President's Information Technology Advisory Committee

August 13, 1998

Honorable F. James Sensenbrenner, Jr.
Chair, House Science Committee
House of Representatives
2332 Rayburn House Office Building
Washington, DC 20515-4909

Dear Chairman:

Thank you for your letter of April 12 and your interest in the activities of the President's Information Technology Advisory Committee (PITAC). For the past year, the Committee has been reviewing the High-End Computing and Computation (HECC) component of the Computing, Information, and Communications (CIC) Research and Development (R&D) program in the context of the entire Federal R&D investment in information technology. Although the PITAC is concerned that changes in the structure of both corporate and government support are putting US leadership in high-end computing at risk, we see this as only one aspect of a larger problem—the erosion of support for long-term fundamental research in information technology. In the past, this kind of research has provided the foundation for dramatic growth in the computing and communications industries, helped address problems of national importance, and contributed to the education of an entire generation of information technology professionals. If the nation is to continue the rapid pace of advances experienced over the past decade, we must reverse the drift away from fundamental research.

Before exploring this issue in more detail, let us respond to the specific questions raised in your letter:

(i) *[Are] the overall level of resources for the HECC component [of the CIC R&D program] adequate and are they allocated appropriately?*

Broadly speaking, there are two components to the HECC program, funding for fundamental and applied research in high-end computing systems and software and infrastructure to support high-end computation by the national community of scientists and engineers. In our view, neither of these components is receiving adequate resources.

One of the original goals of the Federal High Performance Computing and Communications (HPCC) Initiative, the predecessor to CIC R&D, was to extend U.S. technological leadership in high performance computing and communications. Over time, many participants in that program came to believe that the way to maintain leadership was to effect a paradigm shift from vector supercomputers to scalable parallel computers, sometimes called massively parallel processors (MPPs). Since scalable parallel processing could build on commodity technologies, it could ride the technology

curve to achieve enormous processing power. Furthermore, if applications were written to efficiently exploit parallelism, performance would smoothly scale from desktop to supercomputer.

For this goal to be realized, however, there would have to be major shifts in the way these machines were programmed to make effective use of the available parallelism. Furthermore, there would need to be innovations in hardware and software to provide adequate performance of data movement between memory and the processors, and finally there would need to be a new generation of programming tools to make it easier to reprogram applications for scalable machines. To address these needs, the program included a sizeable component of fundamental and applied research support. Unfortunately, the level of research support fell short of what was needed because (1) the problem proved to be much harder than originally thought, and (2) as agency budgets became tight, funds were shifted from long-term research to short-term development efforts so some applications would be available on time. The result is what might be called a "usability gap" for these machines. Although there is every reason to believe that this gap can be bridged, increased investments in research and development will be needed to do so.

The usability gap also affects the availability of high-end computing cycles to the civilian science community. When NSF restructured its supercomputer centers program, a major component was focused on enabling the technologies needed to make the transition to scalable parallel machines. However, this change was made without a corresponding increase in the budget, leaving fewer dollars to purchase increasingly expensive high-end systems. As a result, compared to the National laboratories and university centers in other countries, NSF center users have lost ground in terms of relative high end resources available, which threatens the nation's leadership in computational science and engineering.

Both of these areas need to be addressed if we are to continue the rapid progress in the solution of important societal and scientific problems through high-end computation. The Committee recommends that the Federal government recommit itself to the support of high-end computing through research on new computing architectures, technologies, and software that can overcome the limitations of current supercomputer designs. In addition, the Federal government should increase funding for the acquisition of the powerful high-end computing systems needed by the nation's science and engineering community, both within and outside the government.

(ii) *[Is] the HECC component coordinated effectively with other high performance computing activities of Federal mission agencies?*

To be truly effective, coordination must take place both within the HECC program and between HECC and other federal programs involving high-end computing, such as DOE's ASCI program.

Within HECC, coordinating HECC-designated programs works well. The coordination process, however, has only a limited ability to manage the programs in ways that would lead to greater government-wide efficiencies or to consider national benefits that transcend the specific needs of contributing agencies. Furthermore, since a substantial fraction of the HECC budget lies in mission agencies, resources tend to be refocused on mission priorities over time, contributing to the drift toward

short-term applied research. To remedy this weakness, our report recommends that an agency committed to fundamental research be designated as a lead for coordinating information technology research including HECC, with the mission to ensure that federal IT research is well directed and with the flexibility and authority to carry out that mission.

With regard to programs outside of HECC, we are particularly concerned about the lack of coordination between HECC and the large computer investments (both hardware and software) in the Departments of Energy (ASCI) and Defense (HPC Modernization). To address this problem, these and other similar agency programs – current or under development, such as DOE's Strategic Simulation Program -- should be coordinated with HECC by the same mechanisms currently used within HECC.

(*iii*) *[Will] the overall Federal R&D investment in high performance computing ensure U.S. leadership in computing technology?*

To maintain our leadership in high performance computing technology, the United States must effect a transition from technologies that have application exclusively at the high-end to strategies that are more compatible with the commercial market. With the end of the cold war, the market for the most powerful supercomputers has shrunk to the point that it is no longer viable for a company to support itself exclusively in that market. In the United States, the solution has been to build on commodity technologies by using scalable parallel computation. However, there have been significant problems with the transition to these machines. The usability gap, alluded to earlier, has slowed the translation of many applications to these systems. In addition, there is evidence that some applications will not perform well on these systems until bottlenecks in the memory hierarchy are overcome by improved hardware and software. The committee believes that it is imperative that we develop new, commercially viable computing technologies to help solve the nation's critical problems. However, we cannot move forward without overcoming the impediments to usability of these technologies. Thus, we must recommit ourselves to a program of increased research and development on high performance computing architecture and software along lines similar to that proposed in the recent HECC draft strategy. Such an increase will continue to pay benefits to our nation by making it possible to build computers of the highest performance largely from inexpensive but powerful commodity technologies.

Over the past few years, there has been real erosion of Federal research and development investments in information technology—HECC is only one example. Left uncorrected, this erosion threatens the nation's leadership in computing and communications technology and the economic boom it has fostered. The Brooks-Sutherland report, "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure" (NRC 1995), documented a clear relationship between investment in fundamental research in computer and computational science and commercial impact a decade later. Foundations for the current boom were laid a decade or more ago by the investments from DARPA, NSF, and other government and industrial entities. The erosion of government research funding, coupled with reductions in long-term funding by industry and a continuing drain of university faculty to industry, threaten to interrupt the stream of good ideas that have fueled the information economy and led to the solution of many crucial national problems in the 1990's.

To address this problem, the Committee recommends that funding for fundamental research in information technology increase dramatically over the next several years. A portion of that increase should help reinvigorate research in high-end computing, including architecture, software, and applications. In addition, the nation's non-defense high-end computational infrastructure should be brought back up to levels competitive with defense computing and with scientific computing facilities in other nations, notably Japan and Europe. However, the investment in computational infrastructure should be viewed as an investment in instrumentation for science and engineering as well as an investment in computing. In particular, it must be carefully balanced against the goal of increasing fundamental research support in all areas of information technology.

Thank you once again for your interest in our activities. The comments in this letter are excerpted and adapted from our Interim Report to the President, a copy of which is attached. Please do not hesitate to contact either of us or any member of the Committee if you wish further information or clarification.

Respectfully,

Bill Joy
Co-Chairman

Ken Kennedy
Co-Chairman

President's Information Technology Advisory Committee

Attachment

cc: Honorable George E. Brown, Jr.