for this reason that we believe that no more that twenty to twenty five percent of any increase in information technology funding should go to acquisition of high-end facilities.

The April letter from the House Science Committee to the Advisory Committee posed three specific questions:

- [Are] the overall level of resources for the HECC component [of the CIC R&D program] adequate and are they allocated appropriately?
- [Is] the HECC component coordinated effectively with other high performance computing activities of Federal mission agencies?
- [Will] the overall Federal R&D investment in high performance computing ensure U.S. leadership in computing technology?

Our answer to all three of these questions is "no." However, the Advisory Committee's interim report proposes strategies that can overcome all of these difficulties. A dramatic program of investment in research and development on high-end computing technologies and software, substantive investments in the purchase of facilities for science, and a well-coordinated management strategy can ensure that the United States continues to lead the world in the effective application of high-end computing and communications.

In short, we cannot overcome the shortage of high-end computing cycles by a one-time infusion of funding for computer facilities. The usable lifetime of any facility purchased today will two to four years. To ensure that our scientists continue to have access to computing facilities that will permit them to look far into the future, we must invest in the research that will produce those facilities five, ten, and fifteen years from now.

Conclusions and Next Steps

Subsequent to our final report we have initiated a process for refining its recommendations with a goal of finishing the process by February 1999, when the two-year terms of most Advisory Committee members expire. At the same time, many of us have been working with Dr. Lane, the Office of Science and Technology Policy, and the Office of Management and Budget to help them respond to the President's call for an increased investment in information technology R&D in the FY 2000 budget.

The Advisory Committee believes that increasing the investment in information technology research and development, with an emphasis on fundamental, long-term research, is the best way to ensure that the economic and social benefits of the information revolution will continue to be enjoyed by the Nation in the decades to come. On behalf of my co-chair, Bill Joy, and the entire Advisory Committee, I thank Chairman Pickering and all the members of the Basic Research Subcommittee for the opportunity to address you on these concerns. We welcome your comments and questions on our report, and look forward to working with you to make its vision a reality.