United States Senate
Committee on Commerce, Science and Transportation
Subcommittee on Technology, Innovation and Competitiveness

Congressional Testimony

July 19, 2006

"High-Performance Computing"

Prepared Statement of Mr. Christopher Jehn

Good morning, Mr. Chairman and distinguished Members of the Committee. I am Christopher Jehn, Vice President, Government Programs of Cray Inc. I commend you for holding this hearing on high-performance computing, and I want to thank you for this opportunity to testify on behalf of Cray.

Cray's rich history began in 1972, when the legendary Seymour Cray, the "father of supercomputing," founded Cray Research. The first supercomputer the company built, the Cray-1, broke the world record for computational speed at the Los Alamos National Laboratory.

Cray continues that tradition today. We are a global leader in high performance computing, and we are the only company in the world solely focused on designing, building, and supporting the world's most powerful supercomputers.

Our computers are purposely built to address the most demanding scientific and engineering problems. We give scientists and engineers the ability to not only get answers faster but to ask new questions at the frontiers of scientific discovery.

Today, Cray's high performance computers are addressing key national security missions, helping to predict severe weather, fight forest fires, build safer cars, discover new medicines and uncover the secrets to fusion power and superconductivity.

As we discuss high performance computing today, I want to emphasize four points.

First, supercomputing is vitally important to the federal government. Federal agencies tell us this everyday. As the largest user of supercomputing, the federal government understands how necessary supercomputers are to fulfilling the requirements of government missions – from national defense and homeland security to scientific leadership. Agencies need supercomputing to help maintain military superiority, enable scientific research, advance technological development, and enhance industrial competitiveness. For decades, supercomputing has paved the way for real progress for federal agencies.

Supercomputing is also important to academic researchers and industry. As a key enabler for furthering science and technology, supercomputing has helped advance U.S. productivity and ability to compete in the global economy and to ultimately drive long-term economic growth.

In all these areas, the need for supercomputing is growing, and to sustain progress as it has for decades, the federal government, academic researchers and industry must have access to increasingly more capable supercomputers.

The second point I want to make is that progress in advancing supercomputing technology has slowed considerably. Over the last decade, the computer industry has standardized on commodity processors. With high volume low-cost processors, supercomputer clusters consisting of commodity parts held out a promise to users of ever-more powerful supercomputers at much lower cost. At the same time, the federal government dramatically reduced investments in supercomputing innovation, leaving the future of supercomputing in the hands of industry. But from industry's perspective, the supercomputing market is not large enough to justify significant investment in unique processor designs and custom interconnects—as the supercomputer market is less than 2% of the overall server marketplace, according to International Data Corporation. To advance supercomputing, industry has relied on leveraging innovation from the personal computer and server markets.

Today, it has become clear that the promise of commodity-based supercomputers has not materialized. Because supercomputers are based on technology optimized for other purposes, they are exceedingly complex and extraordinarily difficult to use and administer. Computational scientists now spend enormous amounts of time, effort and cost modifying software algorithms to run efficiently across homogeneous processors. In many cases, as soon as the task is complete, these scientists have to repeat the process for the next-generation supercomputer.

Future trends in supercomputing will only exacerbate this problem. Because engineers are running into physical limits trying to speed up individual processors in supercomputers, they are resorting to increasing the overall number of processors in a given system to get better speed. We work with hundreds to thousands of processors in supercomputers today. In a couple of years, we will have to work comfortably with tens of thousands to hundreds of thousands of identical processors. Since all of the processors are of the same architecture, further performance gains from other types of processors that exploit different processing models are lost. Further, as commodity-based supercomputers add more processors, these systems become less balanced as their internal commodity network becomes overloaded thus resulting in decreased efficiency. These systems will often run real-world scientific and engineering applications at only a small fraction of their theoretical peak capability. Most of the resource is wasted.

While cheap supercomputer clusters still prove adequate for some applications, more and more science and engineering applications need better-balanced systems. That means systems with far more bandwidth and better reliability, cooling and power utilization, packaging, systems software, programming models, tools and other features than are available on mass-market system architectures.

The lack of advancement in supercomputing technology not only puts our nation's leadership in supercomputing at risk, but it also creates significant technology gaps that threaten our lead in national security, science and engineering, and economic competitiveness. This impacts the scientific and engineering community in such a way that many critical computational problems remain unsolvable in a timely and efficient manner.

The third point I want to make is that the U.S. Government recognizes the importance of a healthy domestic supercomputing industry. A series of recent U.S. government-commissioned studies on supercomputing unanimously argue for increased federal government support for supercomputer research and development. In fact, the Defense Department's integrated high-end computing report states, "… many of the advantages the U.S. enjoys in technologies critical to national security depend to a substantial degree on the relative strength and diversity of its domestic commercial sources for high-end computing" and recommends quadrupling federal funding for R&D on supercomputing over the next five years.[1] The report highlights that the U.S. advantage in advanced aircraft designs, ballistic missile defense systems, cryptanalysis, biological sciences, stealth materials, and many other technologies are at risk without additional federal

---

[1] Department of Defense IHEC Report - "High Performance Computing for the National Security Community." July 1, 2002.
http://www.hpcmo.hpc.mil/Htdocs/DOCUMENTS/04172003_hpc_report_unclass.pdf

support for supercomputing R&D. The other government-sponsored reports[2] delivered over the last few years also describe in more detail the difficulty the federal government faces effectively running applications of national importance on most of today's supercomputers. All of these reports call for increased federal support for supercomputing.

The government is doing more than just writing reports. The Department of Energy's Office of Science has proposed funding the development and deployment of a petascale computer, one capable of performing 1,000 trillion calculations per second. So has the National Science Foundation. The Department of Defense, most notably through DARPA's High Productivity Computing Systems (HPCS) program and the National Security Agency, supports research to help reinvigorate the advancement of supercomputing technology. For example, the goal of the HPCS program is to provide economically viable next-generation petascale supercomputing systems for the government and industry user communities in the 2010 timeframe. HPCS will significantly contribute to DoD and industry superiority in areas such as operational weather and ocean forecasting, analysis of the dispersion of airborne contaminants, cryptanalysis, military platform analysis, stealth design, intelligence systems, virtual manufacturing, nanotechnology, and emerging biotechnology.

---

[2] National Science Foundation report, "Revolutionizing Science and Engineering Through Cyberinfrastructure: report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure." January 2003.
http://www.nsf.gov/od/oci/reports/atkins.pdf

Classified JASON's Report examining the requirements for supercomputing which derive from DoE's classified weapons research. Fall 2003.

Interagency High End Computing Revitalization Task Force Report - "Federal Plan for High-End Computing." May 10, 2004.
http://www.nitrd.gov/pubs/2004_hecrtf/20040702_hecrtf.pdf.

National Research Council Report - "Getting Up to Speed: The Future of Supercomputing." November 2004.
http://www7.nationalacademies.org/cstb/project_supercomputing.html.

The President's Information Technology Advisory Committee (PITAC) report - "Computational Science: Ensuring America's Competitiveness." June 2005.
http://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf.

Joint U.S. Defense Science Board/UK Defence Scientific Advisory Council Task Force report – "Defense Critical Technologies." March 2006

My final point is that Cray is acutely aware of the current crisis in supercomputing. We believe we have a vision for overcoming this crisis. We call it adaptive supercomputing—have the machine adapt to the user, not the user to the machine. But we need federal government support for this vision to reach its fullest potential in a timely manner, as the market is not large enough to fund the risky, leading-edge research and development that is required.

Cray's vision of adaptive supercomputing grew out of its partnership with the DARPA HPCS program. Cray, in collaboration with AMD, has proposed a paradigm shift in the supercomputing industry that will enable the building of much more powerful, yet significantly easier to use supercomputers than are built today. Like all previous Cray computers, the new supercomputers will be designed from the bottom up rather than be based on a collection of PC and server commodity parts.

Using revolutionary technology, future Cray supercomputers will employ diverse microprocessor architectures that can dynamically adapt to scientific requirements in a transparent, scalable, robust and optimized way. This will allow computational scientists to focus on their unique scientific problems and application requirements instead of being forced to conform to the supercomputer. Systems will be radically easier to program, much more broadly applicable, and more resistant to failure. They will give scientists and engineers the tool they need to solve the multi-scale, multi-physics problems of the future. Computational scientists will experience a tremendous productivity boost saving government and industry time and money while enhancing competitiveness.

Recommendations

Our recommendation to this committee and the Congress is to fully fund the Administration's proposed government investments in supercomputing. This includes funding supercomputing programs in the Department of Energy, the National Science Foundation, the National Aeronautics and Space Administration, and within the Department of Defense. To continue international leadership in science, industry and national security, the United States government must fully fund the continued evolution of supercomputers and give scientists access to the computational capability for a wide range of scientific and engineering disciplines. This investment will be justified by an array of future breakthroughs from more efficient, quieter planes and space vehicles to improvements in digital imaging and drug discovery. The promises of supercomputers are limited only by our imagination.

For its part, the Administration should build on its recent initiatives and develop and fund a coordinated research and development program for supercomputing as recommended in the many reports cited above. The Administration should also take a stronger leadership role in persuading other federal agencies to make use of supercomputing and computational science to carry out agency missions. Many agencies have realized only limited scientific progress because they are reluctant to complement experiment-based science with computational science. The Administration should identify gaps in computational science usage and develop programs to close these gaps.

We also want to express our support for H.R. 28, the High Performance Computing Revitalization Act of 2005. We worked with the House Science Committee on this bill. It not only updates current law, but it reemphasizes the need for continued advances in supercomputing.

In conclusion, I would like to laud both the Administration's and Congress's leadership with respect to high performance computing. Recent developments have been very encouraging. Both Congress and the Administration are seeing high performance computing as a key enabler, even a catalyst for pushing out the frontiers of science and technology. In the report "Rising Above the Gathering Storm"[3] the National Academies of Science stated: "The committee is deeply concerned that the scientific and technical building blocks of our economic leadership are eroding at a time when many other nations are gathering strength." Supercomputing is one of those key building blocks. The Japanese and Chinese governments recognize this and have taken significant steps to boost supercomputing activities domestically. They see what we see. What supercomputers have done for us today will pale in comparison to what supercomputers will do for us tomorrow. Now is the time to invest.

Biography for Christopher Jehn

Chris Jehn is Vice President, Government Programs, of Cray Inc. Responsible for the various government relations of Cray, Mr. Jehn represents Cray in Washington and helps Cray leadership and Cray's sales and marketing staff communicate with, and better understand, Cray's public sector customers. Before joining Cray in 2001, Mr. Jehn was

---

[3] National Academies of Science, "Rising Above the Gathering Storm: Energizing and Employing America for a Brighter Economic Future," 2006, http://newton.nap.edu/catalog/11463.html

the Assistant Director for National Security of the Congressional Budget Office. From 1989 to 1993, he was Assistant Secretary of Defense (FM&P). Mr. Jehn was responsible for policy and oversight of the recruiting, training, compensation, support, and management of the then more than three million military and civilian personnel in the Department. He planned the substantial personnel reductions for the post-Cold War Defense Department and managed their initial implementation.

Mr. Jehn has also been an executive with the Center for Naval Analyses, the Institute for Defense Analyses, and ICF Kaiser International, Inc. An economist, Mr. Jehn was educated at Beloit College and the University of Chicago, where he was a University Fellow. Among his other awards and honors are the Benjamin Hooks Distinguished Service Award from the NAACP, the Distinguished Public Service Medal from the Department of Defense, and the Meritorious Police Cross from the government of Spain. In 1996, the Senate Armed Services Committee appointed him to the Commission on Servicemembers and Veterans Transition Assistance.