S. HRG. 108–674

HIGH-PERFORMANCE COMPUTING

HEARING

BEFORE THE

SUBCOMMITTEE ON ENERGY

OF THE

COMMITTEE ON
ENERGY AND NATURAL RESOURCES
UNITED STATES SENATE

ONE HUNDRED EIGHTH CONGRESS

SECOND SESSION

ON

S. 2176

THE HIGH-END COMPUTING REVITALIZATION ACT OF 2004, WHICH WOULD AUTHORIZE THE SECRETARY TO CARRY OUT A PROGRAM OF R&D TO ADVANCE HIGH-END COMPUTING THROUGH THE OFFICE OF SCIENCE

AND

TO RECEIVE TESTIMONY REGARDING THE DEPARTMENT OF ENERGY'S HIGH-PERFORMANCE COMPUTING R&D ACTIVITIES IN BOTH THE NATIONAL NUCLEAR SECURITY ADMINISTRATION AND THE OFFICE OF SCIENCE

JUNE 22, 2004

Printed for the use of the
Committee on Energy and Natural Resources

U.S. GOVERNMENT PRINTING OFFICE

96–630 PDF WASHINGTON : 2004

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) 512–2250 Mail: Stop SSOP, Washington, DC 20402–0001

COMMITTEE ON ENERGY AND NATURAL RESOURCES

PETE V. DOMENICI, New Mexico, *Chairman*

DON NICKLES, Oklahoma
LARRY E. CRAIG, Idaho
BEN NIGHTHORSE CAMPBELL, Colorado
CRAIG THOMAS, Wyoming
LAMAR ALEXANDER, Tennessee
LISA MURKOWSKI, Alaska
JAMES M. TALENT, Missouri
CONRAD BURNS, Montana
GORDON SMITH, Oregon
JIM BUNNING, Kentucky
JON KYL, Arizona

JEFF BINGAMAN, New Mexico
DANIEL K. AKAKA, Hawaii
BYRON L. DORGAN, North Dakota
BOB GRAHAM, Florida
RON WYDEN, Oregon
TIM JOHNSON, South Dakota
MARY L. LANDRIEU, Louisiana
EVAN BAYH, Indiana
DIANNE FEINSTEIN, California
CHARLES E. SCHUMER, New York
MARIA CANTWELL, Washington

ALEX FLINT, *Staff Director*
JUDITH K. PENSABENE, *Chief Counsel*
ROBERT M. SIMON, *Democratic Staff Director*
SAM E. FOWLER, *Democratic Chief Counsel*
PETE LYONS, *Professional Staff Member*

SUBCOMMITTEE ON ENERGY

LAMAR ALEXANDER, Tennessee, *Chairman*
DON NICKLES, Oklahoma, *Vice Chairman*

JAMES M. TALENT, Missouri
JIM BUNNING, Kentucky
CRAIG THOMAS, Wyoming
LISA MURKOWSKI, Alaska
LARRY E. CRAIG, Idaho
CONRAD BURNS, Montana

BOB GRAHAM, Florida
DANIEL K. AKAKA, Hawaii
TIM JOHNSON, South Dakota
MARY L. LANDRIEU, Louisiana
EVAN BAYH, Indiana
CHARLES E. SCHUMER, New York
MARIA CANTWELL, Washington

PETE V. DOMENICI and JEFF BINGAMAN are Ex Officio Members of the Subcommittee

JONATHAN EPSTEIN, *Legislative Fellow*
ADAM ROSENBERG, AAAS FELLOW

CONTENTS

STATEMENTS

HIGH-PERFORMANCE COMPUTING

TUESDAY, JUNE 22, 2004

U.S. SENATE,
SUBCOMMITTEE ON ENERGY,
COMMITTEE ON ENERGY AND NATURAL RESOURCES,
Washington, DC.

The subcommittee met, pursuant to notice, at 2:45 p.m., in room SD–366, Dirksen Senate Office Building, Hon. Lamar Alexander presiding.

OPENING STATEMENT OF HON. LAMAR ALEXANDER, U.S. SENATOR FROM TENNESSEE

Senator ALEXANDER. Good afternoon. The hearing of the Energy Subcommittee will come to order.

Senator Bingaman is here, our ranking member of our full committee, former chairman of the Energy and Natural Resources Committee, and we look forward to today's discussion about high-performance computing.

Excuse us for being a little late. We both had to attend our class photo. They have that once a year in the U.S. Senate, so there we all were.

The purpose of this hearing is to examine how the United States can recapture worldwide leadership in high-performance computing. To that end, we are here today to consider S. 2176, the High-End Computing Revitalization Act of 2004, which Senator Bingaman and I co-sponsor.

Until March 2002, 2 years ago, the United States was the undisputed leader in high-speed computing. That advantage has played a significant role in our ability to compete in the global marketplace and our standard of living. Sometimes we overlook the fact in the United States we have 5 to 6 percent of the people in the world and better than a third of the gross national product. There is a reason for that, and one of the reasons is, according to the National Academy of Science, half our job growth since World War II can be attributed to our investments in science and technology.

In 2002, however, Japan introduced its Earth Simulator, which is currently two and a half times more powerful than any other high-performance computer in the world. So Japan is the king of high-performance computing today. When Japan first introduced the Earth Simulator, it was nearly five times more powerful than any other high-performance computer in the world.

Senator Bingaman and I both recently visited Japan. We have both been there in the last 6 or 8 months, and we have both been briefed by the Japanese on their investment in the Earth Simulator.

Japan's development of the Earth Simulator meant that the United States no longer was the clear leader in high-performance computing, and for the first time, American researchers were looking overseas to obtain access to the latest computing tools.

Recapturing the lead in high-speed computing is one of the top priorities of the Secretary of Energy's 20-year facility plan. This bill, the one we are talking about today, and the companion bill that was reported out of the House Committee on Science last week will help the United States do just that.

High-performance computing is important to this country for several reasons. First, it will allow us to address a variety of scientific questions. For example, there is a lot of debate around here about global warming and climate change. We make a lot of decisions about clean air regulations, decisions that cost us money, that conceivably could limit our economy, that affect our health. High-end advanced computing will help us simulate the earth's climate and have better science upon which to base these very important policy decisions.

Second, high-performance computing is required to examine whether fusion power might become a reality. Fusion could provide low-cost energy for people around the world. We all are dramatically reminded today in the United States of how important that could be. Also, nanoscience has the possibility of revolutionizing chemistry and materials sciences. The full benefit of nanoscience may not be reached without detailed simulation of quantum interactions.

And third, there is a large concern and much debate in the U.S. Senate about our keeping jobs from moving overseas. Advanced computing would enable us to lower our manufacturing costs and improve our technologies. That means better jobs here in the United States. If you go to Europe, you do not see headlines about jobs outsourcing. You see headlines about brains outsourcing, brains being attracted to the United States from Europe by our research universities and our great laboratories. Investing and recapturing the lead in high-performance computing would attract more of the most talented scientists and students to the United States, which will help fuel our economy.

Last month, the Department of Energy took an important step toward putting America back in the forefront of high-speed computing. DOE announced that the Oak Ridge National Laboratory in Tennessee was selected as the winner of its competition to develop a leadership class computational facility. ORNL will lead an effort that includes many of the brightest minds in our country to try to reassert our leadership in high-speed computing.

Today we will hear firsthand how reestablishing our leadership will enable us not only to address grand scientific challenges, but to advance our manufacturing industry to enhance our U.S. competitiveness in the world marketplace.

We will also hear about the need for a commitment by the Federal Government to develop high-performance computing systems and the clear signal that this commitment sends to our computer manufacturers and our universities.

We have a distinguished panel of witnesses. I will introduce them in just a few minutes, but first I wanted to invite Senator Bingaman, the ranking member of the Energy Committee and someone who helped encourage my interest in this subject, if he has an opening statement.

STATEMENT OF HON. JEFF BINGAMAN, U.S. SENATOR FROM NEW MEXICO

Senator BINGAMAN. Thank you very much, Mr. Chairman. Thank you for your leadership on this issue and for holding this hearing.

I do think this is a very important subject. We had a chance to visit on this as we were coming over here from the Capitol just a few minutes ago. The point I made is that this is one of the long poles in the tent, as the saying goes, as far as the ability of the United States to remain a world leader in science and technology. I believe very strongly that leadership in high-end computing is an essential part of leadership overall in science and technology, and S. 2176 is the legislation that we have introduced to try to help us in this regard. It is based very much on the Office of Science's plan in its well-conceived "Facilities for the Future" report that was issued last November.

A lot of discussion around the Senate, of course, around the Congress generally is that much of this investment might better be left for the private sector. This is an area where no single company can plan on capturing the full value or a substantial portion of the value of the investment that is required here. This has to be an area where the Government steps in and provides assistance. We have done that in the past. We have been the leader in this area. Our leadership is not there today, and we need to reinstate that. So I feel very strongly that we should move ahead.

Again, I thank you, Mr. Chairman, for your interest and leadership, and I hope that this hearing will help us in that effort to move ahead. Thank you.

Senator ALEXANDER. Thank you, Senator Bingaman.

Let me now introduce the witnesses that we have. We have two panels. We asked Dr. James Decker, Deputy Director of the Office of Science in the Department of Energy, to be here. Dr. Decker is here on behalf of Dr. Ray Orbach, who could not attend because of personal reasons. We understand that and we hope you will convey to him our best wishes.

After Dr. Decker's testimony, Senator Bingaman and I will ask questions, and then we will go to the other witnesses. The other witnesses I will introduce at that time.

Let me suggest, Dr. Decker, we have your full statement. We have read it. If you could summarize your statement—and I would ask the other witnesses to be thinking of that too—if you could do that in about 5 minutes, then that would leave Senator Bingaman and me and any other Senators who might come the opportunity to go back and forth with questions. Dr. Decker.

STATEMENT OF DR. JAMES F. DECKER, PRINCIPAL DEPUTY DIRECTOR, OFFICE OF SCIENCE, DEPARTMENT OF ENERGY

Dr. DECKER. Thank you, Mr. Chairman. Mr. Chairman and Senator Bingaman, I certainly commend you for holding this hearing, and I appreciate the opportunity to testify on behalf of the Department of Energy's Office of Science on a subject of importance to science in this Nation, advanced scientific supercomputing capability.

Dr. Orbach, who was originally scheduled to appear, asked me to convey his regrets to the committee, that he is unable to be here today.

Computational modeling and simulation on today's supercomputers is already an important tool for scientific discovery. For example, simulation validated by experimental observations has played a key role in understanding energy transport due to complex turbulent processes in magnetic fusion devices.

In climate modeling, where it is impossible to do controlled experiments, computational modeling is essential. In fact, modeling has given us very successful forecasts of seasonal and inter-annual climate variability. For example, we now have quite reliable predictions of the onset and duration of El Niño's southern oscillation climate phenomenon.

With potential advances in computer capability that will increase our computing power by factors of a hundred or thousand in the next few years, researchers will be able to attack larger, more complex scientific questions that will make computational science an even more important tool for scientific discovery.

The advent of Japan's Earth Simulator 2 years ago gave us a glimpse of the potential that can be achieved using computer architectures that are optimized for scientific problems. Coupled with models developed by integrated multidisciplinary teams of researchers, computer scientists, and mathematicians, such computers offer the promise of discovery and design of advanced materials, development of catalysts that dramatically reduce energy costs and emissions, understanding of the dynamics of combustion systems, dramatically better understanding of climate change, integrated simulation of fusion experiments, optimization of the design and technology of future accelerators. Each of the above examples—and there are many more—will have a significant effect on the missions of the Department of Energy, the missions of other Government agencies, and the economy.

The Bush administration has developed a coordinated multiagency approach to revitalizing U.S. high-end computing. An interagency study by the High-End Computing Revitalization Task Force identified our critical needs and, in a report released in May of this year, proposed a game plan to improve U.S. capabilities. The Office of Science and other Federal agencies are working to implement the recommendations of the High-End Computing Revitalization Task Force report and develop the next generation of leadership class computational capability, as well as the networks needed to allow widespread access to these new supercomputers.

On May 12 of this year, Secretary Spencer Abraham announced that the Department of Energy will provide $25 million in this fiscal year to a team led by Oak Ridge National Laboratory to begin to build a new supercomputer for scientific research. In addition to Oak Ridge, the team includes the Argonne National Laboratory, Pacific Northwest Laboratory, and others. This is an important step toward achieving our leadership goals.

Mr. Chairman, you captured the importance of that leadership very well in your floor statement on the Oak Ridge facility when you said it is one of the critical science fields in which we need to be the world's leader. This is because high-performance computing produces scientific discoveries that were once thought only possible through experimentation. I would add in some cases experimentation is not practical or possible, for example, climate change.

Mr. Chairman, high-performance computing provides a new window for researchers to understand the natural world with a precision that could only be imagined a few years ago. It is clear that in combination with our computing industry, we can build the necessary tools. The administration has developed a clear path forward for revitalizing U.S. high-end computing, and with vital support from Congress and the administration, I am confident that we will succeed.

Once again, thank you for the opportunity to testify before this committee on this important matter.

[The prepared statement of Dr. Decker follows:]

PREPARED STATEMENT OF DR. JAMES F. DECKER, PRINCIPAL DEPUTY DIRECTOR, OFFICE OF SCIENCE, DEPARTMENT OF ENERGY

Mr. Chairman and members of the Committee, I commend you for holding this hearing—and I appreciate the opportunity to testify on behalf of the Department of Energy's (DOE) Office of Science, on a subject of central importance to this Nation: advanced supercomputing capability for science.

The Bush Administration has recognized the need for the U.S. to emphasize the importance of high-end computing and is working as a team to address it. The Administration commissioned an interagency study by the High End Computing Revitalization Task Force (HECRTF). The HECRTF report (*http://www.itrd.gov/pubs/2004*_hecrtf/20040510_hecrtf.pdf) reinforces the idea that no one agency can—or should—be responsible for ensuring that our scientists have the computational tools they need to do their job, but duplication of effort must be avoided.

Through the efforts of DOE's Office of Science and other federal agencies, we are working to implement the recommendations of the HECRTF Report by investing in the development of the next generation of supercomputer architectures, as well as the networks to enable widespread access to these new supercomputers.

On May 12th of this year, Secretary Spencer Abraham announced that the DOE will grant Oak Ridge National Lab (ORNL), Argonne National Lab, Pacific Northwest National Lab and its development partners, Cray, IBM and SGI, $25 million in funding to begin to build a new supercomputer for scientific research. The Department selected ORNL from four proposals received from its non-weapon national labs. The Department is in the final stages of completing this award and expects to start the project before the end of this fiscal year.

Computational modeling and simulation rank among the most significant developments in the practice of scientific inquiry in the latter half of the 20th century and are now a major force for discovery in their own right. In the past century, scientific research was extraordinarily successful in identifying the fundamental physical laws that govern our material world. At the same time, the advances promised by these discoveries have not been fully realized, in part because the real-world systems governed by these physical laws are extraordinarily complex. Computers help us visualize, test hypotheses, guide experimental design, and most importantly determine if there is consistency between theoretical models and experiment. Computer-based simulation provides a means for predicting the behavior of complex systems that can only be described empirically at present. Since the development of digital computers in mid-century, scientific computing has greatly advanced our understanding of the fundamental processes of nature, e.g., fluid flow and turbulence in physics, molecular structure and reactivity in chemistry, and drug-receptor interactions in biology. Computational simulation has even been used to explain, and sometimes predict, the behavior of such complex natural and engineered systems as weather patterns and aircraft performance.

Within the past two decades, scientific computing has become a contributor to essentially all scientific research programs. It is particularly important to the solution of research problems that are (i) insoluble by traditional theoretical and experimental approaches, e.g., prediction of future climates or the fate of underground contaminants; (ii) hazardous to study in the laboratory, e.g., characterization of the chemistry of radionuclides or other toxic chemicals; or (iii) time-consuming or expensive to solve by traditional means, e.g., development of new materials, determination of the structure of proteins, understanding plasma instabilities, or exploring the limitations of the "Standard Model" of particle physics. In many cases, theoretical and experimental approaches do not provide sufficient information to understand and predict the behavior of the systems being studied. Computational modeling and simulation, which allows a description of the system to be constructed from basic theoretical principles and the available experimental data, are keys to solving such problems.

We have moved beyond using computers to solve very complicated sets of equations to a new regime in which scientific simulation enables us to obtain scientific results and to perform discovery in the same way that experiment and theory have traditionally been used to accomplish those ends. We must think of computation as the third of the three pillars that support scientific discovery, and indeed there are areas where the only approach to a solution is through high-end computation.

Combustion is the key source of energy for power generation, industrial process heat and residential applications. In all of these areas, combustion occurs in a turbulent environment. Although experimental and theoretical investigations have been able to provide substantial insights into turbulent flame dynamics, fundamental questions about flame behavior remain unanswered. Current limitations in computational power do not allow combustion scientists to address the range of conditions needed to have environmental and economic impact. Leadership class computers should enable us to model more complex fuels with emission chemistry under conditions typical of industrial settings. These computations should make it possible to design new low-emission burners that could dramatically reduce NO_X emissions.

The Fusion Program must be able to model an experiment the size of the International Thermonuclear Experimental Reactor (ITER) through the duration of a discharge that may last on the order of hundreds of seconds. Current codes are able to model a variety of the physical phenomena that occur in small experiments operating on a millisecond time scale. Leadership class computers should enable scientists to simulate burning plasmas in ITER and include new physics such as more realistic treatment of electron dynamics and multiple species of fusion products such as high energy alpha particles.

High-end computing must be coupled with high-performance networks to fully realize its potential. These networks play a critical role because they make it possible to overcome the geographical distances that often hinder science. They make vast scientific resources available to scientists, regardless of location, whether they are at a university, national laboratory, or industrial setting. We work with the National Science Foundation and university consortia such as Internet 2 to ensure that scientists at universities can seamlessly access unique DOE facilities and their scientific partners in DOE laboratories. In addition, the emergence of high performance computers as tools for science, just like our light sources, accelerators and neutron sources, has changed the way in which science is conducted. Today and in the future, large multidisciplinary teams are needed to make the best use of computers as tools for science. These teams need access to significant allocations of computer resources to perform leading edge science. In the Office of Science we are building on the experience of the National Nuclear Security Administration's Office of Advanced Simulation and Computing program to build and manage these teams.

The astonishing speeds of new high-end machines, including the Earth Simulator, should allow computation to inform our approach to science. We are now able to contemplate exploration of worlds never before accessible to mankind. Previously, we used computers to solve sets of equations representing physical laws too complicated to solve analytically. Now we can simulate systems to discover physical laws for which there are no known predictive equations. We can model physical structures with hundreds of thousands, or maybe even millions, of "actors" interacting with one another in a complex fashion. The speed of our new computational environment allows us to test different inter-actor relations to see what macroscopic behaviors can ensue. Simulations can help determine the nature of the fundamental "forces" or interactions between "actors."

The ASCR program mission is to discover, develop, and deploy the computational and networking tools that enable scientific researchers to analyze, model, simulate, and predict complex phenomena important to the Department of Energy—and to the U.S. and the world.

Advanced scientific computing is central to DOE's missions. It is essential to simulate and predict the behavior of nuclear weapons and aid in the discovery of new scientific knowledge.

As the lead government funding agency for basic research in the physical sciences, the Office of Science has a special responsibility to ensure that its research programs continue to advance the frontiers of science. This requires significant enhancements to the Office of Science's scientific computing programs. These include both more capable computing platforms and the development of the sophisticated mathematical and software tools required for large-scale simulations.

Existing highly parallel computer architectures, while extremely effective for many applications, including solution of some important scientific problems, are only able to operate at 5-10% of their theoretical maximum capability on other applications. For most vendors, today's high performance computer market is too small a fraction of the overall market to justify the level of R&D needed to ensure development of computers that can solve the most challenging scientific problems or the substantial investments needed to validate their effectiveness on industrial problems.

Therefore, we are working in partnership with the National Nuclear Security Administration (NNSA), the National Security Agency (NSA), and the Defense Advanced Research Project Agency (DARPA) to identify architectures which are most effective in solving specific types of problems; to evaluate the effectiveness of various different existing computer architectures; and to spur the development of new architectures tailored to the requirements of science and national security applications.

This partnership is working to ensure the development of computers that can meet the most demanding Federal missions in science and national security. We are also working to transfer the knowledge we develop to U.S. industry to enable a vibrant U.S. high performance computing industry, which can provide the impetus for economic growth and competitiveness across the nation. The Office of Science plays a key role in providing these capabilities to the open science community to support U.S. scientific leadership, just as we do with other facilities for science.

Advanced scientific computing will continue to be a key contributor to scientific research in the 21st century. Major scientific challenges in all Office of Science research programs will be addressed by advanced scientific supercomputing. Designing materials atom-by-atom, revealing the functions of proteins, understanding and controlling fusion plasma turbulence, designing new particle accelerators, and modeling global climate change; are just a few examples.

In fact, in fulfilling its mission over the years, the Office of Science has played a key role in maintaining U.S. leadership in scientific computation and networking worldwide. Consider some of the innovations and contributions made by DOE's Office of Science:

- helped develop the Internet;
- pioneered the transition to massively parallel supercomputing in the civilian sector;
- began the computational analysis of global climate change;
- developed many of the computational technologies for DNA sequencing that have made possible the unraveling of the human genetic code.

Various computational scientists have said that discovery through simulation requires sustained speeds starting at 50 teraflops to examine a subset of challenging problems in accelerator science and technology, astrophysics, biology, chemistry and catalysis, climate prediction, combustion, computational fluid dynamics, computational structural and systems biology, environmental molecular science, fusion energy science, geosciences, groundwater protection, high energy physics, materials science and nanoscience, nuclear physics, soot formation and growth, and more.

The Office of Science also is a leader in research efforts to capitalize on the promise of nanoscale science and biotechnology. This revolution in science promises a revolution in industry.

To develop systems capable of meeting the challenges faced by DOE, universities, and industry, the Office of Science invests in several areas of computation: high-performance computing, large-scale networks, and the software that enables scientists to use these resources as tools for discovery. The FY 2005 President's Request for the Office of Science includes $204 million for ASCR for IT R&D and approximately $20 million in the other Offices to support the development of the next generation of scientific simulation software for SC mission applications.

As a part of this portfolio the Office of Science supports basic research in applied mathematics and the computer science needed to underpin advances in high performance computers and networks for science.

In FY 2001 the Office of Science initiated the Scientific Discovery through Advanced Computing (*www.science.doe.gov/SciDAC/*) effort to leverage our basic research in mathematics and computer science and integrate this research into the scientific teams that extend the frontiers of science across DOE-SC. We have assembled interdisciplinary teams and collaborations to develop the necessary state-of-the-art mathematical algorithms and software, supported by appropriate hardware and middleware infrastructure, to use terascale computers effectively to advance fundamental scientific research at the core of DOE's mission.

All of these research efforts, as well as the success of computational science across SC, depend on a portfolio of high performance computing facilities and test beds and on the high performance networks that link these resources to the scientists across the country. DOE and the Office of Science have been leaders in testing and evaluating new high performance computers and networks and turning them into tools for scientific discovery since the early 1950s. The Office of Science established the first national civilian supercomputer center, the Magnetic Fusion Energy Computer Center, in 1975. We have tested and evaluated early versions of computers ranging from the first Cray 1s to the parallel architectures of the 1990s to the Cray X1 at ORNL. In many cases these systems would not have existed without the Office of Science as a partner with the vendors. Our current facilities and test beds include:

- The Center for Computational Sciences (CCS) at Oak Ridge National Laboratory, has been testing and evaluating leading edge computer architectures as tools for science for over a decade. The latest evaluation is on a Cray X1 formed the basis for ORNL's successful proposal to begin developing a leadership class computing capability for the U.S. open scientific community. In his remarks announcing the result of this competition, Secretary of Energy Spencer Abraham stated, "This new facility will enable the Office of Science to deliver world leadership-class computing for science," and "will serve to revitalize the U.S. effort in high-end computing." This supercomputer will be open to the scientific community for research.
- The National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory, which provides leading edge high-performance computing services to over 2,000 scientists nationwide. NERSC has a 6,000 processor IBM SP3 computer with a peak speed of 10 TeraFLOPS. We have initiated a new program at NERSC, Innovative and Novel Computational Impact on Theory and Experiment (INCITE), to allocate substantial computing resources to a few, competitively selected, research proposals from the national scientific community. Last year, I selected three proposals for INCITE. One of these has successfully simulated the explosion of a supernova in 3-D for the first time.
- The Energy Sciences Network (ESnet), which links DOE facilities and researchers to the worldwide research community. ESnet works closely with other Federal research networks and with university consortia such as Internet 2 to provide seamless connections from DOE to other research communities. This network must address facilities that produce millions of gigabytes (petabytes) of data each year and deliver these data to scientists across the world.

We have learned important lessons from these test beds. By sharing our evaluations with vendors we have enabled them to produce better products to meet critical scientific and national security missions. Our spending complements commercial R&D in IT which is focused on product development and on the demands of commercial applications which generally place different requirements on the hardware and software than do leading edge scientific applications.

The Office of Science coordinates with other federal agencies to avoid duplication of efforts. In the areas where the Office of Science (DOE-SC) focuses its research—High End Computing and Large Scale Networking—DOE-SC co-chairs the relevant federal coordinating group. In addition to this mechanism, DOE-SC has engaged in a number of other joint planning and coordination efforts.

- DOE-SC participated in the National Security community planning effort to develop an Integrated High End Computing plan.
- DOE-SC and DOD co-chaired the HECRTF.
- DOE-SC and NSF co-chair the Federal teams that coordinate the engineering of Federal research networks and the emerging GRID Middleware.
- DOE-SC is a partner with DARPA in the High Productivity Computing Systems project, which will deliver the next generation of advanced computer architectures for critical science and national security missions through partnerships with U.S. industry.
- DOE-SC works closely with NNSA on critical software issues for high performance computing.
- DOE-SC, DOE-NNSA, DOD-ODDR&E, DOD-NSA, and DOD-DARPA have developed a Memorandum of Understanding to jointly plan our research in high performance computing. This MOU will enable us to better integrate our substantial ongoing collaborative projects.

High end computing is a key tool in carrying out Federal agency missions in science and technology, but the high end computer market is simply not large enough to divert computer industry attention from the much larger and more lucrative commerce and business computing sector. The federal government must perform the research and prototype development on the next generation of tools to meet those needs. This next generation of computers, however, might also serve to benefit industry.

Mr. Chairman, high-performance computing provides a new window for researchers to understand the natural world with a precision that could only be imagined a few years ago. Research investments in advanced scientific computing will equip researchers with premier computational tools to advance knowledge and to help solve the most challenging scientific problems facing the Nation.

With vital support from this Committee, the Congress and the Administration, we in the Office of Science hope to continue to play an important role in the world of scientific supercomputing.

Thank you very much.

APPENDIX

OFFICE OF SCIENCE: WHO WE ARE

The Office of Science is the single largest supporter of basic research in the physical sciences in the United States, providing more than 40 percent of total funding for this vital area of national importance. It oversees—and is the principal federal funding agency of—the Nation's research programs in high-energy physics, nuclear physics, and fusion energy sciences.

The Office of Science manages fundamental research programs in basic energy sciences, biological and environmental sciences, and computational science. In addition, the Office of Science is the Federal Government's largest single source of funds for materials and chemical sciences, and it supports unique and vital parts of U.S. research in climate change, geophysics, genomics, life sciences, and science education.

The Office of Science manages this research portfolio through six interdisciplinary program offices: *Advanced Scientific Computing Research, Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences*, and *High Energy Physics* and *Nuclear Physics*.

The Office of Science also manages *10 world-class laboratories*, which often are called the "crown jewels" of our national research infrastructure. The national laboratory system, created over a half-century ago, is the most comprehensive research system of its kind in the world. The 10 Office of Science laboratories are: Ames Laboratory, Argonne National Laboratory, Brookhaven National Laboratory, Fermi National Accelerator Laboratory, Thomas Jefferson National Accelerator Facility, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory, Pacific Northwest National Laboratory, Princeton Plasma Physics Laboratory and the Stanford Linear Accelerator Center.

The Office of Science oversees the construction and operation of some of the Nation's most advanced R&D *user facilities*, located at national laboratories and universities. These include particle and nuclear physics accelerators, synchrotron light sources, neutron scattering facilities, supercomputers and high-speed computer networks. Each year these facilities are used by more than 18,000 researchers from universities, other government agencies and private industry.

The Office of Science is a principal supporter of graduate students and postdoctoral researchers early in their careers. About 50 percent of its research funding goes to support research at 250 colleges, universities, and institutes nationwide.

For more than half a century, every President and each Congress has recognized the vital role of science in sustaining this Nation's leadership in the world. According to some estimates, fully half of the growth in the U.S. economy in the last 50 years stems from federal funding of scientific and technological innovation. American taxpayers have received great value for their investment in the basic research sponsored by the Office of Science and other agencies in our government.

Ever since its inception as part of the Atomic Energy Commission immediately following World War II, the Office of Science has blended cutting edge-research and innovative problem solving to keep the U.S. at the forefront of scientific discovery. In fact, since the mid-1940's, the Office of Science has supported the work of more than 40 Nobel Prize winners, testimony to the high quality and importance of the work it underwrites.

Office of Science research investments historically have yielded a wealth of dividends including: significant technological innovations; medical and health advances; new intellectual capital; enhanced economic competitiveness; and improved quality of life for the American people.

Senator ALEXANDER. Thank you, Dr. Decker.

Senator Bingaman, let me suggest I will take 5 minutes, you take 5, and we will go back and forth for a little while. I would like to perhaps aim that we end the hearing by 4 or 4:15. Would that be all right with you?

Senator BINGAMAN. I do not know that I can stay that long, Mr. Chairman, but I will stay as long as I can.

Senator ALEXANDER. We will make sure you have plenty of time to ask questions while you are here because I am glad that you are here.

Dr. Decker, I cannot speak for both Senator Bingaman and myself, but I think I can perhaps to this extent. We are trying to take a look a long way down the road here. For myself, I compliment the Department for its 20-year plan. Chet Atkins used to say in this life you have to be mighty careful where you aim because you might get there. So we have a 20-year plan for science. That is very helpful.

That is the purpose of this legislation that we are introducing. We have in front of us a situation, as you have said and our other witnesses say, where the United States, which has relied upon science and technology for our standard of living to a great degree, has lost the lead in high-performance computing and we need to get it back and we know how to get it back. So we have developed a piece of legislation here called the High-End Computing Revitalization Act of 2004 that we believe would authorize the steps and authorize the funding, which Congress would then have to decide whether it had the money or not, along with the President. We believe these are the right steps.

So I guess my first question to you is this. Does the administration support this legislation? Or if you do not, can you suggest improvements or changes that would make it a better path toward recapturing our lead in high-performance computing?

Dr. DECKER. Mr. Chairman, we certainly very much appreciate the support that is indicated in that bill for the Office of Science and for fixing this important issue. I think the activities that are laid out in the bill are definitely the right ones. There is not an administration position on this bill, to my knowledge, at this time, so I am not able to comment on specifics.

Senator ALEXANDER. Well, what I would like to do, as just one Senator, is to suggest to the Department and to the administration that this would be a good subject to be specific about. We know—and we will hear from other witnesses today—that we can recapture the lead in high-performance computing. It is going to take specific goals. It is going to take some money. We have all been around long enough to know that the budget-setting priority has to begin somewhere and we are hoping to begin it here.

A very important step was the $25 million that you pointed out, which the Congress added and the administration is now spending to begin to do this, but this legislation would authorize the appropriations for the Secretary of Energy for $150 million in the year

2005 on up to $170 million a year for the year 2009, some of that for ultra-scale scientific computing and $10 million for a software development center. I would like to see the administration add to its 20-year plan a budget for this year, for the next year, for the following year that would permit us to go forward.

NOTE: S. 2176 has not yet been reported out of the Energy and Natural Resources Committee. As a matter of policy, OMB does not issue Statements of Administration Policy (SAPs) prior to the bill being reported out of committee because the reported version may differ from the introduced version. The DOE will request a SAP once the bill is reported.

We are in a Presidential year and there will be a lot of back and forth over which political party deserves the most credit or blame for funding for research and development. I happen to think that as a Nation, both parties have done pretty well in some areas over the last several years, including the Bush administration. R&D funding for the National Institutes of Health is up 44 percent over the last 3 years, and we can go down through, the National Science Foundation, up 27 percent over the last 3 years.

But as I have tried to point out, as others on this committee, we need to begin to do for the physical sciences what we have done in the health sciences. The physical science funding has been relatively flat or a little worse in the Bush years and in the Clinton years. So I think there is blame to go around and credit to go around on both sides of the aisle.

What I would like to see us do is to say this is a very specific area in which it is extremely appropriate for the United States to be involved, for the U.S. Government to fund. We have these secret weapons in our country called research universities and national laboratories. No other country in the world has anything like it. They have a few, but it is one of the clear advantages we have. And it is remarkable, in fact, that we could fall behind in high-performance computing and then lay out a plan that within a few short years, by the year 2008, clearly recapture—everyone concedes we can recapture—that lead for a relatively modest sum.

So you may be limited by OMB or Presidential budgets or other priorities in what you might be able to say today. It would be my hope that soon the administration could say that it fully supports this legislation, not just the objectives, which you said it did support, but that we could agree on some goals for authorization levels, or if these are not the right goals, maybe the administration could suggest other goals so we could be on a clear path and so that we, in a bipartisan way, can support implementation. This, after all, was No. 2 I believe on the Secretary of Energy's 20-year plan for where we hope to go with the physical sciences.

Dr. DECKER. Mr. Chairman, I can certainly take that message back to the Department and to the administration.

Senator ALEXANDER. Thank you, Dr. Decker.

Senator Bingaman.

Senator BINGAMAN. Thank you very much, Mr. Chairman.

Let me just underscore what the chairman basically said on the importance of this. When I was in Japan, we did get a briefing by the director of the Japanese Earth Simulator. My strong impression—I believe my recollection is right—is that he said that they were doing some computing on that machine for various companies and others in this country, and in particular, I think he said Lawrence Livermore Lab had contracted with them to do some calculations, some computing. Are you familiar with that?

Dr. DECKER. I was not aware that Lawrence Livermore Laboratory was doing that. I know that they said that they would provide some opportunity for our researchers to use the Earth Simulator, but I do not know how much of that has been done. I certainly can find out and get back to you.

[The information follows:]

None of the DOE laboratories has contracted with the Japanese Earth Simulator Center for scientific calculations. There have been some visits by individual scientists from these laboratories, including one from an Earth Scientist at Lawrence Livermore National Laboratory, to evaluate the capabilities of the Earth Simulator for their particular classes of applications. In addition, there is a Memorandum of Understanding between the Earth Simulator Center and the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory, which is focused on joint activities in performance evaluation and benchmarking to improve our understanding of the factors that affect application performance on large computers.

Senator BINGAMAN. I would appreciate that.

Obviously, I commend the Japanese for the initiative they have shown and the leadership that they have demonstrated in this area. I also appreciate very much their willingness to take on advanced computing work for the United States, our own laboratories, and our own companies.

But if you put this in the larger context, we have had a lot of debate around here about outsourcing. I am not opposed to outsourcing in all its various forms, but this is one area where I would prefer us not to have to outsource. I think it would be much better if we had the capability to do whatever computing we determine we need to do right here. I know that is your view, so I appreciate that.

One other area I wanted to question you about is if we are successful and we go ahead and are able to increase funding in this area, make the investment necessary, and develop the computing capability necessary, how would that be accessed by a professor in my home State if we had a professor at New Mexico State or the University of New Mexico or a researcher or engineer in a private company? How would they access that computing capability if we are going to be paying for this with the taxpayer dollars? It is my view that it should be readily accessible to those who have a legitimate need for it and have a legitimate purpose to pursue with it.

Dr. DECKER. Senator, I agree with that. Certainly it is our intent with a leadership class machine to make it available on a peer-reviewed competitive basis. As you know, we operate a number of large scientific facilities in our national laboratories primarily. The access to those facilities is on the basis of proposals that are submitted by researchers, reviewed, and a decision made based largely on scientific merit. I think that is a model that applies probably with some modification to a leadership class machine.

Senator BINGAMAN. Mr. Chairman, I could ask a series of questions, but I think we have made a good record here with Dr. Decker. I think he is clearly a strong proponent of doing more in this area, and clearly that is our intent with this legislation. So I will stop with that. Thank you.

Senator ALEXANDER. Thank you, Senator Bingaman. I agree with that.

Dr. Decker, thank you very much for your presentation.

We have five other witnesses from whom we would like to hear and we will now invite them to come to the table. We have five witnesses whose resumes are so distinguished, it would take most of our remaining time if I properly introduced them all. So let me give them a brief introduction in the order in which I will ask them to testify.

Dr. Jeff Wadsworth is director of the Oak Ridge National Laboratory. Dr. Wadsworth, thank you very much for being here, and good to see you again.

Dr. David Turek, vice president of Deep Computing for IBM. Thank you very much for coming.

Dr. Daniel Reed is director of Renaissance Computing Institute, University of North Carolina at Chapel Hill. Dr. Reed, thank you for being here.

Mr. Dimitri Kusnezov, director of Advanced Simulation and Computing of the National Nuclear Security Administration. Thank you very much for coming.

Mr. Vincent Scarafino, manager of Numerically Intensive Computing of Ford Motor Company.

You are in the right order. I got a little bit out of order there. So thanks to each of you for coming.

Let me ask again, starting with Dr. Wadsworth and simply going across the row, in about 5 minutes each, can you give to Senator Bingaman and me and to our colleagues in the Senate as the Senator says, as we build a record and develop understanding of the importance of this, a picture of where we have been, what we are capable of doing, what we need to do to recapture the lead in high-performance computing, and what it will cost to get there? I am delighted that this brings us a perspective from a variety of areas in our country, from our laboratories, from our universities, from our private institutes, from other parts of the Federal Government, including national security. So, Dr. Wadsworth, we will begin with you.

STATEMENT OF DR. JEFFREY WADSWORTH, DIRECTOR, OAK RIDGE NATIONAL LABORATORY, OAK RIDGE, TN

Dr. WADSWORTH. Thank you, Mr. Chairman, Senator Bingaman. Thank you for the opportunity to join you today. My name is Jeffrey Wadsworth and I am the Director of the Department of Energy's Oak Ridge National Laboratory. I am particularly pleased to be able to provide this testimony on the role of high-performance computing in addressing major scientific challenges. It is a subject I care deeply about.

For many of us, it has become clear that computational simulation has joined theory and experiment as the third leg of science, and as with theory and experiment, we need increasingly powerful tools to deal with the ever increasingly difficult problems we want to solve. There are at least four types of problems that we need computing for. At least.

One of them is the type of problem that just cannot be solve experimentally. Predicting climate change is the premier example.

There is a second class where we may choose not to do the experiment for policy reasons; underground nuclear testing being the prime example. And that led to the development of the first teraflop class computers in this country as we solved that problem without doing those experiments.

A third class of problems is our desire to design large, complicated structures for economic benefit, and I think we will hear about that, but certainly the Boeing 777 was designed using a large amount of computing capability rather than building prototypes. So that is a third class of problem.

A fourth class is that we can accelerate scientific discovery. If we can accurately simulate structures at the atomic level, this opens the way to solving and designing new materials, solving biological problems with a confidence we did not have before using computing. I am pleased to tell you that in our own at Oak Ridge National Lab in certain industrial materials, computing is leading experiment. Our simulations are now leading the experiments we choose to do because of the accuracy of the simulations, and that can lead us in new directions.

But in high-performance computing, it is well known that if you are standing still, you are falling behind. If you are standing still, you are falling behind. And the Nation has invested in powerful supercomputers for classified work but that similar investment in computing for unclassified work has not happened and we have fallen behind. And as described earlier, in 2002 the Japanese surprised the world with a computer that at that time had more power than our Nation's 20 top unclassified computers. Those 37 trillion calculations per second surpassed that total, and America no longer leads in high-performance computing.

We want to regain that leadership, as do you, and the foundation for addressing this issue is in place. Last month the Secretary of Energy announced that Oak Ridge National Lab and its partners had been selected to establish the National Leadership Computing Facility and to reinvigorate our country's ultrascale computing program.

This facility will bring together world-class researchers. It will bring an aggressive, sustainable path for hardware, an experienced operational team, a strategy for delivering capability computing, and modern facilities connected to the national computing infrastructure through state-of-the-art networking.

As we just heard, this new facility will be open to the scientific community. We will place the world's best scientific application codes and computational specialists at the service of researchers who are attacking problems that can only be solved with this large computing capability. And these teams will be selected through a competitive peer review.

We have made investments at the laboratory that support the Nation's need for this type of computing. We used private funding to build a new computational facility which has 1 acre—that is 40,000 square feet—of world-class computing space to house the next generation supercomputers. In all of these areas, we are partnering not only with the Federal Government, but with industry, with universities, and with other laboratories. And I would like to mention that the State of Tennessee invested $10 million in a joint institute for computational sciences at Oak Ridge, and this building anchors a partnership between the laboratory and the University of Tennessee that is being expanded to include other universities and industry. Every dollar received from now on will be devoted to developing the supercomputer, using it for scientific research because the facility is in place.

This new machine should be larger and more powerful than the Japanese Earth Simulator. Being the largest is not the only goal, but it certainly is a measure of our progress that we expect and we expect this computing power to help revolutionize our scientific research and solve some of our most challenging technical problems. We have heard about some of them: climate change at the local, regional level, energy security through fusion plasmas and the delivery of electrical power, and new avenues of research in biology, pharmaceuticals, chemicals, industrial materials, and so on.

We cannot afford to miss out on these opportunities. Half of our economic growth in the past few decades can be traced to our advances in science and technology. High-performance computing played a critical role and will increase in its importance in the next several decades.

So I would like to commend the committee for putting in the proposed bill, and I am happy to discuss the levels of funding that would be needed to compete with the best computers in the world.

[The prepared statement of Dr. Wadsworth follows:]

PREPARED STATEMENT OF DR. JEFFREY WADSWORTH, DIRECTOR, OAK RIDGE NATIONAL LABORATORY, OAK RIDGE, TN

Mr. Chairman and Members of the Committee, thank you for the opportunity to join you today as you consider a topic that many believe is critical to America's ability to retain world leadership in science and technology.

My name is Jeffrey Wadsworth, and I am director of the Department of Energy's Oak Ridge National Laboratory. I am pleased to provide this testimony on the role of high-performance computing in addressing grand scientific challenges.

In many areas of science, computational simulation—a means of scientific discovery that employs a computer to simulate a physical system—has attained peer status with theory and experiment. Scientific computing has advanced our understanding of the fundamental processes of nature (e.g., fluid flow and turbulence, molecular structure and reactivity, drug-receptor interactions) and of complex natural phenomena (weather patterns) and engineered systems (aircraft and automobiles). Computers are essential for the advanced signal and image processing that underpin modern communications and medical diagnostic systems.

As the complexity of the system being simulated increases, however, so does the computing power needed for an accurate simulation. Just as we have built larger experimental devices and developed more complex theories to understand the most demanding scientific problems, we find that we need high-performance computing to deliver solutions.

This need is particularly acute for those problems that simply cannot be solved experimentally. Climate change is a classical example.

There are also problems that we choose not to solve experimentally, for ethical or policy reasons. The most familiar example of such a challenge emerged after the decision to suspend underground testing of nuclear weapons. Deciding not to "experiment" with actual weapons meant that we needed to find another way to measure and understand forces and reactions of enormous magnitude. Part of the solution required supercomputing at a previously unimaginable scale, and to meet this need we have constructed supercomputers that can simulate a nuclear device by performing literally trillions of calculations per second.

A third class of problems involves the economical design of large structures by using a computer to avoid costly experimentation. During the development of the Boeing 777, for example, it was both physically and financially impossible to build and test prototypes. The solution was a computer simulation that provided a safe and cost-effective new product for American industry.

Finally, we can use supercomputers to accelerate scientific discovery. It is now feasible to accurately simulate structures at the atomic level in a way that can lead to the design of new materials and solve biological problems such as protein folding and cell signaling. In recent work at ORNL on silicon nitride, a ceramic used in a number of industrial applications such as turbochargers and ball bearings, simulation has led experiment—that is, our ability to model the behavior of this material at the atomic level is driving the structural engineering required to develop the next generation of ceramics.

In the field of high-performance computing, however, there is a saying that if you are standing still, you are really falling behind. Our defense laboratories in America have done a marvelous job of developing supercomputers for classified weapons research, but as a nation we have not made a similar investment in supercomputing for unclassified scientific research. Not surprisingly, our international competitors took advantage of our stagnation.

In the spring of 2002, the Japanese surprised the world with the announcement of a supercomputer that could perform at a peak power of 37 teraflops, or 37 trillion calculations per second. Put in perspective, the Japanese machine was more powerful than the 20 largest unclassified computers combined in the United States. Without question, America had surrendered our leadership in high-performance computing. The potential consequences to our nation's prestige, to our economic vitality, and to our historic leadership in the international scientific community were profound.

Mr. Chairman, our discussion today addresses America's opportunity to regain our leadership in high-performance computing. We commend the Chairman and the Committee for recognizing this issue of national importance.

The foundation has already been laid for this initiative. Last month the Secretary of Energy announced that a team led by Oak Ridge National Laboratory was the winner of a competition to establish the National Leadership Computing Facility (NLCF), with the mission of reinvigorating America's ultrascale computing program.

The NLCF brings together world-class researchers from national laboratories, universities, and industry; a proven, aggressive, and sustainable hardware path; an experienced operational team; a strategy for delivering true capability computing; and modern computing facilities connected to the national infrastructure through state-of-the-art networking to deliver breakthrough science. Combining these resources and building on expertise and resources of the partnership, the NLCF will enable scientific computation at an unprecedented scale.

As is the case for other large-scale experimental facilities constructed and operated by DOE's Office of Science, the NLCF will be a world-class resource open to the international research community. At typical experimental facilities, scientists and engineers make use of "end stations"—best-in-class instruments supported by instrument specialists—that enable the most effective use of the unique capabilities of the facilities. At the NLCF, we will organize "computational end stations" that offer access to best-in-class scientific application codes and world-class computational specialists. Multi-institutional, multi-disciplinary teams will undertake scientific and engineering problems that can only be solved on the NLCF computers. These computational end stations will be selected through a competitive peer review process.

We are delighted to have been selected to attack this extraordinarily important problem. Oak Ridge has been a leader in scientific computing throughout its history, and during the past several years our Center for Computational Sciences has addressed the challenges of scientific computing through the evaluation of new architectures and the development of the system software, communications protocols, visualization systems, and network interfaces that must work together with the hardware in solving problems. The Center is a principal resource for DOE's Scientific Discovery through Advanced Computing program, which has created partnerships between computing professionals and researchers throughout the nation to build applications software that makes the most efficient use of the available computing power. Many of these partnerships involve the more than 200 computational scientists who work at ORNL.

We have also made a substantial investment at ORNL that provides a unique national resource for attacking the challenges of high-end computing. Using private funding, we have constructed a brand-new, 130,000-square-foot state-of-the-art computational facility in Oak Ridge. This facility contains a full acre of floor space designed to accommodate next-generation supercomputers and their requirements for electric power and cooling.

To make our computing resources available to the scientific community and to enhance the sharing of data among the nation's leading research institutions, we have developed a variety of high-speed networks, and we are playing a lead role in establishing DOE's Science UltraNet.

In all of these areas, we are working with a number of partners in industry, at universities, and at other national laboratories. Of particular note, the State of Tennessee invested $10 million to construct a facility at ORNL that houses the Joint Institute for Computational Sciences. This new 52,000-square-foot building anchors a unique partnership between the Laboratory and Tennessee's flagship university that is being expanded to include the broader university community.

Thanks to these efforts, we have in place the infrastructure and personnel at Oak Ridge National Laboratory to build a 100-teraflops machine by 2006 and a 250-teraflops machine by 2008 and to use these machines to deliver scientific computation at an unprecedented scale.

To stress what may already be apparent, thanks to the investment of Federal, State, and private resources at ORNL, no funds will have to be spent on building an expensive new facility. Every dollar can be devoted to the development of a supercomputer and the mission of scientific research.

While we anticipate that the size and efficiency of this American supercomputer will surpass the Japanese machine, merely being the largest is not and should not be our only goal.

Just as surely as information technology revolutionized America's economy in the 1990s, high-performance computing could help revolutionize basic scientific research in ways that were unimaginable just a few years ago.

If time permitted, I could share with the committee a lengthy list of potential scientific breakthroughs directly related to the kinds of policy issues that confront the Senate every day. As you discuss clean air, we will be able for the first time to manage the data needed to understand climate changes on global, regional, and local scales.

As you discuss America's energy challenges, we can build models that help us determine how best to control a fusion plasma and reliably deliver power across the national electric grid. In similar fashion, high-performance computing can open up avenues of research for pharmaceuticals, chemicals, industrial materials, and a host of other areas vital to the health of our citizens and the strength of our economy.

Indeed, as I noted earlier, we have already reached the point at which computation is integral to research in virtually every field of endeavor. The two principal tools of scientific discovery—theory and experiment—have been joined by a third: modeling and simulation.

As a nation, we have done a great job in investing in the health sciences. I want to thank Senator Alexander and Senator Bingaman for their leadership in calling for a comparable investment in the physical sciences, which underpin many of the remarkable advances in the life sciences achieved during the last century.

The importance of high-performance computing to both the physical and the life sciences cannot be overstated: the convergence of nanoscale science and technology, computing and information technology, and biology at the "nano-info-bio" nexus affords remarkable opportunities for discovery that we cannot afford to miss out on. It is now generally accepted that half of our economic growth over the past few decades can be traced directly to advances in science and technology. High-performance computing played a critical role in these advances, and it will continue to do so as we extend the frontiers of science.

An investment in high-performance computing would enable the Department of Energy to move forward with plans to attain the ultrascale scientific computing capability needed to realize its goals in nanoscience, biology, fusion science, physics, chemistry, climate simulation and prediction, and related fields.

In summary, these investments would make it possible for America to regain our leadership role in high-performance computing and lay the groundwork for addressing some of the nation's greatest scientific challenges.

Mr. Chairman, I commend you and your colleagues for your vision and your understanding of the challenges facing the nation's research community.

Thank you. I would be happy to answer any questions that you or other members of the committee may have.

Senator ALEXANDER. Thank you, Dr. Wadsworth.

Mr. Turek.

STATEMENT OF DAVID TUREK, VICE PRESIDENT, DEEP COMPUTING, IBM CORPORATION, POUGHKEEPSIE, NY

Mr. TUREK. Mr. Chairman, Senator Bingaman, thank you for inviting me here today. My name is David Turek. I am vice president of Deep Computing for IBM Corporation.

I commend the committee to helping to ensure the continued leadership of the United States in high-performance computing, and I would like to thank you personally for sponsoring S. 2176, which demonstrates the Federal commitment to supporting high-end computing research and development.

I would like to make two points today. One, high-performance computing is an essential ingredient for U.S. scientific and economic competitiveness, and second, the role of government in facilitating partnerships between the Government and industry is critical to further advancing high-performance computing. The Federal Government has had a long and outstanding tradition of support for the advancement of high-performance computing. This has clearly well served diverse agency and departmental missions directly. Federal funding in high-end computing has also provided a stimulus for innovative computing design which has diffused more broadly into the commercial marketplace over time as well.

The tangible benefits that have accrued have been significant. Today our consumer products are better designed and more abundant. Our medical diagnostics and therapeutics are superior. Our ability to analyze the risk of financial instruments takes place at a pace never before imagined. Our understandings of the origins of the universe have developed to an extraordinary extent, and even our movies employ fantastic synthetic images and scenes that entertain an amaze in ways unimaginable even a decade ago.

In essence, then commercial deployment of high-performance computing has become a vehicle for competitive advantage. As a consequence, demand for this level of technology has grown dramatically, creating the success that underlies a considerable level of research and development performed by leading high-performance computing companies.

Today we are also beginning to witness the emergence of small, highly creative and skilled companies that are choosing to compete using high-performance computing technology. IBM has implemented a number of supercomputing on-demand facilities, accessible to customers for short periods of time over the Internet to meet this new need. This accelerates the diffusion of technology into some of the most competitive enterprises in the economy, the small and medium business. This is a proposition that we would not have readily imagined a decade ago, but it has elevated the competitiveness of U.S. industries on the international stage.

As the Government outlines its strategy for high-performance computing, I am sure you realize the enormous impact that you have on the entire Nation in dealing with the changes and challenges facing us in science, business, and homeland security. The careful choices we make through our partnerships and initiatives can significantly enhance our competitiveness and preparedness on all fronts.

It is through the partnership between the Federal Government and computer manufacturers that many of the key advances in high-performance computing have become ubiquitous, and it is one of the principal ways that IBM and other companies achieve and maintain technological leadership. For example, the Department of Energy has contracted with IBM to build the two fastest supercomputers in the world, the Advanced Simulation and Computing Project Purple, based on IBM's power technology, and Blue Gene/L, which together have a combined peak of 460 trillion calculations per second, or teraflop, at Lawrence Livermore National Laboratories, effectively 10 times the power of the Earth Simulator today. Recently the DOE has also announced that IBM will work with the ASCR, or the Advanced Scientific Computing Research program, to build a Blue Gene/L system at the Argonne National Laboratory.

These projects are shaping our approach to system design in terms of systems scaling, tools, system availability and usability to a degree never before imagined. The rate and pace of improvement is truly unprecedented, and much of the credit must go to the demanding requirements of customers like Lawrence Livermore National Laboratory and Argonne National Laboratory.

It is also important for me to address the state of the U.S. supercomputing industry and its ability to deliver on the promise of enhanced scientific and commercial competitiveness. Earlier this week, the semiannual report from the top 500 organizations was published. The publication listed top 500 supercomputers in the world. It is important to note that out of 500 systems, 456 come from U.S. companies. IBM supplies 224 of those. U.S. computer companies account for 89 percent of the total compute power embodied in that list, and the U.S. economy consumes more than 55 percent of the aggregate compute power. This is five times greater than the aggregate compute power consumed by any other country in the world from that list of 500. Our industry is alive, well, and serving the needs of the United States to an unmatched degree.

Finally, as we look out in time over the next 5 years, we expect certain trends to continue. Prices will continue to decline, and the community of potential customers in scientific, commercial, and research enterprises and institutions for high-performance computing will expand as a result. Evolved models of delivery based on on-demand principles will become more prevalent. Systems will become progressively more physically compact, easy to use, and manage, and new applications will stretch our thoughts on systems architecture in currently unanticipated ways. We look forward to the Federal Government's continued role in advancing high-performance computing.

Thank you for the opportunity to speak today.

[The prepared statement of Mr. Turek follows:]

PREPARED STATEMENT OF DAVID TUREK, VICE PRESIDENT, DEEP COMPUTING, IBM CORPORATION

Good morning, Chairman Alexander and members of the Energy Subcommittee. My name is David Turek and I am Vice President, Deep Computing for the IBM Corporation. I have responsibility for providing the products, solutions and services offerings designed to meet the high performance computing needs of customers in market segments as diverse as financial services, business intelligence, scientific research, medical imaging, petroleum exploration, pharmaceuticals, manufacturing and industrial design and digital media.

Thank you for inviting me here today. I commend you and the committee for helping to ensure the continued leadership of the U.S. in high performance computing.

First, I'd like to thank Senators Alexander and Bingaman for sponsoring S. 2176. IBM is fully supportive of the basic tenets of this bill: 1) advancing high end computing in the U.S.; 2) advancing hardware and software development through an ultrascale computing program for scientific research and development; and 3) supporting the DoE's role in advancing high performance computing, especially in the area of nonclassified scientific discovery.

I believe that it is critical to extend U.S. leadership in high performance computing—it is an increasingly important tool facilitating scientific discovery, business competitiveness, and homeland security in a rapidly changing world. Indeed, the, scientific and engineering research communities are increasingly accepting the two main supercomputing activities—*simulation and data analysis*—as two new pillars for discovery, expanding beyond the traditional activities of *theory* and *experimentation*. Through the pursuit of a computing technology to serve diverse agency missions, the federal government has provided a stimulus for innovative computing design that has often, over time, diffused more broadly into the commercial marketplace. The process of innovation and diffusion has been active for decades and the results have been cumulative and profound. We can all remember a time when the concept of supercomputing was restricted to a narrow community of users, extraordinarily skilled and extraordinarily financed to support the operation and acquisition of expensive and exotic technology. Over time, as the inexorable decline in cost of computing progressed, the financial impediments to supercomputing also declined and the community of potential users expanded. Financial accessibility enabled exploration and experimentation with supercomputing in applications that were unanticipated and novel in many wonderful ways. People, enterprises, and institutions which had previously been unable to afford access to this type of technology became able to do so. Creativity blossomed and we began to see the deployment of supercomputers in a broad array of industries outside the domain of the classic large scale research institutions. Commercial deployment of supercomputing became a vehicle for competitive advantage, generating significant commercial demand for supercomputing and creating the economic circumstances that drive the considerable level of research and development prevalent among the leading supercomputing companies we observe today.

Proliferation of supercomputing, enabled in part by affordability, has created cadres of sophisticated users across the entire portfolio of industries served. Many of these people have followed their entrepreneurial instincts and have started or joined new companies, some of modest size, to which they have brought their knowledge of the value and application of supercomputing. The consequence is that today we are beginning to witness the emergence of small, highly creative and skilled companies that are choosing to compete by developing applications based on supercomputing technology. While it may be true that many of these companies still find conventional access to supercomputing limited by concerns of affordability or limited inhouse operational expertise there are new ideas being deployed in the marketplace that are beginning to ameliorate these difficulties.

IBM has implemented a number of on demand supercomputing facilities accessible to customers for short periods of time via the internet. We call this Deep Computing Capacity on Demand. The aggregate compute power in one facility in New York is roughly equivalent to the 4th most powerful supercomputer in the world in terms of the recently published TOP500 list. Yet customers with less than 100 employees in total can access this system for short periods of time to compete with large companies in areas like therapeutic drug design, animation, and petroleum exploration. The ability to proliferate supercomputing into small and medium size companies through mechanisms like IBM's on demand centers enhances the competitiveness of entire industries in ways never before imagined.

As government outlines its strategy for high performance computing, I am sure you realize the enormous impact that you can have on the entire nation in dealing with the ongoing changes and challenges that we face in leveraging economic development and spurring free markets, growth and innovation. The U.S. is experiencing increasing competition from nations worldwide. Our innovativeness can establish our continued competitive standing in the world and assure the advancements necessary to maintain our standard of living for generations to come. High performance computing is an essential element in our effort to compete worldwide. While IBM and many other companies have strong research programs, the federal government is the key to making certain that basic research is done today to ensure tomorrow's inventions.

The High-End Computing Revitalization Act of 2004

The High-End Computing Revitalization Act of 2004 demonstrates that the federal government would like to extend its commitment to support high-end computing research and development.

This is critically important because in addition to meeting its own agency mission requirements, federal funding has traditionally seeded high risk research and enabled the critical university research necessary to advance high performance computing and other important areas in information technology. This investment in research has complemented the financial risks taken by the firms in our industry. It has enabled the development of technologies at a faster pace than could be accomplished by the risk capital of private industry by itself. As a result, innovation has accelerated and new technologies which provide competitive advantage on a national scale across private industry and research institutions are introduced much more quickly than would be possible without federal funding.

The partnership between the Federal government and the computer manufacturers.

The partnership between the federal government and computer manufacturers has been a key driver in advancing high performance computing and making it more ubiquitous. I would, therefore, like to address this in three ways: First, why high performance computing is important; second, the importance of the partnerships that exist between IBM and the DoE; and third, the five year outlook for high performance computing.

IMPORTANCE OF HIGH PERFORMANCE COMPUTING

High performance computing (or supercomputing) provides the means to solve problems that appeared to be unsolvable by conventional means, to solve hard problems with extraordinary speed, and to plumb the depths of complex problems to provide insights never before realized. IBM supercomputers, for instance, have been platforms for analysis in areas such as modeling transportation routes through congested urban areas for the purpose of efficient delivery of goods and services, identity theft prevention, pharmaceutical development, weather forecasting, disease research, petroleum discovery, digital animation, financial services, and basic research on materials and scientific phenomena.

The consequence of such supercomputing applications are manifold: our consumer products are better designed, cheaper and more abundant, our medical diagnostics and therapeutics are superior, our ability to analyze the risk of financial instruments takes place at a pace never before imagined, our understanding of the origins of the universe is developed to an extraordinary extent, and even our movies employ fantastic synthetic images and scenes that entertain and amaze in ways unimaginable even a decade ago. To a substantial degree, these types of benefits have accrued as a result of the relentless decline in computing costs and have enabled a broader community of users to get access to high performance computing capabilities. But we must take into account that not all companies or institutes have equivalent financial or business circumstances: if access to supercomputing is an important ingredient to maintaining or amplifying scientific or business competitiveness, we must contemplate a variety of mechanisms by which access to supercomputing can be made available.

As previously mentioned, we have a service called IBM Deep Computing Capacity on Demand, which enables customers to access IBM supercomputing power over the Internet without the costs and management responsibilities of owning their own supercomputer. Customers can:

- easily tap into massive amounts of supercomputing power that could be otherwise unaffordable
- rapidly deploy supercomputing capacity in response to urgent business opportunities
- pay for supercomputing capacity on a variable cost basis, avoiding large upfront capital outlays and long term fixed IT cost commitments
- lower overall supercomputing ownership and operating costs
- take advantage of a scalable, highly secure and highly resilient on demand operating environment

This approach to providing access to supercomputing resonates with many customers because they pay for what they use, they do not have to worry about technological obsolescence nor housing a supercomputer. This is an important example of how supercomputing, as a means to competitiveness, can be more broadly propagated throughout the marketplace.

But access is not solely a function of affordability; skill within an enterprise or institution also plays a critical role in terms of the ability to exploit the power of supercomputing. To that end, IBM has begun the Productive, Easy to use, Reliable Computing Systems (PERCS) project, one of three projects under Phase II of DARPA's High Productivity Computing Systems (HPCS) program. HPCS is a long-term investigation of a range of issues that define the overall value that a user obtains from a computing system, including performance efficiency, scalability, robustness, and ease of use. The HPCS program emphasizes groundbreaking, high-risk/high-reward research with a close eye on commercialization prospects. IBM is partnering with multiple universities and Los Alamos National Laboratory in this project.

I would also like to address the general state of the U.S. supercomputing industry and its ability to deliver on this promise of enhanced scientific and commercial competitiveness. Earlier this week, the semi-annual report from the TOP500 organization was published. This publication lists the top 500 supercomputers in the world, ordered by sustained performance on a standard benchmark. Out of 500 systems, 456 come from U.S. companies with IBM supplying 224 of the total.U.S.computer companies account for 89% of the total compute power ascribed to these 500 systems. The U.S. economy consumes more than 55% of the aggregate compute power generated by the computers on this list which is five times greater than the compute power consumed by any other country in the world. Our industry is alive, well, and serving the needs of the U.S. economy to an unmatched degree. If you inspect this list, you will note that many of the industries I have previously mentioned are well represented.

IMPORTANCE OF PARTNERSHIPS

An important means by which U.S. supercomputing companies maintain technological leadership is through partnerships with some of our most sophisticated customers. For purposes of this hearing, I will primarily discuss our partnerships with the U.S. Department of Energy (DoE) which have been notable in terms of the extent to which DoE computational requirements have impacted our system designs.

DoE has contracted with IBM to build what will soon be the two fastest supercomputers in the world, ASC (Advanced Simulation and Computing) Purple, based on our high end POWER systems, and Blue Gene/L, based on our low power embedded POWER processors, together they have a combined peak speed of 460 trillion calculations per second (teraflops) at Lawrence Livermore National Laboratory. The ASC POWER system will be used for simulation and modeling in the U.S. nuclear weapons mission and Blue Gene/L will be focused on enhancing ASC scientific simulations and providing ASC researchers with a cutting-edge tool for computational science. The ASC program has been extremely beneficial in its mandate to manage the nuclear stockpile as well as in advancing high performance computing.

We will also work with the ASCR (Advanced Scientific Computing Research) program to build a 5-teraflop Blue Gene/L machine at the Argonne National Laboratory. That marks the third announced installation of Blue Gene/L, after the Lawrence Livermore National Laboratory system and ASTRON, a radio telescope project in Netherlands. Two Blue Gene/L prototypes have been ranked among the most powerful supercomputers in the world today, ranking number four and eight in the Top500 list announced yesterday in Heidelberg. The Blue Gene/L at Argonne National Laboratory will be part of the DoE Office of Science Leadership Class RFP.

The projects that we are executing in partnership with the DoE are shaping our approach to system design in terms of system scaling, tools, system availability and usability to a degree never before attempted. At the end of 1999 the most powerful supercomputer in the world was about 3 teraflops; by the middle of 2005 the Blue Gene system at Lawrence Livermore National Laboratory will be 100 times more powerful and it will incorporate a host of novel technologies and design ideas motivated entirely by the desire to build a system of this class of computational capability at an affordable price. The rate and pace of improvement is truly unprecedented and much of the credit is due to the demanding requirements of, and strong partnerships with customers like Lawrence Livermore National Laboratory.

FIVE YEAR OUTLOOK

As we look out in time over the next five years we expect certain trends to continue: prices will continue to decline and a broader community of potential customers will obtain access to supercomputing as a result; evolved models of delivery based on on-demand principles will become more prevalent; systems will become progressively more physically compact, easy to use and manage; and new applications will emerge in importance that will stretch our thoughts on system architectures in currently unanticipated ways. It is imperative that our industry, sustain and amplify the utility of supercomputing as we make technological advances through this period. We must not create obstacles that will block the use of new technologies. While we stretch towards the future we must be mindful of the past, so that the investments our customers have made in training and application development are not wasted. For example, when we set out to design the Blue Gene system in late 1999, one of its goals was that applications written over the intervening years be portable to this system at the time of its debut. Thus the radical improvements in performance and price performance embodied in the Blue Gene system are perfectly accessible to applications written over the last fifteen years on a wide variety of cluster and massively parallel processor (MPP) systems without, for the most part, any modification. The introduction of new technologies must always make accommodations to the burdens levied on users so that the cost of transitioning to the technology does not dominate the benefits of using the technology.

Within IBM we are pursuing multiple design paths built around a handful of guiding principles: First, although the requirements of the industry are extraordinarily diverse, the fundamental approach to supercomputing will remain wedded to principles of parallel computing. Second, from an implementation perspective this need will be accommodated with "scale-out" or cluster models of computing as well as "scale-up" or symmetric multiprocessor (SMP) models of computing. As is the case today, many customers will deploy both models simultaneously to accommodate the diversity of computational needs they encounter. Third, the centerpiece of our strategy is our POWER architecture. It enables models of parallelism at a variety of price and capability points to better accommodate the broad needs of our customers. Fourth, we will complement our product portfolio with offerings based on industry standard commodity technologies. Fifth, we will continue to embrace open standards. And sixth, all of our design decisions will be driven by customers and market based opportunities.

CONCLUSION

High performance computing requires continued advancement to handle the increasing complexity, scale and scope of challenges arising in industry, government, and the scientific community and solve consistently larger and more complex problems more quickly and at lower costs. The application of high performance computing has allowed us to better understand the complexities of scientific discovery and business-responding to the challenges of national security; environmental impacts; designing large aircraft; simulating critical medical procedures; designing new pharmaceutical drugs; and more. In addition, the range of uses of these tools is being extended as they become progressively more affordable and accessible. It is therefore critical for the U.S. government to develop and fund a creative and productive high performance computing environment and strategy to help enable problem-solving tools for the significant challenges that lie ahead.

Senator ALEXANDER. Thank you, Mr. Turek.

Dr. Reed.

STATEMENT OF DR. DANIEL A. REED, DIRECTOR, RENAISSANCE COMPUTING INSTITUTE, UNIVERSITY OF NORTH CAROLINA AT CHAPEL HILL, CHAPEL HILL, NC

Dr. REED. Thank you. Good afternoon, Mr. Chairman and Senator Bingaman. I am Daniel Reed. I am director of the Renaissance Computing Institute, a collaborative venture of the University of North Carolina, Duke, and North Carolina State University. I also chaired the recent community input workshop for the High-End Computing Revitalization Task Force.

In response to your questions, I would like to make a few brief points today.

First, high-end computing systems share many features with other large-scale scientific instruments. However, I think there is one unique aspect that distinguishes them from other instruments, and that is their universality as an intellectual amplifier. Powerful new telescopes advance astronomy but not material science. Powerful new particle accelerators advance high energy physics, but not genetics. In contrast, high-end computing advances all of science and engineering because all disciplines benefit from high resolution model predictions, from theoretical validations, and from experimental data analysis.

Over 2 centuries ago, the English scientist Sir Humphrey Davy could well have been speaking about high-end computing when he said: “Nothing tends so much to the advancement of knowledge as the application of a new instrument.” In a phrase, success accrues to the talented with access to the most powerful tools.

At several recent workshops, researchers have made the case for sustained performance 50 to 100 times beyond that available on current systems to achieve disciplinary frontiers in physics, astronomy, chemistry, biology, and other disciplines. However, beyond these disciplinary frontiers lie even greater interdisciplinary challenges. For example, in hurricane preparedness, multidisciplinary computations must fuse models of ocean and atmosphere, of transportation and telecommunications systems, and of social dynamics.

Today, computing pervades all of science and it is only slightly hyperbolic to say that today science is computational science.

This brings me to my second point, the need for ongoing balanced investment in high-end architectures to continue to advance this frontier. The explosive growth of commodity clusters has reshaped high-end computing. However, not all applications map efficiently to this model. We substantially under-invested in my judgment in the research needed to develop a new generation of high-end architectures. The consequence of this is limited support for many important scientific and national defense applications.

This leads me to my third point, the critical importance of software and the centers necessary to make these systems useful. Today scientific applications are developed with software tools that are often crude compared to those used to develop many desktop applications. We need new programming models to simplify application development and to reduce software maintenance costs.

Hence, I was pleased that S. 2176 includes support for a high-end computing software development center. Such a center is an institutional mechanism for evaluating new approaches and supporting valuable software tools over the decade or more often needed to maximize their efficacy.

How then in this context do we maintain competitiveness and sustain communities for the long term? High-end computing, as many have noted, is an increasingly international activity with all the associated competition for intellectual talent. To attract and retain the best and brightest talent, we must recognize that computational science requires long-term coordinated support, and that means funding for the staff, the students, the post-doctoral research associates, and the faculty and laboratory researchers that use these systems.

Finally, in this context, what is the appropriate role for the Federal Government? Many of the non-recurring engineering costs necessary to design high-end systems specifically targeted to scientific and government needs are not necessarily repaid by commercial sales. Hence, I believe we must rethink some of our support for models of high-end computing as part of a strategic plan that includes at least four features.

First, support for the long-term R&D necessary to create new generations of high-end systems.

Second, the sustained support for the grand challenge application teams that will develop the next generation applications to use these systems.

Third, regular deployment of the world's fastest computing facilities as part of a broad infrastructure that sustains and supports them.

And finally and equally importantly, vendor engagement to ensure technology transfer and economic leverage.

In summary, the opportunities afforded by high-performance computing are great if we continue to commit to the balanced investment in hardware, software, and applications. Thank you very much.

[The prepared statement of Dr. Reed follows:]

PREPARED STATEMENT OF DANIEL A. REED, DIRECTOR, RENAISSANCE COMPUTING INSTITUTE, UNIVERSITY OF NORTH CAROLINA AT CHAPEL HILL, CHAPEL HILL, NC

Good afternoon, Mr. Chairman and Members of the Committee. Thank you very much for granting me this opportunity to comment on the future of high-end computing I am Daniel Reed, Director of the Renaissance Computing Institute (RENCI), a collaborative activity of the University of North Carolina at Chapel Hill, Duke University and North Carolina State University. I also chaired the recent community workshop[1] for the interagency High-End Computing Revitalization Task Force (HECRTF). I am also a researcher in high-performance computing, with collaborations in both technology and applications.

I would like to begin by commending Senators Bingaman and Alexander for their sponsorship of S. 2176, the High-End Computing Revitalization Act of 2004. In response to your questions regarding high-end computing and S. 2176, I would like to make five points today.

1. SCIENTIFIC COMPUTING: THE ENDLESS FRONTIER

Often, the phrase high-end computing (HEC) is used without adequate definition. This impreciseness has often confused discussion about the unique capabilities of high-end computing, its intended uses and the impact of market forces on access to high-end computing systems. Evolving technology continues to change the quantitative lower bound on the definition of high-end computing—today's desktop computer was yesteryear's supercomputer. However, at any moment, high-end computing is most accurately defined by its impact—those computing systems with transformative power to enable breakthrough scientific discoveries, ensure defense preeminence and maintain international competitiveness.

At the highest level, HEC systems share many features with large-scale scientific instruments, whose national and international deployments are also funded by the Federal government. Each new and more powerful scientific instrument allows us to probe further into the unknown, whether it is deep field images from the Hubble telescope and insights into the origins of the universe, the high energy detectors of Fermi Lab's Tevatron and refinements to the Standard Model of subatomic particles, or large-scale genetic sequencers and an understanding of the deep biological basis of life and disease.

Similarly, each new and more powerful generation of high-end computing systems has enabled validation of theoretical predictions, particularly when circumstances prevent experimental testing (e.g., in cosmology). Where experiments are possible, high-resolution computational models allow researchers to shape those experiments more efficiently. High-end computing also allows experimentalists to capture and analyze the torrent of data being produced by a new generation of scientific instruments and sensors, themselves made possible by advances in computing and microelectronics.

However, one aspect of high-performance computing distinguishes it from other scientific instruments—its universality as an intellectual amplifier. Powerful new telescopes advance astronomy, but not materials science. Powerful new particle accel-

[1] The HECRTF community workshop report is available at *www.hpcc.gov/hecrtf-outreach/20040112__cra__hecrtf report.pdf*.

erators advance high energy physics, but not genetics. In contrast, high-end computing advances all of science and engineering, because all disciplines benefit from high-resolution model predictions, theoretical validations and experimental data analysis.

The English scientist Humphrey Davy could well have been speaking about high-end computing when he said:

> Nothing tends so much to the advancement of knowledge as the application of a new instrument. The native intellectual powers of men in different times are not so much the causes of the different success of their labors, as the peculiar nature of the means and artificial resources in their possession.

In a phrase—success accrues to the talented who have access to the most powerful tools.

Although incremental advances in computing continue to bring research advantages, there are transition points, where major advances in computing have qualitatively changed the range of problems that can be solved. In the 1970s, the emergence of vector computing first made it possible to construct realistic models of many phenomena. In the 1980s and 1990s, massively parallel systems based on commodity processors opened new doors to computational modeling. *However, realistic three-dimensional models of many time varying phenomena remain out of reach with today's HEC systems.*

Two recent workshops, the interagency HECRTF community workshop and the DOE Science Case for Large-scale Simulation (SCALES) workshop, researchers from multiple disciplines made the quantitative case for speedups in sustained performance of 50-100 over current levels to reach new, important scientific thresholds. For example, in quantum chomodynamics (QCD), HEC systems with a sustained performance of 20-100 teraflops (TF) would enable calculations of sufficient precision to serve as predictions for ongoing and planned high-energy physics experiments. In magnetic fusion research, sustained execution at 20 TF would allow Tokamak simulations that resolve the natural length scales of micro-turbulence. Finally, 50 TF was identified as an important threshold for creation of new catalysts that are more energy efficient and generate less pollution.

However, beyond these opportunities lie scientific and public policy problems of even greater complexity—ones that will require the coupling of models from multiple disciplines to understand the complex interplay of many forces, all subject to real-time constraints. For example, in hurricane preparedness, multidisciplinary computations must fuse models of the ocean and atmosphere (for weather prediction and damage assessment), transportation systems (for evacuation and recovery), telecommunication system structure and use (for public and government usage patterns) and social dynamics (for predicting social response).

Similarly, multilevel models of biological processes will be necessary to understand the complex interplay of disease heritability and environmental impact. Constructing a first principles, predictive model of a biological organism is multiple orders of magnitude beyond our current capabilities. However, an accurate computational model of even a single cell could save trillions of dollars in drug testing and would allow us to accelerate the development of new drugs that could be tailored to maximize efficacy and minimize toxicity.

At the end of the World War II, Dr. Vannevar Bush famously noted in his report, *Science: The Endless Frontier*, ". . . without scientific progress no amount of achievement in other directions can insure our health, prosperity, and security as a nation in the modern world." Today, high-end computing is the enabler for scientific progress of all types; it has become the third pillar in the triad of theory, experiment and computation. Indeed, it is only slightly hyperbolic to say that all science is now computational science.

Given the deep interdependence of computing and science, the university community could readily exploit access to new generations of high-end computing systems. Indeed, the community eagerly awaits such access. However, without continued investment in high-end computing capabilities, our rate of scientific discovery will be limited, not by our insights or our imagination, but by the ability to develop and evaluate complex computational models.

This brings me to my second point: the need for investment in new high-end architectures.

2. ARCHITECTURES, SOFTWARE AND INTEGRATED SYSTEMS

The explosive growth of scientific computing based on clusters of commodity microprocessors has reshaped the high-performance computing market. Although this democratization of high-performance computing has had many salutatory effects, including broad access to commodity clusters across laboratories and universities, it is not without its negatives. Not all applications map efficiently to the cluster programming model of loosely coupled, message-based communication. Hence, many researchers and their applications have suffered due to lack of access to more tightly coupled high-end systems. Second, an excessive focus on peak performance at low cost has limited research into new architectures, programming models, system software and algorithms. The result has been the emergence of a high-performance "monoculture" composed predominantly of commodity clusters and small symmetric multiprocessors (SMPs).

We have substantially under-invested in the research needed to develop a new generation of high-end architectures. The result is a paucity of new approaches to managing the increasing disparity between processor speeds and memory access times (the so-called von Neumann bottleneck). Hence, we must target exploration of new systems that better support the irregular memory access patterns common in scientific and national defense applications. In turn, promising ideas must be realized as advanced prototypes that can be validated with scientific codes.

Finally, although high-end hardware is necessary, it is by no means sufficient. Scientific discovery also requires access to large-scale data archives, connections to scientific instruments and collaboration infrastructure to couple distributed scientific groups. *Any investment in high-end computing facilities must be balanced, with adequate investments in hardware, software, storage, algorithms and collaboration environments. Simply put, scientific discovery requires a judicious match of computer architecture, system software, algorithms and software development tools.*

This leads me to my third point: software and the importance of centers.

3. THE CRITICAL IMPORTANCE OF SOFTWARE AND ALGORITHMS

Without appropriate software, the full potential of HEC systems will remain unrealized. In the 1990s, the U.S. high-performance computing and communications (HPCC) program supported the development of several new computer systems. In retrospect, we did not recognize the critical importance of long-term, balanced investment, particularly in software and algorithms.

Today, scientific applications are developed with software tools that are crude compared to those used in the commercial sector, or even available on a personal computer. Low-level programming, based on message-passing libraries, means that application developers must provide deep knowledge of application software behavior and its interaction with the underlying computing hardware. This is a tremendous intellectual burden that, unless rectified, will continue to limit the usability of high-end computing systems, restricting effective access to a small cadre of researchers.

New programming models and tools are needed that simplify application development and maintenance. The current complexity of application development unnecessarily constrains use of high-performance computing, particularly for commercial use. Finally, increases in achieved performance over the past twenty years have been due to both hardware advances and algorithmic improvements; we must continue to invest in basic algorithms research.

Hence, I was pleased to see that S. 2176 includes support for a high-end computing software development center. Indeed, several community workshops and reports have advocated creation of just such a software development center. The limited market for high-end systems means, concomitantly, that software tailored for them also has limited markets. This makes long-term government sustenance of software tools critical to the success of high-end systems.

Given the unique software needs of high-end computing and the importance of long-term research, development and deployment, a software development center provides an institutional mechanism for evaluating new approaches and developing and supporting valuable software tools. Experience has also shown that effective software tools are developed over periods of a decade or more, as experience with applications and architectures is used to rectify software shortcomings and enhance software strengths. The Japanese Earth System Simulator is an exemplar of this experience; it relies on software ideas originally developed by the U.S. high-performance computing program, but later abandoned before they could be fully implemented and proven.

This brings me to my fourth point: competitiveness and community sustainability.

4. COMPETITIVENESS AND RETAINING TALENT

Not only has high-performance computing enriched and empowered scientific discovery, as part of a larger information technology ecosystem, it has also been responsible for substantial economic growth in the United States. Because of this success, information technology and high-performance computing are increasingly international activities, with associated competition for intellectual talent and access to world-class computing resources. *Today, we are in danger of losing our international competitive advantage in high-end computing, with serious consequences for scientific research and industrial competitiveness.*

Investment in high-end computing has advanced a broad array of computing technologies, with associated enhancement of industrial competitiveness. However, today's HEC systems are too difficult to use and often fail to deliver sufficiently high performance on important industrial applications. Multidisciplinary manufacturing optimization, high-speed data mining, virtual prototyping and rational drug design are all targets for industrial application of HEC.[2]

To attract and retain the best and brightest talent, we must create an environment where students and practicing researchers believe, and experience shows, that computational science can catalyze scientific discovery of the first order. Concomitantly, we must sustain the level of investment needed to educate multiple generations of students and allow them to reap the benefits of scientific discovery via computational science. *In the past, the uncertain and highly variable support for high-end computing has led many of these researchers to focus their efforts on theoretical or experimental studies where funding was perceived to be more stable and where access to experimental facilities was assured.*

We must recognize that creating a leading edge computational science code is a multiyear project that requires coordinated effort by professional staff, students, post-doctoral research associates and faculty or laboratory researchers. The research rewards are reaped only after a multiyear, upfront investment. In contrast to many other scientific instruments, whose operational lifetimes are measured in decades, the 2-3 year lifetimes of high-end computing facilities means that new systems must be procured and deployed regularly, as part of a long-term, strategic plan that includes coordinated investment in people and infrastructure.

Science is a "learn by doing" enterprise where excellence begets excellence; computational science is no different. Support is needed for computational science grand challenge teams that can address large-scale problems. The opportunity for students and other researchers to apply their talents using the world's best tools will, as Sir Humphrey Davy famously remarked, yield the competitive advantage.

We must also encourage risk taking and innovation, both in high-end system design (hardware, software and applications) and in scientific applications. A balanced research portfolio includes both low risk, evolutionary approaches and higher risk, revolutionary approaches. By definition, many of the latter fail, but a few will have transforming effects. The opportunity to explore new ideas within an environment that embraces innovation and provides access to the world's highest end computing systems is the clarion call that will continue to attract the best talent.

Finally, my fifth point concerns the role of the Federal government.

5. FEDERAL GOVERNMENT ROLES

The dramatic growth of the U.S. computing industry, with its concomitant economic benefits, has shifted the balance of influence on computing system design away from the government to the private sector. *Given their unique attributes, the very highest capability computing systems have a very limited commercial market, nor is it likely a broad market will ever develop. The high non-recurring engineering costs to design HEC systems matched to scientific and government needs are not repaid by sales in the commercial market place.*

Hence, we must rethink our support models for research, development, procurement and operation of high-end systems. Just as certain capabilities. are supported by the Federal government for the common good—Interstate highways for transportation, national parks for protecting our natural heritage and ships and aircraft for the national defense—so too must high-end computing be sustained by the Federal government. This new approach may well require 10-20 year commitments to strategic vendor partnerships, just as is common in defense procurements. The Federal commitment to fund research and development, together with many years of procurements, can provide the long-term economic incentives needed by the computing industry to justify HEC development.

Hence, ongoing Federal investment, as part of a strategic, long-term computing plan, is critical to ensuring that HEC systems remain accessible for scientific dis-

[2] Many of these topics will be discussed at the upcoming High-Performance Computing User's Conference; see *www.hpcusersconference.com/home.html*.

covery, industrial development and national needs. This strategic plan should include at least five features:

1. Support for the long-term research and development to create new generations of HEC systems matched to the needs of scientific, government and critical industry needs.
2. Sustained support for computational science grand challenge teams to create and use leading edge computational codes and to educate new generations of HEC users.
3. Regular deployment and support of the world's highest performance computing facilities for scientific use, as part of a broad ecosystem of supporting infrastructure, including high-speed networks, large-scale data archives, scientific instruments and integrated software.
4. Coordination and support for national priorities in science, engineering, national security and economic competitiveness.
5. Vendor engagement to ensure technology transfer and economic leverage

The opportunities afforded by high-end computing and computational science are great. However, continued U.S. leadership and the associated scientific benefits can be reaped only by sustained investment in long term strategic plans. We must not waiver in our commitment.

Thank you very much for your time and attention. I would be pleased to answer any questions you might have.

Senator ALEXANDER. Thank you, Dr. Reed.

Mr. Scarafino.

STATEMENT OF VINCENT SCARAFINO, MANAGER, NUMERICALLY INTENSIVE COMPUTING, FORD MOTOR COMPANY, DEARBORN, MI

Mr. SCARAFINO. Thank you. I appreciate being able to discuss the importance of government leadership in advancing the state of high-end computing. My name is Vincent Scarafino and I am manager of Numerically Intensive Computing for Ford Motor Company.

Ford has a long and proud history of leadership in advancing engineering applications and technologies that covers our 100 years of operations. Today we spend billions of dollars every year on worldwide engineering, research and development, reflecting our ongoing commitment in technology to bring innovative products to markets around the world.

The effect government decisions have on the direction of high-end computing has been well demonstrated. Up until the mid 1990's, the Federal Government played an active role in funding the development of high-end machines with faster, more powerful processing capability and matching memory bandwidth. Built to meet the needs of government security and scientific research, their development spurred new applications in the private sector.

The mid 1990's, however, brought an embracement of parallel processing as the holy grail for harnessing computing power to solve the next generation of intractable problems. What followed were significant advances in computer science in the area of parallel processing. Nevertheless, an unfortunate and unintended consequence was that scientists and engineers who, for the most part, did not have the necessary computer science expertise, were not in a position to participate in these advances.

I am encouraged by this committee's interest in advancing the fundamental speeds and capabilities of high-end computers and reestablishing U.S. leadership in the field of supercomputing. There are still difficult problems waiting to be solved and many of them may not be parallel in nature. A parallel approach is effective in many instances, but there are limitations. We are at a level for many applications where further development requires higher levels of individual processor performance.

For example, the current state-of-the-art in simulation programs used by industry apply a single type of computational analysis. Some examples are heat transfer, physical deformation, vibration, and fluid flow. The ability to apply more than one of these fundamentals simultaneously is one of the evolutionary directions that will move science forward. This is referred to as multi-physics simulation and is very computationally demanding. An example is computational aero-acoustics where the characteristics of fluid flow and structural behavior are modeled. This provides a virtual wind tunnel that can potentially predict the wind noise characteristics of a vehicle, which is among the most cited customer issues. Another automotive application could be the design of exhaust systems for effective noise management. Ford is planning to work with Oak Ridge National Laboratory to evaluate the feasibility of this with current available software on a very large capability platform at the lab.

Advances in vehicle safety analysis, which currently depends on finite element models, could be enhanced with improved high-performance computers. New element formulations have been created that have the potential to provide improved fidelity but at a cost of needing significantly more computing power. Also, more detailed material property modeling will expand the application to new levels. Accurate prediction of human injury waits for the arrival of faster processors. Predicting the behavior of composite materials in impact situations is also too difficult for today's machines.

Computing capabilities allow Ford to accelerate the design cycle. More powerful high-end computing systems will help engineers balance competing design requirements. Performance, durability, crashworthiness, occupant and pedestrian protection are among them. These tools are necessary to be competitive in today's technology driven and intensely competitive markets. The United States is the largest and most open market in the world and the battleground for the world's global auto makers.

The competitive impact of government policies and technological support from other countries is easily noted. Germany provides its industries access to high-end computers through universities that have a core mission objective to support industry. The United Kingdom and France provide supercomputer resources to European aerospace industries. Japan produces high-end computers and makes them available to its industries for research through universities.

U.S. leadership in the area of supercomputing is needed to promote technologies and scientific advancements that provide the basis for economic growth and competitiveness. The Federal Government cannot rely solely on market-based economic forces with fragmented and relatively low volume applications to advance high-performance computing capability.

I would also like to mention the importance of software development as an integral part in achieving high-end computing capability. Many of the application codes used by the automotive industry have their roots in government-funded development projects.

NASTRAN from NASA and DYNA3D from Lawrence Livermore Labs provided the solid background. Languages and programming environments need to allow scientists and engineers to express their problems in terms they are familiar with.

Advancing high-performance computer capabilities will enhance U.S. manufacturing competitiveness. Our experience over the past 100 years in product development and manufacturing has shown that continued investment in technology is needed in order to provide cleaner, safer, more efficient, and more affordable products to our customers. Technology will play an increasingly important role moving forward as a key competitive driver for U.S. industry and the economy as a whole.

Once again, I applaud the focus of this committee on ensuring that we can meet the competitive challenges of the future by promoting funding initiatives at the National Science Foundation and at the Department of Energy in the area of high-performance computing.

Again, thank you for this opportunity.

Senator ALEXANDER. Thank you, Mr. Scarafino.

Dr. Kusnezov.

STATEMENT OF DR. DIMITRI KUSNEZOV, DIRECTOR, OFFICE OF ADVANCED SIMULATION AND COMPUTING, NATIONAL NUCLEAR SECURITY ADMINISTRATION

Dr. KUSNEZOV. Thank you. Mr. Chairman and Senator Bingaman, it is an honor for me to be here and be afforded the opportunity to provide you an overview of the advanced simulation and computing program.

The central problem this program addresses is the replacement of underground testing with the more rigorous scientific methodology with which to assess and maintain our confidence in our nuclear stockpile.

The first point I would like to make is that ASC deliverables are time sensitive. Supporting national policy with respect to the maintenance of our nuclear stockpile requires that we be able to certify annually to the Secretaries of the Departments of Energy and Defense that the stockpile is safe, reliable, and secure.

The stockpile is aging and refurbishment of some parts is essential. This drives a sense of urgency on our part to have the tools, both the codes and the supporting computer infrastructure, in place and tested so that they can be applied and provide answers to stockpile questions. This is our mission, to provide leading edge, high-end simulation capabilities needed to meet weapons assessment and certification requirements. We cannot achieve this mission without the multidisciplinary scientific underpinnings critical to this major computational effort. Computation underpins all we do.

Second, simulating the time evolution of the behavior of an exploding nuclear device is not only amount of the scientific enterprise from a computational perspective, it probably represents the confluence of more physics, chemistry, and material science, both equilibrium and non-equilibrium, at multiple length and time scales than almost any other scientific challenge. Both our legacy and our modern codes must be able to reproduce the data taken in

Nevada and in the Pacific, and with the exception of some anomalies that remain to be explained, they do.

However, now we are calling on the simulations to evaluate phenomena that result from changes to the devices from the way they were originally designed and built. The systems, most of which are decades old, are not aging gracefully. The radioactive environment in the interior of a nuclear device causes uncertain changes in the material properties and their subsequent behavior. We rely on our ability to predict the burning of high explosives, the fission properties of critical metals, and the stability of various inert materials. The physics and chemistry of aging is far from understood and will require increasingly microscopic descriptions to characterize their effects accurately.

Surveillance activities regularly open existing devices and examine them for these kinds of changes. Now we have to understand how much these changes matter, how critical they are. We can only do this through detailed simulations that include the necessary physical representations. These stockpile effects, almost all of which are three dimensional, currently require heroic, nearly yearlong calculations on thousands of dedicated processors. It is essential that we provide the designers with the computational tools that allow such simulations to be completed in a reasonable timeframe for systematic analysis. This is one of the requirements that drives us well into the pedascale regime for our future platforms. An ingredient of this landscape is that most of the work that we do is and must remain classified, which limits the kinds of collaborations we are able to do with various other agencies and academia.

My last point is that there is a broad and fertile ground for serious collaborations. Today scientific enterprise is enabled through large supercomputers. Clearly one cannot just buy such machines and plug them in. There are complex operating systems, compilers to translate human written code into machine language, sophisticated debugging tools to find the inevitable errors in any large programming enterprise, and evaluation techniques such as that which enables three-dimensional visualization of the results that we get from the codes. Each of these is essential for our success and does not need to be invented here. We can share ideas, share implementations, and provide serious peer review of approaches we are taking.

I support the work of the committee to inject energy, resources, and commitment to strengthening the scientific enterprise of this Nation. It is essential for our national security in all its manifestations from defense to economic competitiveness to the quality of individual life. At NNSA our focus has been and must continue to be to support national policy in the arena of nuclear competence. I choose the word "competence" carefully because it implies many things. It implies a powerful scientific underpinning to a most complex enterprise and it implies the infrastructure to support that science. Most of all, it demonstrates to our adversaries that we know what we are doing. That is our first and foremost responsibility.

In closing, we in the Department of Energy are charged with two disparate missions: one of scientific exploration and the other of national security. I would like to emphasize that we cannot afford to exchange one for the other. We are mutually stronger because of the commitment and the dedication to innovative science that the basic and applied work of the two parts of the Department respectively bring together to their missions. The country is stronger as a consequence.

Thank you.

[The prepared statement of Dr. Kusnezov follows:]

PREPARED STATEMENT OF DR. DIMITRI KUSNEZOV, DIRECTOR, OFFICE OF ADVANCED SIMULATION & COMPUTING, NATIONAL NUCLEAR SECURITY ADMINISTRATION

I thank the committee for the opportunity to address the Members and to express my support for computation as a major underpinning of the scientific enterprise. As it is in many contexts, within my sphere of NNSA, computing is making possible, things previously thought to be impossible.

INTRODUCTION

Within the Stockpile Stewardship Program, the National Nuclear Security Administration and the Department of Energy and its three weapons laboratories are responsible for assuring the President, annually, that each nuclear weapon system in the existing stockpile is safe, secure and reliable, without the need to resume underground testing. This is a scientific and engineering challenge that many have likened to the Manhattan Project and the Apollo Project. One of the most important elements of the Stewardship Program is the Advanced Simulation and Computing Program (ASC, formerly ASCI).

In the post cold war world many have asked why the United States still needs to maintain a nuclear stockpile. As international events have proved since the fall of the Berlin Wall and the collapse of the Soviet Union, the world remains a dangerous and unpredictable place. A safe, secure and reliable nuclear deterrent reassures our allies that the security umbrella which helped secure the peace during the cold war remains effective; it deters potential adversaries, and advances non-proliferation goals. We approach our mission with these ends in mind.

Achieving the necessary credibility, both internally and externally, reflects our commitment to the nation to ensure that it can continue to depend on the reliability of the stockpile. The simulation tools we develop to this end rely for their credibility on a combination of non-nuclear experiments, comparisons with analytic solutions where possible, rigorous analysis of the scientific data gathered from over 1000 nuclear tests and extraordinary computing.

Since the dawn of the nuclear age, computation has been an integral part of the weapons program and our national security. With the cessation of testing and the advent of the science-based Stockpile Stewardship Program, ASC simulations have matured to become a critical tool in stockpile assessments and in programs to extend the life of the nation's nuclear deterrent. Using today's ASC computer systems and codes, scientists can include unprecedented geometric fidelity in addressing issues specific to life extension. They can also investigate particular aspects, such as plutonium's equation of state, scientifically and in detail heretofore impossible, and then extend that understanding to the full weapons system. The results of these simulations, along with data from legacy testing and ongoing experimental activity, improve the ability of weapons designers to make sound decisions in the absence of nuclear testing. Given the critical role that numerical simulations play in the Stockpile Stewardship Program, the credibility of our simulation capabilities is central to the continued certification of the nuclear stockpile.

ASC STRATEGY

Simulating the time evolution of the behavior of an exploding nuclear device is not only a mammoth scientific enterprise from a computational perspective, it probably represents the confluence of more physics, chemistry and material science, both equilibrium and non-equilibrium at multiple length and time scales than almost any other scientific challenge.

Changes that we must make in nuclear weapons to extend their lifetime, under the Life Extension Program to compensate for unavoidable corrosion and chemical decomposition also require the application of sophisticated engineering modeling to enable us to replace components and to perform refurbishments of existing weapons without altering weapon performance. Moreover, understanding the consequences of aging, evaluating the effects of corrosion and oxidation, folding into our calculations the inevitable changes in material properties in self-irradiating environments, all require a deeper understanding and the ability to model the relevant physical phenomena.

The ASC Program must be a balance of short-time-line deliverables, like the annual assessment, and longer-term research activities. The latter are essential to reduce the uncertainties in our simulations and to better model aging effects outside of the parameter space defined by the nuclear test base.

As regards weapons simulations, there are many areas of classified research that we must perform in a secure manner, for example, understanding specific properties of special nuclear material as well as analyzing the behavior of systems under a particular set of extreme conditions (stockpile to target sequence). For this we must maintain a strong, in-house scientific capability. While much of what we do can and does benefit greatly from work with others, "outside the fence", our core mission and the rationale behind our structure and activities has been and will continue to be the support of the Stockpile Stewardship Program.

To deal with the complex needs of Stockpile Stewardship, ASC has developed as a comprehensive ten-year program tuned to deal with the schedule of deliverables. It includes the development of two- and three-dimensional weapons codes and physics models built on a validated scientific/engineering base, the scientific resources necessary to develop better models, the acquisition of powerful computing platforms and the creation of the supporting hardware and software infrastructure. A balanced allocation of resources across these components is essential for program success. For example, platform costs represent about 15% of the overall ASC budget—the greatest investment is in the people, particularly those focused on scientific applications, physics and model development.

The FY 2005 request now before the Congress provides a total of $435M to pay for people at the weapons labs; this is an increase of 3.6% over FY 2004. Recent action by the House Energy and Water Development Committee to cut $75M places at risk not only these critical people but also the next generation of machines that are needed at the laboratories to tackle the ever-increasing demands of the weapons designers and engineers. A recent study by the JASONS highlighted both the capability and capacity constraints.

Weapons code development and computing infrastructure have evolved together in complexity and sophistication. At the very beginning of the ASC Program, we looked at the kinds of problems we would need to solve, when we needed to be able to solve them, and how quickly we would need to get results from calculations. This analysis determined both the size of the computers we set out to acquire through partnerships with computer industry leaders and their acquisition schedule. In 1995 our computing platform goal was to obtain a computer system by 2004 that could process 100 trillion floating-point operations per second (a trillion floating-point operations per second is one teraflop or TF)—the "entry-level" capability for high fidelity, 3D, full system weapon simulations. Clearly, major innovations in massively parallel computer systems and computing infrastructure would be required to meet this goal. At the same time, highly scalable weapons simulation codes that could make effective use of these computers had to be developed.

The ASC platform strategy is to provide robust production level capability to the program today, while staying abreast of recent advances in computer technology to prepare for the future. Each platform, which necessarily pushes the current state-of-the-art, requires a close partnership between the weapons laboratories and industry to bring to fruition. ASC has produced four generations of powerful platforms having impacts on stockpile decisions code-named: Red, Blue, White, and Q. Today, the ASC platforms of highest capability are LLNL's "White" at 12.3 TF and LANL's "Q" at 20 TF. The present acquisitions are SNL's "Red Storm" projected to be 40 TF and LLNL's "Purple" at 100TF, arriving in mid 2005. The 100 TF platform was sized during original program planning activities to provide a reasonable turn-around time for 3-dimensional weapons calculations, taking into account the minimal resolution and physical models required. A one-week calculation was estimated to require roughly a 100 TF supercomputer. This represents an entry-level calculation since it begins to make 3D calculations more of a tool than 476-de-force, with sufficient resolution and science to render the simulations of value to the designers. In the interim, as the Stockpile Stewardship mission has progressed, new issues and questions have come to light. As we address these emerging needs through improved science and resolution, we balance the program planning to evolve accordingly.

The acquisition of Purple is the fulfillment of the original ASC 100 TF goal.

Nearly 9 years after the original plan, it should be delivered within a few months of the anticipated date. But this is only the capability demonstration. There is a clear need, well supported by distinct technical requirements, for almost equal amounts of capability and capacity, leading up to but not stopping with a petaflop (PF = 1000 TF) class computer by the end of the decade.

To meet the broader, evolving computing needs of the future, ASC is now acquiring Blue Gene/L, a 360 TF platform that will be used extensively to improve physics models in ASC codes starting in FY05. This platform will also be used to evaluate the technology for suitability to a broader workload. Blue Gene/L represents a very positive benchmark for high performance computing in the United States. The system represents a substantial R&D investment by IBM in a "computer for science". This investment was initiated and encouraged by NNSA and the Office of Science long before the Japanese Earth Simulator was widely discussed in American circles. This technology demonstrates that American industry and government partners have never wavered from focusing on the very difficult issues associated with scientific computing. Considering that Blue Gene/L in 2006 will be running problems ten times more demanding than are currently possible on the Earth Simulator and that it will cost less than 1/6 as much as the Earth Simulator, demonstrates the vitality and imagination of American industry and the forward-looking planning and commitment of resources by NNSA and the ASC Stockpile Stewardship Program.

Although our current acquisition model meets our present programmatic needs, we remain supportive for additional investments in innovative architectures that will carry us to the next generation of computing architectures. As an integral part of the NNSA ASC Program, we fund targeted efforts to study advanced architectures and a program we call "PathForward" that looks to the future in both hardware and software components. Additionally, we seek opportunities to capitalize on the work of others through formal structures, such as the HEC Revitalization Task Force and the DARPA HPCS Program, as well as less formal collaborations, many of which are with Office of Science principle investigators.

THE FEDERAL ROLE IN HIGH-PERFORMANCE COMPUTING

Due to programmatic requirements, NNSA has historically been the owner of the largest high-end machines in the world. This has created-an expectation on the part of the open science community that some fraction of these resources would be available for basic research modeling, computing and analysis. Consistent with our responsibility to deliver on our mission, we have always made a large number of cycles available to the scientific community, taking great care with the restrictions imposed by maintaining the security of our classified workload and paying attention to export control issues.

However, the demand has historically outstripped the availability and resulted in a tension between open and secure needs. This is alleviated to some extent today by the advent of inexpensive, terascale Linux clusters at many centers, particularly in the academic communities. Comparing the top 500 list 5 years ago with today's list, one finds today over 100 machines with greater than. one teraflop peak, compared 5 years ago when there were only four. Clearly we are entering a time of a greatly enhanced capacity of cycles for science, spread throughout the world's scientific community. A large fraction of these cycles have become available outside our borders. In fact, in 1998, 290 of the top 500 most capable machines were U.S. machines. In 2003, that number had dropped to 248. Although the total teraflops in the top 500 available in the U.S. has increased from 28 teraflops to 531 teraflops, the numbers overseas has increased from 16 to 391 teraflops. The challenge to American success in this endeavor is obvious.

In November 2002, the Secretary of Energy, Spencer Abraham, announced the ASC Purple contract between IBM and LLNL, for the 100 teraflop Purple platform and the 360 teraflop Blue Gene/L system. Last month, Secretary Abraham announced the ORNL procurement, which will deliver even more computing to the open scientific community. This commitment to computing from the Department of Energy demonstrates the leadership role the Department has taken in overseeing the development of computational science in the U.S.

In order for the country to move forward effectively, it is essential that multiple architectural approaches and technologies be explored systematically. For the past decade, the NNSA ASC Program, working with first tier vendors, has demonstrated that very large systems can be built successfully on accelerated timescales and at reasonable cost to meet extraordinary programmatic objectives. In recent years, the DARPA High Productivity Computing Systems (HPCS) Program has invigorated U.S. vendors through its unprecedented investments to build innovative high-end computing solutions. Even so, for there to be long-term, sustainable paths in multiple technologies to reduce risk, additional investments are essential beyond those possible by NNSA and DARPA, and so the DOE's Office of Science Leadership Class computing effort represents a welcome development.

In addition to the most capable high-end computing platforms, advanced applications require a powerful supporting infrastructure that includes integrated systems of compilers, debuggers, visualization tools, and secure computing and data transmission over long distances. For many of these support activities we rely on an industrial sector that we believe must be motivated to continue to work with us on our problems of such national significance.

ASC AND SC IN PARTNERSHIP

The Secretary's announcement of the ASC Purple contract between IBM and LLNL, for the 100 teraflop Purple platform and the 360 teraflop Blue Gene/L system along with last month's, announcement of the ORNL procurement highlights a major source of commonality in our goals, in this case for high-performance tools to enable our scientific endeavors.

Additionally, we have collaborated on and jointly issued a policy with the Office of Science that directs that software developed under contracts from the Department will be licensed as open source. This will make available the fruits of our joint labors to the academic community and to the industrial sector. On the hardware side our procurements of the Cray Red Storm and the IBM Blue Gene/L machines not only include Office of Science, but also involve other agency and academic leaders in peer reviews, and allows these partners to weave first-available technologies into their activities.

To accomplish our mission, now and in the future, the program must rely upon scientific progress in many fields of physics and engineering, as well as innovative advances in computer science and modern architectures. We cannot do this in isolation but must continue to remain connected to the broader science community as a whole. Although the nation's nuclear weapons program has a long history of leadership in driving the supercomputer industry and in using the largest capability machines to inform design and maintenance decisions, the enormity of the problems we face today are beyond NNSA's ability to go it alone. We are actively partnering with other agencies, industry and academia to develop tools and techniques of applicability to our programmatic challenge.

We are committed to maintaining the country's scientific strength. To that end, we nurture computation at every level, particularly at the high end, and we support recruitment and the training of the next generation of computational physicists and engineers to whom we will eventually entrust our national security responsibilities. One example in this respect is our funding contributions to the Computational Science Graduate Fellowships Program, which we do in conjunction with the Office of Science.

With that goal in mind, the Computational Science Graduate Fellowship program, jointly funded by the DOE Office of Science and NNSA/DP, is administered by the Krell Institute to support highly capable individuals pursuing doctorates in applied science or engineering disciplines with applications in high-performance computing. The fellowship program requires completion of a program of study that provides a solid background in three areas: a scientific or engineering discipline, computer science, and applied mathematics.

U.S. COMPUTING IN A GLOBAL CONTEXT

We have heard much in the past two years on the Earth Simulator, the Japanese supercomputer primarily focused on climate modeling. With roughly five years in the planning, the delivery of the Earth Simulator was not a surprise. Neither is the performance of the particular set of applications chosen to run on it. We have not ceded super computing leadership to the Japanese as a result of their fielding of the Earth Simulator. To achieve the results they exhibit, they spent two years tuning a climate code to run on that particular architecture and the government invested well over $350M, three times the amount we spent on bringing the ASC White and Q machines up. Their success does demonstrate the power of governmental will and commitment.

It is fair to say, however, that the debut of the Japanese Earth Simulator has revived the debate about the role of vector computing, whether ASC should reconsider the role of vector processing in its future machines. Although vector supercomputers provide large performance gains in certain applications, they are not well suited to ASC applications and, in particular, do not provide sufficient performance gains to outweigh their increased costs. The large, multi-physics applications that dominate the Stockpile Stewardship workload display a relatively large scalar fraction since the algorithms that provide the shortest time to solution are often not the ones most amenable to vectorization.

In the past two years, NNSA platforms and their performance have been measured against the Earth Simulator and other vector-based architectures. We take the issue of performance very seriously and actively model our applications across architectures, paying close attention to the cost vs. performance and to the time to solution of our codes and the platforms on which they run. A metric that has received wide currency is the 'efficiency' ratio of floating point operations to peak floating point potential. This metric does not account for many of the details of our applications (e.g. memory fetches, integer arithmetic, logic operations). One cannot separate the specifics of physics models and their implementation from machine architectures; some applications will run better on platforms better suited to the details of their problem suite. One can increase performance as measured by percent of peak floating point operations and significantly increase the total time it takes to complete the calculation. However, this exchange of making an improvement in an arbitrary metric may discourage the users of our codes.

In a recent analysis, it was shown that for ASC applications, vector machines were approximately 3 times less cost-effective than commercial-off-the-shelf processing nodes. This follows because ASC codes have a relatively small (0.1 to 0.75) vector fraction compared to some other codes of interest to the scientific community. These are the technical and financial considerations that drive different programs to seek different computational architectures.

CLOSING REMARKS

In the realm of collaborations, it is important to recognize that the most fruitful collaborations take place on the scientist-to-scientist level. Agency management can foster an environment in which such collaborations can flourish, and they do so even today, but we cannot force them. We have many collaborations with many agencies, most especially our, sister agency, the Office of Science. These are good and productive collaborations, often focused on computer science solutions and ideas for new solvers, in the general sense, that benefit us both.

In addition to our own intra-agency and interagency-focused efforts, we continue actively to work with the broader community engaged in promoting high-end computing and the development of a supporting infrastructure. Our recognition of the need for a vigorous partnership between agencies and government sponsors as well as for interagency collaborations demonstrates that commitment. Further, the ASC Program supports the Council on Competitiveness' Initiative in Supercomputing and Productivity, along with our colleagues from DARPA and the Office of Science.

I hope it is clear from my comments and the actions of our program that we recognize the importance of sustaining a broad scientific community. In addition to the work performed at the Defense Programs laboratories to develop key models that reflect the physical reality encompassed by our mission, we must and do rely upon the work of our colleagues in other agencies. In particular it is the responsibility of the ASC Program to turn the sum of our understanding into high-fidelity computer representations that are the crucial underpinnings of our ability to respond to the nation's policy decisions with respect to the nuclear deterrent. Our substantial investments are sized and balanced against our need for experimental facilities and our support of the ongoing workload across the weapons complex.

A healthy and vital U.S. High End Computing industry is crucial to our continued success in Stockpile Stewardship. We recognize that we cannot go it alone but must engage and even rely upon the technical achievements of our colleagues in all aspects of scientific computation and in the development of the supporting infrastructure. This is a massive enterprise from which we all gain, especially as we partner and build productive relationships for the greater benefit of this country and its people.

Senator ALEXANDER. Thank you, Dr. Kusnezov, and all of you.

I will ask a few questions and then turn to Senator Bingaman.

Dr. Kusnezov, the National Nuclear Security Administration, which you describe, has historically been the owner of the largest high-end computing machines I guess in the world. What fraction of these machines has been available for unclassified scientific computing?

Dr. KUSNEZOV. Thank you, Mr. Chairman. That is a very good question. We have a number of restrictions with our largest platforms mainly because of the nature of our work. It is classified. And we put it behind the fence and it is largely unavailable to the open scientific community.

During the stand-up period, as we introduce these machines into the complex, they are in the open environment. This is because it facilitates the work of the vendors in standing these up and implementing the environment to make these usable. During that period, we traditionally have made the machines available to some leading edge scientific work, but this is not an overall commitment to open science mainly because we do not have the resources to support that.

We do have a fair amount of open scientific work through our university partnerships. To support that, we have leveraged scientific resources within the country. In particular, we use now the scientific computing at the University of California at San Diego because this allows us not to worry about the export control and classification issues of having foreign nationals use our platforms.

Senator ALEXANDER. Thank you.

Dr. Wadsworth, let me go back to you with some basic things. You are fairly precise in your testimony about where you believe this project can go by the year 2008. Could you just, in shorthand, describe in summary the teraflops or the calculations, where we are today with the kind of high-performance computing we are talking about studying at Oak Ridge and where you hope to go and where that will put the United States at that time in comparison with the rest of the world?

Dr. WADSWORTH. Yes, I will be happy to do that. We prepared these estimates for the proposal that we submitted to the Department of Energy. At a substantial investment of the kind contained in S. 2176, we believe we can be at 270 teraflops in 2007.

Senator ALEXANDER. Today we are where?

Dr. WADSWORTH. Maybe 10. At a lower level of investment, then we would get to about 100 teraflops in 2007, a lower being at the current level of investment of $25 million or so. So at $100 million a year, you can get up to a number like 270; at a lower number, you would get to about 100.

But one has to remember that the rest of the world does not stand still. So we would advocate a very aggressive investment. That aggressive investment would not be out of line with the kind of investments for world class facilities in other fields of science.

Senator ALEXANDER. Let me ask one more question before I go to Senator Bingaman. In my conversations with the managers of the Oak Ridge program before the competition was conducted, some of your colleagues felt like one of your advantages there was your ability to provide an easier access for other scientists, other business people. Talk a little bit about the focus that you are putting not just on developing this capacity, but then on making it useful and available to those who might apply it in ways like Mr. Scarafino, for example, was talking about.

Dr. WADSWORTH. Yes, indeed. First of all, we built a facility, which is a beautiful building which can house a world class computer. And this is important when we are recruiting. Having a program that is sustainable, world class, cutting edge in a facility that looks like the world's leading capability is an important tool for bringing in the best minds in the country and from around the world. So part of the plan was to build a facility that has the ability to be expanded, that can allow different contractors to compete for the next generation of machine, and we also adopted a notion from the large scale scientific facilities where we would have end stations or user stations.

So our model is to have seven or eight different scientific problems formulated and competed by the scientific community and those people, industry, university students, would come into the facility and execute their research on these so-called end stations of the computer. So our notion from day one was to have an open environment where we would attract people from all walks of the scientific community into Oak Ridge in a facility that was modern and was able to sustain change not only in a scientific agenda but also in the type of computing that would come along in years to come.

Senator ALEXANDER. Thank you.

Senator Bingaman.

Senator BINGAMAN. Thank you very much.

What we focused on, in this legislation and this hearing, is the capability that we are developing and already have in our Nation to do high-end computation. Clearly the extent of that capability is one indicator of how well we are doing in competition with others and in leadership in science and technology. I would think another good indicator of how well we are doing is who the people are who are standing in line waiting to use this new computing capability. I just wonder if any of you have any insight into that.

Are U.S. companies actually anxious to or interested in using this capability if we go ahead and develop this very advanced capability? Are foreign companies interested, more interested than U.S. companies, or is this strictly an academic kind of a thing or a national security kind of a enterprise that we are looking at here?

Dr. Wadsworth, maybe you have a view.

Dr. WADSWORTH. Yes, I can certainly attest to the degree of interest in the laboratory since we won this competition. We are engaged with numerous universities, numerous industries, and numerous other laboratories from around the world. Our challenge will be to find the most effective peer review process to get the best possible teams together to use the computer.

Senator BINGAMAN. But you are not concerned about any lack of interest by U.S. researchers.

Dr. WADSWORTH. Absolutely not. No. We are engaged with over 25 U.S. universities right now and many different industries, as well as computer companies themselves. There is no lack of interest at all.

Senator BINGAMAN. Let me ask another question. One of the big problems that we have created for ourselves—and maybe it is built into the real world environment we are in—is this distinction that we have built into all structures between defense-related research and non-defense-related research. Of course, NNSA is focused on the defense-related research and as Dr. Kusnezov just indicated, their work is of a classified nature and therefore they are not able to open up their computing capability for the use of others.

It seems like, though, in developing this tool that we are talking about, this high-end computing capability, we need to have very good cooperation and communication between the defense side and the non-defense side. I mean, if we have got the greatest concentration of high-end computing in NNSA, presumably there are some people within NNSA who know something about high-end computing. Of course, I am particularly interested because of Los Alamos Lab and Sandia Lab in my State.

To what extent can we be sure that there is a cooperative effort between the NNSA labs and the rest of the DOE labs in the development of this capability, and not only the development of it, but the use of it?

Dr. WADSWORTH. Not to take all the questions, but Los Alamos, Livermore, and Sandia are part of our proposal at Oak Ridge National Lab. I was at Livermore for 10 years, and our colleagues from Livermore visited us last week, as a matter of fact, at Oak Ridge. So we are sharing very much in that capability.

Senator BINGAMAN. Yes, Mr. Scarafino.

Mr. SCARAFINO. We had visited Los Alamos a number of years ago in order to get information on what kind of advancements they have been making, specifically in the parallel environment. We learned a lot from that. In fact, I think that probably gave us a 9 month or so advantage over our competition in being able to get a parallel processing environment up and running at Ford. So the information was very helpful and actually directly applied.

Although my emphasis here was pushing for high-end computers, faster unit processors, we do use a significant amount of the parallel type, the commodity. They are very difficult to manage because of just the high numbers of units and stuff like that. And we did learn a lot from Los Alamos, and it provided us a very useful and very helpful interchange of information.

Senator BINGAMAN. I just wanted to make the point, which I am sure everyone here is aware of, that when we established the NNSA as a separate unit within the Department of Energy, several of us expressed a concern that this might cordon off the laboratories that were going to be part of NNSA from the other scientific work that the Department was pursuing through the Office of Science and others. I am encouraged to hear that is not happening in this case, and I hope that is still the case.

Dr. Kusnezov.

Dr. KUSNEZOV. Thank you, Senator Bingaman.

I would like to comment a little bit on that. I think there are very good relations between the Office of Science and the NNSA, both in Washington and in the field, and there are some good examples about how people work together.

One thing to keep in mind is the research communities are typically pretty small and irrespective of where the people are found, whether in industry or in universities or at the labs, they tend to run into each other everywhere. So there is a very good communication network at really all levels.

With respect to examples of good collaboration, I think you could consider, for example, platform architectures. Part of the leadership class proposal now at Oak Ridge is going to include one of the machines that was developed in part with Sandia, the Red Storm architecture. Following that, the Sandia people, together with Oak Ridge and other labs, are working together with Cray for the next generation beyond that for the 1906 timeframe, the Black Widow. So there is very good work together of these people to push the architectures forward.

The types of communication networks we use on our computers as well, these 10,000 processor machines, require a certain type of communication. You have essentially 10,000 different computers or processors calculating something, and they have to send information back and forth to give you the final result. The message passing interface—the MPI it is called—is developed at Argonne in collaboration with the defense program labs. So you find it in many places and there are many success stories about how we work together.

Senator BINGAMAN. Mr. Chairman, I have one other question. I will just ask that if I could, and then I am going to have to leave.

Senator ALEXANDER. Go ahead.

Senator BINGAMAN. Dr. Reed, you referred to these strategic plans that should include at least five features. The second one you list here is sustain support for computational science grand challenge teams to create and use leading edge computational codes and to educate new generations of HEC users. Do those computational science grand challenge teams exist today?

Dr. REED. There are certainly some of them, and this touches on the interplay across the community. They are from the academic side as well as from both sides of the Department of Energy. There are lots of collaborations. Those teams have been funded from many sources. One of them has been funded out of defense programs at some of the university agencies. There are several examples there.

I think the message I would leave you with is that investment in high-end computing is a balanced process. The software, the architecture, and deployment of systems are critical, but as is the investment in people. Developing a large scale computational science code, one that will yield new scientific results either in an individual discipline or increasingly in an interdisciplinary world is a large scale enterprise, the development time to create these codes is measured in years. It is no longer a case that an individual researcher can create one in his or her laboratory. So the sustenance for that community is really critical if we want to use the machines. We can build a highway, but we need the cars to drive on it as well. The human component is the part that is renewable that allows us to understand the strengths and weaknesses of particular machines to develop the next generation of systems that will be more effective.

So there are some of those teams, for sure. We could benefit from additional investment in that, and that goes hand in glove with the investment in software and systems.

Senator BINGAMAN. Thank you, Mr. Chairman, very much for having the hearing and I thank all the witnesses.

Senator ALEXANDER. Thank you, Senator Bingaman. We will keep talking about this. I have maybe one more question. Then we will bring the hearing to a conclusion.

When I was in Yokohama a couple or 3 months ago being briefed on the Earth Simulator, the Japanese computer, my sense of things was that it was sold to the Japanese people and Japanese government primarily as a way to understand climate change, that that was the major use for it. That kind of high-end computing, as I understand, is not the only kind of high-end computing. There are different kinds of architecture.

Mr. Scarafino, there was some skepticism there that that sort of architecture would be very useful in manufacturing, in other words, that other kinds of architecture which already existed and might not require such an accelerated investment as we are talking about might be fine for designing automobiles, while we might need to catch the Earth Simulator to figure out climate change.

Now, it sounds today like you might not agree with that. What is your view on these different types of architecture?

Mr. SCARAFINO. Actually the Earth Simulator is made up of a classical design. They are made up of NEC vector computers, processors that are very similar to the C series and T series Crays that were made in the mid 1990's. They are very good general purpose processors. They can run at high utilization rates. Some problems run on these machines run at utilization rates are in the mid 30's percent-wise, which is a little over three times the type of efficiencies you can get in a typical off-the-shelf commodity-based cluster. So what the Japanese built was a machine capable of basically solving general purpose problems.

In addition to the climate aspect, they also were studying earthquake simulation too. But as far as it being a specialized machine only for climatology, I do not see that——

Senator ALEXANDER. So the effort we are describing you believe has a real relevance to our manufacturing and competitiveness in the United States.

Mr. SCARAFINO. Yes. The processors are expensive. They have a very good balance between processor speed and their access to memory, the memory bandwidth and latency. Also, being vector processors, vectors are kind of the first level of parallelism that is very highly efficient. So they did not invent a new architecture at all. They basically refined an old one and put together a very large machine. It has got over 5,000 processors in it.

Senator ALEXANDER. Dr. Wadsworth, would you have anything to add to that?

Dr. WADSWORTH. I think that was a good summary.

Senator ALEXANDER. Well, let me thank each of the five of you, as well as Mr. Decker for coming earlier. Senator Bingaman and I intend to continue to press to provide the support from the Federal Government to help the United States regain the lead in high-performance computing. We want to do that intelligently and we want to spend whatever money Congress appropriates as wisely as we can.

This hearing today has defined specific goals. It has given us a perspective from a broad variety of sectors. It has suggested that we can reach those goals and that the benefits would have broad implications, not narrow implications in America's society.

We have heard also that the Oak Ridge effort may be centered there but it is in partnership with other major laboratories, universities, and major businesses in the country and that a focus is being paid on making sure that whatever the results are they are broadly available in an easy way. And the facility is already built to help do that.

Dr. Kusnezov has said to us that the very important national security work we are already doing in high-end computing is a very busy operation, already using much of our capacity and that we need more. At least, there is not enough there to meet the demand that we have in the unclassified world, and there is no conflict with this effort and the effort that you manage. In fact, the two would work in parallel.

So this has been a very useful hearing. I thank you for your time.

The hearing is adjourned.

[Whereupon, at 4:03 p.m., the hearing was adjourned.]

○