S. HRG. 109–1150

# HIGH-PERFORMANCE COMPUTING

## HEARING

BEFORE THE

SUBCOMMITTEE ON TECHNOLOGY, INNOVATION, AND COMPETITIVENESS

OF THE

## COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION UNITED STATES SENATE

ONE HUNDRED NINTH CONGRESS

SECOND SESSION

JULY 19, 2006

Printed for the use of the Committee on Commerce, Science, and Transportation

U.S. GOVERNMENT PRINTING OFFICE

71–158 PDF WASHINGTON : 2011

For sale by the Superintendent of Documents, U.S. Government Printing Office
Internet: bookstore.gpo.gov Phone: toll free (866) 512–1800; DC area (202) 512–1800
Fax: (202) 512–2104 Mail: Stop IDCC, Washington, DC 20402–0001

SENATE COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION

ONE HUNDRED NINTH CONGRESS

SECOND SESSION

TED STEVENS, Alaska, *Chairman*

JOHN McCAIN, Arizona
CONRAD BURNS, Montana
TRENT LOTT, Mississippi
KAY BAILEY HUTCHISON, Texas
OLYMPIA J. SNOWE, Maine
GORDON H. SMITH, Oregon
JOHN ENSIGN, Nevada
GEORGE ALLEN, Virginia
JOHN E. SUNUNU, New Hampshire
JIM DEMINT, South Carolina
DAVID VITTER, Louisiana

DANIEL K. INOUYE, Hawaii, *Co-Chairman*
JOHN D. ROCKEFELLER IV, West Virginia
JOHN F. KERRY, Massachusetts
BYRON L. DORGAN, North Dakota
BARBARA BOXER, California
BILL NELSON, Florida
MARIA CANTWELL, Washington
FRANK R. LAUTENBERG, New Jersey
E. BENJAMIN NELSON, Nebraska
MARK PRYOR, Arkansas

LISA J. SUTHERLAND, *Republican Staff Director*
CHRISTINE DRAGER KURTH, *Republican Deputy Staff Director*
KENNETH R. NAHIGIAN, *Republican Chief Counsel*
MARGARET L. CUMMISKY, *Democratic Staff Director and Chief Counsel*
SAMUEL E. WHITEHORN, *Democratic Deputy Staff Director and General Counsel*
LILA HARPER HELMS, *Democratic Policy Director*

---

SUBCOMMITTEE ON TECHNOLOGY, INNOVATION, AND COMPETITIVENESS

JOHN ENSIGN, Nevada, *Chairman*

TED STEVENS, Alaska
CONRAD BURNS, Montana
TRENT LOTT, Mississippi
KAY BAILEY HUTCHISON, Texas
GEORGE ALLEN, Virginia
JOHN E. SUNUNU, New Hampshire
JIM DEMINT, South Carolina

JOHN F. KERRY, Massachusetts, *Ranking*
DANIEL K. INOUYE, Hawaii
JOHN D. ROCKEFELLER IV, West Virginia
BYRON L. DORGAN, North Dakota
E. BENJAMIN NELSON, Nebraska
MARK PRYOR, Arkansas

# CONTENTS

# HIGH–PERFORMANCE COMPUTING

**WEDNESDAY, JULY 19, 2006**

U.S. SENATE,
SUBCOMMITTEE ON TECHNOLOGY, INNOVATION, AND COMPETITIVENESS,
COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION,
*Washington, DC.*

The Subcommittee met, pursuant to notice, at 11:05 a.m. in room SR–253, Russell Senate Office Building, Hon. John Ensign, Chairman of the Subcommittee, presiding.

## OPENING STATEMENT OF HON. JOHN ENSIGN, U.S. SENATOR FROM NEVADA

Senator ENSIGN. I call the Subcommittee to order. I want to welcome everybody to today's hearing on high-performance computing, and I want to thank my colleague, Senator Cantwell from Washington State, who was the inspiration for this hearing. I am very excited to listen and learn today.

In an increasingly interconnected global economy, high-performance computing plays an important role in maintaining the United States' economic and scientific competitiveness and national security. From aircraft and automotive design to weather prediction to advanced medical research to financial modeling on Wall Street, high-performance computing accelerates the innovation process by shrinking time to insight and time to solution for both discovery and invention.

In the 21st century, together with theory and experimentation, computational science now constitutes the third pillar of scientific inquiry. High-performance computers enable researchers to build and test models of complex phenomena such as multi-century climate shifts and multi-dimensional flight stresses on aircraft. Without high-performance computers, these phenomena cannot be replicated effectively in the laboratory.

High-performance computers also enable organizations to manage huge volumes of data rapidly and economically. As the Council on Competitiveness recognized at its High-Performance Computing Users Conference last year, companies that leverage high-performance computing tools realize a range of competitive benefits from shortened product development cycles and faster time to market to reduced research, development, and production costs, all of which improve a company's bottom line and the country's competitiveness.

Moving forward, high-performance computing will continue to facilitate innovation in our Nation's industries, improve our research capabilities, and enhance our national security. For example, high-performance computing holds the promise of recovering 75 percent or better of an oil reservoir's capacity, up from 50 percent today, through more accurate seismic modeling.

In addition, high-performance computing can help the United States to explore and maximize usage of alternative energy technologies, such as hydrogen fuel cells. High-performance computing can help optimize the design of a wide range of consumer products that we use every day, from cars to Pringle's potato chips, to make sure they are safe and reliable.

High-performance computing can also support breakthroughs in medical research and treatments for disease. For example, it has been used to help figure out how Alzheimer's affects the brain. It has also been used to improve treatment for various forms of cancer.

Finally, high-performance computing can assist in research undertaken across all scientific disciplines at our Nation's universities.

I am eager to hear about the progress that is being made in high-performance computing by both the public- and private-sectors. In addition, I look forward to discussing future challenges and opportunities that may impact the development and application of high-performance computing. I look forward to the expert testimony of our distinguished witnesses and I want to thank everyone for attending and participating in today's hearing.

Before we hear from our witnesses, I'd like to welcome any opening statement by Senator Cantwell.

## STATEMENT OF HON. MARIA CANTWELL, U.S. SENATOR FROM WASHINGTON

Senator CANTWELL. Thank you, Mr. Chairman, and thank you for conducting this important hearing on high-performance computing, and thank you for your enthusiasm and interest in this particular area. I also want to thank the witnesses today and particularly the two with ties to Washington State. Michael Garrett, Director of Airplane Performance at the Boeing Commercial Aircraft Division, made a long trip from Everett, Washington, to join us today and he will describe how Boeing uses high-performance computers to design revolutionary aircraft such as the Boeing 787.

I also want to welcome Mr. Christopher Jehn, Vice President for Government Programs from Cray Computers, based in Seattle. In my mind Cray is synonymous with high-capability, high-performance computing, and Mr. Jehn is standing in for the CEO, who could not be with us today.

This is an exciting time for high-performance computing in the State of Washington and across the country. Companies, universities, and research institutions are taking full advantage of our expertise in both high-performance computing and in life sciences. Nobel Prize-winning scientist Dr. Leroy Hood, Founder of the Institute of Systems Biology in Seattle, is leveraging the power of grid computing, for example, to accelerate the development of predictive, preventive, and personalized medicine.

Meanwhile, the University of Washington, with one of the top computing science departments in the Nation, continues to apply cutting-edge research to advanced networking and distributed computer systems. And through an innovative collaboration with the Fred Hutchison Cancer Research Center, we are training the next generation of scientists in the nascent field of computational molecular biology.

Other research institutions in the state are making use of high-performance computing assets through the Pacific Northwest Gigapop, and Microsoft is working to make high-performance computing technology more mainstream by developing software to network desktops and clusters of computer servers.

Over the past decade, the high-performance computing industry has grown beyond the national security interests that first created it and drove it, and today researchers at universities, national laboratories and corporations with access to supercomputers are using their tremendous computational powers to model complex natural phenomena, to test expensive systems through simulation, and to replace experiments that are hazardous, illegal, or forbidden by official policies and treaties.

So America has always been at a technologically-advanced stage when it comes to this kind of innovation. But now is not the time for us to fall behind, and to stay competitive as a Nation we must maintain computer leadership in high-performance computing and computational sciences.

The U.S. Government still remains a primary user of high-performance computing and we use it to maintain our military superiority, to achieve goals and to defend in other areas of national security. I also want to make sure that we continue to look at this important role as it relates to the Department of Energy and the various missions that the Department of Energy carries out.

I look forward to hearing from many of the individuals here today, and particularly about some of those relevant applications as it relates to the Pacific Northwest National Laboratories in Richland, Washington.

In 1991 when Congress passed the High-Performance Computing Act, the Act that was first established under the first President Bush, we had some goals and standards for Federal high-performance research. I think it is important now that we look at what kinds of changes and upgrades need to be made to that policy. That is why in 2004 I co-sponsored the High-End Computer Revitalization Act of 2004, which became law, and focused really only on the Department of Energy.

So today I hope that we can discuss how we need to broaden that to focus on other areas. So I look forward to working with the Chairman as we move forward on this important issue of high-performance computing and the needs for our Nation and what other additional language, whether that is H.R. 28, the High-Performance Computing Revitalization Act of 2005, or other language that helps us maintain our effectiveness and our advantage in high-performance computing as a Nation.

So, I thank the Chair.

Senator ENSIGN. Thank you, Senator Cantwell.

Our first witness is Dr. Simon Szykman. Dr. Szykman is the Director of National Coordination Office for Networking and Information Technology Research and Development, and we welcome your testimony and welcome you to the Subcommittee.

**STATEMENT OF DR. SIMON SZYKMAN, DIRECTOR, NATIONAL COORDINATION OFFICE FOR NETWORKING AND INFORMATION TECHNOLOGY RESEARCH AND DEVELOPMENT**

Dr. SZYKMAN. Thank you very much, Mr. Chairman, Senator Cantwell. I am pleased to have been invited here today to discuss the role of the government in funding high-performance computing, or HPC, research and development.

The Networking and Information Technology Research and Development Program, which I will refer to as the NITRD program, represents the coordinated efforts of many Federal agencies that support R&D in the areas of networking and information technology. I am the Director of the National Coordination Office for the NITRD program, the office which is responsible for supporting interagency technical and budget planning and assessment for the NITRD program.

Today I would like to discuss three different aspects of high-performance computing: its place as a priority in the Federal Government R&D portfolio, the impact of successful interagency coordination in this area, as well as U.S. leadership in HPC technologies.

As the NITRD program has evolved over the years, HPC has not only remained the dominant element of the program, but has been cited on a recurring basis as a priority within the Federal R&D portfolio. This has led to significant investments in HPC. In Fiscal Year 2002, funding for HPC in the NITRD program was less than $800 million. In 2007, next year's budget request, the budget has grown over 65 percent since then to a budget request of over $1.3 billion for HPC.

Development of high-end computing capability and capacity is also a priority research area for the American Competitiveness Initiative which was announced by the President earlier this year. The Fiscal Year 2007 HPC budget requests for the ACI agencies—NSF, DOE's Office of Science, and NIST—are collectively over $135 million above 2006 levels. DARPA, although not part of the ACI, is also a strong supporter of HPC R&D and is expected to have a budget increase of $23 million next year.

The release of the "*Federal Plan for High-End Computing*" in 2004 represented the start of a renewed emphasis on HPC R&D within the Federal NITRD program. Interagency coordination with strong leadership from the Office of Science and Technology Policy and OMB in the White House have resulted in unprecedented cooperation on HPC issues across the Federal Government.

Programs such as DARPA's HPCS, High Productivity Computing Systems program, and the NSF-led High-End Computing University Research Activity have garnered support from all of the Federal agencies involved in HPC R&D within the NITRD program. More importantly, recognizing the importance of these efforts for next-generation technologies, several agencies are providing their own funding to support these programs in addition to the funding provided from DARPA and NSF.

In other noteworthy policy developments, addressing the issue of accessibility of HPC resources DOE and NASA have opened up their resources to communities beyond their traditional research communities. This has enabled millions of hours of supercomputing time to be made available to industry projects as well as other government agencies that in the past would not have had access to these HPC resources. HPC system procurement practices are also being influenced through the sharing of best practices, performance metrics, and benchmarks developed through interagency collaboration.

In 2002, HPC gained high visibility when Japan announced the bringing online of a new supercomputer called the Earth Simulator. Although this garnered some attention within certain policy circles, the government research community had been aware of this system being developed. three weeks ago, a new version of the Top500 Supercomputer Sites list was released. The list, which surveys the world's 500 fastest supercomputers, clearly confirms that the U.S. continues to hold a leadership position in HPC technologies.

Some interesting statistics drawn from the latest version of the list: The Earth Simulator which I just mentioned now sits at the number ten position, not the number one position. Of the nine machines in front of it, six of them are inside the United States, including the top four machines. The U.S. dominates the list as a whole, with over 60 percent of the world's 500 fastest machines being in the United States. U.S. vendors are dominant suppliers of HPC technologies. The top three vendors account for approximately 75 percent of the world's 500 fastest systems. Even those outside of the United States rely strongly on U.S.-developed and U.S.-sold technologies.

Looking back, we can confirm that the launch of the Earth Simulator did not represent a crisis for U.S. competitiveness in the context of HPC technologies. This is very important to note in the context of recent announcements from Japan indicating that they are undertaking the development of a new next-generation supercomputing system over the next few years as a successor to the Earth Simulator.

The fact that the U.S. currently holds the title of world's fastest supercomputer does not herald a new era in HPC leadership any more than the loss of that number one position represented a loss of leadership several years ago. HPC has been and will continue to be a priority within the Federal R&D portfolio. The clearest demonstration of progress over the past 4 years should not be viewed in terms of the raw speed of the world's fastest machine, but rather in the context of a growing focus on HPC technology policy in the government, unprecedented interagency coordination and collaboration on planning and implementation of plans within the government, and the increasingly cooperative ties between the government and the private sector.

The progress that has taken place has been the result of concerted efforts aimed at fostering a vibrant government research community, as well as the work of many dedicated individuals working collaboratively toward shared objectives and goals.

Once again, I would like to thank you for the opportunity to be here today and I am happy to answer any questions.

[The prepared statement of Dr. Szykman follows:]

PREPARED STATEMENT OF DR. SIMON SZYKMAN, DIRECTOR, NATIONAL COORDINATION OFFICE FOR NETWORKING AND INFORMATION TECHNOLOGY RESEARCH AND DEVELOPMENT

Mr. Chairman and members of the Subcommittee, I am pleased to have been invited here today to discuss with you the role of the Federal Government in funding high-performance computing research and development (R&D), and to place these investments in the broader context of global competitiveness.

The Federal Networking and Information Technology Research and Development (NITRD) Program, was established by the High-Performance Computing Act of 1991 (Pub. L. 102–194) and further elaborated upon by the Next Generation Internet Research Act of 1998 (Pub. L. 105–305). Federal networking and information technology research and development, which launched and fueled the digital revolution, continues to drive innovation in scientific research, national security, communication, and commerce to sustain U.S. technological leadership. The NITRD Program, now in its 15th year, represents the coordinated efforts of many Federal agencies that support R&D in networking and information technology.

I am the Director of the National Coordination Office (NCO) for Networking and Information Technology Research and Development. The NITRD National Coordination Office is responsible for supporting technical and budget planning and assessment activities for the NITRD Program. The interagency coordination of NITRD activities takes place under the auspices of the National Science and Technology Council (NSTC), and more specifically through the NSTC's Networking and Information Technology Research and Development Subcommittee, and several interagency working groups and coordination groups that operate under this Subcommittee. The collaborative efforts of the interagency NITRD community increase the overall effectiveness and productivity of Federal networking and information technology R&D investments.

Today I would like to discuss three different aspects of high-performance computing: (1) high-performance computing as a priority in the overall Federal R&D portfolio, (2) the impact and success of interagency coordination in the area of high-performance computing, and (3) U.S. leadership in high-performance computing in the context of global competitiveness in information technology and its applications.

**High-Performance Computing as a Priority in the Federal R&D Portfolio**

Fifteen years ago, what is now the NITRD Program was established in legislation as the National High-Performance Computing and Communications Program, having at that time a narrower focus on R&D in high-performance computing technologies and high-speed networks. Today, investments in high-performance computing support a variety of important Federal agency missions, including national security; climate modeling and weather prediction; modeling and simulation in biology, chemistry, materials science, nanoscale science and technology, and physics; and others. Over the years, the program evolved in scope into one that covers information technologies more broadly, including not only high-performance computing and advanced networking, but also cyber security and information assurance, human computer interaction and information management, software design, high confidence software and systems, and other important areas. Through this evolution, high-performance computing not only remains the dominant element of the NITRD Program, but has been cited on a recurring basis as a high priority within the Federal R&D portfolio.

The Office of Management and Budget (OMB) and the Office of Science and Technology Policy (OSTP) annually issue a joint memorandum on the Administration's R&D budget priorities. In the past 4 years, high-end computing has been identified as one of those priorities. These memoranda set the stage for significant focused interagency coordination by Federal agencies, which I will discuss further shortly, from the establishment in 2003 of the High-End Computing Revitalization Task Force that led to the development of the *Federal Plan for High-End Computing* in 2004, to directing agencies to "aggressively focus on supercomputing capability, capacity and accessibility issues," in accordance with that plan.

The Administration's support has led to significant investments in high-performance computing. In 2002, the funding for high-performance computing in the NITRD Program was less than $0.8 billion. In 5 years, that budget grew by over 65 percent to a Fiscal Year 2007 budget request of over $1.3 billion for high-performance computing R&D, R&D infrastructure, and applications. The National Science Foundation, the Department of Defense, and the Department of Energy together account for over $1 billion of that investment (see *Table 1*). In the Administration's FY 2007 budget, high-performance computing accounts for over 40 percent of the $3.1 billion

NITRD Program budget request, and accounts for more than half of the increase in the NITRD Program budget from the previous year.

The President's emphasis on science and technology, which is in part embodied in the American Competitiveness Initiative (ACI), is further contributing to the development of world-leading high-end computing capability and capacity, which is identified as a key goal for ACI research. The three agencies that are part of the American Competitiveness Initiative—the National Science Foundation (NSF), the National Institute of Standards and Technology (NIST), and the Department of Energy's (DOE's) Office of Science—are all members of the NITRD Program, and all fund high-performance computing investments.

Table 1: Largest Government Funders of High-Performance Computing R&D, R&D Infrastructure, and Applications

| NITRD Agency | FY 2007 Budget Request ($M) |
| --- | --- |
| *NSF* | *337* |
| *DOD:* | *375* |
| OSD and DOD Service organizations | 195 |
| DARPA | 118 |
| NSA | 62 |
| *DOE:* | *329* |
| Office of Science | 296 |
| NNSA | 33 |

As a result of the ACI, the high-performance computing budget at NSF is expected to increase by more than $53 million above its FY 2006 level, enabling NSF to pursue the goal of a petascale computing environment and resources by 2010. Similar investments at DOE's Office of Science are expected to increase by more than $82 million above their FY 2006 levels due to the ACI, which will make possible upgrades and diversification of existing high-performance computing platforms and the acquisition of a next-generation platform, at various DOE National Laboratories. NIST's investments are supporting the development of high-performance computing tools, standards, and algorithms, as well as research on quantum computing and secure quantum communications. The Defense Advanced Research Projects Agency (DARPA), though not part of the ACI, is another key supporter of high-performance computing R&D, an area in which its budget is increasing by over $23 million above FY 2006 levels.[1]

High-performance computing has been and continues to be a funding priority within the Federal R&D portfolio. Together, the guidance, leadership, and past and future investments in high-performance computing have demonstrated and solidified the Administration's commitment to U.S. leadership in this area.

**Impact and Success of Interagency Coordination of High-Performance Computing**

I would now like to take the opportunity to highlight some of the success stories that have emerged from the interagency coordination activities of the government's high-performance computing research community.

Until 2003, interagency coordination of high-performance computing activities took place through the NITRD Program's High-End Computing Coordinating Group. It was then that a decision was made within the Administration to increase the government's focus on high-performance computing. In April 2003, Dr. John H. Marburger III, Science Advisor to the President and Director of OSTP, established the High-End Computing Revitalization Task Force (HECRTF) and charged this group to develop a Federal plan that covered high-performance computing R&D; capability, capacity, and accessibility of high-performance computing resources; and procurement issues. The release of the *Federal Plan for High-End Computing* in May 2004, and the increase in visibility through elevating the High-End Computing Coordinating Group to an Interagency Working Group under the umbrella of the National Science and Technology Council, represented the start of a renewed emphasis on high-performance computing within the NITRD Program.

This cooperation, along with strong leadership from the OSTP and OMB, has resulted in unprecedented coordination on high-performance computing issues among Federal agencies. A few examples follow:

[1] Additional detail about high-performance computing budgets, technical activities, and coordination activities can be found in the *FY 2007 Supplement to the President's Budget for the Networking and Information Technology Research and Development Program* (*http://www.nitrd.gov/pubs/2007supplement/*).

- *DARPA High Productivity Computing Systems (HPCS) Program:* DARPA's HPCS program was established in order to develop a new generation of economically-viable high productivity computing systems for national security and industrial user communities by the end of this decade, producing substantial advances in the performance, programmability, portability, and robustness of these systems. Although initiated by DARPA, this program has garnered the support of over a half dozen Federal agencies which have contributed to HPCS technical planning and coordination, and also of the broader multi-agency research community.

  More importantly, as a result of recognition that this program is the government's primary effort directed at next-generation high-performance computing architectures, several of these Federal agencies have contributed their own funding to the program, thereby increasing the leverage of DARPA's investments. The HPCS program is close to entering its third phase, which is aimed at development and prototype demonstration. It is expected that the additional funding provided by other agencies will make it possible to fund more projects in Phase III than would have been possible with DARPA funding alone. This will increase the diversity of architectures that will be explored through this program, thereby expanding the pool of concepts available on which to build next-generation systems in the future, and helping to cement U.S. leadership in this critical technology area.
- *High-End Computing University Research Activity (HEC–URA):* HEC–URA is a program for funding university research that has been supported by interagency planning. A group of NITRD agencies has been collaborating since 2004 to identify research needs for high-performance computing, and to develop programs to meet those needs. Most recently, following a pair of workshops held last year, a solicitation was released by NSF this year to fund university research in file systems and storage technologies for high-performance computing systems. Though led by NSF, three other Federal agencies contributed funding to support HEC–URA file systems and storage projects that have direct relevance to their agency missions, helping to ensure the availability of research results that would not necessarily have emerged from their own agencies' research programs.
- *High-performance computing benchmarks, performance metrics, and performance modeling:* The use of benchmarks, performance metrics, and performance modeling are key to a variety of high-performance computing issues, ranging from guiding decisions on which architectures to invest in at research stages, supporting procurement decisions by providing consistent bases for comparing alternative systems, and predicting the performance of various types of systems on different classes of computing applications. Because of the importance of these issues and their broad relevance to needs that are shared by multiple agencies, over a half dozen Federal agencies have been collaborating on the development of performance metrics, measurement tools, and benchmarks, with several of these agencies providing funding to support related research.

In my preceding discussion, I have highlighted several examples of high-impact results of interagency coordination. These are just a few of the many instances of the cooperation that is taking place across Federal agencies and the positive effects that these collaborative efforts have produced. Numerous other examples are identified in the *FY 2007 Supplement to the President's Budget for the NITRD Program,* which I referred to earlier. I would now like to close my remarks with a brief discussion of U.S. leadership in high-performance computing technologies.

**U.S. Leadership in High-Performance Computing in the Context of Global Competitiveness**

I described earlier the establishment of the High End Computing Revitalization Task Force that led to the development of the *Federal Plan for High End Computing.* Agencies are now working together to implement that plan, focusing on R&D programs for hardware, software, and systems, the different technical elements of the roadmap laid out in the plan. Distinctions between different classes of machines (capability machines, also referred to as leadership class machines, *versus* capacity machines intended to provide the high-performance computing capacity needed to meet government agency needs), and collaborative funding of programmatic activities such as those I described earlier, have helped make better use of Federal R&D investments in high-performance computing.

The focus of the government research community on issues that extend beyond technical program planning is as noteworthy as the level of collaboration on R&D that I have described previously. In the area of benchmarking and performance metrics that I discussed earlier, agency sharing of technical results and best practices is already productively influencing the procurement of high-performance systems. The issue of accessibility of high-performance computing resources as a new Administration priority represents another important evolution in thinking within the government research community outside the direct scope of R&D investment. This issue emerged with the realization that the use of government high-performance computing resources should not be restricted only to the community of researchers directly funded by a given agency. With support from OSTP and OMB, agencies are now working to ensure that the use of computing resources they fund can also be used meet the needs of broader constituencies.

Two notable examples of the impact of this policy change are the DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, and NASA's National Leadership Computing System (NLCS). Both of these agencies opened up the use of their systems to users outside of their traditional user community, while still maintaining the high standards of merit-based peer review. As a result, DOE awarded millions of processor hours of supercomputing time to four industry research projects in the latest INCITE program cycle, and NASA awarded a million hours on a NASA supercomputer to the National Institute of Standards and Technology (NIST), an agency that has important problems that require high-performance computing to solve, but which does not procure its own dedicated high-performance computing systems.

In addition to making high-performance computing resources available to support private sector R&D, the government is working more generally to foster the use of high-performance computing in the private sector. Several government agencies are providing funding for the Council on Competitiveness's High-Performance Computing Initiative. This initiative, undertaken by this well-known nonpartisan and nonprofit organization, is funding studies, conferences, and educational activities to stimulate and facilitate wider usage of high-performance computing across the private sector, in order to propel productivity, innovation, and competitiveness.

In March 2002, issues of innovation and competitiveness in the context of high-performance computing gained high visibility when Japan brought online a new supercomputer called the Earth Simulator, which became at that time the world's fastest supercomputer. Although people in some policy circles were caught by surprise by this development, the system had been publicly announced long in advance and its existence was known by experts in the research community. Many in the research community called for a tempered reaction, arguing that the leap-frogging by a Japanese supercomputer to the position of world's fastest machine was simply a result of the natural march of progress.

Three weeks ago, a new version of the Top500 Supercomputer Sites list was released.[2] This list, which surveys the world's 500 fastest supercomputers (excluding classified systems) as ranked by a well-known benchmark clearly confirms that the United States continues to hold a strong leadership position in the world of high-performance computing technologies. Some interesting statistics drawn from the latest Top500 list:

- The Earth Simulator, which held the number one position 4 years ago, now sits in the number ten position. Six of the nine machines above it are in the United States, including all of the top four machines.
- The U.S. dominates the list as a whole; 60 percent of the world's 500 fastest supercomputers are installed in the United States.
- U.S. vendors are the dominant suppliers of supercomputing systems in the world. The top three vendors of systems on the Top500 list are all U.S. companies, and account for nearly 75 percent of the systems on the list, including those outside the U.S.
- Even foreign systems rely overwhelmingly on U.S. technologies. Of the top 20 non-U.S.-based systems, 15 were sold by U.S. companies. Of the remaining five that were built by foreign companies, a majority were built using high-performance microprocessors supplied by U.S. companies.

Looking back, we can now confidently say that while the clamor surrounding the launch of the Earth Simulator 4 years ago brought to the attention of policymakers the importance of supercomputing, it did not represent a pivotal crisis to U.S. global competitiveness. This is important to note in the context of recent announcements from Japan regarding an undertaking to develop a successor to the Earth Simulator, which will take place in phases over the next few years.

[2] See *http://www.top500.org/*.

**Conclusion**

The fact that the U.S. currently holds the title of world's fastest supercomputer does not herald a *new* era in U.S. leadership in high-performance computing any more than the loss of the number one position implied a loss of leadership. High-performance computing has been—and will continue to be—a cornerstone in the government's networking and information technology R&D portfolio.

The clearest demonstration of progress over the past 4 years, however, should not be viewed in terms of the raw speed of the world's fastest machine, but rather in the context of the growing focus on domestic high-performance computing policy, the unprecedented interagency coordination and collaboration on technical planning and implementation taking place within the government research community, and the increasingly cooperative ties between the government research community and the private sector. These latter attributes are not simply due to the march of technological progress, but are the result of focused efforts aimed at policy development, budget and technical planning, and the fostering of a vibrant government research community consisting of dedicated individuals with shared priorities committed to working toward common objectives.

Once again, I thank you for the opportunity to be here today and would be happy to answer any questions.

Senator ENSIGN. Thank you, doctor.

Let me just start with this question, because you have talked about the United States and its commitment to high-performance computing. When we compare ourselves with other countries, you said when Japan came online with the fastest high-performance computer and now we have the fastest computers again—at any one point in time it is probably not that important. What seems to be more important is the commitment to high-performance computing. Which country has the commitment to high-performance computing? How do we compare with other countries, other major industrialized countries? How do we compare to Europe, because it is hard to just pick one country there with the EU? How do we compare with China, EU, India, Japan, Korea, those types of countries?

Dr. SZYKMAN. Globally, I would say that Japan is probably at a leadership position for making a national commitment to high-performance computing. They recently, earlier this year announced a national technology area of importance in the context of high-performance computing. That is a high-level policy commitment. I would say that that commitment is shared within the United States.

The Office of Science and Technology Policy and the Office of Management and Budget annually issue a memo, a guidance memo to agencies, that directs agencies to focus on certain R&D priorities, and high-performance computing has been identified as a priority in the memo in each of the past 4 years. So there is certainly very high-level White House commitment to high-performance computing in this country.

Senator ENSIGN. How does that translate down into other countries dollarwise *versus* our country dollarwise, investment by the government?

Dr. SZYKMAN. I would probably have to check some figures and I would be happy to get back to you.

Senator ENSIGN. Could you get that for us?

Dr. SZYKMAN. I can say that in terms of countries in Asia besides Japan, there certainly is an interest in high-performance computing, but not a local, national capability for developing very high-end systems in, for example, Japan—I am sorry, in for example China or India.

Senator ENSIGN. Thank you.

You mentioned performance metrics and I have always been interested in performance metrics throughout our government. I think it is very important that when we are putting money into something we try to at least measure what we are getting for our investment as much as possible. It is not always possible, but we should make every attempt to measure that.

In your testimony you talk about a 65 percent increase since 2002 for funding for high-performance computing. Do you have any metrics that explain the increase? That is a fairly large increase, 65 percent in a very short period of time. Do you have any metrics? You talk about metrics. Now do you have any metrics to show us what we are getting back as a Nation for that investment? Sorry to use simple terms, but hopefully you understand the essence of the question.

Dr. SZYKMAN. Certainly. The work that is being done collaboratively in the area of metrics has in fact been aimed at being able to make better decisionmaking about investments in the area of high-performance computing. Developing better predictions of different alternative architectures helps people decide what areas to invest in, both at the research level as well as informing procurement decisions for more advanced systems.

I could probably do some research and come up with particular examples, but certainly if we look at the fastest machine that is online today funded by the U.S. Government, it is considerably more advanced than the one that was in place in the past, and the performance metrics and benchmarks that are being developed today are being incorporated in calls for proposals as well as procurement issues for future systems.

Right now the most advanced program in the area of developing next-generation technologies will be making use of significantly improved benchmarks over what was available just a few years ago.

Senator ENSIGN. A lot of this information on high-performance computing is new to me. So let us just take even the top five or top ten computers in the world today as far as performance capabilities, can you give the Subcommittee an idea what it costs to put one of those computers together? I realize there is a lot of R&D, but from start to finish what does it cost? If the computer is going to be put in a university someplace, what kind of investment are we looking at for one of these, especially the very, very fast ones, and maybe also for computers that are within the second hundred in terms of speed and processing ability, how the costs would compare between these two types of computers?

Dr. SZYKMAN. I would say that that would probably be an ideal question for some of the industry members of the second panel. They are the ones who sell these machines and could give a more informed answer than I could.

In terms of the R&D that leads up to the development of these systems, though, it is on the order of hundreds of millions of dollars of R&D investment on the private sector side to support next-generation architecture development.

Senator ENSIGN. Senator Cantwell.

Senator CANTWELL. Thank you, Mr. Chairman.

Mr. Szykman, I have a couple questions. First of all, did you support the legislation that moved through the House of Representatives? Did the Administration support that language?

Dr. SZYKMAN. Yes, the Administration was supportive. You are referring to H.R. 28?

Senator CANTWELL. Yes.

Dr. SZYKMAN. Yes. And we did provide a couple of comments and suggestions, but in general we did support that legislation.

Senator CANTWELL. Why do you think that it is important to upgrade the coordination? What specifically do you think that that legislation gets at?

Dr. SZYKMAN. I would say one of the benefits of that legislation is to revise the statutory descriptions of the NITRD Program. If we go back 15 years to when the program was initially established, it was focused on high-performance computing and the advanced networking to support those high-performance computing centers. The program over the past 15 years has broadened considerably into new program component areas and certainly having an expansion of the scope of the program in legislation was helpful, as well as the rearticulation of the priority of some of these technology areas.

Senator CANTWELL. What do you think those priorities are in the technology areas?

Dr. SZYKMAN. High-performance computing, as I mentioned in my testimony, remains one of the main priorities and in fact accounts for 40 percent of the program budget, even though there are eight different program component areas in the program. Other R&D priorities in the area of IT R&D include advanced networking, large-scale networking, which is clearly needed to support interconnectivity between high-performance computing centers as well as more general research facilities and the connectivity needed for users to access those facilities.

Senator CANTWELL. What areas do you think the United States right now has a lead in in the area of high-performance computing research?

Dr. SZYKMAN. I would say the United States has a strong lead in most or all areas. High-performance computing is not an area that really can be looked at as a collection of individual areas. It needs to be looked at holistically at the level of architectures and systems, and the most advanced systems are the ones that are being developed through R&D in the United States. Those systems include aspects of hardware, aspects of software, aspects of storage systems, as well as the overall architectures needed to bring these together into functioning systems.

Senator CANTWELL. So there is not an area of concern that you have where the United States may be losing an R&D advantage because dollars are going to another country because of their particular focus in an area of supercomputing?

Dr. SZYKMAN. I would say through interagency coordination the government research community is able to identify weak areas in the overall portfolio and put funding in those areas. In fact, one of the areas that was not highlighted as a high priority in the "*Federal Plan for High-End Computing*" was the area of storage systems for high-performance computing. That was mentioned in there, but was not highlighted as a strong priority. However, over the couple of years since the release of that plan interagency coordination has identified that as an important need. Agencies have come together to fund university research in that area with funding from multiple agencies to help fill in some of the gaps.

So I would say in summary that the interagency coordination mechanisms that are in place are very effective at identifying needs for the future of high-performance computing R&D and are able to direct funding and technical planning against those needs.

Senator CANTWELL. So you do not have a particular area? For us in Washington State, since we see so much that goes on with computing in general and supercomputing and systems biology, I would say that we are doing pretty well there as a country and having an advantage in driving R&D investment. But there are probably some other areas that I would say I am not so sure. I mean, the Japanese have taken some lead on various models as it relates to weather and climate; is that not correct?

Dr. SZYKMAN. In the area of applications, I would say that most likely agencies themselves could provide clearer information than I can from my office. We are focused more typically on the R&D and less on the particular applications that are being done, even though we collect funding information on those applications. But certainly the most advanced high-performance computing models that are in place today are being supported by high-end capabilities funded by the U.S. Government.

In the area of climate modeling and weather prediction, for example, NOAA, the National Oceanographic and Atmospheric Administration, has recently over the past few years very significantly upgraded its capabilities for doing modeling and prediction in ways that allow longer term modeling of climate and weather, and these are things that are having a direct influence on people every day when they turn on the weather in the morning.

So the U.S. is, I believe, maintaining its leadership in the applications areas as well.

Senator CANTWELL. I know you mentioned advanced networking. Do you believe that there is an under-investment in software for high-performance computing?

Dr. SZYKMAN. That is one of the other areas that I think that agencies recognized as being an important area that in the past, if we look back perhaps 3, 4 years, had been somewhat underfunded, and there is a renewed interest in putting funding in those areas and software programs within different agencies, including the National Science Foundation and the Department of Energy.

Senator CANTWELL. Mr. Chairman, since I am on a time limit and I have two witnesses I would like to hear from, I could ask Mr. Szykman more questions, but I think I would just file those for the record and thank him for his testimony and allow us to hear from some of the other individuals that are here with us today.

Senator ENSIGN. I would like to thank you for your testimony. We will have other questions for the record, and we appreciate your being here.

Let us call the second panel to the table. I am just going to introduce all six of you at once and we will go in the order of your testimony. All of you can please come up now, and I will introduce all of you at the same time. Then we will hear from you in the order in which I introduce you.

Our first witness on this panel will be Dr. Irving—and this is going to be a tough name——

Dr. WLADAWSKY-BERGER. "Vla-DOW-skie."

Senator ENSIGN. "Vla-DOW-skie"–Berger. How is that?

Dr. WLADAWSKY-BERGER. Perfect.

Senator ENSIGN. He is IBM's Vice President for Technical Strategy and Innovation.

After Dr. Wladawsky-Berger will be Mr. Christopher Jehn. He is the Vice President of Government Programs at Cray Incorporated. The next witness after that will be Mr. Jack Waters. Mr. Waters is the Executive Vice President and Chief Technology Officer of Level (3) Communications.

After that, we will hear from Joseph Lombardo. Mr. Lombardo is the Director of the National Supercomputing Center for Energy and the Environment at the University of Nevada, Las Vegas, my *alma mater*.

Our next witness after that will be Mr. Michael Garrett, who is the Director of Airplane Performance, Boeing Commercial Airplanes, for The Boeing Company. And our final witness will be Dr. Stanley Burt. Dr. Burt is the Director of the Advanced Biomedical Computing Center in Frederick, Maryland.

So we can make sure to get in all the questions that we possibly can, because we have all of your written testimonies, if you could sum up your testimonies in about 5 minutes, that would give us a chance to hear from each one of you and then allow plenty of time for further discussion. So I appreciate each one of you being here and look forward to your testimony. Let us start with Doctor Wladawsky-Berger.

## STATEMENT OF DR. IRVING WLADAWSKY–BERGER, VICE PRESIDENT, TECHNICAL STRATEGY AND INNOVATION, INTERNATIONAL BUSINESS MACHINES CORPORATION (IBM)

Dr. WLADAWSKY-BERGER. Thank you. On behalf of IBM I want to thank you, Mr. Chairman, for the opportunity to testify. I am Vice President for Technical Strategy and Innovation at IBM and have been involved with supercomputing initiatives for over 20 years. With your permission, I will simply summarize my written testimony.

Today IBM leads the industry with the world's top three supercomputers and almost half of the world's top 500 supercomputers. We were first to deliver over 100 teraflops, or 100 trillion operations per second, $10^{12}$ of peak performance, to the DOE's Lawrence Livermore National Lab, where we also first demonstrated the practicality of using well over 100,000 microprocessors on a single problem.

Likewise, we have been working with DARPA to help them make very high-end systems more productive, and are investing in advanced hardware and software that will culminate in a system capable of more than one sustained petaflop, $10^{15}$.

In the process of all this, we have learned many lessons, but two are especially significant. First, it is vital to work closely with lead partners in research labs and universities to push the envelope of performance, applications and discovery. I can not overemphasize enough from my personal experience the importance of this pilot to developing working systems for real research on important applications, and the fact that it is the Federal Government that is instrumental in creating them.

Second, the marketplace is all-important. Many supercomputing companies have failed because they relied solely on government-based projects and were heedless of marketplace requirements that go beyond leading performance to competitive prices, energy efficiency, and sophisticated software and applications. While we are very proud of IBM's leadership in supercomputers, we are equally proud that it is an actual, viable business for us with clients around the world.

Why is a national supercomputing capability vital to the U.S.? Supercomputing systems and applications push the envelope in multiple dimensions. They analyze huge amounts of information. They accurately simulate both the natural world and the world of manmade objects, and they present the results in highly visual and realistic ways so we can interact with them.

Additionally, supercomputing architectures and applications foreshadow the future. If one is removed from the advanced research, new ideas, and creative minds in supercomputing, one will inevitably misread the major trends in computing.

Finally, supercomputers enable scientists to make discoveries that would be difficult, perhaps impossible, to accomplish experimentally.

The supercomputing market has changed radically in the last decade. It was once a niche market because the hardware was so expensive, but that has all changed with the introduction of workstation and PC-based technologies. Today we are even working to build supercomputers with technology from the worlds of consumer electronics and video games. All these approaches use components from high volume markets, thus the costs are significantly lower than in the early days and the potential markets are much, much bigger.

My basic point is this: A commercial business model has reduced costs and enabled us to address the vast spectrum of public- and private-sector applications. One-off machines built for a single mission are usually very expensive, impractical in the marketplace, and not viable in the long term.

Progress in supercomputing hardware has been outstanding. The real challenge, however, lies in both application software and system software, as has been widely recognized in a number of studies. Software is so consequential because supercomputing's value is not in the technology, important as it is, but in its applications, which makes software critical. My formal testimony reviews the progress and promise of some key applications. There is enormous promise both in classic or more mature applications, such as defense and national security, science, engineering, and weather and climate, and in the newer applications that are so vital to the national interest, like energy, healthcare and bioinformatics, learning and training, and business in general.

In civilian nuclear energy, for example, the GNEP, or Global Nuclear Energy Partnership, and ITER, International Thermonuclear

Energy Research, programs are excellent examples of government-industry academic collaboration on matters of national and international importance, market relevance, and timeliness. They deserve support.

Supercomputing today is essential for innovation in both the world of science and the world of commerce. It is an indispensable tool if our country is to thrive in a global economy that grows more competitive by the day. It is therefore essential to pass an innovation authorization bill this year, as in S. 2802, as you, Senator Ensign, Senator Stevens, and others on this committee know.

The Federal Government funds basic research and establishes priorities for research in the pursuit of innovation and competitiveness. That makes wise policy choices critically important to a national supercomputing capability. To realize that capability, Congress should clearly outline and invest in a long-term strategy. For example, the President's budget request for Fiscal Year 2007 includes high-performance activities that range from biomedical computing to Earth and space science research and many others. Clear direction and consistent funding will foster investment by industry and academia, so that together we can address the challenges that face our country and grow the capabilities of our knowledge. This is the kind of joint effort between government, universities, and private industry for which there is no substitute.

Thank you.

[The prepared statement of Dr. Wladawsky-Berger follows:]

PREPARED STATEMENT OF DR. IRVING WLADAWSKY-BERGER, VICE PRESIDENT, TECHNICAL STRATEGY AND INNOVATION, INTERNATIONAL BUSINESS MACHINES CORPORATION (IBM)

**Introduction**

On behalf of IBM, I would like to thank the Subcommittee, and especially Chairman Ensign and Senator Kerry, for the opportunity to address the evolution of supercomputing during the 1990s, the priorities we should focus on, and the challenges we face.

First, by way of introduction, I am currently Vice President for Technical Strategy and Innovation at the IBM Corporation, and am responsible for identifying emerging technologies and marketplace developments critical to the future of the IT industry, and organizing appropriate activities in and outside IBM in order to capitalize on them.

My association with computers began in the Summer of 1962, when prior to entering the University of Chicago, I was employed at the computation center where I worked part-time through my college years doing scientific programming. I later went on to get a Ph.D. in Physics, and did my research on computational atomic and molecular physics. After finishing my Ph.D. at the University of Chicago, I decided that I was better suited for computing than for physics, switched fields and joined the Computer Sciences Department at IBM's Thomas J. Watson Research Center in June 1970.

By then, IBM had decided, for a variety of reasons, to exit the scientific computing market, where it had been a leader for a number of years. We re-entered the market in the second half of the 1980s by adding what is called a vector feature to our mainframes, and a few years later in the early 1990s we became a leader in the emerging area of parallel supercomputing with our SP system. I was the General Manager of both these efforts.

Today, according to the authoritative Top500 Supercomputing rankings, IBM:

- Leads with the world's top three supercomputers: BlueGene/L for the U.S. Department of Energy (DOE) with 280.6 sustained teraflops (trillion floating point operations per second), BlueGene/W at Watson Research with 91.3 sustained teraflops (or 114.7 teraflops peak) and DOE's ASC Purple at Lawrence Livermore National Laboratory with 75.8 sustained teraflops (or 92.8 teraflops peak).

- Has supplied 240 of the world's top 500 supercomputers—more than any other vendor.
- Accounts for over 1.5 petaflops of aggregate performance in the TOP500 list (from a total of 2.79 petaflops).
- Has supplied more supercomputing systems than any other vendor in the Top10, Top20, Top100 and Top500.
- Has supplied the most cluster systems with 177 of 364 (48.6 percent).
- Has built the largest university supercomputer in the U.S.—Big Red—a Cluster rated at 15 teraflops and installed at the University of Indiana.

IBM has held the number one spot in the Top500 list since June 2005. Japan's Earth Simulator held the number one spot for the previous 3 years. IBM was the first to deliver a system that achieved over 100 teraflops—that is a system that could perform over 100 trillion operations per second—of peak performance; in fact BlueGene/L has tested at 360 teraflops (peak). And we have achieved actual sustained performance of from 100 to 200 teraflops on a number of applications of real importance to the National Nuclear Security Administration. The BlueGene/L supercomputer has been measured at 10-times the energy efficiency (measured by Watts of electricity needed to attain a particular level of performance) of any of the top 20 supercomputers, and it is similarly efficient in its space requirements. In the process, we have demonstrated the practicality of using well over 100,000 microprocessors and then leveraging their computational capability efficiently on a single problem.

Along the way, we have learned many lessons, but I believe that two are especially significant. First, it is vitally important to work closely with lead partners in research labs and universities in order to "push the envelope" in terms of performance, applications and discovery. I cannot overemphasize the importance of these leading-edge pilots in propelling us forward and bringing together all the elements needed to develop working systems that can be used for real research in important application areas.

The second key point is the importance of the marketplace in guiding our actions. Through my professional career I have seen many supercomputing companies fail because they relied solely on government-based projects and were heedless of marketplace requirements for, not only leading-performance, but competitive prices and sophisticated software and applications as well. While we are very proud of IBM's leadership in the Top10, Top100 and Top500, we are equally proud that supercomputing is a viable business for us with many clients around the world in both the private and public sectors.

**What Is Supercompting?**

Supercomputing is defined by three key characteristics. First, the applications are information-intensive; second, they deal with computation-intensive simulations—both in the natural world of physics, chemistry and biology, and in virtual worlds, such as engineering objects and entertainment; third they enable the visualization of information and simulations so people can interact with the results—as exemplified by scientific visualization and the more recent emergence of video games played between myriad participants.

**Why Is Supercomputing so Important?**

Through the years, we have come to realize that supercomputing architectures and applications foreshadow the future of computing itself. Indeed, if one is removed from the advanced research, new ideas and creative minds in supercomputing, one will inevitably misread the major trends in computing. This is among the main reasons IBM re-entered the market in the late 1980s.

Beyond its role as a precursor, supercomputing has become essential to the pursuit of scientific inquiry. To quote the June 2005 report by the President's Information Technology Advisory Committee (PITAC), "Computational science has become the third pillar of the scientific enterprise, a peer alongside theory and physical experiment." Supercomputers enable scientists to either make discoveries that would be difficult (perhaps impossible) to accomplish experimentally or to point researchers in new directions.

Examples abound. They include (but are certainly not limited to) developing insight into the behavior of materials under extreme conditions that cannot be reproduced experimentally, enabling scientists to make reliable predictions about the behavior of our nuclear stockpile or the safety of aging nuclear reactors, for example. Supercomputers can also find previously undiscovered sequences in so-called "junk DNA" that may lead to new insights into its "function." They can also discover "docking sites" for new drugs, *i.e.,* receptors on molecules where a drug can potentially attack a disease. Or supercomputers can perform multi-century simulations to understand trends in the Earth's climate.

**Growth of the Supercomputing Market**

Supercomputing was once confined to a niche market, because the hardware was so very expensive. That changed over time with the introduction of workstation and PC-based technologies, the latter becoming immensely popular in Linux clusters during the late 1990s. Today, we even use low-power, low-cost micros—consumer-based technologies—to attain very high degrees of parallelism and performance, as in our Blue Gene system, which has reached a peak of 360 trillion calculations per second. Now, we are seeking to build supercomputers using technologies from the gaming world, such as the Cell processor.

All these approaches leverage components from high-volume markets, and aggregate them using specialized architectures; thus the costs are significantly lower than in earlier days and the potential markets are consequently much bigger.

Progress in supercomputing hardware has been nothing short of astounding. The real challenge, however, is software, both application software and systems software. In fact, both the 1999 PITAC report, with which I was personally involved, and the June 2005 PITAC report made precisely that point.

**Key Application Areas**

But the real value of supercomputing to society is not in the technology, architecture and software, important as they are. The value of supercomputing is best appreciated by considering its application, so let me review the recent progress and the promise in a few key application areas, starting with the "classic" or more mature ones and then moving on to some of the newer opportunities.

*Defense and National Security*

Let me start by discussing Blue Gene/L and ASC Purple, two of the world's top three supercomputers, residing at Lawrence Livermore National Laboratory (LLNL). They are vital to the National Nuclear Security Administration's (NNSA) Advanced Simulation and Computing Program (ASC), which in turn is an essential element of our Nation's Stockpile Stewardship Program. ASC provides the integrating simulation and modeling capabilities and technologies needed to combine new and old experimental data, past nuclear test data, and past design and engineering experience into a powerful tool for future design assessment and certification of nuclear weapons and their components.

Already, the simulation and modeling tools are improving the assessment of stockpiles far in advance of schedule. Indeed, weapons designers, scientists, and engineers now rely on ASC simulation and modeling capabilities and technologies to assess changes occurring in aging stockpiles of nuclear weapons and to assess and certify planned refurbishments of weapons system components.

On March 9, 2006, Lawrence Livermore National Laboratory and IBM announced a fundamental breakthrough using ASC Purple. IBM and LLNL demonstrated over 102 gigabytes per second of sustained read and write performance to a single file using specialized software that orchestrates thousands of processors and thousands of disk storage devices. The breakthrough is expected to stimulate development of data-intensive applications in areas like customized medicine, online gaming, entertainment, and homeland security, as well as in traditional high-performance computing applications.

Then on June 22 of this year, the NNSA announced that it had achieved an unprecedented level of performance using our Blue Gene/L. This world record for a scientific application was set by achieving a sustained performance of 207.3 teraflops, running "Qbox" computer code for conducting materials science simulations critical to national security.

*Science*

In addition, the unmatched cost-effective computational capability of Blue Gene has already resulted in new insights in biology.

The scientists at the T.J. Watson Research Center have applied supercomputing to demonstrate that "junk DNA" could have very startling ramifications on cell regulation and species evolution. In another computational experiment, they have shown that a single mutation in a protein can render it unstable, causing it to misfold. Similar techniques and computational models can be applied to better understand fatal diseases.

ASTRON, in the Netherlands, is using the Blue Gene supercomputer to develop a new type of radio telescope capable of looking back billions of years. This research project will enable scientists to examine the beginnings of the earliest stars and galaxies after the formation of the universe in the wake of "the Big Bang."

Blue Gene/L will give ASTRON the flexibility and unparalleled speed it needs to gather and analyze information from its Low Frequency Array (LOFAR) "software telescope" network. Unlike current observatories that use big optical mirrors or radio dishes to point to distant galaxies, ASTRON will harness more than 10,000 simple radio antennas spread across the northern Netherlands and into the German state of Lower Saxony and interpret them using high-speed calculations.

In many domains, theory, experimental capabilities, and computational advances are coming together in a manner that will significantly accelerate scientific discovery.

*Weather, Climate*

Supercomputing is also taking weather forecasting, modeling and research to new levels. Research groups at several government agencies and research laboratories are moving traditional models to scalable supercomputer systems. These models are then used to test the validity of our current understanding of the physics of weather and to develop more detailed, robust, high-resolution models. When the models are considered trustworthy, they are used for operational forecasting by the National Weather Service and by environmental analysts to assess air quality.

In addition, there is emerging a generation of localized, high-resolution weather prediction capabilities customized for application one to 2 days ahead of time by businesses with weather-sensitive operations. Industries that would benefit are as diverse as aviation, agriculture, broadcasting, communications, energy, insurance, sports, entertainment, tourism, construction and others in which weather is a crucial factor. Extremely fast, ultra precise weather forecasts would be invaluable to these businesses' day-to-day decisionmaking. Such forecasts could be used for competitive advantage or to improve operational efficiency and safety.

*Engineering*

Automobile companies run virtual car crashes using complex supercomputer simulations to ascertain how different designs react in collisions. This reduces the number of costly prototypes, and speeds the delivery of new models. With new regulations on safety in the auto industry and buyer preferences for safer cars, keeping this competitive advantage is of paramount importance to manufacturers.

Supercomputing is also being used to create more fuel-efficient automobile designs. Exa Corporation, a global provider of wind tunnel design simulation software uses our supercomputers to help major automotive manufacturers and smaller suppliers solve larger, more complex aerodynamic, acoustic and thermal engineering problems. With virtually unlimited amounts of compute capacity available as needed, Exa's clients can perform more analysis in less time—improving quality and time-to-market and overall competitiveness.

Seismic imaging is an application critical to our energy future. Seismic imaging is the process by which acoustic waves are generated and their reflections off the Earth's subsurface are collected. Seismic imaging applications then convert the reflected waves into a 3D image of the subsurface, revealing an image of a petroleum reservoir. This process is used by all major oil and exploration companies. Good quality seismic imaging is critical since dry holes can cost millions—in the deep waters of the Gulf of Mexico as much as $100 million.

IBM and Compagnie Générale de Géophysique (CGG), a world leader in geophysical services, recently announced deployment of Europe's most powerful seismic supercomputer to respond to growing global demand in the petroleum industry. The system is expected to significantly reduce processing times from the moment the geophysical data is collected to the point when it generates a seismic image. This clustered supercomputer will also allow CGG to boost its worldwide computing capacity to a maximum of 113 teraflops, and give them an unprecedented ability to respond to the extremely high performance requirements of the oil industry.

This supercomputer installation is a result of CGG's need to continually improve its performance in response to the demands of a highly competitive market by optimizing the quality and speed of processing in specific applications. The new system is being deployed at the company's premises in Massy (France), London, Kuala Lumpur and Houston, Texas.

Let me now focus on some of the newer application opportunities.

*Energy*

*GNEP* (The Global Nuclear Energy Partnership) is a Presidential initiative to establish nuclear energy as the preferred emissions-free alternative source of electric power. By reprocessing spent nuclear fuel and recycling it for reuse in nuclear power plants we can control the process and share recycled fuel and technology with developing countries that need inexpensive energy. The United States has been encouraged in this effort by China, France, Japan, Russia, and the United Kingdom as well as the International Atomic Energy Agency (IAEA).

Computer simulation will be essential to the success of GNEP, allowing us to rapidly test innovative approaches and improve our ability to understand and control very sensitive materials. The President has requested $250M in FY07 for the program, while the House has recommended $150M and the Senate $250M. We support the effort and the concept of "making nuclear energy a renewable source of power".

*ITER* is a large international fusion experiment aimed at demonstrating the scientific and technological feasibility of fusion energy and at trying to answer the question: Can we produce practical amounts of fusion power on Earth? In fusion, heavy forms of hydrogen are fused at high temperatures with an accompanying production of heat energy. ITER is a step beyond the study of plasma physics and toward the possibility of fusion power plants actually producing electricity and hydrogen.

The international project is made up seven partners including the United States, China, the European Union, India, Japan, Russia and South Korea. The facility will be housed at a site in Cadarache, France. We support the United States' participation in ITER and the funding requested by the President in the Fiscal Year 2007 budget.

*Bioinformatics and Computational Biology*

These involve the use of techniques from applied mathematics, informatics, statistics, and computer science to solve biological problems. Genomes (an organism's complete information set) are sequenced and assembled, and then become candidates for data mining. This data mining is often referred to as bioinformatics.

The objective is a better understanding of the relationship between specific genes and diseases, an understanding that is essential to the development of therapies. The point is to develop drugs that will target specific genes and focus on a specific disease. With the volume of genetic data proliferating, it long ago became impractical to analyze DNA sequences manually. Today, computer programs search the genome of thousands of organisms, containing billions of nucleotides.

Bioinformatics has great potential for expediting delivery of new, individualized therapies to patients.

Brain research is another promising scientific pursuit utterly dependent on supercomputing. It is also the purpose of a joint research initiative between the Ecole Polytechnique Fédérale de Lausanne (EPFL) and IBM. Nicknamed the "Blue Brain Project," it is intended to take brain research to a new level.

Scientists from EPFL and IBM are working together using the huge computational capacity of Blue Gene to create a detailed model of the circuitry in the neocortex—the largest and most complex part of the human brain. By expanding the project to model other areas of the brain, scientists hope eventually to build an accurate, computer-based model of the entire brain.

Relatively little is actually known about how the brain works. Using the digital model, scientists will run computer-based simulations of the brain at the molecular level, shedding light on internal processes such as thought, perception and memory. Scientists also hope to understand more about how and why certain microcircuits in the brain malfunction—a failure thought to be the cause of psychiatric disorders such as autism, schizophrenia and depression.

*Health Care*

Medical science increasingly relies on advanced information systems to share information, mine that information for trends and insights, and use those findings to head off disease or improve treatment. This takes sophisticated computer hardware and software, and the technology has advanced to the stage where truly wondrous things that yesterday were only wishful thinking can now be tackled.

For example, The Scripps Research Institute and IBM researchers are working on new technology to anticipate, manage and contain infectious diseases like avian flu. Using Blue Gene, they are trying to devise a way to track the emergence of new virus strains and map human and animal responses to them. This capability will help scientists and governments to better understand viruses and respond effectively to potential pandemics. It could also enable vaccines to be created quickly enough to prevent massive outbreaks.

Likewise, QuantumBio Inc., a provider of software tools for drug, biotechnology, and pharmaceutical companies, uses the Blue Gene supercomputer to help satisfy its testing needs. With Blue Gene, QuantumBio is able to provide users with the opportunity to study molecules of interest over a secure and integrated system on an as-needed or on-demand basis.

*Business*

SmartOps, a leading provider of enterprise-class supply chain optimization solutions for the manufacturing and distribution industries, used the Blue Gene supercomputer to port and test their Multistage Inventory Planning and Optimization (MIPO) solution in preparation for offering a large-scale hosted solution for their clients.

Likewise, in the Finance Industry the most competitive firms are those that can maximize returns and minimize risk, all in the shortest time possible. Key to success is the ability to apply computational power to increasingly complex and demanding business processes. Workloads such as risk management, portfolio analysis, derivatives pricing and actuarial simulations can all benefit from the application of supercomputing's greater computational power.

*Learning*

Highly realistic, visual interfaces first appeared with scientific applications as well as with flight simulators used to train pilots and with war game simulators used to train military personnel. These visual interfaces (along with the accompanying sounds) have been increasingly enhanced with digital animation and video games. Video games are particularly important because in addition to their very realistic visual images and sound effects, they are also highly interactive and increasingly collaborative, and thus a good launch pad for thinking about how people can best interact with all kinds of computer applications as well as with each other in the future. Furthermore, the success of video games with millions of people has stimulated the introduction of very inexpensive and powerful technologies, such those around Microsoft's Xbox and Sony's upcoming Play Station 3.

The new highly visual, realistic, and interactive interfaces now hold the promise of sparking a major round of innovation for computer applications in general, both in rethinking how to best integrate these new kinds of visual interfaces with existing applications, as well as inspiring whole new categories of applications that we cannot even envision today.

One application area that holds great promise is learning across the broad spectrum of needs, from K–12 all the way to the introduction of sophisticated new procedures for professionals. After all, since our brains are wired for sight and sound, these new applications should be able to approach humans on human terms, and thus significantly facilitate the learning process.

**Conclusion**

Clearly, supercomputing has advanced to the point of being essential in myriad endeavors, in the laboratory certainly but most assuredly in the commercial world as well. It is indispensable to the process of innovation and to the ability of the United States to thrive in a globalized economy that grows more competitive by the day—something The National Innovation Act of 2005 (S. 2109), which we support, is meant to foster.

The Federal Government has significant influence in setting the agenda for basic research and in turn the use of high-performance computing in pursuit of innovation and competitiveness. We, in our industry participate in that agenda as partners. In order to realize the full benefits for our country, Congress, in partnership with the industry, should clearly outline and invest in a long-term strategy. For example, the President's budget request for Fiscal Year 2007 includes high-performance computing activities funded by the Networking and IT Research and Development (NITRD) agencies including the National Science Foundation, Department of Energy and the National Aeronautics and Space Administration. These activities range from biomedical computing to Earth and space science research to weather modeling frameworks.

Clear direction and consistent funding will prompt industry and academia to invest as well, and in partnership we can address many of the serious challenges that face our Nation. In the process, we will expand and deepen our knowledge of much of the world around us and our ability to influence it. These kinds of efforts unite government, universities and private industry in a productive collaboration—a partnership for which there is no substitute.

Senator ENSIGN. Thank you. Our next witness will be Mr. Jehn.

## STATEMENT OF CHRISTOPHER JEHN, VICE PRESIDENT, GOVERNMENT PROGRAMS, CRAY INC.

Mr. JEHN. Thank you, Mr. Chairman, for inviting me here on behalf of Cray. And I would like to thank Senator Cantwell too for those generous words about Cray. I also would like to thank you for holding a hearing addressing this important subject.

I have submitted a written statement for the record and will briefly summarize it here. I have just a very straightforward story that contains only four key points. First, supercomputing is vitally important, as both you and Senator Cantwell recognized. It is key to many critical national security missions and it is essential for the country's scientific leadership and our global economic competitiveness.

Second, progress in supercomputing technology has slowed. As we have increasingly relied on commercially-available parts for supercomputers, we have come to realize that those solutions are not always the best for the most demanding technical and scientific applications in government, industry, and academia.

Third, the Federal Government has recognized this reality. In a series of recent reports, the government has recognized that a vital supercomputer industry is important and that U.S. Government support is necessary to achieve that end. These reports all cite the need for a systematic research and development program that supports R&D in the supercomputer industry. Industry cannot do it alone because the market for supercomputers is simply not deep enough to justify the kind of investment, the amount of investment, necessary to sustain progress in this area.

Fortunately, the government is doing more than just writing reports. The Department of Energy has recently announced its intention to develop and deploy a petascale computer, that is, one capable of performing 1,000 trillion calculations per second. The National Science Foundation has announced a similar intent.

Meanwhile, the Defense Advanced Research Projects Agency and the National Security Agency are supporting R&D. For example, DARPA's high-performance computing systems program is aimed at developing a commercially-viable system by the end of this decade, a system that can deliver sustained petaflops of performance. It would be more productive than today's computers, but also, equally important, more robust, use power more efficiently, be much easier to program, and be available and applicable to a much wider range of applications.

Fourth, at Cray we understand all these problems and believe we have developed a vision, a plan that we are now acting on, to develop what we call adaptive supercomputers. These will be supercomputers that will combine in one system multiple processor technologies, so that the computer can adapt to the scientists' requirements, rather than demanding that the scientist, adapt their science to the available supercomputer.

I would like to conclude by urging the Congress to fully fund the current Administration initiatives in this area. I would also encourage the Administration to build on these initiatives and develop and fund an R&D program like those described in the reports I cited above.

In conclusion, I would like to thank you again for holding this hearing, and also thank the Congress and the Administration for their leadership in supercomputing over the past several years. There has been a lot of progress and we need to build on that momentum. The time to invest is now.

[The prepared statement of Mr. Jehn follows:]

PREPARED STATEMENT OF CHRISTOPHER JEHN, VICE PRESIDENT, GOVERNMENT PROGRAMS, CRAY INC.

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 percent 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 5 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 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.

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.

---

[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.*

[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.

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 U.S. 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.

## STATEMENT OF JACK WATERS, EXECUTIVE VICE PRESIDENT/ CHIEF TECHNOLOGY OFFICER, LEVEL (3) COMMUNICATIONS

Mr. WATERS. Mr. Chairman, Senator Cantwell, my name is Jack Waters and I'm Executive Vice President and Chief Technology Officer of Level (3) Communications. Thank you for the opportunity to testify today on behalf of Level 3. We believe that high-perform-

[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.*

ance networking is an essential element of high-performance computing, critical to our Nation's competitiveness, and should be a central part of Federal policy regarding the Nation's cyber-infrastructure.

With that, I am going to turn to a presentation that summarizes my testimony.


I would like to make two points. The first is that high-performance networking is a key component of high-performance computing and the second is that the government has played, and should continue to play, a leadership role in both networking and computing.

Next slide.

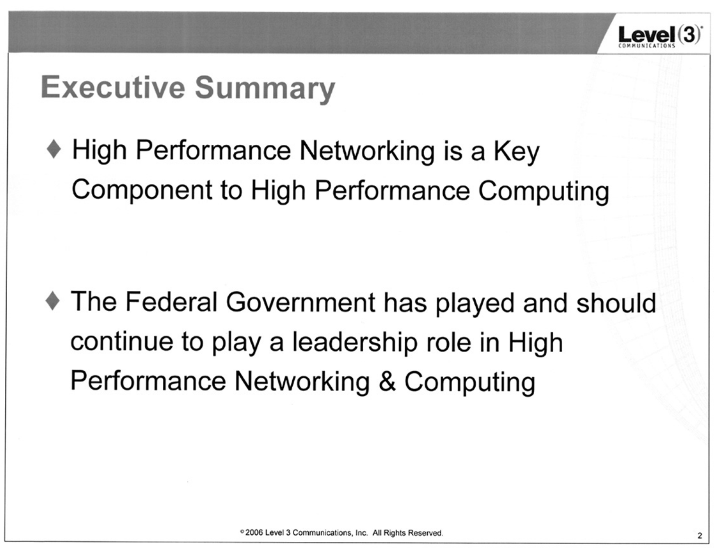

I will tell you very briefly about Level 3. Level 3 is a little bit new, a new company to the telecom world. We were founded about 9 years ago and we constructed a network in the U.S., across the U.S., and in Europe. We had the idea in mind at the time that, what seems obvious today, Internet technology and optical technologies would be the waves of the future and we should optimize our company around those technologies.

Next slide.

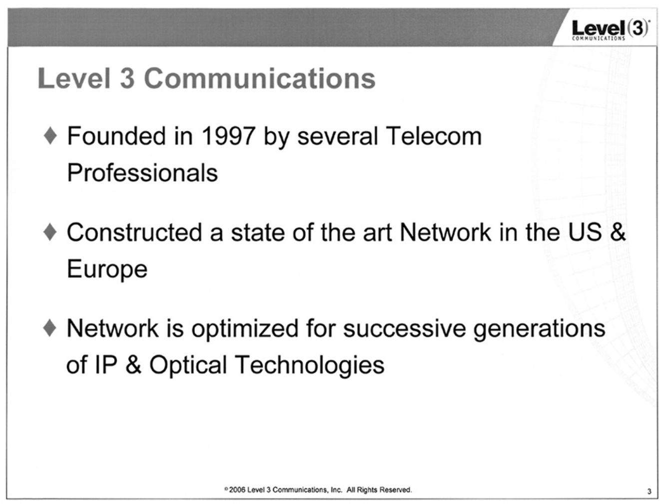


The point of this slide is really to depict that for years we have considered computing information and storing information as the only variables in how we handle information as an industry, and recently the economics of moving information around on a network have become compelling, and that is why networking should be such a key component of high-performance computing.

If you turn to the next slide.

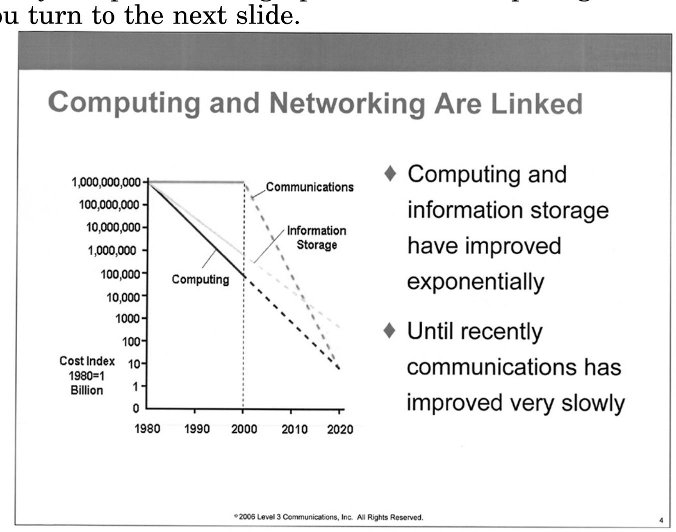

The Federal Government has played a pretty key role and has made some important investments over the last couple of decades, and I would like to point out three: the original investment in the ARPAnet and its communication protocol, TCPIP, the language that computers speak to each other over networks; the original funding in the NSF supercomputer centers and the network that interconnects those supercomputer centers.

Well, what did we get out of those investments? It is pretty daunting when you think those original investments led to the commercial Internet as we know today in both our personal and professional lives, and, even more basic, the browser. The browser actually came from one of the NSF-funded supercomputer centers, NCSA. It was called Mosaic and it is probably something that everyone in this room and around the world uses every day as part of their Internet use.

If you turn to the next slide.

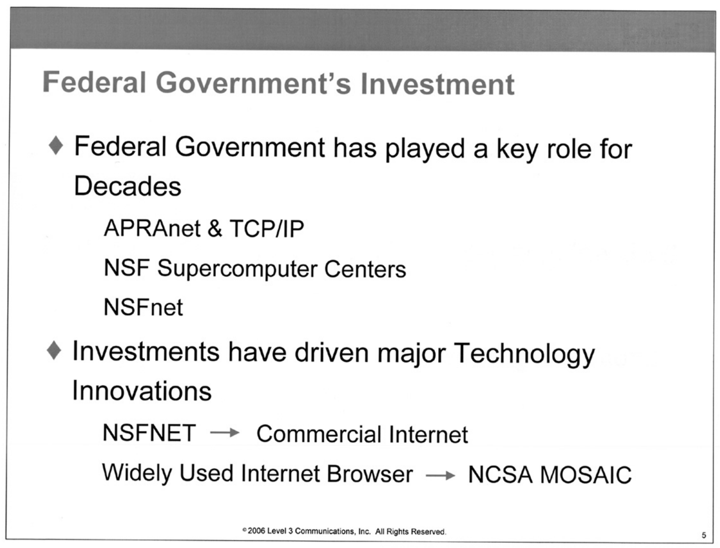


Those investments have led to the growth that I know that we all have seen. This picture actually shows how incredible it has been and frankly continues to be.

If you turn to the next slide.

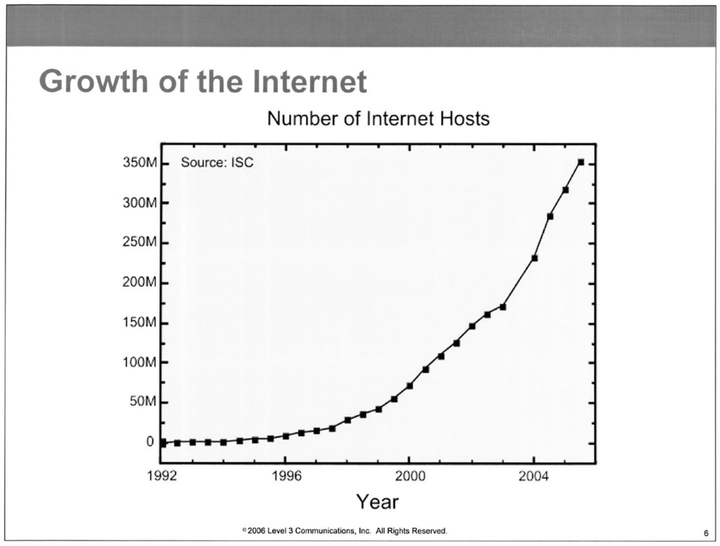


I would like to point out a couple of particular research projects. CERN is a particle physics laboratory in Switzerland and they are building a machine called the Large Hadron Collider. It costs about $8 billion. After it is finished, it is going to produce about 1,500 times the information held in the Library of Congress, so a fair amount of information, and 2,000 scientists around the world in 31 countries, including our country, need access to that information. The only way to provide access to that information is through high-performance networking.

If you turn to the next slide.

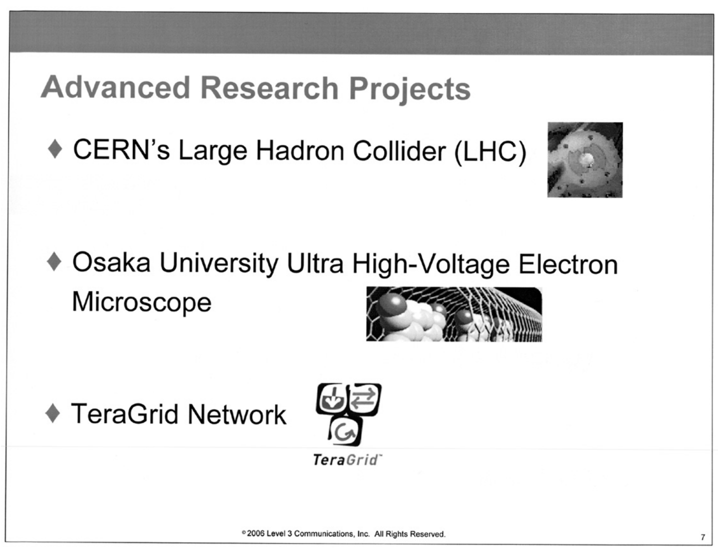


There is another initiative closer to home, called "Internet2," a not-for-profit company that is focused on providing networking to the R&E community. They have recently announced plans to design and deploy a new nationwide infrastructure. Once completed, this infrastructure will allow researchers an unparalleled access to computing and storage facilities across our country.

If you turn to the next slide.

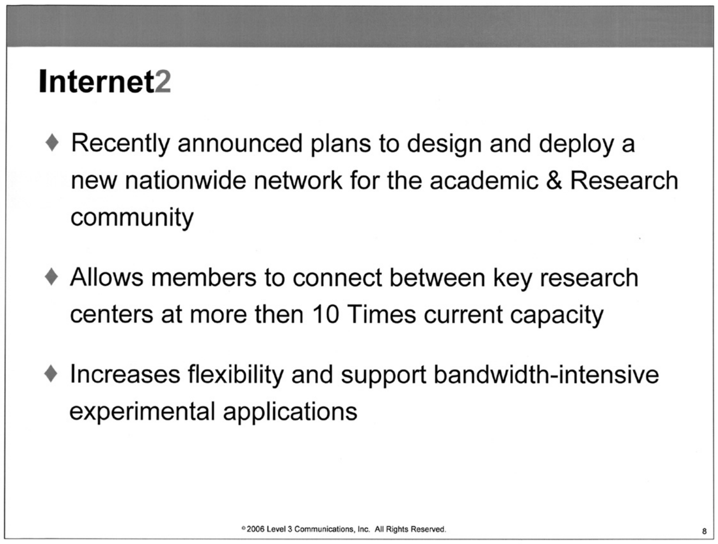


You will see what we get from this investment and Internet2's broad reach which touches many institutions across the country. At the end, users will be able to dynamically tailor their networking resources and needs to suit their research initiatives.

If you turn to the next slide——

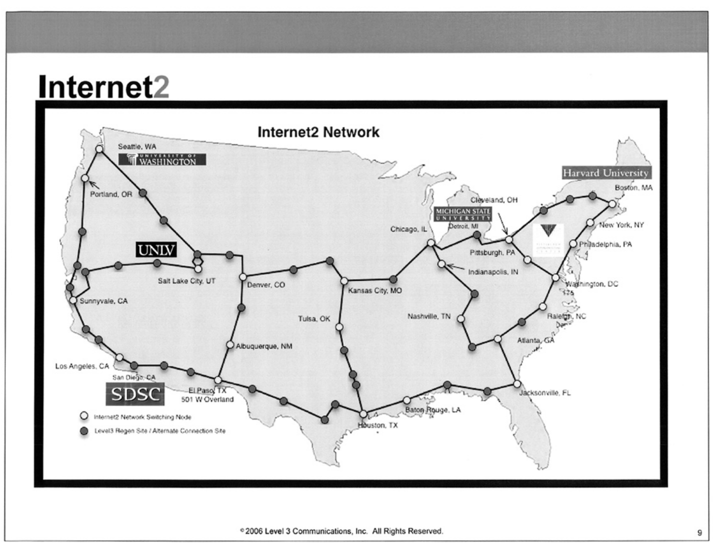

—I thought I would take a moment and give you an idea of what these investments have given us as far as a figure of merit. The figure of merit we networking people love to use is the amount of information in the Library of Congress and how long it would take to transfer it. If you looked 20 years ago, at the original NSFNet, it would have taken about 50 years on that network to transfer that information. If you looked at the investment that Internet2 is making this year in the infrastructure that is being built, it would take about 15 minutes to transfer that information. So the investments have garnered great capability in networking.

**National Research Network Growth**

| Network | Approximate Time to Transmit Contents of the Library of Congress (10 Terabytes) |
|---|---|
| NSFnet – 56K | 50 Years |
| NSFnet – T1 | 2 Years |
| vBNS – 155M | 7 Days |
| Abilene – 10G | 2 Hours |
| New Internet2 – 100G | 15 Minutes |

 10

In summary, we believe that Federal policy should achieve a balance and focus on three key elements: computing, software, and what I have spoken about today, networking; a balanced approach will lead to American innovation, facilitate advanced research, contribute to our homeland security and national defense, and fortify our competitive position.

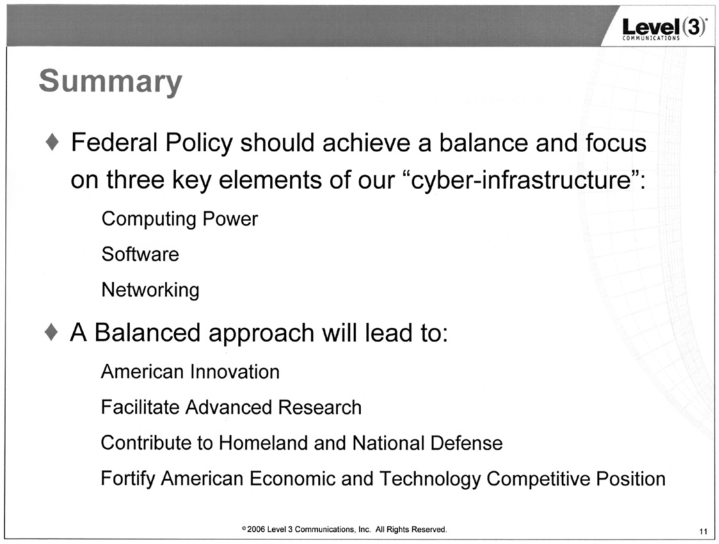


Thank you very much.

[The prepared statement of Mr. Waters follows:]

PREPARED STATEMENT OF JACK WATERS, EXECUTIVE VICE PRESIDENT/CHIEF TECHNOLOGY OFFICER, LEVEL (3) COMMUNICATIONS

Mr. Chairman and Members of the Subcommittee, my name is Jack Waters, and I am Executive Vice President and Chief Technology Officer of Level (3) Communications. Thank you for the opportunity to testify today on behalf of Level (3) Communications. We believe that high-performance networking is an essential element of high-performance computing, critical to our Nation's competitiveness and should be a central part of Federal policy regarding the Nation's cyber-infrastructure.

First, let me commend the Subcommittee for its work in approving the American Innovation and Competitiveness Act (S. 2802). I share the Subcommittee's view that a renewed commitment to basic research will go a long way to ensuring the competitiveness of the United States and to maintaining and improving the United States' innovation in the 21st century by: increasing research investment, increasing science and technology talent, and developing the Nation's "innovation infrastructure." I believe that high-performance computing and the high-performance networks that interconnect and facilitate information-sharing between the high-performance computing centers are key elements of the Nation's "innovation infrastructure," and are essential to ensuring our homeland security, the strength of our national defense, and ultimately, our continued economic competitiveness in the global economy.

Level (3) Communications is a U.S. company focusing on international communications and infrastructure services. Our network has more than 36,000 fiber route miles and provides high-bandwidth services in 15 countries. Level (3) Communications constructed and now operates one of the largest Internet Protocol (IP) and optical transport backbone networks in the United States and Europe, utilizing the latest fiber and optical technologies. Level 3 is regarded as one of the most technologically advanced carriers in the world, recognized by the Smithsonian Institution for building "The world's first upgradeable international fiber optic network to be completely optimized for Internet Protocol technology . . ."

The Federal Government has played a vital role in both high-performance computing and high performance networking for several decades. In 1979, after the successful deployment of the ARPAnet (originally a military network funded by the Department of Defense) the National Science Foundation saw the need to link computer science researchers across the Nation and funded a basic network called CSnet. In 1984, with several advancements in high-performance computing occurring, the NSF funded the construction of five supercomputer centers across the country and connected these centers with a network called the NSFnet. Although more than twenty years ago, these investments, along with subsequent others by the Federal Government, have helped drive many technology innovations that all of us use every day. A few examples will help me illustrate my point.

Today, we all know and use the Internet in both our personal and professional lives. The NSFnet mentioned previously, was a key piece of the early infrastructure that started it all. This network which interconnected 5 supercomputer centers in 1985 and 50,000 networks in 1995, the time of its decommissioning, was the platform on which the commercial Internet that we know today was founded.

Another key piece of Internet technology came from one of the five supercomputer centers that the NSF funded. Although Tim Berners-Lee is quite rightly credited with the idea of the World Wide Web, the first widely used Internet browser was developed at the University of Illinois' National Center for Supercomputing Applications. This browser, named Mosaic, became an overnight success allowing early Internet users the ability to find information across the vast global network. All of this happened in 1993, many years before the world really understood what the Internet would be, through our government's foresight and financial support.

Increasingly, advanced research in the United States and around the globe is accomplished collaboratively by researchers and data sources which are geographically distributed. The quantities of empirical and higher-order data used in this research are also increasing at an incredible pace. As such, the need to share large quantities of information in a timely manner among geographically distributed research centers becomes an essential part of accomplishing the objectives of these advanced research programs. Let me use several examples to illustrate this point:

The Large Hadron Collider (LHC), located at the European particle physics research center, CERN (*Conseil Européen pour la Recherche Nucléaire*), cost approximately $8 billion to construct and is planned to begin operation in 2007. Once online, the Collider will produce an output stream of data approaching a Terabit (one trillion bits) per second which will be shared with 34 research centers around the world. The existing network infrastructure is not sufficient to handle this demand.

In the field of medical research, the newest 1.25 MeV (Mega-electron volts) ultra high-voltage electron microscopes, which allow detailed structural studies of biological specimens, produce network bandwidth requirements that approach 100 Gigabits per second—a requirement equivalent to the capacity planned for the largest American research network, Internet2 now under construction.

Today the Tera-grid network, which recently received increased funding from the National Science Foundation, links seven U.S.-based supercomputing and research centers. Tera-grid has 200 teraflops (one trillion floating point operations per second) of computational capacity, 20 terabytes of storage and will reach sustained data flows between these centers approaching and eventually exceeding 1 terabit per second.

As the first two examples illustrate, the basic instrumentation in advanced research can be so costly that simple economics mandate that these essential elements be shared by the many research centers and scientists rather than duplicating the basic functions. Further, all of the examples demonstrate the trend toward distributed use of enormous quantities of basic research data. Increasingly, refined specialty and inter-disciplinary research initiatives also create an increasing need for collaboration among various research centers and inter-disciplinary research teams. These two factors, cost-efficiency and the need for research collaboration among geographically distributed centers, underlie and motivate the need for efficient, high-performance networks to interconnect these various research centers.

A final case in point is the National Science Foundation's (NSF) Major Research Engineering Facility Construction (MREFC) Program, which provides funding for complex research instruments at 10 centers across the United States, plus one in Antarctica. Each of these centers has an instrumentation and discipline-specific focus—such as ecology, physics, magnetism, etc. The basic data produced by these instruments are shared among the scientific community by manual transference of data or, more efficiently, across networks which can speed the researchers' access to these basic data streams.

It is clear that the Federal Government has historically recognized the need to fund both high-performance computing and high performance networking. The investments made two decades ago have left a proud legacy for the benefit of the entire world. It is also clear that this Subcommittee and the Federal Government recognize the need for continued funding and research in the network component of the Nation's "cyber-infrastructure" and have taken important steps to address these issues. Examples include:

In 2003, the NSF Blue Ribbon Advisory Panel, published a report entitled "Revolutionizing Science and Engineering through cyber-infrastructure" in which it stated, "High-speed networks are a critical cyber-infrastructure facilitating access to the large, geographically distributed computing resources, data repositories, and digital libraries. As the commodity Internet is clearly not up to the task for high-end science and engineering applications, especially where there is a real-time element (*e.g.,* remote instrumentation and collaboration), a high-speed research network should be established and the current connections program extended to support access to this backbone as well as to provide international connections."

The National Science and Technology Council in 2004 called for achieving aggressive networking goals such as:

- networks with 1,000 times existing capabilities,
- with better security and trust mechanisms, and
- development of inter-optical networks and middleware to couple networks with software.

The National Science Foundation's recently announced plans for the Global Environment for Network Innovations (GENI) with the primary goal to enable the research community to invent and demonstrate a global communications network and related services that will be qualitatively better than today's Internet.

In addition to these important Federal initiatives is the work of our Nation's research universities. Recently, the nonprofit consortium known as Internet2, serving more than 200 research universities, took an important step toward meeting the growing bandwidth requirements of many of the United States' top research centers.

On June 15, 2006, Level 3 and Internet2 announced an agreement to design and deploy a new nationwide network and new services to enhance and support the advanced needs of the academic and research community. Internet2's new network will provide its members with 100 gigabits per second (Gbps) of network capacity between key research centers, more than 10 times that of its current backbone. Even with this big step forward, Internent2's members have asked Level 3 to provide a network platform capable of handling even larger bandwidth demands. Accordingly, a key design characteristic of this network is the ability to quickly scale to add capacity as members' requirements evolve.

[Visual Aid]: This map represents a small fraction of the Institutions who are members of Internet2. This illustration also shows a number of the federally-funded research and development centers which will directly or indirectly benefit from the Internet2 backbone.

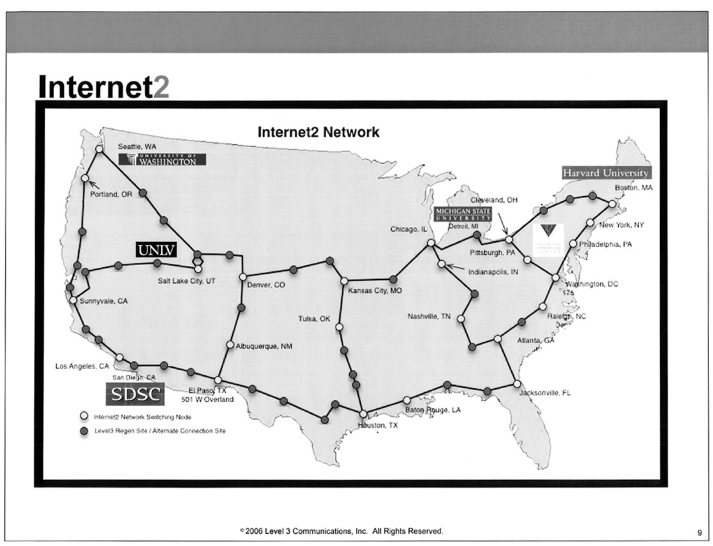


Under the agreement with Internet2, Level 3 will deploy leading edge digital optical technologies to provide multiple ten (10) Gbps wavelengths and enable rapid bandwidth provisioning across the entire network. These new optical services will allow researchers and scientists to obtain dedicated one (1) Gbps sub-wavelengths or entire ten (10) Gbps wavelengths and optimize the utilization of the network to suit the information-sharing needs of the researchers.

In addition to providing high capacity, scalable bandwidth, achieving efficient utilization of the network is critical to ensuring that researchers have the bandwidth they need when they need it. Optimal utilization of network resources improves the economic efficiency of the research, allowing more robust and dynamic use of the network. Internet2's network and the flexible optical services it provides, will increase flexibility and support bandwidth-intensive experimental applications which have direct impact on the United States' research agenda, homeland security, national defense and our economy. Like the Federal initiatives cited earlier, Internet2 demonstrates that the network is an essential component of the Nation's "cyber-infrastructure" and essential to achieving the objectives of the most advanced research being conducted in the United States and abroad.

**Summary and Recommendations**

In summary, I believe that a Federal policy that achieves a balance of investment and focus on the three key elements of the Nation's "cyber-infrastructure"—computing power, software, and networking—is likely to yield the greatest benefits. A balanced approach will: (1) contribute to the attainment of the goals of the American Innovation Act; (2) work to ensure that all of the essential elements of the Nation's "innovation infrastructure" are available to facilitate advanced research; (3) contribute to homeland security and national defense; and (4) fortify the United States' economic and technological competitive position.

I thank you for your time and I am happy to answer any questions you have.

Senator ENSIGN. Thank you.

Mr. Lombardo.

## STATEMENT OF JOSEPH M. LOMBARDO, DIRECTOR, NATIONAL SUPERCOMPUTING CENTER FOR ENERGY AND THE ENVIRONMENT, UNIVERSITY OF NEVADA, LAS VEGAS

Mr. LOMBARDO. Thank you for the opportunity to appear before the Committee and offer observations on the role of Federal policy in the area of high-performance computing as it relates to academic research. My name is Joseph Lombardo, Director of the National Supercomputing Center for Energy and the Environment. The Center is a mature high-performance computing center located at the University of Nevada, Las Vegas. The Center was established in 1989 and has played an important role in the high-performance computing community by providing a resource for academic researchers in the fields of energy and the environment and has an impressive track record of sponsored research from a range of Federal agencies, including the Department of Energy, the Department of Defense, Interior, EPA, Health and Human Services, NOAA, NSF, and others.

I would like to make the following observations. The history of Federal support for high-performance computing has been tied to perceptions that high-end computing is crucial to a broad definition of national security. That is, that the strength of the U.S. is tied not only to military hardware, but to scientific and technological preeminence. High-performance computing is crucial to that preeminence, as it is a basic tool for the advanced research across many fields.

Initial Federal support came with the Lax report issued in 1983 by the National Science Foundation, which perceived that the Japanese sixth generation computer would give Japan a large lead over the U.S. in high-end computing. The Lax report recommended Federal funding for supercomputing centers in open environments such as universities and for training, software engineering, and related activities.

This report led to the formation of the NSF centers as well as other Federal and state-funded centers across the country. In the late 1980s and early 1990s, high-performance computing funding was directed through collaborative relationships between government, corporate sector, and academic research consortia, leading to the formation of policies that established a national priority list of Grand Challenges in research, addressed through high-performance computing tools. This era saw broadening of the high-performance computing manufacturing base, as well as significant software and tool development. This might be considered the highwater mark of Federal interest and funding for high-performance computing.

Beginning in 1993, Federal policy reversed and deemphasized Grand Challenge problems. Grand Challenge problems are extremely difficult to solve, requiring several orders of magnitude improvement in computational capability. The focus then shifted to distributed computing and moved toward commercial off-the-shelf technology. Such initiatives led to a broader range of individuals working in the scientific computing, but basically starved the high end of the high-performance computing field.

At the end of the 1990s, DARPA and other organizations began to see that foreign countries such as the Asian groups were overtaking the U.S. position in high-performance computing once again and recommended policies that would fund and support the high end of the field again. The DARPA High Productivity Computing Systems program is a good example of this shift back toward an emphasis on high-end capability. The DARPA program is focused on providing a new generation of economically viable high-productivity computing systems for the national security and industrial user community in the 2010 timeframe.

This trend has continued with the High Performance Computing Revitalization Act, the President's 2006 State of the Union Address, and with the Fiscal Year 2007 budget, which increased DOE's high-performance computing programs by almost $100 million.

This brief recounting of the history of Federal support for high-performance computing demonstrates that national interest, academia, and high-performance computing communities are joined at the hip. Scientific and technological preeminence for the U.S. is related directly to high-performance computing. Support for Federal funding of high-performance computing has ebbed and flowed as a result of perceived foreign competition. Collaborations of Federal laboratories and agencies, academic institutions, and corporate interests are key to advancing both technologies and applications, but require Federal funding to do so.

Based on the above, I would like to make the following observations: One, Federal policy should recognize high-performance computing as vital to the scientific and technological strength of the U.S. and as such should be considered as crucial to national security.

Two, Federal funding for high-performance computing should encourage development of cutting edge high-end technologies capable of addressing the Grand Challenge problems as well as the midrange problems.

Three, Federal policy should encourage expansion of applications in fields where high-performance computing is not yet a core research tool, as an example agriculture, many of the biomedical areas, and transportation.

Four, Federal policies and funding should be allocated to encourage a diverse industry with a range of companies given opportunity to develop and deploy their technologies. Such broad applications and procurements are crucial to sustain a viable high-performance computing manufacturing community not dominated by a single corporate interest.

Thank you for your interest and for the opportunity to share my thoughts with the panel.

[The prepared statement of Mr. Lombardo follows:]

PREPARED STATEMENT OF JOSEPH M. LOMBARDO, DIRECTOR, NATIONAL SUPERCOMPUTING CENTER FOR ENERGY AND THE ENVIRONMENT, UNIVERSITY OF NEVADA, LAS VEGAS

**Introduction**

Thank you for the opportunity to appear before the Committee and offer observations on the role of Federal policy in the area of High-Performance Computing as it relates to academic research.

My name is Joseph Lombardo, Director of the National Supercomputing Center for Energy and the Environment—the center is a mature High-Performance-Computing Center located at the University of Nevada Las Vegas. The center was established in 1989 and has played an important role in the High-Performance Computing community by providing a resource for academic researchers in the fields of Energy and the Environment, and has an impressive track record of sponsored research from a range of Federal agencies, including the Department of Energy, Department of Defense, Interior, EPA, Health & Human Services, NOAA, NSF and others.

I'd like to make the following observations:

The history of Federal support for High-Performance Computing has been tied to perceptions that high-end computing is crucial to a broad definition of national security—that is, that the strength of the U.S. is tied not only to military hardware but to scientific and technological preeminence. High-performance computing is crucial to that preeminence, as it is a basic tool for advanced research across many fields.

Initial Federal support came with the Lax Report, issued in 1983 by the National Science Foundation, which perceived that the Japanese 6th generation computer would give Japan a large lead over the U.S. in high-end computing. The Lax Report recommended Federal funding for supercomputing centers in open environments, such as universities, and for training, software engineering, and related activities. This report led to the formation of the NSF centers as well as other Federal and state-funded centers across the country.

In the late 1980s and early 1990s, High-Performance Computing funding was directed through collaborative relationships between government, corporate sector, and academic research consortia, leading to the formation of policies that established a national priority list of "Grand Challenges" in research, addressed through High-Performance Computing tools. This era saw broadening of the High-Performance Computing manufacturing base, as well as significant software and tool development. This might be considered the "highwater mark" of Federal interest and funding for High-Performance Computing.

Beginning in 1993, Federal policy reversed and deemphasized "Grand Challenge" problems. Grand Challenge problems are extremely difficult to solve, requiring several orders-of-magnitude improvement in computational capability. The focus shifted to distribute computing and moved toward "commercial off-the-shelf" technology. Such initiatives led to a broader range of individuals working in scientific computing, but basically starved the high end of the High-Performance Computing field.

At the end of the 1990s DARPA and other organizations began to see that foreign countries, such as Asian groups, were overtaking the U.S. position in High-Performance Computing once again, and recommended policies that would fund and support the high end of the field once again. The DARPA High Productivity Computing Systems program is a good example of this shift back toward an emphasis on high-end capability. The DARPA program is focused on providing a new generation of economically-viable high productivity computing systems for the national security and industrial user community in the 2010 timeframe. This trend has continued with the High-Performance-Computing Revitalization Act, the President's 2006 State of the Union Address, and with the FY07 Budget which increased DOE's High-Performance Computing programs by almost $100 million.

This brief recounting of the history of Federal support for High-Performance-Computing demonstrates that national interest, academia, and the High-Performance-Computing community are joined at the hip.

Scientific and technological preeminence for the U.S. is related directly to High-Performance-Computing. Support for Federal funding of High-Performance-Computing has ebbed and flowed as a result of perceived foreign competition. Collaborations of Federal laboratories and agencies, academic institutions and corporate interests are key to advancing both technologies and applications, but require Federal funding to do so.

Based on the above, I would make note of the following observations:

1. Federal policies should recognize High-Performance Computing as vital to the scientific and technological strength of the U.S. and as such, should be considered as crucial to national security.
2. Federal funding for High-Performance Computing should encourage development of cutting edge, high-end technologies, capable of addressing "Grand Challenge" problems as well as mid-range projects.
3. Federal policies should encourage expansion of applications in fields where High-Performance Computing is not yet a core research tool—*e.g.,* agriculture, many bio-medical areas, and transportation.
4. Federal policies and funding should be allocated to encourage a diverse industry, with a range of companies given opportunity to develop and deploy their technologies. Such broad applications and procurements are crucial to sustain a viable High-Performance Computing manufacturing community not dominated by a single, corporate interest.

Thank you for your interest and for the opportunity to share my thoughts with the panel. I would be pleased to answer questions the members may have.

Senator ENSIGN. Thank you.

Mr. Garrett.

## STATEMENT OF MICHAEL GARRETT, DIRECTOR, AIRPLANE PERFORMANCE, BOEING COMMERCIAL AIRPLANES

Mr. GARRETT. Mr. Chairman, Senator Cantwell, Good morning. My name is Michael Garrett, Director of Airplane Performance for the Boeing Commercial Airplane Company in Seattle, Washington. In that role I am responsible for the performance characteristics of all our commercial products, including the new product development such as on the 787. That includes the mission performance capabilities such as fuel burn and range, as well as the environmental performance and capability of our products, such as noise characteristics and emissions.

Today I am going to give you a brief summary of the role of supercomputing, of high-performance computing at the Boeing Company. Let me get this up real quick.

Senator CANTWELL. It is hard when you have to testify and run your own demo.

Mr. GARRETT. Thank you.

The first one I want to talk about is scope. High-performance computing is not only used on our commercial transport aircraft, but on our military aircraft, launch vehicles, and our space vehicles as well. It plays a significant role and has been in the development and design of these products.

With respect to the impact of performance computing, high-performance computing, when it is connected with our computational tools and methods, it kind of supercharges the design process. They provide solutions much faster than we have ever been able to do before, to much more complex problems, more accurate solutions in that we are able to predict better and lower the risk of the development of our products because we can predict how the airplanes are going to perform when they deliver to our customers better than we ever had before.

We also are able to enhance the safety and the environmental aspects of our airplanes, so that they deliver to our customers what they are looking for.

On the business side, we have significantly reduced the cycle time as we develop our products and lower the costs. Basically, high-performance computing allows us to validate technology and get it to the market faster, which results in lower costs and better performance of our products for our customers.

So I want to give one current example of high-performance computing and what it has done. This is in the area of noise reduction. The picture on the lower right shows an application of technology on a 777 flight test airplane. It shows what are called chevrons, these saw-tooth like structures, which is typically a straight leading edge of the nacelles. We have looked at a technology that allows us to reduce noise on the 787 and by doing so and making it the lowest noise airplane in its class, from an environmental standpoint lowering the noise in the areas we live and work.

We did not get there without high-performance computing and our analysis tools, which are shown in the upper left and in the lower left. Those are simulations that allow us to run a number of different configurations very quickly before ever going into the acoustic tunnel to test and before we ever go to flight test. So it reduces the cost because the cost of the wind tunnel goes up and the cost of flight test obviously is going up with the price of fuel. The more we can do to simulate, the better we are.

The future. What we show here is the ability that we need in the future to look at the acoustics of the entire airplane, not just of the engine. The model you see there that is represented from our computational fluid codes uses high-performance computing to do a very complex type of problem with separated flows, looking at the gear down, at the flaps down. It is time-based because it is separated flow, so it is time-dependent as well as running the engines at the same time.

This solution takes hours. We need to get this down to seconds because this is only one solution. We need multiple solutions to handle all the flight conditions that are needed, and this is where the future is.

So in summary, four major points. High-performance computing is used throughout the Boeing Company from a product development standpoint and we will be using it more and more in the future as well. It allows us to meet our business case conditions by reducing our cycle time and getting our products to the market faster. We are investing in high-performance computing now primarily because of the 787. On an annualized basis, we have increased our investment 50 percent year over year the last 7 years. We have invested tens of millions of dollars in the development of our new products.

Last, the continued improvements in high-performance computing through faster and more efficient computers will enable Boeing to provide more efficient and more capable products to our customers at reduced cost.

Thank you.

[The prepared statement of Mr. Garrett follows:]

PREPARED STATEMENT OF MICHAEL GARRETT, DIRECTOR, AIRPLANE PERFORMANCE, BOEING COMMERCIAL AIRPLANES

**Introduction**

Good morning, Mr. Chairman and Members of the Subcommittee. I am Michael Garrett, Director, Airplane Performance for the Commercial Airplane Division of the Boeing Company. I have worked at Boeing and McDonnell Douglas for 27 years with a broad range of experiences in product development, program management and marketing. In my current position, I have responsibility for the overall performance characteristics of all Boeing Commercial Airplanes, including new products, such as the 787.

At Boeing, we pride ourselves for understanding our customers' requirements and then developing, designing, building and delivering airplanes that meet or exceed those requirements.

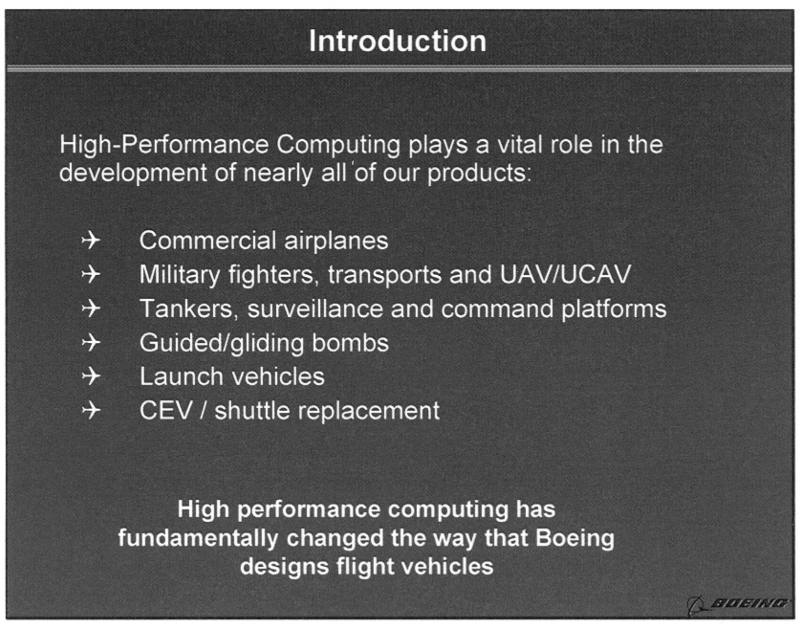


High-performance computing has fundamentally changed the way that Boeing designs flight vehicles, whether it be commercial transports, military fighters, unpiloted aircraft, guided bombs, launch vehicles, or crewed-space exploration vehicles. Computational tools are being used to create numerical simulations to assess system performance—replacing the more costly and time consuming requirements for physical testing. For example, these new tools are being used to determine the aerodynamic performance of entire airplanes, the optimum structural layout to minimize weight, and the radar cross-section of a stealthy vehicle. It is the evolution of computing hardware that has enabled more efficient simulations with reductions in overall design cycle times.

**High-Performance Computing Is Good Business**

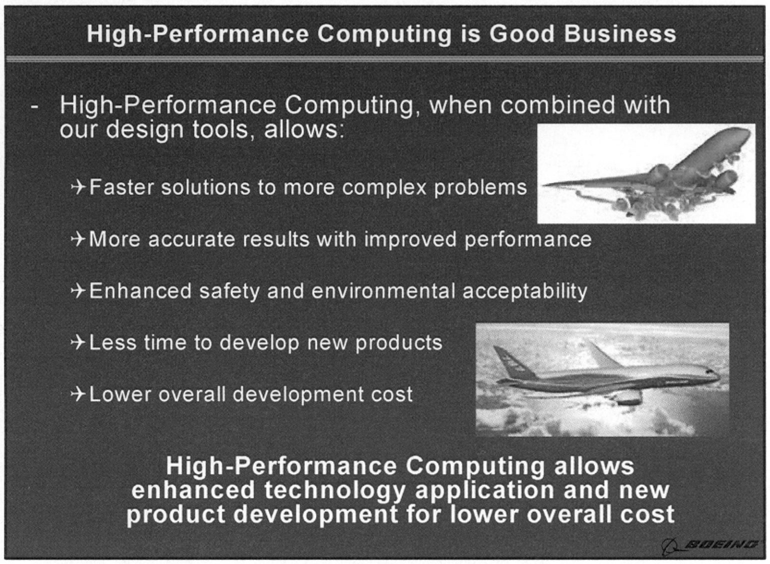


When we combine our computational design tools with the high-performance computing resources, we obtain incredible efficiencies in the design processes we use to develop our airplanes. Complex processes and simulations, such as computational fluid dynamics, can be run much more quickly, at lower cost, and at a level of fidelity and accuracy that is equal to that achieved in physical vehicle testing. While we will never eliminate wind tunnel and flight testing, more powerful computing tools allow us to better predict results, therefore reducing technical risk, while reducing physical testing costs.

**High-Performance Computing in Computation Fluid Dynamics Applications**

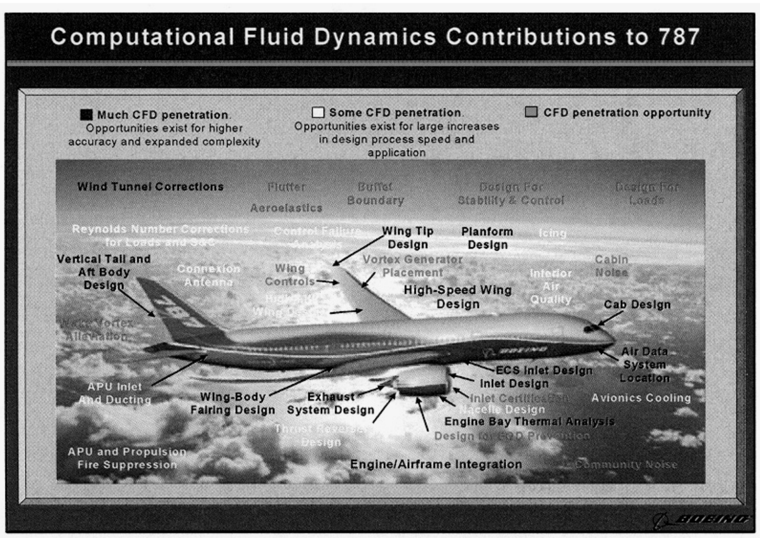

One of the best utilizations of high-performance computing is in the development of computational fluid dynamics (CFD). While CFD has been in use at Boeing for 30 years, the most extensive application has been on our newest commercial aircraft, the 787 Dreamliner. The use of CFD tools has allowed Boeing engineers to address a wide range of design challenges, including traditional wing design, the even distribution of cabin air and reduction in overall airplane noise.

**High-Performance Computing in Aircraft Noise Reduction**

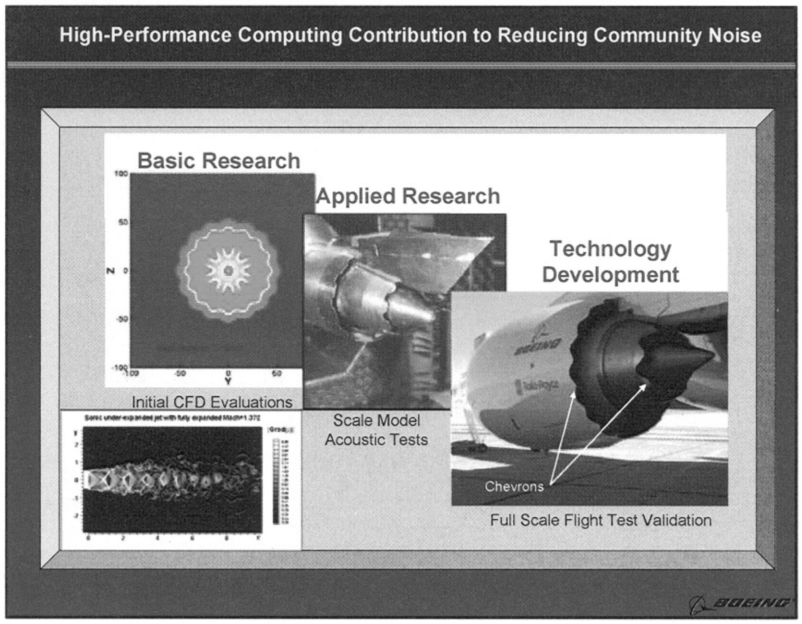


High-performance computing, together with our CFD tools, has also played a significant role in reducing airplane noise. The example above shows how an engine noise-reduction feature called "chevrons" was developed for application to our commercial airplanes. The 787 will be the first Boeing airplane with this technology. We were able to simulate the noise reduction characteristics of numerous chevron configurations and determine the best configuration for noise reduction before ever testing in the acoustic tunnel or in actual flight test. This means the 787 will be a quieter aircraft, making it more environmentally-friendly for those who live and work near airports.

It is the knowledge gained from the this process that reduces the product development life-cycle allowing our customers to get our products faster while meeting our commitment to improve the environmental performance of our airplanes.

**High-Performance Computing in Wing Design**

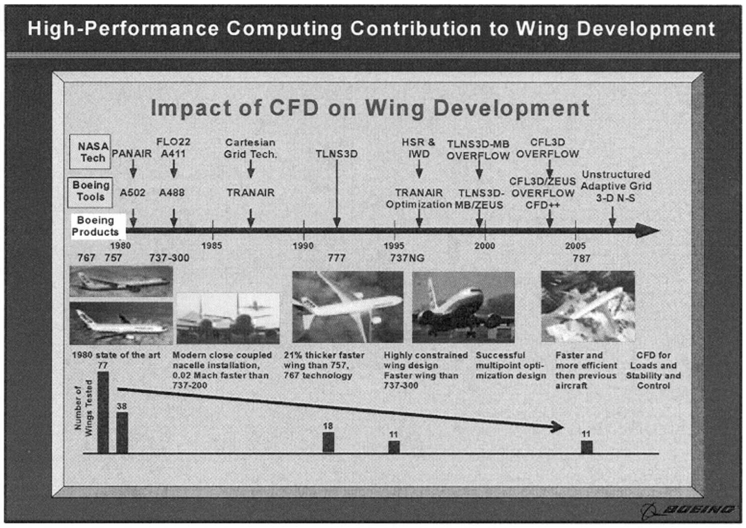


Another example of the benefit of high-performance computing and improved computational capability has been in our wing development over the last 25 years. In 1980, Boeing tested 77 wings in wind tunnels to arrive at the final configuration of the 767. Just 25 years later, we built and tested 11 wings for the 787—a reduction of over 80 percent. Those 11 wings took less people to design, less time to design, and the wind tunnel results matched the CFD predictions.

While our CFD tools today are very good, there are still some flight conditions that cannot be simulated very well. These conditions will continue to require significant wind tunnel testing. As more advanced computer hardware is developed, the computational tools and processes should improve and we will one day be able to calculate the airplane's characteristics everywhere in the flight regime with high fidelity.

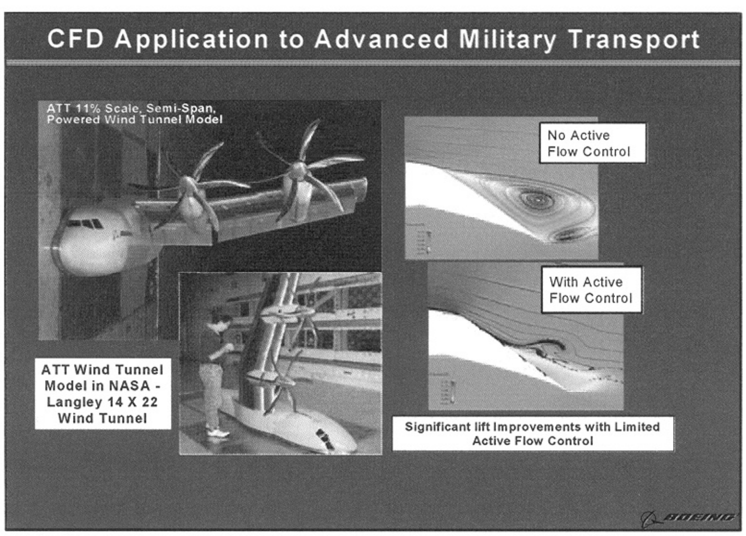

The chart above shows another wing design application which resulted in a configuration that could not have been designed without CFD design tools and high-performance computing. The Air Force has a requirement for a battlefield delivery transport. It must operate out of unimproved, very short runways—runways shorter than C–17s can use today. Meeting these challenging specifications requires a new and innovative wing with performance never previously demonstrated. A new active flow control technology was evaluated to achieve that performance using CFD. This computation, shown in the pictures above, demonstrated that when air remains attached to the wing, as in the picture in the lower right hand corner, lift is increased. This additional lift enables the aircraft to take-off and land in shorter distances. After computer simulation, this concept was then successfully demonstrated on the Advanced Tactical Transport model at the NASA Langley wind tunnel.

**Future of High-Performance Computing**

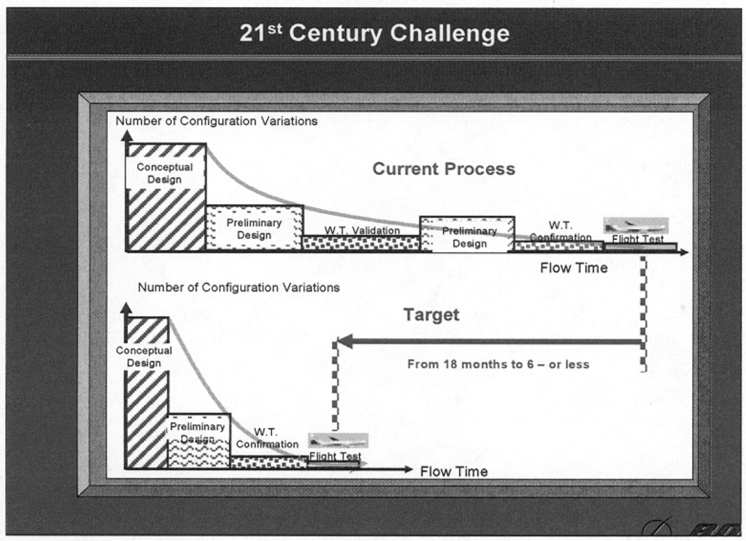


As previously stated, we have reduced the amount of wind tunnel testing required for new product development. Our vision for a future design environment would be that all simulation work would be done computationally, enabled through more powerful high-performance computing tools. This would allow us to test only two or three wings in the wind tunnel versus the 11 for the 787. This will not only dramatically reduce the non-recurring cost to develop an airplane but also reduce the time it takes to bring a new product to market. Instead of developing a new airplane once a decade, we can potentially develop one in significantly less time, allowing us to be more responsive to market demand.

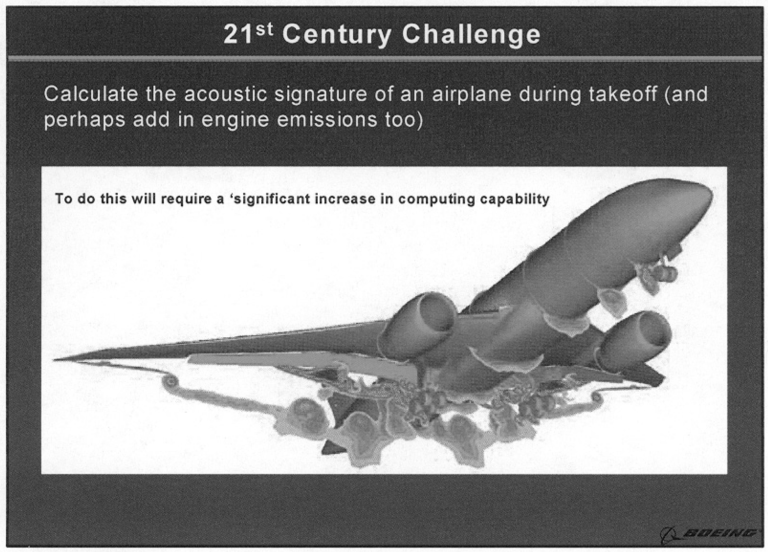


An even greater challenge lies in the area of acoustics. Today laboratory and/or flight tests must be conducted to determine the acceptability of candidate airplane configurations to meet community noise requirements. In the future we hope to do all the simulations within the computer.

This is a problem that is probably decades away from being addressable because noise covers a wide frequency. The numeric grid to capture the shorter wavelengths drives up the size of the problem dramatically—as does the requirement to model the landing gear, all flap and slat details, the engine (running!)—and it is all time-dependent. As the hardware continues to improve, we will incrementally work our way up to meeting this challenge.

**High-Performance Computing Usage at Boeing**

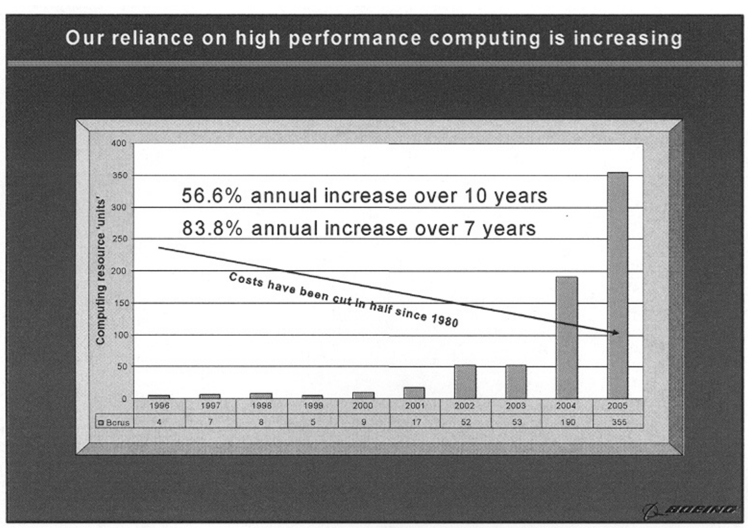

Boeing is committing large amounts of resources to provide the necessary computing capability we require. During the development of the 787, we have nearly doubled the capacity of our high-performance computing data center year after year. This is a big investment of capital, but one that we are willing to make because there is a measurable return for that investment. While our high-performance computing usage has increased, the cost per unit has been dramatically reduced by 50 percent making our development tools more and more cost effective.

**Conclusions**

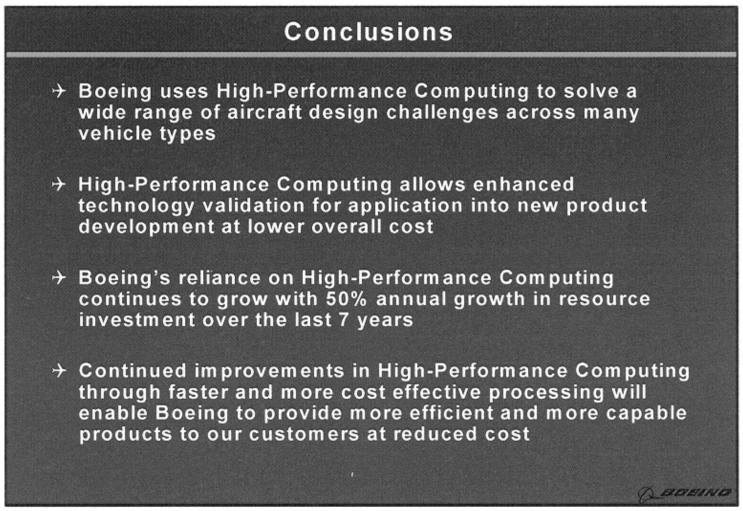


Boeing has made extensive use of high-performance computing in addressing a wide range of issues across all of its products. While high-performance computing is a valuable tool across the entire product cycle, its primary contribution has been in technology validation and its application into new product development.

Our reliance on High-Performance Computing continues to grow as better, faster and more cost effective processing is available. This will enable Boeing to deliver better value to our customers through products that are more efficient and capable at significantly lower cost.

Again, Mr. Chairman, I appreciate this opportunity to testify before the Subcommittee.

Senator ENSIGN. Thank you.
Dr. Burt.

## STATEMENT OF STANLEY K. BURT, PH.D., DIRECTOR, ADVANCED BIOMEDICAL COMPUTING CENTER

Dr. BURT. Thank you, Mr. Chairman, for allowing me to testify today. My written report has many examples of high-performance computing applied to biology and computational bottlenecks. I will limit myself to just a few remarks relating to biology, high-performance computing, and biomedical research.

I am the Director of the Advanced Biomedical Computing Center, ABCC, which is the principal high-performance, high-capacity computing resource for the National Cancer Institute of the National Institutes of Health. The ABCC was founded in 1986, when its first supercomputer came in. We provide researchers high-performance computing tools for their research in the complex field of cancer. The goal of the ABCC is to provide the cyberstructure for data-intensive computing. And unlike some other supercomputing centers, the ABCC only supports biological research.

Both molecular biology, which is the basic science, and oncology research in high-performance computing have advanced dramatically over the past decade. In fact, high-performance computing has emerged as a basic tool to address very complex issues, commonly referred to as "Grand Challenges," in most areas of research. Computer simulations in fields such as physics and chemistry have become the third leg of the stool along with theory and experiment.

As molecular biology has advanced rapidly, there has not been the adoption of HPC as a necessary tool and the simulation of biological problems has not in general been carried out by true molecular biologists. Applications of HPC in medicine and the biosciences have lagged behind the physical sciences. While Boeing designs its new airplanes using HPC technology, all too many molecular biologists are still at the bench using familiar approaches to address the very complex issues and identifying and assessing normal and abnormal cell structures that are the basis of cancer research.

There is the need for biology to change to a hypothesis-driven research field. This is the perfect role in which HPC can be used. However, computational biology in my estimation is today where computational physics was in the late 1960s and 1970s. The use of HPC will be required if biologists are not to be overwhelmed by the amount of data being generated. This is a particular feature of biology. The sequence of the human genome has been completed. Another 100 genomes have been sequenced and published and another 700 are in the works. The amount of information that can be derived from these genomes is immense.

In addition, experimental methods and chip designs have increased the enormity of the data. For example, we now have a chip about the size of a microscope slide that is able to detect every known mammalian and bird virus. This is how SARS was detected. These chips have the potential for producing hundreds of thousands of data points and even millions of data points for altered disease and up-regulation, down-regulation, and appearance of proteins in cancer.

In addition, the National Cancer Institute has invested in other experimental techniques, such as mass spectroscopy, imaging, nanotechnology, et cetera, all of which will contribute more data. The data must be annotated, curated, and analyzed in order to extract information, but it is not just information that we are after; it is actually knowledge.

The goal is not to create just databases, but the goal is to create knowledge databases. These bases must be usable by nonexperts in computer science, especially by clinicians, if we are going to have translation into medical benefit. This is a challenge for HPC and the ability to store data.

The ABCC has been in the forefront of the effort the expand the application of HPC into the field of molecular biology and cancer research. We provide advanced computational technologies, high-performance computer software, scientific expertise, and scientific support to both NIH internal researchers, NCI, at the universities and research institute scientists around the country.

One example of this is in the discovery of cancer biomarkers, which is extremely important to the NCI because early detection is key for treatment of cancer. The ABCC has been involved with cancer researchers at Texas A&M, M.D. Anderson, Eastern Virginia, and others to analyze data for biomarkers for different cancers, and we have helped identify markers for bladder and colorectal cancers which could lead to inexpensive screens for cancer.

The ABCC offers or provides training and related outreach programs to scientists in computational biology technology, and I am pleased to say that we have begun offering classes in molecular modeling and our classes are filled within 30 minutes of being announced.

We also reach out to other agencies to adapt technology from other fields to molecular biology. For example, we adapted technology developed by the National Security Agency in cryptological pattern recognition to the study of, recognition and study of human genome work, especially tandem repeats. Tandem repeats are pieces of DNA bases that are expanded in many disease states. We were able to use technology developed by the National Security Agency and we could scan 150 million bases on a chromosome in 2 seconds and find all the tandem repeats, and we could do the entire human genome in 2 minutes, something that had never been done before.

We have also applied field programmable gate arrays to certain biological algorithms to exceed thousandfold increases, but programming field programmable gate arrays is tremendously difficult.

The ABCC has an ongoing funding line in the National Cancer Institute and the NIH budget, which enables our staff to provide the missionary work of enhancing cancer research protocols with the advanced tools of HPC. This is useful and I think a critical element in the fight to cure cancer. The Congress and the Administration, the NIH and the NCI, should be commended for developing policies and a program that fund these activities.

I do have a couple of recommendations to make. As I mentioned before, cancer research and molecular biology scientists are substantially behind physical scientists in the application of HPC to their research procedures and protocols. Accordingly, the outreach program of ABCC and other institutes that promote HPC as a basic research tool in their biosciences should be encouraged and supported in their work. The National Cancer Center's program should embrace and promote the use of HPC tools and approaches as they credential and support programs of the external centers, such as Anderson, Sloan-Kettering, San Francisco, and others. The center's program should also consider funding HPC applications at new emerging centers across the country, as opposed to the low level of support and sometimes the non-forethought of needing high-performance computing for their analyses.

I recommend that the Congress and the Administration encourage a fundamental change in the education of biologists in our university community to provide much greater emphasis on computational biology, which I believe will be the foundation of future biology and medical research. As a matter of fact, I just finished a RAND report to the National Cancer—National Intelligence Committee, on technology for the future, in which the statement is made that “The intersection of high-performance computing and biology and the fields that are developed through this effort, such as nanotechnology and biomedical materials, will be the future.”

The third thing I recommend is that the Congress and the Administration raise funding for applications of HPC in biomedical research to that of an NSF program and that agency’s track one and track two supercomputing program, and that a major flagship national center similar to the model of the National Center for Atmospheric Research in Boulder should be commissioned and funded on the main campus of the NIH to create a critical mass for computational medical research, with emphasis on cancer, heart disease, and other genetic diseases. This new center should possess a variety of computer architectures, should be added to the NIH Teragrid and the DOD research and engineering network to make its resources widely available for biological researchers around the country.

Biology is a particularly complex and difficult field. It requires the synergy of physicists, computer scientists, mathematicians, biologists to all work together and to be able to speak the common language. One of the difficulties in translating HPC into action in biology is that the person who is performing the simulations and helping the biologists analyze their data must be able to understand the language that the biologists are speaking.

I recommend that the ABCC outreach program that brings HPC tools and resources to cancer researchers across the country be extended and expanded, especially to new cancer centers that are coming online.

That concludes my remarks and I appreciate the opportunity to appear before the Committee today and welcome questions regarding my testimony or the field of HPC applications in biomedical research.

[The prepared statement of Dr. Burt follows:]

### Prepared Statement of Stanley K. Burt, Ph.D., Director, Advanced Biomedical Computing Center

Over the past couple of decades many important milestones in biology have been obtained. These include completing the genomic sequences from several mammalian genomes, including human, and producing draft sequences for several additional genomes. Also, new technologies allowing for simultaneous measurement of mRNA expression levels for thousands of transcripts and application of this method to RNA samples from tumors and normal tissues have identified many genes whose expression is influenced by cancer and other disease states. Further improvements in this technology and other discoveries have led to chip technologies capable of simultaneously monitoring the entire genome for person-to-person variations including both single nucleotide polymorphisms (SNPs) and other larger alterations such as deletions and duplications. Other data derived from experiments that measure microRNA expression levels and transcription factor binding studies have also contributed to the extensions of this technology. Finally, measurements of protein expression levels using high throughput mass-spectrophotometry and chip-based tissue and antibody arrays have given biologists the ability to correlate changes in mRNA expression with changes in protein expression levels that may contribute to the disease process or at least be markers for these altered physiological states.

New instrumentation in microscopy has allowed for simultaneous monitoring of cells responses to drugs or other agents in parallel. Further, con-focal imaging techniques have allowed for multiple slices of the same fields to be examined in detail so that a three dimensional image of a specimen can be reconstructed. Whole-animal imaging is also being used to study drug distribution throughout an animal’s tissue.

Other methodologies allowing for higher levels of protein expression and purification are being leveraged to allow for more direct biochemical metrics of an enzymes function to be collected. High throughput binding technologies can be used to determine affinities of proteins with cofactors and drugs. Better docking software applications now exist to screen some of these interactions *in-silico.* Newer, more sensitive and reliable methods have been developed out to identify protein-protein interactions. The biomedical literature has also grown dramatically as all of these new methods and the data associated with each of them has increased.

Taken together, these new methods and the need to process and analyze the data produced by them have resulted an explosion in the need for high-performance computing in biology and medicine. This need requires both increased capacity, as the sheer volume of data generated is considerable, and also increased capability. One of the confounding problems associated with the needs analysis of this problem is that there does not appear to be any single solution to the problem. Because of the diversity in the algorithmic requirements for analysis of each of these data types, no particular computer hardware seems suited for all of the problems. Thus, some of the problems are embarrassingly parallel, meaning they are ideally suited to a cluster environment. Good examples of this type of application would be comparing fragments of one genome to another, where each computation is entirely independent of the other. Advances in microarray plating technology now allows for increased spot density. This translates into a tremendous increase in the amount of data from a single experiment, at a significantly reduced cost. In addition, since these experiments only produce useful results when they are run for many samples (*e.g.,* tumor and normal tissues) a greater volume of data is produced. This is leading to the need for the biologist to have access to computers with more memory and higher processor speeds to allow the data to be analyzed in a reasonable time. Already, the ABCC has received requests from cancer biologists for help with genomic analysis of promoters, control regions, miRNAs, better annotation of the genome and comparison of genomes, understanding of fragile sites, sites of chromosome translocation, and the relationship to cancer of segmental duplications. In addition, the new 500K SNP chips are flooding researchers with data that requires big computers to process, store, and interpret. Cancer biologists want new methods to look at the data and estimate haplotypes and look for interaction among many loci. This and all of the abovementioned challenges require the use of HPC resources.

Another area in which cluster computing can be useful is in biomarker discovery. Aside from prevention, diagnostic tools to detect cancer at an early stage are of great benefit to patients. Great efforts are being made to identify biomarkers from gene or protein expression profiles. One tool being used to find biomarkers is mass spectroscopy, which can identify proteins and their fragments based on their size and electrical charge. In mass spectroscopy experiments thousands of spectral peaks are produced. These peaks are then used to find biomarkers for proteins. Because there are so many data points that are trying to be fit to few markers, this can lead to false results because the problem is over-determined. In order to avoid this mistake, one needs to perform thousands of calculations to develop a consistent set of models to find the proper biomarker. The ABCC does this by using methods that converge on a model that has the same biomarkers in each solution, thereby guaranteeing a biologically relevant answer. This procedure can benefit from hundreds of processors, but large memory is not needed since each calculation is independent of the others. The ABCC has been successful in finding biomarkers for bladder cancer and colorectal cancer. Hopefully, these markers, which are derived from urine and human serum, will translate into efficient, inexpensive screens that can be used for early detection of these cancers.

Another problem that confronts biological computing and cancer research in particular is the sheer volume of data that must be collected, analyzed and compared. Data already exists in older databases in many places and in different formats. Part of the problem is already being approached by the NCI through its caBIG (Cancer Bioinformatics Grid) initiatives of NCICB and it involves identifying and leveraging information technologies that facilitate data interconnectivity, amongst other goals. In this regard, the development and enforcement of data exchange standards through caDSR and caCore are designed to bridge the gap between a clinicians and a bioinformaticists perspective of a set of genomic data. In order to analyze and house this data, there needs to be a computational infrastructure and visualization capabilities. Furthermore, while distributed databases are convenient for data maintenance, the National Security Agency has found that having all the data reside locally, where it can be called into computer memory, is essential for rapid data scanning. This will require HPC resources with large memory resources. Also, database consolidation is not enough. There needs to be development of methods for the construction of a knowledge base in which nonexperts, especially clinicians, can query data from various sources. This will require a serious research effort in knowledge base development area, although some manufactures have obtained preliminary results in this area. In addition, because there are problems suited for both hardware configurations mentioned above, the data I/O infrastructure must also be able to be connected to both of these scenarios. Again, because of bandwidth issues resulting from the sheer volume of the data, this results in a need for new technologies in computer architectures.

Another complicating factor in data combination and analysis for biological research is that while massive storage and bandwidth have become relatively cheap and abundant, the data can not only be from different sources but it can represent experiments in different scales, from years to femtoseconds—time-scales that go across orders of magnitude. This is a problem that is referred to as multiscale modeling, and it is a profound problem in computational science. Solving this problem will require a commitment of resources to advanced architecture development, more efficient algorithms, and clever data reduction.

I will now address some computational bottlenecks for a few areas that the National Cancer Institute has identified for their roadmap.

**Nanotechnology**

Nanoparticles typically have dimensions smaller than 100 nanometers, which are smaller than human cells. Nanometer devices smaller than 50 nanometers can easily enter most cells. Nanoscale devices can interact with biomolecules on both the cell surface and within the cells. Despite their small size, nanoscale devices can also hold tens of thousands of small molecules such as a contrast agent or a multi-component diagnostic system capable of assaying a cell's metabolic state. This can provide a mechanism for detecting cancer at its earliest stages. Nanoscale constructs, such as dendrimers and liposomes, can provide customizable drug delivery to targeted cancer cells or tissues. This has already been demonstrated experimentally.

While nanoparticles have great promise, it also has to be demonstrated that they are not toxic to normal tissue. The ABCC is supporting the NCI's Nanoparticle Characterization Laboratory through modeling of bulk properties and calculation of atomistic properties. At the nanoscale, the physical, chemical, and biological properties of matter differ fundamentally and often unexpectedly from those of corresponding bulk material because of the quantum mechanical properties of atomic interactions which are influenced by material variations on the nanometer scale. Modeling of bulk properties such as surface charge or shape is not difficult. The calculation of atomic level quantities is a huge computational issue, even atomistic calculations on quantum dots are beyond our current capability, and will require large increases in HPC.

**Drug Design**

Over the years there has been great success in drug design using HPC. Drug design is usually done against a protein target, such as an enzyme whose function one wants to inhibit. A great recent example is the discovery of Gleevec, an inhibitor of protein kinase activity, which brings about complete and sustained remission in nearly all patients in the early stages of chronic myeloid leukemia. If the structure of the protein is known, docking calculations can be performed. This usually involves docking thousands of molecules into an active site and scoring the resultant interaction. If the docking is done with rigid molecules, the calculations are fairly trivial. If, however, flexibility is allowed, and most proteins and ligands do flex, then the problem becomes enormously computationally expensive.

If the protein structure is not known, and the protein is not similar to another one, then one must perform *ab initio* structure determination. David Baker's group at Illinois took approximately 150 CPU days to determine the structure of the CASP6 target T0281. Also to do a docking interaction between two proteins took 15 CPU days. He makes particular note that his group is limited by computational power. Our group has studied the enzyme mechanism of many enzymes involved in cancer. For an enzyme named Ras, which is mutated in over 30 percent of known cancers, we modeled 1,622 atoms of the protein by molecular mechanics and only 43 atoms by quantum chemistry. These studies took several years and were bound by computational power. To calculate reaction surfaces normally takes several months of time on HPCs. Luthey-Schulten's group at Illinois did molecular dynamics simulations of Imidazole Glycerol Phosphate Synthase, an enzyme involved in making DNA and RNA. It took 10 hours, 12 hours, and 40 hours to animate one nanosecond on three cluster machines (with different processor speeds). It takes many nanoseconds of simulation to just relax the systems to prepare for further simulations. It has been estimated that to go from nanoseconds to milliseconds will require an increase in computer capacity of approximately 1,000,000. This can only be achieved by the combination of improved hardware and software.

**Integrative Biology**

Computer aided design of HIV protease inhibitors remains one of the most successful stories in modern biology. Although this was a remarkable achievement, the complexity of a single viral particle pales in comparison to characterizing the complete catalog of the cell (the proteome) and the full map of the interactions of the members of the proteome. For a subset of interactions of the proteome, the immunome, the combinatorial problem of treating all possible pairs in the immunome (1,000,000 of them) escapes the capacity of current computers.

**Synzymes**

There is great interest both in academia and industry for the creation of artificial enzymes that are much smaller but duplicate the enzymatic activity of the large natural ones. Because they are smaller, they can be tethered to other molecules or nanopartices, such as dendrimers or liposomes, and delivered to a particular targeted area such as a tumor cell. The ABCC staff has experience in this area. We modeled a particular inorganic catalyst know as Mn-salen, which is used commercially in the chemical industry for epoxidation reactions. After studying this reaction, we were able to convert this catalyst into one having biological activity and could act as a free radical scavenger. This could be useful for traumatic injuries, strokes, or even for cancer. However, this falls into the same category as enzyme mechanism studies and the calculations take months and months to perform. Complete characterization of these reactions took several years running on fast HPCs.

**Specialized Hardware**

Being able to take advantage of specialized HPC resources and software written for those resources can lead to dramatic increases in time to solution. In one instance, the ABCC staff in a research parternership with several NCI biologists investigated how to rapidly scan for microsatellites (tandem repeats). Tandem repeats are groups of DNA nucleotides ranging from two to sixteen bases that are expanded in several diseases. For example, in normal people there is a pattern of DNA nucleotides, CAG that is expanded 10–35 times. In Huntington's this same pattern is expanded between 36–121 times. In the past finding these repeats were found in a heuristic and probabilistic manner on conventional computers.

Using specialized hardware such as bit matrix multiply and pop count, which had been requested by the NSA to be incorporated in the machines they were using in order to perform rapid pattern matching, we, along with industrial programmers, were able to drastically reduce the time to find all tandem repeats on chromosomes and the entire human genome. To scan a chromosome of approximately 150 million bases took 2 seconds. To scan the entire human genome took 2 minutes. We discovered 47 potential disease sites, 8 of which could be associated with cancer, and we more than doubled the known numbers of repeats. We also used this specialized hardware search for another genomic feature, known as segmental duplications, which are associated with diseases. This involved finding clusters of DNA bases approximately tens of thousand bases long that are separated by approximately 1 million bases from another cluster of bases that are the complimentary complement of the original DNA base cluster. When these complimentary clusters find each other during replication they combine and huge sections of the genome are excised. We could not have done this without these specialized hardware features.

We have also used FPGAs, which are reprogrammable hardware and support the custom computing needs that are characteristic of data-intensive problems. We programmed the FPGA for a powerful sequence alignment algorithm known as Smith-Waterman. The Smith-Waterman alignment method is a powerful algorithm for aligning sequences in which there may be gaps and one is trying to find the "best" alignment. This algorithm is widely used in the biological community but is particularly computationally demanding. We obtained speed-ups of over a thousand fold. However, the difficulty is that the programming of FPGAs is not a trivial task, and one that would not be normally within the expertise of a biologist. However, FPGAs offer great promise because there are expected to be huge increases in performance on these types of machines.

**Recommendations**

It has been said that biology will be the science of the 21st century. Due to the complexity of biology, the sheer volume of data, the fact that the environment of a cell, (particularly for cancerous cells) must be taken into account means that biology must be tackled using a systems biology approach. This means that teams of scientists such as biologists, computer scientists, mathematicians, physicists, and chemists should work on these problems in conjunction. In order to do this, it will require cross-training to have a meaningful dialogue. I believe that in order for the United States to remain competitive we should devote funding to education and training in the above disciplines. We also need to find mechanisms to encourage young people to enter the scientific field. I have seen for several years the lack of U.S. citizens applying for jobs in the ABCC. I believe that this reflects the national trend.

As biology matures the use of HPC in biological research will grow. There is clearly a need for large memory assets coupled with fast processors. I believe that cluster computing will still have its place, but as the problems grow in size and complexity, the need for HPC resources will be inevitable. One can already see this trend in Europe where several national centers have made purchases of Blue Gene machines, and others have made investments in large memory machines.

There needs to be funding of new computer architectures, specialized hardware, faster interconnects, etc. One area of funding that is especially important is software development. One thing holding back HPC development is that the software available today is not written for HPC machines. Sometimes software engineers spend considerable time to port nonparallel applications to parallel machines without much increase in speed or efficiency. We are running "old" software on newer architectures. Along with developing new software, research into new compilers must be encouraged.

I also recommend that the United States fund several centers for Integrative Computational Technology for Systems Biology. These centers would provide for the integration of biology with strong computational infrastructures and analytic tools. These centers need to provide intuitive, visual interfaces for biologists with real-time interactive data analysis. These centers could also serve as training facilities and facilitate communications between scientists of diverse backgrounds, disciplines, and expertise within a common framework. These centers would also facilitate the interplay between discovery and hypothesis-driven science. Several other countries are already creating such centers.

Maintaining a leadership role is vital for the economic health of the United States. We need to maintain our leadership in HPC in order to have the advantage in intellectual property, which is connected to our economic well-being. Support for our HPC industry is vital. Countries such as Japan, China, and India are making substantial investments in HPC. We need to do the same.

The need for supporting HPC extends across all of the hard science disciplines. I hope that I have been able to show in this statement that the increased need for this support is arising from biology. A recent RAND report entitled "*The Global Technology Revolution*" was prepared for the National Intelligence Council. In this report it summarizes how the future will be determined by the intersection of IT and biology, and the industries such as nanomaterial, materials, and biotechnology that are spun from this intersection. Clearly, the future is in this area. We should make the investment now.

Senator ENSIGN. I want to thank all of you for some very fascinating testimony.

I am going to take about 5 minutes and cover one topic and then I am going to turn it over because Senator Cantwell has to leave and let her spend some time, and then I am going to come back and ask a few other questions.

I want to explore this because at this hearing we have academia represented, we have the private sector and industry represented, and I have always believed that it is a fundamental role of government to conduct basic research. Applied research is more the responsibility of the private sector, and sometimes there is that little nebulous area in between where sometimes we use government programs to try to bridge the gap. But when we are looking at the idea of basic research and funding basic research in the area of HPC—now I am using that acronym; I did not know the acronym before we got it today, but I will join right in with it.

Can you take a minute each and talk about where the line between basic and applied research is when it comes to HPC?

Dr. WLADAWSKY-BERGER. If I may start, in the past the bulk of the basic research was in the actual technology component. While some of that, especially pushing the leading-edge, might be needed, what we are finding is that the really complicated problem is how do you put the whole system together, including software and applications. That goes well beyond the components and that is where the testbeds, the pilots focusing on applications, are so important.

For example, what Dr. Burt talked about. Focusing a problem on advanced cancer research or brain mapping research, which is a major area, brings all these ingredients together. The testbeds are extremely important that become a kind of boundary between researching the components and beginning to get it into the marketplace, and then the private sector takes that and then they themselves bring it to a lot more applications, a lot more business areas, and make it less expensive and so on.

Senator ENSIGN. Thank you.

Mr. Jehn.

Mr. JEHN. Yes. I would like to just echo what you said a moment ago, and that is to say that this is really a classic example of a place where Federal support for R&D is justified, and the government itself has recognized this in the reports I have cited. I would simply echo those as well, that in addition to basic research that would develop the building blocks that companies like IBM and Cray could use to develop next-generation supercomputers, support for research is also required in the area of basic architectures and ultimately to build prototypes.

In fact, the Federal plan for high-end computing and the similar report that the Defense Department released the so-called IHEC report, proposed a four-component R&D program: one, basic and applied research; two, applied developmental work; three, building prototypes; and four, establishment of a handful of laboratories in the government that would consolidate this. A program like that should support R&D in government, in industry, and in academia. That systematic approach is what is missing today in Federal policy.

Senator ENSIGN. Mr. Waters.

Mr. WATERS. When I think of basic research, I think a little bit back in the past when it all used to sort of happen in one place, whether it was a lab or a campus environment. I actually came from Dr. Burt's lab. About 20 years ago, it was my first job out of school. What happened was that the basic research occurred on one campus.

I think what is occurring is that basic research is getting so expensive in some examples that the instruments themselves cannot be duplicated. So that the need is for a distributed architecture that we have not had for basic research, and I think that is where a lot of emphasis needs to be put, whether it is grid computing, whether it is distributed storage, the middleware or the software that sits and controls how people use distributed computing. That is not an area where I think we have focused a whole lot.

Senator ENSIGN. Is that part of the government role or is that more the role of the private sector?

Mr. WATERS. I think the government does play a role and can play a role in that basic research, because it is not an area that I think industry has focused on.

Senator ENSIGN. Mr. Lombardo?

Mr. LOMBARDO. From my perspective, the basic research really chimes in with the Grand Challenge problems. We are looking at answers and solutions that may be a decade away, maybe 20 years away. So it fits very nicely in that realm, and I can see where the Federal Government would want to fund that research. And the opposite is also true. If there is a solution that is ready in 3 months, 6 months, 18 months, that may be better left for the corporate side or the corporate sector to handle.

Senator ENSIGN. Thank you.

Mr. Garrett?

Mr. GARRETT. Sir, from a user perspective, most of our research and development is going into the tools that we use to do the analysis, not into the high-performance computing area, though our tools are running neck-in-neck with the high-performance computing as far as which is getting further ahead. So that is part of it.

Just as an example, the tools that we used 15 years ago which were using high-performance computing are today run on laptops. And we still use those tools, but now they have migrated down to where the user can use them on a laptop. So we develop the capacity and as the capacity comes online we need it, but there are cases where we need more.

Senator ENSIGN. Does Boeing, as a private company needing HPC, use or have access to supercomputers at an academic institution?

Mr. GARRETT. We do as needed. We have our own, for example, Cray X–1 in our Bellevue campus, which we do use, which is supplied by Cray, obviously. And it has been upgraded tremendously in the last 5 years. But we also have access to other companies' usage, as well as sometimes, based on government contracts, to government computers to be used on the defense side. We do not use those very much on the commercial side, but in support of defense contracts.

Senator ENSIGN. Dr. Burt, Please hold off on your answer until after Senator Cantwell is finished with her questions, because I have some even more detailed questions for you. Being a veterinarian, I am very interested in how high-performance computing can impact medical treatments as well. Now I will turn it over to Senator Cantwell.

Senator CANTWELL. Thank you, Mr. Chairman, and thank you for your indulgence.

I want to follow-up on where the Chairman went as far as the next phase of supercomputing and what we need to do as it relates to the House legislation and Senate legislation that we have been considering. Mr. Jehn, you talked about this stage of adaptive computing so that you can basically—just the diverse processing and scientific requirements. Mr. Waters kind of alluded to it in the sense of networks with a thousand times capabilities recommendation.

So what is it specifically that we need to do in that area? And I know that the Chairman just asked in the sense of what role we should play. But are you saying that this is an area that is the next phase and that we ought to be more specific in outlining this as far as our competitiveness?

Mr. JEHN. Well, a government-sponsored R&D program would be much more specific and systematic, than the government has funded to date. However, I again go back to the report of the High-End Computing Revitalization Task Force of 2 years ago. They laid out a roadmap, a technological roadmap, that would be a candidate for a systematic R&D program. The government over the last 3 or 4 years has made enormous strides in this direction.

As you may remember, I joined Cray a little less than 5 years ago and was relatively new to this industry and this set of issues, and particularly Federal policy in this area. Frankly, there has been a very significant change in Federal policy and activity over the last 3 or 4 years that Dr. Szykman referred to as well. I think building on that momentum is what is required right now. At the moment individual agencies like DARPA and the National Security Agency are supporting R&D in specific areas that they feel are most appropriate and most applicable to their requirements, but we need similar energy and direction elsewhere throughout the Federal Government. And we need better coordination among the various agencies to ensure that nothing slips through the cracks and that we are not duplicating things unnecessarily.

But there is plenty of room for everybody to contribute and I would advocate them doing so.

Senator CANTWELL. Well, when you talk about systematic R&D and Mr. Waters talks about the implications of TCPIP having been developed by DARPA, it sounds like we are talking about a similar focus to really bring together the great computing science, but with an easier use for the individual researchers in that application; is that correct?

Mr. JEHN. Well, certainly in the field of high-performance computing there is no mistaking that as capable and powerful as these systems are, they are very, very difficult to use. Imagine programming a system that contains 100,000 laptops, basically, 100,000 microprocessors. It is a daunting challenge, and the focus of the DARPA program has been and is to develop systems that are far easier to use, more accessible to a wider range of users, and I think this is a trend, a bit of momentum, again to use the same word, that we need to build on. I think Mr. Waters' example of the Internet is a great example or a great analogy here, where a very narrow, focused bit of technology was expanded to——

Senator CANTWELL. Explosive.

Mr. JEHN. Yes, it has been explosive, with wide availability and applicability to enormously different areas. That is what we need to promote in this area as well.

Senator CANTWELL. Mr. Garrett, I think that when you were using the aviation applications that you were talking about, I do not know if that is what Mr. Lombardo would talk about, the midrange applications, but it certainly has great applications for your competitiveness today.

Mr. GARRETT. Yes.

Senator CANTWELL. What are your competitors doing in this area?

Mr. GARRETT. Our competitors, Airbus specifically, is doing the same thing. They are investing, making those same investments.

Senator CANTWELL. Where do they get their research?

Mr. GARRETT. Well, they get that either through themselves or academia or through the government.

Senator ENSIGN. If you could elaborate on that point, because we talked about the U.S. *versus* other nations with Dr. Szykman on the first panel, and Airbus is obviously your competitor. On the applications side, do we have an advantage? Does the farther we stay ahead in the United States on high-performance computing give Boeing an advantage in its competition with Airbus?

Mr. GARRETT. Yes. In the long run, we have had the issues of our competitiveness and what has played out in the marketplace has been of benefit to us in the last year especially. At Farnborough that is playing out. But the bottom line is the performance. Our ability to meet our performance commitments is getting more and more attention because of the lack of that in many campaigns with our competition.

That does not happen by accident. It happens because of our ability to predict the airplane's capability from a weight, engine, and drag standpoint has been significantly improved and actually has been our competitive advantage for years. It is just coming out now in some of these key campaigns. As they get more and more aggressive, this is where we start to shine, in our ability to use these tools.

It is HPC is one element, but it is the step ahead that we have in our proprietary analysis tools and the development of those computational methods which is reducing the risk. It is basically risk reduction and our ability to hit our target where we think we are going to be, and our customers are acknowledging that and we are getting credit for that in the marketplace now, which is nice to see.

Senator CANTWELL. Thank you, Mr. Chairman. I will leave you to query Mr. Burt on a lot of areas. But I certainly agree with his philosophy that the intersection of supercomputing and biology is where the future is, and we would like to think a piece of that is in Seattle. So thank you.

Dr. BURT. Senator Cantwell, if I may just make one comment. When I mentioned data-intensive computing, I have been for several years on the scientific advisory board for Pacific Northwest National Laboratories in their computer science initiative. One of their major concerns is how do you do this data-intensive computing and how do you gather it, because they of course are concerned across many fields, not only biology but national security, homeland security, et cetera, et cetera. So it is a constant. So I serve on a panel that is devoted to looking at how they are approaching this problem.

Senator CANTWELL. Thank you.

Thank you, Mr. Chairman.

Senator ENSIGN. Thank you.

Let me start with Dr. Burt. You mention that because of the complexity of biology and the life sciences, a cell is infinitely more complex than an airplane. It just is.

Dr. BURT. Yes, he is only doing airplanes.

[Laughter.]

Senator ENSIGN. So it would seem to me that supercomputing—and I remember when I was in veterinary school in the 1980s, the talk about pharmacology and the idea of some day replacing animal research with computer modeling. With a lot of the research, without high-performance computers you are not going to even come close to doing some of the things that Mr. Garrett is talking about in predictability. I think you called it hypothesis-driven.

Dr. BURT. Hypothesis-driven. It means that you can use the computer to help you generate various models, which will then help direct where the experiments need to be run.

Senator ENSIGN. Right, and that would seem to me, similar to the airplane example, how it shortened the life-cycle down and the development cycle of new products, the same thing with new drugs, the same thing with new treatments, the same thing with a lot of different things in the field of medicine.

The question is on the Federal Government's role. We doubled the budget for NIH and it continues to go up from there. One of the goals that we have is to develop the budget for the physical sciences in the same way. Do we need to do something differently up here? Since we doubled the budget, we hate to get into telling the scientists what to do with their money. It has worked fairly well with Congress not politicizing that too much. But are there other areas that need focus or should we at least say we need, as a national priority to invest for the life sciences into high-performance computing? Do we need to start directing NIH and some of their funding to do that?

Dr. BURT. I believe so, Senator. The ABCC happens to be probably the only example of a truly integrated high-performance computing center of the NIH involved in biomedical research. The Center for Research Resources has made grants to other universities for computing and et cetera and I think they have been successful in that. But I think that we are just now beginning to realize that you really need, because you have to take a systems approach to biology because you must take into account not only the cell, not only the proteins, but the pathways, and especially in cancer it has been shown that the environment of the cell plays a large role.

So we really need to fund something like I proposed in order that we can get more people from different disciplines involved. As I said, one of the things that makes the—I did not say this, but one of the things that makes the ABCC successful is that the bioinformaticists in my group all have experience in the labs as well as with computer scientists. The people who do the modeling for me also are physicists and quantum chemists and mathematicians who have now received enough experience that we can translate. So that is a real key to applying HPC in biology, is the synergy, this collaboration, and we need cross-training so that people can speak the languages.

Senator ENSIGN. Well, I would like to work with you, if you could make yourself available to work with my staff, and try to come up with some ideas along these lines. I do not like to just have hearings. I like hearings to lead to actually policy. So I would like to work with you on some of the things that we are talking about today.

Sorry, doctor. You wanted to comment?

Dr. WLADAWSKY-BERGER. Yes, if I may add. Dr. Burt said something really, really important, which is more and more we need to take a system approach to these problems. If I could use that to link back to the innovation authorization bill, because in the past the bulk of the fundamental research was components in the laboratory, and then we threw it over the fence and then people built things with them. The problem now is that the things we are talking about building are systems of such extraordinary complexity that we need to actually do the research in how do you build those systems effectively. And since those systems have to be usable by human beings—otherwise why build them—how do you, for example, visualize the results so that the physicians or the veterinarians or whoever is using them can actually work with them?

So the change that has happened is it is not just in the laboratory; it is almost more, there is a lot of marketplace innovation, if I may use that term, in how do we build these systems and make them usable by human beings?

The Internet, by the way, is probably exhibit A of what we are talking about, which is we did not just do the TCPIP in the lab. NSFNet actually built the Internet, and then the World Wide Web came out of that. So that is something that is very different that links HPC to the innovation bills that you have been working so hard to authorize.

Senator ENSIGN. Mr. Waters, along those lines, because you folks are more involved obviously with networks than anybody else here—you talked about the Internet2. If we are linking Mr. Lombardo with other places around the country, what are the requirements? We are talking about more broadband in the country. Are we just talking fiber? Are we talking compression technology? What is necessary for that Internet2—I would imagine when you are talking teraflops—or what is the next one beyond teraflops?

Dr. WLADAWSKY-BERGER. Petaflops.

Senator ENSIGN. Petaflops.

Dr. WLADAWSKY-BERGER. Ten with 15 zeroes.

Senator ENSIGN. Right, I knew the $10^{15}$. I just wanted to confirm the name.

When we are talking about transferring that type of data, I would imagine that over a traditional phone line the transfer might be a little slow. I am just guessing. But what kind of network capabilities are required? And on Internet2, how long into the future before Internet2 is available?

Mr. WATERS. The networking requirements that I believe are coming out of the HPC environment are measured in the hundreds of gigabits per second. So in your home line you may have a meg and a half. A gigabit is $10^9$, so you can kind of get a feel for the orders of magnitude that we are talking about.

That is just the start. That is the initial Internet2 infrastructure at 100 gigabits per second. We believe that that will grow to four or five times over the next couple of years as research needs are——

Senator ENSIGN. What is the physical——

Mr. WATERS. The physical is the infrastructure that has been invested over the last decade. The fiber itself is frankly, the physical fiber, is probably already there to link-up the institutes where important. It is really the electronics that have to be put on the ends of the fiber, and that is where some investment needs to be made; and also the control systems, the things that allow a researcher in San Diego to have 100 gigabits at a particular time of day and then perhaps a researcher in Pittsburgh, being able to switch that bandwidth to that researcher in Pittsburgh, because you cannot necessarily provide bandwidth everywhere at the peak capacity at all times, and we want to be able to shift that bandwidth to individual users and applications over time.

Senator ENSIGN. Mr. Lombardo, could you just make a comment on what that would do to centers like your own, being able with your researchers to communicate with other researchers, not just working on your computer there but actually, just like you join PCs together, you can actually collaborate a lot more?

Mr. LOMBARDO. That is exactly right. In fact, for us, we jumped onto the Internet2 backbone about 6, 7 years ago, and the amount of data that we can currently download is incredible. Having access to such data enhances the research outcomes in modeling and simulation problems. We are currently in a planning stage to develop a center for the simulation and modeling of brain disorders, such as schizophrenia and Alzheimer's. But what is key to this research—and this is one of the rules I always like to mention about supercomputing: we need computers, storage and, high-performance networking capability. Without which we would not have the ability to move these enormous amounts of data without waiting—in some cases you would wait weeks and on Internet2 it may take you an hour.

Internet2 access is both critical and an essential component of HPC. And the fourth element of HPC, which I still think is the most important, is the access to very smart people.

Dr. BURT. Senator, I would like to just speak to that just very briefly. Since you are a veterinarian, there has been a big emphasis in Frederick on doing animal imaging. Now the goal is to be able to share those images with the people in Bethesda. So we have had to put in a big pipe so that people can stand in what we now have is a wall, so that you can look at this thing on the big screen and in three dimensions, but at the same time it can actually be seen by pathologists and other people like that. We have them at Frederick, too. But other people at the NCI in Bethesda, a distance of only 30, 35 miles, but we need to have it to where you can see it; as it moves, each person can see it move in real-time.

Senator ENSIGN. Well, thank you all. We have rules in the Senate about the length of Subcommittee hearings, and I think they did that just because Senators' attention spans are often not that long. But I thank all of you. It has been a fascinating hearing today, and once again compliments to my colleague for inspiring this hearing.

The hearing is adjourned.

[Whereupon, at 12:35 p.m., the hearing was adjourned.]

# APPENDIX

PREPARED STATEMENT OF TOM WEST, CEO, NATIONAL LAMBDARAIL

Dear Mr. Chairman,

On behalf of the member organizations of the National LambdaRail, we thank you for the opportunity to provide a prepared statement for the hearing record on high-performance computing.

We applaud you and your distinguished colleagues for their dedication to sustaining and stimulating investments in technology for this Nation's innovation and competitiveness in the 21st century. One of the most effective ways to advance our Nation's research capacity to lead the world in innovation is ubiquitous access to high-performance computing resources (HPC). High-capacity optical networks are the means by which the scientific community broadly harnesses HPC resources for innovation and competitiveness.

Today a new global network infrastructure owned and operated by the research and education (R&E) community has been deployed, and is being utilized. In the United States, National LambdaRail (NLR) owns a nationwide networking infrastructure that leverages regional and local efforts to provide a flexible infrastructure capable of supporting multiple, advanced research and education networks—a "network of networks." NLR is available to all researchers in academe, Federal agency laboratories and non-profit and for-profit research organizations. It serves researchers in all scientific disciplines, providing the critical advanced network infrastructure to access the Nation's high-performance computing facilities for advancing big science initiatives. The NLR infrastructure is the result of over 3 years of work and nearly $120 million in funding by its members.

The mission of the NLR is to build an advanced, nationwide network infrastructure to support many types and levels of networks for research, clinical, and educational fields. This infrastructure consists of 15,000 miles of fiber and optical networking equipment, all of it owned by NLR. NLR's potential capacity is 40 10-gigabit (10 billion bits per second) nationwide networks. By comparison, a 10-gigabit network is roughly 10,000 times faster than today's commodity networks such as cable modem and DSL. The infrastructure supports both experimental and production networks, fosters networking research, promotes next-generation applications, and facilitates interconnectivity among regional and international high-performance research and education networks. Furthermore, NLR is scalable to accommodate the ever-increasing computing demands of the future.

The hallmark of 21st century big science applications is multi-disciplinary, multi-investigator research collaborations across time and space. This distributed approach can lead to more rapid and systematic solutions to society's most intractable challenges. High-capacity optical networks are critical to leveraging innovation across these worldwide assets.

Moreover, there is a growing urgency to develop new network technologies that scale to the growing needs of the worldwide R&E community and, later, to commodity Internet users. We are encouraged that the Administration's FY 2008 research and development (R&D) budget guidelines prioritize R&D in advanced networking technologies. NLR's high-performance network infrastructure enables the next generation of technologies, protocols, and services. This enabling infrastructure is critical to progress in essentially every interagency R&D priority for FY 2008—from homeland security to energy security, and from nanotechnology to complex biological systems and the environment.

The focus on network researchers is a distinguishing characteristic of NLR. Fifty percent of NLR capacity is devoted to support network research projects at the forefront of developing and testing revolutionary, not just evolutionary, networking technologies and capabilities not possible in the laboratory or any other national-scale network.

Undertaking this R&D requires an experimental testbed where network researchers can experiment with new approaches to all levels of networking technology. The results of this research will enable networks capable of supporting scientific projects in fields such as high-energy nuclear physics and radio astronomy, which require real-time collaboration among scientists and manipulation of enormous data sets. Already, individual projects in these fields can usefully consume a majority of the largest network links available. Together, even a few of them could potentially overwhelm existing advanced research and education networks. And, these kind of bandwidth-hungry applications are spreading. Applications in almost every discipline are now emerging with the same need for big, broadband networks.

While regional optical network (RON) infrastructure development emerged a few years ago, the formation of NLR spurred numerous new regional efforts. Now 14 NLR members operate 21 regional optical infrastructures to serve as the pillars of connectivity to NLR. Across the United States, an additional 15,000 miles of fiber-optic cable controlled by RONs significantly enhances access to rich high-performance computing capabilities; unique, expensive research resources; and linkage of enormous amounts of data through federated databases.

Importantly, NLR's diverse membership includes RONs as well as many of the Nation's premier research and education organizations, private sector technology corporations, and Federal agencies. Today, NLR's members include—

- Case Western Reserve University
- Cisco Systems
- Committee on Institutional Cooperation
- Cornell University/Northeast LambdaRail
- Corporation for Education Network Initiatives in California (including the University and Community College System of Nevada)
- Duke University (representing a coalition of North Carolina universities)
- Florida LambdaRail
- Front Range GigaPop/University Corporation for Atmospheric Research
- Internet2
- Lonestar Education and Research Network
- Louisiana Board of Regents
- Mid-Atlantic Terascale Partnership/the Virginia Tech Foundation
- National Aeronautics and Space Administration
- Oak Ridge National Laboratory
- Oklahoma State Board of Regents
- Pacific Northwest Gigapop
- Pittsburgh Supercomputer Center/University of Pittsburgh
- Southeastern Universities Research Association
- Southern Light Rail
- University of New Mexico (on behalf of the State of New Mexico)

An excellent example of the computational environments enabled by NLR's infrastructure is the Extensible Terascale Facility (ETF) supported by the National Science Foundation. The ETF is a multi-million dollar, multi-year effort that has built and deployed the TeraGrid, a world-class networking, computing and storage infrastructure designed to engage the science and engineering community to catalyze new discoveries. The Pittsburgh Supercomputing Center, one of the original TeraGrid participants, was the first organization to use NLR to connect its facilities to the nationwide TeraGrid facility. More recently, the Texas Advanced Computing Center acquired a 10 Gigabit wave from NLR to connect Austin to Chicago. Oak Ridge National Laboratory is also using NLR for back-up waves between Atlanta and Chicago as part of ETF.

Today more than ever, growth in our economy is increasingly linked to the investments made in fundamental research to advance computing and communications technologies. We urge your continued support for strengthening investments in America's future with a strong national research infrastructure for advancing discovery, innovation, and education.

Thank you.

RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. JOHN ENSIGN TO DR. SIMON SZYKMAN

*Question.* How do other countries—particularly those in Asia—rate relative to the United States in the context of high-performance computing research?

Answer. According to the most recent list of Top500 supercomputer sites* released in June, over 90 percent of the world's top 100 fastest machines are located in the United States, Europe and Japan. Representation by other nations among the elite computing facilities is minimal. The locations of the top 100 machines break out geographically as follows:

United States: 57 machines
Europe: 18 machines

Japan: 16 machines

Canada: 3 machines

Korea: 2 machines

China: 2 machines

Australia: 1 machine

Russia: 1 machine

As the most recent snapshot highlights, more than half of the world's top 100 machines are in the United States, about three times as many machines as Europe or Japan. Eighty-eight of the world's top 100 machines, including more than two-thirds of those machines that reside outside of the U.S., were built by U.S. companies. Within Asia, the critical mass of supercomputing capability is clearly in Japan. In relative terms, representation by other parts of Asia is minimal; Korea and China each have two machines in the top 100, and India has none.

The Top500 Supercomputers list provides a picture of investments in high-performance computing infrastructure. We can also consider related R&D investments as an indicator of how leadership may change with time. With the rapid growth of the economies of certain nations, it is important to consider not only a current snapshot, but to look forward as well.

The European Union (EU) has taken a fundamentally different approach to high-performance computing than the U.S. has. The bulk of EU R&D investments in recent years has been through the EU's Framework Programme (FP), Both of the most recent Programmes (FP5 and FP6) have emphasized grid computing environments for high-performance computing. In this context, it is very important to note that grid infrastructure is not a substitute for tightly-coupled supercomputing architectures or centralized computing facilities.

The U.S. Government's High End Computing Revitalization Task Force (HECRTF) recognized this limitation and deemed grid computing to be out of scope during the planning and roadmapping that resulted in the *Federal Plan for High End Computing*. The EU characterizes their grid computing investments as "high-performance computing," but has no substantial R&D investments in high-end computing as characterized by the *Federal Plan for High End Computing*. Although the EU's Seventh Framework Programme is still in planning, current indications are that investments will continue expanding grid-based high-performance computing infrastructure rather than being aimed at R&D for new high-end computing technologies.

In the HECRTF context of high-end computing R&D (*i.e.,* viewing grid computing technologies as out of scope), Japan is the only competitor to the U.S. in the global playing field. I mentioned in my testimony that 4 years ago the Japanese Earth Simulator System became the world's fastest supercomputer. It is now ranked tenth. Japan's *3rd Science and Technology Basic Plan (FY 2006–FY 2010),* a Japanese national policy document, identified supercomputing as a key technology of national importance to the national infrastructure. Japan has announced plans for the development of a successor to the Earth Simulator System called the Next Generation Super Computer (NGSC). While the overall investment required for this effort is estimated at about $1 billion from 2006–2012, to date the level of funding that has been approved is only about ¥30 million in (Japanese) Fiscal Year 2006. The NGSC is expected to be a ten petaflops (10 quadrillion floating point operations per second) computing facility. The U.S. expects to have several petascale computing facilities by early next decade. As I pointed out in my testimony, global leadership should not be defined solely by the speed of the world's fastest machine. The NGSC will likely be a highly capable machine, but if the U.S. continues along its current trajec-

* See *http://www.top500.org/*.

tory of high-end computing R&D and infrastructure investments, the NGSC will not pose a threat to U.S. leadership in the high-end computing arena.

At present, R&D capabilities and investments are significantly weaker elsewhere in Asia. Although China is emerging as an important global competitor as a producer of information technology (IT) equipment and products, China does not yet have a strong capability for innovation in the area of IT. Their core technologies and key equipment still rely on imports of technology and intellectual property. China is advancing their domestic IT innovation capabilities by focusing on integrated circuit and CPU technology and system and applications software, and not on high-performance computing technologies.

There is some Chinese government support for high-performance computing, with around a dozen high-performance computing centers in China. However, these computing centers are not at the highest levels, as demonstrated by the fact that China has only two machines on the Top100 list (#35 and #53). Two Chinese IT companies (Lenovo and Dawning Group) have announced long-term objectives of building petaflop machines. but it is not clear how much R&D will go into supporting these efforts (as opposed to simply building machines through large investments technology that is within the state-of-the-art). In the long-term, China's domestic R&D capabilities are expected to improve and should be monitored. But at present and in the short- to medium-term future, Chinese R&D in high-performance computing is not expected to be a threat to U.S. leadership.

India has a strong IT sector, but one that is primarily limited to software development and IT services. India does not have a strong hardware sector, and is still working to establish a basic electronics industry. India does not yet have a world-class computer science R&D community. A national grid computing effort is only just beginning, and is still at a proof-of-concept stage. Physical infrastructure in India also lags behind that of China. India has pledged to increase R&D investments in the future, but given their current standing in the computing arena, they are not viewed as a significant threat at this time, and will need to make substantial strides in their ability to innovate from multiple perspectives (government commitment, workforce, infrastructure) before they are able to compete globally in the high-performance computing arena.

I hope that I have addressed your inquiry to your satisfaction. I would be pleased to respond to any additional questions you may have.

---

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. JOHN ENSIGN TO DR. IRVING WLADAWSKY-BERGER

*Question 1.* Please explain what features of present-day supercomputing do you feel will have the greatest impact on computers being used by the American public in the future?

Answer. Hybrid programming models (such as the Los Alamos/DOE "Roadrunner" which will use these software models to coordinate different chip architectures simultaneously) for supercomputers will become applicable to smaller systems—perhaps even PCs and game systems. This will greatly improve personal systems' performance and resulting capabilities—as Moore's law becomes less and less applicable over time. Chips can't just keep getting smaller to get faster (because each reduction in size brings a commensurate increase in heat) so new methods to combine multiple chips and varying chip architectures for even the smallest systems will be required. This innovation is happening today at the forefront of supercomputing.

- Fundamental technology enhancements are in the areas of multi-core architectures with latency hiding techniques.
- Very fast interconnects that have the ability to make remote memory access seem like local memory access

*Question 2.* In your testimony you point out that "while progress in supercomputing hardware has been astounding, both applications software and systems software remain a real challenge." Why do you believe this is so?

Answer. Boosting system performance over the years—especially relative to price—has been comparatively easy, given Moore's law which posited a doubling of performance every 18 months. But the scope of these systems requires greater levels of coordination—to, among other things, orchestrate the tens or hundreds of thousands of processors that must work together in the world's most powerful computers. Further, the evolution of hybrid supercomputer architectures (The Los Alamos "Roadrunner" supercomputer, for example, will use both Cell and AMD Opteron processors) requires new software to take complex problems and divide their mathematical components for routing to different chip architectures simultaneously. On the application side, software development is required to extend the capabilities of high-performance computing into commercial markets, as well as deeper into the academic and scientific arenas. The objectives of supercomputing hardware innovation are relatively uniform over time: more performance in less space requiring fewer/less resources. In the case of software, each new application requires new thinking, methods and skills.

*Question 3.* Please explain how PC-based technologies have impacted the way that IBM develops its high-performance computers.

Answer. Some of the world's most powerful computers are built using processors that are common in today's personal computers and video game systems. The emergence of parallel computing has allowed smaller, cooler, less powerful chips—such as those in PCs and game systems—to be grouped in large pools to split complex problems into smaller pieces.

---

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. JOHN ENSIGN TO CHRISTOPHER JEHN

*Question 1.* In your testimony, you state that, "the supercomputing market is not large enough to justify significant investment in unique processor designs" and that as a result, "to advance supercomputing, industry has relied on leveraging innovation from the personal computer and server markets." Why is this development a bad thing?

Answer. Industry's reliance on leveraging innovation from the personal computer and server markets has led to the production primarily of large collections of commodity processors linked together with commodity interconnects. This architecture dominates supercomputing. Unfortunately, these processor clusters are notoriously difficult to program. As a result, fewer programmers are capable of programming them, and those that do need much more time to program. The total cost of ownership goes up.

Even more problematic is the divergence problem. See the following chart[1] from a 2004 Federal interagency report on supercomputing.

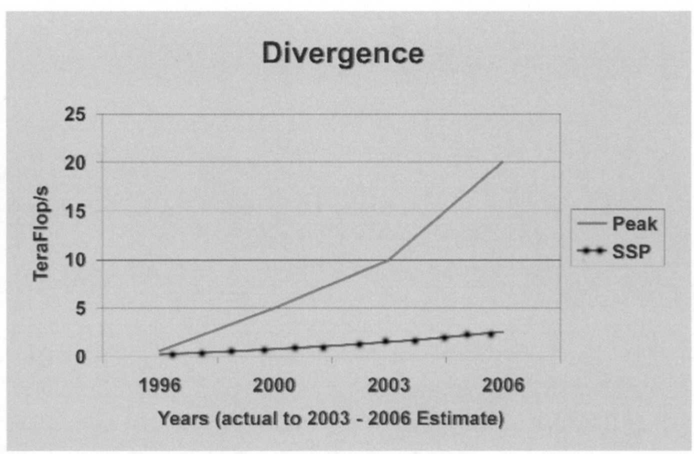


**Figure 1: Divergence Problem for HEC Centers (SSP = Sustained System Performance)**

In looking at the chart, you can see that the theoretical peak performance, the highest performance achievable by a system performing at 100 percent efficiency, has diverged from the sustained system performance (SSP), which is what is actually usable by the applications. As you can see, the gap between the peak performance and the SSP has grown dramatically, meaning most of the system capability is wasted. As a result, supercomputers are far less efficient today than they were in 1996. And, the problem is getting worse. This development is at the core of the current crisis in supercomputing.

*Question 2.* In your testimony, you state that there is currently a crisis in supercomputing. Several other witnesses seem to suggest that the situation is not quite that dire. Please elaborate on why you think that the situation is so problematic and what you believe needs to be done?

Answer. Progress in advancing high-end computing technology has slowed considerably since 1996, when Federal Government R&D funding for supercomputing began to decrease. Supercomputers are now 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 large numbers of microprocessors with relatively weak interconnects. Programming and administrative costs often exceed the costs of the actual machine. Future trends in supercomputing will only exacerbate this problem.

The lack of advancement in supercomputing technology has wide-ranging ramifications. It 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 our economic competitiveness. The most recent report on HEC, which comes from the U.S. Defense Science Board and the U.K. Defence Scientific Advisory Council, had this to say, "There is great concern that lack of investment is eroding U.S. leadership in [supercomputing]; as well as, negatively impacting our ability to meet defense mission requirements . . . technology that is developed in the context of high-performance computing 'flows down' to help advance mass market computers. Thus, if the United States does not aggressively pursue very high performance computer technology, then innovation in mass market computers will slow."[2]

Parallel programming is a prime example. Much has been discussed about how today's software for personal computers and workstations are not ready to run efficiently on multicore processors, even though computer vendors will soon stop selling computers with single cores. Had we continued Federal investments in supercomputing over the last 10 years, it is likely large numbers of programmers would be able to program in parallel efficiently.

The lack of advancement in supercomputing comes at a bad time. While HEC has been vitally important to the Federal Government for many decades, the need for supercomputing is greater today than ever. Federal agencies tell us this everyday. Agencies need supercomputing to help maintain military superiority, enable scientific research, advance technological development, and enhance industrial competitiveness. Just as they have in the past, Federal agencies rely on supercomputing to pave the way for real progress well into the future.

As such, countries with the best high-end computing capabilities will enjoy a significant advantage in scientific competitiveness. Meanwhile, a number of other countries are aggressively pursuing HEC programs. Both Japan and China have announced programs to build multi-petaflops systems in the near future. While little is known publicly of the Chinese plans, the Japanese government's proposed "Keisoku" project would spend US$1billion to build a 10 petaflops supercomputer by 2011 and another US$300 million for a new national HPC center.[3] This aggressive government initiative is reminiscent of the Japanese Earth Simulator project whose performance took this country by surprise just 4 years ago.

Our recommendation is that Congress fully fund the Administration's proposed government investments in supercomputing. This includes funding supercomputing programs in the Department of Energy (DOE), the National Science Foundation (NSF), and within the Department of Defense (DOD). Proposals from DOE's Office of Science and the NSF to fund the deployment of petascale computers are steps in the right direction. DOD, most notably through DARPA's High Productivity Computing Systems program and the National Security Agency (NSA), is supporting research and development 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.

While these are important steps, we recommend the Administration build on these recent initiatives and develop and fund a coordinated research and development program for supercomputing, as many recent government sponsored reports have strongly recommended.[4] The U.S./U.K. Task Force on Defense Critical Technologies' recommendation to make HPCS a recurring program with multiple overlapping waves, each lasting seven to 8 years, is a sensible example of what such a program would include.

We also recommend the Administration 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.

Numbers speak for themselves. While the NSF plans to spend more than $200 million on HEC out of its proposed $4.7 billion annual research budget, the National Aeronautics and Space Administration (NASA) plans to spend only about $30 million out of a proposed $10.5 billion science, aeronautics and exploration budget. The National Institutes of Health (NIH) funding rate is even worse. NIH plans to spend significantly less then $30 million on HEC out of its $25.1 billion annual research and development budget. The science and engineering requirements for NIH and NASA are not any less, dollar for dollar, than NSF's. Excluding NSA and DARPA, DOD plans to spend only $170 million on HEC out of its proposed $73.2 billion research, development, test and evaluation budget. That amounts to spending less than three tenths of 1 percent on high-performance computing, even though DOD's science and engineering requirements are enormous.

Other agencies requiring significant science and engineering, such as the Environmental Protection Agency, the Department of Transportation and the Department of Agriculture, use practically no HEC at all. These agencies would benefit from the increased use of HEC. So, my final recommendation would be for the Federal Government to identify gaps in computational science usage across all of the agencies and develop programs to close these gaps where appropriate.

Had the government continued investing as it had prior to the late 1990s, we would more than likely have such promising technologies such as superconducting multiprocessors, processor-in-memory (PIM), multithreading, streams, and holographic storage today. We would also have seen similar advances in software.

**ENDNOTES**

[1] 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*.

[2] Joint U.S. Defense Science Board/U.K. Defence Scientific Advisory Council Task Force report—"*Defense Critical Technologies*." March 2006.

[3] IDC report—"*The Keisoku Project: Reestablishing Japan's Leadership in Supercomputing?*" June 2006.

[4] 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.

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*.

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/U.K. Defence Scientific Advisory Council Task Force report—"*Defense Critical Technologies*." March 2006.

---

RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. JOHN ENSIGN TO JACK WATERS

*Question.* Please explain how the manner in which research is conducted today reinforces the need for high-performance networks.

Answer. Scientific research occurs across multiple disciplines, *i.e.,* high-energy physics, biology, astronomy, etc. to name but a few, require significant super computing resources to process the extraordinary amounts of data. The research instruments, whether they are supercolliders, radio telescopes, electron-microscopes or laboratories, represent billions of dollars of public capital investment and continued operational investments by the public-sector.

In order to leverage both the significant investment in research instruments and laboratories as well as the computational resources, it is essential that high-speed, high-capacity, high-performance networks be available to connect the disparate entities.

To a very large extent, the research funded by the U.S. Government through various agencies, whether DOE, DOD, NIH, NSF, NASA, *et al.,* is accomplished with multi-collaborating entities. Some, due to their highly sensitive nature and uniqueness are done with fewer collaborators, but even these tend to have multiple entities in disparate locations engaged in the research project. Simply put, supercomputers need to transmit enormous amounts of data amongst one another, hence the need for super networks.

When you move into a research realm which is not as sensitive or unique, be it biomedical program with national defense implications like studying and preparing for a potential bird pandemic, a high-energy physics program or astronomical research program, the number of collaboration partners can grow substantially to include hundreds of individual disparate entities.

Networks are a key to facilitating research in order to truly leverage the investment made into widely dispersed physically entities. High-performance networks allow the sharing of computation resources minimizing the need for continued investment in supercomputers and accelerating the pace of scientific discovery.

---

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. JOHN ENSIGN TO JOSEPH LOMBARDO

*Question 1.* How can a university like UNLV maximize the benefits from high-performance computing for its students and academic programs?

Answer. A comprehensive university such as UNLV can maximize benefits from HPC for its students and faculty through interactive applications of real-time problems across the academic disciplines. Complex problems in research are struggling with the ever increasing amount of data to be processed (such as molecular biology in cancer research, difficult problems in dealing with nuclear waste, and emerging opportunities in the nano sciences). These problems are best addressed via management and analysis of data through applications of high-end computing. Developing an understanding of these research tools early in undergraduate education and hands-on applications by graduate students and faculty in real problems should enhance the capabilities of human resources and lead to advances in technological resources in the field of HPC across the board.

University research programs could advance the applications of HPC through collaborative research programs with Federal agencies whose missions would benefit by use of HPC. Examples are the National Cancer Institute, the U.S. Department of Transportation (materials sciences/new airport construction/development of sophisticated security systems) and the Department of Energy. The Committee might consider introducing resolutions or legislation encouraging such collaborations with various incentives to develop such consortia.

*Question 2.* How does UNLV plan to update its high-performance computing capabilities as the next generation of hardware and software become available? Especially given the rapid pace of technological innovation in the high-performance computing area, how does a leading university like UNLV stay ahead of the curve?

Answer. UNLV has embarked upon a program of developing symbiotic relationships with public-sector and private-sector groups that have need of HPC technology and the human resources that support the hardware and software infrastructure. These resources are sustained and enhanced by grants, contracts and other sources of funding that derive to the National Supercomputing Center for Energy and the Environment. These relationships and attendant contracts allow UNLV to stay current with the research needs and the requisite technology. Should opportunities emerge that would require a quantum leap in technology, UNLV might draw upon technological resources from other research universities (*e.g.,* Ohio State University, University of California at San Diego, University of Alaska, *et al.*) as available and appropriate. Absent ability to access these extant resources, UNLV would seek Federal/state/corporate funding for the large capital outlays to enhance the technological and human resources.

*Question 3.* In your testimony, you mention that, "Support for Federal funding of high-performance computing has ebbed and flowed as a result of perceived foreign competition." Do you think that increased Federal support for high-performance computing research and development over the past few years is well-suited to meet today's foreign competition? Do you believe that data indicating that the United States possesses and produces a majority of the world's fastest high-performance computers is a sign that we are out-computing and out-competing our foreign competitors?

Answer. The U.S. position in high-end computing is under constant global challenge as technological advances go forward. HPC is a vital tool to preserve U.S. preeminence in science, technology, math and engineering, and as such should be considered as an integral part of comprehensive national security. With the cutting-edge of technology always moving, it is necessary to the national interest that the Federal Government work collaboratively with university researchers, the corporate community and other related R&D interests to develop new approaches to the hardware and software technology infrastructure and to increasingly encourage applications of these research tools to areas that have not yet brought HPC to bear on their Grand Challenge problems. As Dr. Stan Burt, Director of the Advanced Biomedical Computing Center of NCI, pointed out in his testimony before the Committee, the biological sciences have lagged far behind the physical sciences in applying these 21st century tools to their research protocols. With increasing concern over biological warfare, communicable diseases, food safety and the traditional challenges of the fight against cancer, the Committee should make every effort to encourage the use of HPC in the field of biology and consortia fields of chemistry/physics/cellular studies, *et al*. Only with applications and access to HPC by all elements of the research community can the U.S. maintain its global leadership in HPC.

---

RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. JOHN ENSIGN TO MICHAEL GARRETT

*Question.* From the perspective of a consumer of high-performance computing tools, how do you think the market for high-performance computers will evolve over the next 5 to 10 years?

Answer. You have asked how we see the high-performance computing market evolving over the next 5–10 years. As far as the total market for high-performance computing we believe that the market will increase. We believe this for several reasons:

Companies are moving to replace physical simulation with computational simulation (model the issue in a computer rather than build a piece of hardware and test it). The numerical simulation is both faster and cheaper—the issue is the fidelity of the answer. New applications for applying high-performance computing are emerging rapidly—such as genetic engineering, material modeling/creation (at the molecular level), weather prediction, entertainment and astrophysics. And finally, as the cost of high-performance computing comes down it will be more affordable to a broader range of customers.

As for Boeing in particular our usage of high-performance computing will increase. The drivers for this are: (1) the need to get to the marketplace more rapidly. To do this we must shorten the design process, and the only way to do that is with more numerical simulation and less physical simulation. (2) We have a “closet” full of problems that are too big for current computers (at least affordable ones). As computer capability increases we will tackle these more complex problems.

Another view of how the market will evolve is from a hardware perspective. We believe that the emphasis will be on large, cluster parallel-processing machines. Why? Because they use inexpensive commodity chips. Increased throughput will be achieved through faster processors and by applying more processors to the problem. Utilizing more processors presents software challenges as well. How will the application software be written so that it can run across 10, or even 100 thousand processors? How will the operating system efficiently manage the data communications among these processors so as to not slow the process? And how will the operating system gracefully handle failures such that the entire process doesn’t fail because a single element of the job does?

The investment that the country makes in high-performance computing in support of nuclear stockpile verification, reconnaissance and intelligence gathering and processing, and for national defense in general will pay off in years to come as that hardware and software becomes available in the commercial marketplace. It is a model that has been in use for more than 30 years and as yet we do not see a reason to change it.

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. JOHN ENSIGN TO STANLEY K. BURT, PH.D.

*Question 1.* In an era when scientific research is increasingly becoming inter-disciplinary, please discuss how the Advanced Biomedical Computing Center has adapted computational methods developed in other areas of science for use in biology research.

Answer. The Advanced Biomedical Computing Center (ABCC) has, from practically its inception, had the good fortune to be composed of a staff that includes training in systems administration, network administration, bioinformatics, and physical sciences. As biology has become more dependent on computers and computational sciences, the ABCC has increased and diversified its staff in different scientific fields in order to utilize different scientific methodologies to solve biological problems. The diversity of the staff has allowed the creation of specialized tools, new algorithms, custom interfaces, and mathematical techniques. The ABCC is what I consider to be the model for a modern computing center for biology with the integration of different scientific disciplines combined with a strong computational infrastructure. The ABCC staff is knowledgeable about computer hardware and many simulation software packages and is able to adapt many techniques and computer architectures to specific problems. The ABCC also benefits from being in a diversified scientific environment that presents many kinds of biological problems demanding different approaches and solutions. The ABCC has benefited from the willingness of software and hardware vendors, other Federal agencies, and universities to work in a collaborative manner to find solutions for cancer.

*Question 2.* What do you see as the greatest potential future benefits from using high-performance computing to assist in medical treatment of diseases like cancer or AIDS?

Answer. I believe that the main role of HPC will be to tackle the inherent complexity of cellular processes and speed the time to solution for complex diseases and potential biological terrorism. As the increase in biological data accelerates at an incredible fast pace due to high-throughput screening, more genomes being mapped, and higher array chip densities, HPC will be needed to analyze the data, integrate the data, and model the complex processes involved in cancer and AIDS. No single experiment will be able to uncover these complexities. As biology moves toward a systems biology approach, which is necessary to use in order to understand the inter-relationships of complex cellular components in diseases, HPC will be necessary to address multiple levels of complexity at once, discover new data that will lead to more focused data-driven testable hypothesis, allow a better coupling of experiment and theory, and lead to more specific treatments. HPC is the best solution to provide a qualitative leap in understanding and solutions.

○