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

Edward D. Lazowska
Chair, Department of Computer Science & Engineering, University of Washington, and
Chair, Board of Directors, Computing Research Association

Introduction

Mr. Chairman and members of the subcommittee, thank you for the opportunity to testify before you today on the health of the nation's computing research program. My name is Edward D. Lazowska. I am Chair of the Department of Computer Science & Engineering at the University of Washington, and Chair of the Board of Directors of the Computing Research Association. I also chair the National Science Foundation's Advisory Committee for the Computer and Information Science and Engineering Directorate, and sit on the Computer Science and Telecommunications Board of the National Research Council.

I am testifying today in my role as Chair of the Computing Research Association. The Computing Research Association (CRA) represents nearly 200 academic departments of computer science and computer engineering and industrial and non-profit laboratories that engage in fundamental computing research. The major professional societies in the computing fields are also CRA affiliates. Including the members of those affiliates, CRA represents a combined membership of more than 100,000 professionals involved in computing.

It is a pleasure to appear before this subcommittee to discuss the state of computing research and to give a computer science and engineering perspective on the President's Information Technology Advisory Committee (PITAC) Interim Report to the President. The House Subcommittee on Basic Research has a long and distinguished record in encouraging and shaping the nation's efforts to develop computing and computational science and engineering. This commitment dates back to your critical work in building the foundations for and overseeing the original High Performance Computing Act (HPCA), and was most recently reflected in your work on the Next Generation Internet Research Act of 1998 (H.R.

3332), recently passed by the House.

In all of this important work, you have demonstrated an awareness of four important facts:

- **A leading-edge, high-performance computing and communications infrastructure is vital to sustaining the nation's science and engineering effort.** Much of modern science and engineering cannot be done without access to such high-end systems.
- **It's more than just a matter of buying hardware and pulling fiber** in two different ways. First, keeping our infrastructure at the leading edge requires a program of fundamental research in a broad range of computing areas, not all of which have an obvious direct link to the visible challenges in high-end computing. And in addition, other fields of science and engineering need more than just hardware speed; for example, astronomers need research on data mining algorithms to discover patterns in the billions of bytes of observational data collected by telescopes, and astrophysicists need research on simulation algorithms to explore theories about the origins of the universe.
- **Computing and communications play essential roles far beyond serving as enabling technologies for science and engineering.** Fundamental advances in computing and communications are transforming commerce, education, employment, health care, manufacturing, government, and entertainment. Leadership in computing and communications is essential to our economy and to our society.
- **Analogously to the situation with leading-edge infrastructure, a broad-based program of fundamental research in computing and communications is required to support these other essential roles.** Today's applications are possible because of yesterday's innovations. Today's innovations will enable tomorrow's applications. Tomorrow's applications will not just be "elaborations on a theme," but in many cases, applications that we cannot even envision today.

This is why the computing research community has always stressed a balanced research portfolio. The HPCA incorporated such a balance by establishing "Basic Research and Human Resources" as one of its four priorities.

In response to your questions concerning the President's Information Technology Advisory Committee Interim Report, I would like to make two specific points.

Report gets the History and Status Exactly right

The PITAC Interim Report gets the history and status exactly right

The bottom-line messages of the PITAC Interim Report, as I interpret them, are the following:
Leadership in information technology is essential to leadership in the 21st century. As I have just noted, advances in computing and communications are transforming not only the conduct of science and engineering, but also our economic health, the efficiency of our nation's processes, and the lives of our citizens.

The conduct of science and engineering are being revolutionized not only by the application of computing (e.g., for simulation, visualization, and collaboration), but also by a synthesis of discoveries in computing with discoveries in other fields, such as biology.

Our economy is bolstered by the incredible robustness of the information technology sector itself, which now boasts an annual income of roughly $1 trillion. And, even greater economic gains have resulted from infusing the benefits of advances in information technology into all segments of government and business.

And the lives of our citizens are being transformed, in areas ranging from education to health care to entertainment, through advances in information technology.

Today, information technology pervades every activity in our society, making our lives better and our economy stronger.

We have had a spectacular return on past Federal information technology research investments. Many, if not most, of the key ideas that are driving the production and application of information technology today resulted from earlier government-funded research projects. Everyday technologies that we all take for granted the mouse, graphical user interfaces, the Internet, the World Wide Web, and many more have their roots in fundamental research programs funded by DARPA, NSF, and other science agencies.

This point has been well documented. In a small report published by CRA, Computing Research: A National Investment for Leadership in the 21st Century, five essays describe examples where fundamental university research in computing has led to the emergence of major new technology markets: databases, computer graphics, RISC computers, the Internet, and artificial intelligence. We could have selected many other examples as well. The 1995 National Research Council report, Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure, makes the same point more broadly, describing the complex, two-way flow of ideas between university research and industrial research and development.

Federal support for fundamental research in information technology has sorely lagged the growth of the field in terms of scientific, economic, and social importance. This support is "dangerously inadequate," in the language of the PITAC Interim Report despite the spectacular record of return on past investments, as the reports cited above make quite clear. Too many good research ideas are being turned away. As a result, we are losing opportunities to explore exciting new directions and, at the same time, adversely affecting the production of the highly trained specialists who will be in great demand in the next century.

Federal programs of support for computing research have become dangerously short-term. This trend may be due in part to the serious pressures on limited funds, a situation in which short-term projects with a predictable (even if modest) payoff, win out over longer-term and riskier ideas. It is precisely these longer-term and riskier projects that are most deserving of federal support.

Unlocking the Future, the National Science Policy Study, the report recently released by the House Committee on Science under the leadership of Congressman Vernon Ehlers, points out the crucial role of the federal science funding enterprise in supporting long-term "understanding-driven" research. Clearly, computers are vital to the missions of many government agencies, and those agencies will continue to pursue a research and development mission focused on their specific needs. But these shorter-term efforts are drawing from a well of ideas that must be continually replenished by fundamental research, lest it run dry.

As a result, "critical problems are going unsolved, and we are endangering the flow of ideas that has fueled the information economy."

Correctly Proposes a Balanced Program of Research

The PITAC Interim Report correctly proposes a balanced program of research

As my colleagues at this table have indicated in their testimony, high-end computing is of great importance to the research communities that will use it, to other users in industry and government who will apply the techniques developed by the research users, and to the nation that will benefit from the technological advances. We also agree that there is a vital role for the federal government because of the unusual and limited commercial market for high-end machines and their importance to some government missions.

My colleagues also have indicated that high-end computing and communication is much more than hardware and fiber. It is tempting to frame computational researchers' needs starkly faster machines and faster networks. The truth is, however, that even when narrowly construed, advancing the forefront of high-end computing and communications requires enormous and fundamental research advances in software as well as in hardware, and the needed progress in hardware in turn requires continued research advances in areas such as VLSI technology, computer-aided design tools, and computer architecture. (I should note that the programs being carried out by the NSF Partnerships for Advanced Computational Infrastructure are exemplars of this balanced approach to high-end computing and communication.)

And we agree that continued fundamental research advances in computing and communications are essential to many key areas of our nation's agenda, far beyond their critical contributions to the conduct of science and engineering.

Appropriately, then, the PITAC Interim Report identified four areas for particular attention:

Software: Our ability to create reliable and efficient software to do the work we need computers to do, lags far behind our ability to make computers run faster. Even more troublesome, we still desperately need to improve the process of creating complex software.

The need for software research arises both from the desire to make better use of new and faster computers, and in response to a wider array of problems society is facing. This committee has, I know, heard a great deal in recent months about the information technology worker shortage. It is a critical issue that needs addressing on many fronts. One partial solution is to gain a better handle on how to build very complex software systems. We know that simply throwing more person-hours at a large programming task may actually be counter-productive. Similarly, some of our nation's most serious vulnerabilities in its critical infrastructure are software-related, and occur not so much from the sophistication of the "bad guys," but from the inherent fragility of the software. A vastly improved fundamental understanding of software engineering is essential to creating systems that are less vulnerable in the first place.

Scalable information infrastructure: This is not a new point for this committee, as it was frequently raised in the debate over the NGI program.

Simply stated, we don't yet know how to build a global, high-speed information infrastructure that will serve vast numbers of users efficiently and reliably. Yet, that is what must happen if the Internet is to continue to grow and evolve into a true information utility. This is not merely a problem of making the communications "pipes" bigger and more efficient. It involves research in areas such as creating protocols; understanding how to interface very-high-speed networks with computers; and improving our ability to monitor the performance of the network and adapt its operation to changing demands. These are huge, complex problems that must be solved before we can build the networks we will need in the 21st century.

High-end computing: This is, of course, a focus of this hearing the need to improve the research community's access to very-high-speed computing. But even here, the focus must be on research in system architecture and software and algorithms and data structures, not just on providing access to today's best and fastest systems and their direct descendants.

Socioeconomic and workforce impacts: Experience has taught us that computer systems are not simply neutral tools that can be used or not with little consequence. Information technology has become a powerful agent of change with important implications for those touched by it. If that were not so, we would not be arguing with one another about control of cryptography, rewriting intellectual property law, filtering content on the web, or the enormous policy implications of electronic commerce.

These debates in the government policy arena, within private organizations, and even within homes might be resolved more quickly, or at least be better informed, if we understood more about the economic and social dynamics of these powerful technologies that we are creating so rapidly.

The PITAC also identifies the need for well-educated information workers. The more researchers succeed in creating new information technologies that are productive and useful, the more industry will need the services of people who are knowledgeable and skilled in using those technologies. One point I want to emphasize particularly is the tight link between academic research funding and our ability to produce the highly skilled specialists that will be needed.

Conclusion

As spectacular as the rapid growth in information technology has been, the best is potentially yet to come.

Advances in communication, such as low-cost wireless systems and very-high-speed Internet technology, together with advances in computing, such as novel software capabilities and the continued improvement of 50 percent per year in cost-performance that we should be able to achieve, create entirely new possibilities that we could only dream about just a few years ago. Information technology is not only providing new tools with which to do our work, but it is fundamentally transforming the nature of that work. The PITAC Interim Report lists several visions of those changes: new, more powerful forms of human communication; greater and more flexible access by people to the knowledge base of society; on-line, lifelong education available without limitations on age, geography, or accessibility; the transformation of medical care; and new forms of electronic commerce.

These visions may seem far off, but many are, in fact, just on the horizon. Yet, to realize this promise requires more, not less, investment in fundamental research on information technology if we are to build systems that are safe, reliable, efficient, and useable.

In science policy, at the end of all the debates over why and how and where, the bottom-line (and toughest) question always seems to focus on level of funding. The PITAC Interim Report suggests increasing support for computing research by $1 billion over the next five years. Given what we, as well as the PITAC, see as a severe shortfall in support at this time, we concur with that recommendation. The enormous potential for the nation justifies this investment.

We concur in this recommendation while fully understanding that, despite the easing of the past budget crunch, demands for funds are still very tight, and there are many conflicting views on how to direct future government spending. Nevertheless, we believe that our current economic strength rests in no small measure on past investments in computing and communications research, and that the future holds similar promise. As the new kid on the block, so to speak, computing research has had to fight from its inception for a share of the R&D pot. When research budgets were growing rapidly, that was not a problem. In more recent years, though, the zero-sum game has become more difficult, and setting priorities based on societal importance is not a popular concept in some science policy circles.

CRA, together with our colleagues in industry and academia, is committed to working toward the goal set by the PITAC. We know it will be difficult, and we recognize that a heavy burden of proof rests with us. But we feel that a renewed commitment to computing research, as called for in the report, is both an appropriate and necessary step toward assuring the future economic, social, and scientific strength of our nation. We look forward to working with this subcommittee to put substantial flesh on the bones of the recommendations in this report.

Thank you, Mr. Chairman, for the opportunity to testify today. I would be happy to address any questions you may have.