Exascale Workshop Panel Meeting Report

January 19-20, 2010 • Washington D.C.

U.S. DEPARTMENT OF ENERGY

Sponsored by the Office of Advanced Scientific Computing Research

DISCLAIMER

On the cover: Oak Ridge National Laboratory's Cray XT5™ supercomputer. The computer, dubbed Jaguar, is the largest in the U.S. Department of Energy's Office of Science. The Jaguar's computing power makes it among the most powerful open scientific systems in the world. Future reports in the Scientific Grand Challenges workshop series will feature different Office of Science computers on their covers.

EXASCALE WORKSHOP PANEL MEETING REPORT

Report from the Meeting Held January 19-20, 2010

Sponsored by the U.S. Department of Energy Office of Advanced Scientific Computing Research

Exascale Workshop Panel

Jim Decker, Principal
Decker, Garman, Sullivan & Associates, LLC

Tom Everhart, President Emeritus
California Institute of Technology

Anita Jones, University Professor Emerita
University of Virginia

Arthur Kerman, Professor Emeritus
Massachusetts Institute of Technology

Cherry Murray, Dean
Harvard School of Engineering and Applied Science

Aristides A.N. Patrinos, President
Synthetic Genomics, Inc

Justin Rattner, Director and Chief Technology Officer
Intel Labs

Daniel Reed, Corporate Vice President
Microsoft Research, Technology Strategy, and Policy and Extreme Computing Group

Charles V. Shank, Director Emeritus
Lawrence Berkeley National Laboratory

C. Bruce Tarter, Director Emeritus
Lawrence Livermore National Laboratory

Alvin Trivelpiece, Director
Oak Ridge National Laboratory (retired)

Irving Wladawsky-Berger, Innovation and Technology Consultant–
IBM and Citigroup; Visiting Faculty–Massachusetts Institute of Technology and Imperial College

Biographical Information is featured in Appendix 1.

FOREWORD

Computers are pervasive in our daily lives. Without them, there would be no Internet. They are essential elements in home appliances, controlling our vehicles, and tracking personal finances. For many businesses, computers are an integral part of managing financial data, inventories, and other information systems.

Although not generally recognized, computers also play essential roles in all aspects of scientific research, engineering, and technology. Powerful computers and specialized software are used in varied applications such as the following:

- Designing aircraft without reliance on wind-tunnel testing
- Crash-testing simulation to design better automobiles
- Predicting the weather
- Designing new computers
- Designing computer chips
- Formulating everyday household products
- Enabling computer animation for movie and game production.

The initial drivers for development of high-performance computer systems designed for scientific and engineering computations were government research and development needs at the Department of Energy (DOE) and the Department of Defense. The DOE and its predecessor agencies, the Energy Research and Development Administration and the Atomic Energy Commission, were pioneers in using computers as essential tools in solving the problems associated with their missions in national security, energy, and key areas of science and technology. These agencies played a pivotal role in bringing about systematic improvements in the hardware, software, and related mathematics, thereby accelerating dramatically increased computational capability. Many of these government developments have found widespread use in the private sector where industry exploits them for commercial applications.

The use of all levels of computing is ubiquitous in academic research. Many innovations (e.g., in open-source systems) have proven beneficial to scientists and engineers at both the DOE laboratories and universities. This report outlines some of the key elements of this history of the various generations of supercomputers now used by DOE's scientists and engineers to solve the various problems associated with DOE mission responsibilities. In many areas of research, the ability to make effective progress has been limited by the speed at which individual calculations can be made and the amount of memory required for managing large amounts of data.

Historically, supercomputers have been defined as the fastest operating and the largest memory computers available at a given time. Over the past 40 years, the state of the art has advanced to a current level of "petaflops," which is shorthand for "10 to the 15th power (10^{15} or 1 thousand trillion) floating point operations per second." Although this remarkable petaflops capability has greatly advanced computer simulations in a number of fields of science and engineering, there are compelling needs to increase this operational speed by a substantial amount. Each new generation of supercomputers is usually designated by a thousand-fold increase in the speed of computing operations per second. The next generation will be an exascale computer operating at a speed of 10 to the 18th power (10^{18}) flops.

The need for exascale computing capability has been analyzed by groups of scientists and engineers who conduct the research and development related to the DOE's mission responsibilities through a series of workshops and town meetings that were documented in a series of reports. The preponderance of the information in this report is based primarily on these reports that were provided to the Exascale Review Panel (hereafter "Panel"). At its only meeting in Washington, D.C., on January 19, 2010, the Panel heard presentations that summarized key results of these reports. A list of these reports with Internet links is provided in Appendix 2. Information was also presented on the benefits of an exascale initiative to our national economy, international competitiveness, and national security, as well as the consequences of not moving forward with an exascale initiative.

No summary report can truly capture the benefits of pursuing an exascale computational capability initiative and the magnitude of the challenge required to achieve this next step in extreme-computing capability. The Panel was impressed with the quantity and quality of the efforts from scientists and engineers who have studied this need and made recommendations regarding the opportunities for advancement in their fields of investigation. While this effort has focused on DOE and its responsibilities, the Panel believes that exascale computing should be considered as a national challenge and opportunity. As with each previous new generation of supercomputers, a thousand-fold increase in computational capability would certainly bring substantial benefits to fields of science, technology, and engineering beyond those funded by DOE.

The negative value proposition is the risk that failure to proceed with an exascale development initiative on a systematic and sustained basis might cause other nations to build on our decades of developments, thereby enabling them to build a capability that competes with our industries. Nuclear power is an example where the United States failed to continue to capitalize on our dominance in this field and thereby lost our advantage as a supplier of those nuclear goods and services both to the U.S. and other nations. Today, some of these goods and services are purchased from other nations that previously bought them from the United States.

Two uses of the term "computing capability" are used in this report. There is the science and technology of computers, which involves the creation of components and their assembly into a computer system. There is also the science, engineering, and technology that can be accomplished using supercomputers. Both need to be advanced, and to some extent, each depends on the other. This report seeks to make the case that both are essential and interdependent on the other, and that the United States should consider a program to advance the state of the art to a new level consistent with national needs.

EXECUTIVE SUMMARY

In Washington, D.C. on January 19-20, 2010, a panel of 12 scientists and engineers with experience in government, universities, national laboratories, and industries met to review a collection of reports prepared to document the need for a new generation of extreme-computing capability for the U.S. Department of Energy's (DOE) missions. These extensive and comprehensive reports were prepared by groups of scientists and engineers in workshops and town meetings held over the past few years. Representatives involved in preparing these reports presented summaries of their efforts to the Exascale Review Panel (hereafter "Panel") regarding the need for a new generation of extreme computing at the exascale (10^{18} floating point operations per second, or flops).

The Panel was impressed with the breadth and depth of these studies. Based on these presentations and on prior study of the reports, the Panel believes that a compelling case has been made regarding the need for exascale computing capability as an essential asset for advancing the DOE's missions. To meet this need, the Panel believes the DOE should initiate a program to help ensure that exascale capability becomes available to its scientists and engineers in a timely manner.

Achieving exascale computing requires overcoming some significant technical challenges that cannot be accomplished with incremental improvements to current hardware and software. Exascale computer systems will likely have between 10 million and 100 million processing elements, which is 100 to 1000 times the number in today's petaflops machines. A simple model using a large collection of petaflop computers would lead to an electric power requirement that would likely exceed a gigawatt. Therefore, to achieve a realistic level of power consumption for an exascale computer, a substantial reduction in power consumption from today's petaflops computers would have to be achieved. This would require a significant reduction in the power consumption of the various individual components that comprise a modern supercomputer. New computer architectures will be required. Software for exascale computers will require new approaches to efficiently use such a massive number of processing elements.

Historically, the development of high-performance computing for government programs has been driven primarily by the need to solve the research and engineering problems related to national security requirements and DOE missions. In the past, market forces alone have not been sufficient to attract the private investments necessary to bring about major advances in supercomputer development and significant investment has been necessary.

The most recent TOP500 list (www.top500.org) shows that U.S. companies developed over 90% of the 500 largest supercomputers in the world. This is largely a result of the close partnerships these U.S. companies have had with the DOE in developing advanced supercomputers. On its own, the U.S. private sector might develop hardware technologies to assemble an exascale computer sometime in the 2020s. This would be at least 5 to 10 years later than what would be possible with a major DOE-led program. Maintaining U.S. leadership without a DOE-led initiative assumes that other nations will not mount their own aggressive effort to develop exascale capability. Past hubris-shattering experiences include the belief that no other nation would be capable of launching a satellite such as Sputnik before the United States. This is a good example to keep in mind.

Without a government initiative, computer vendors would likely focus on developing systems for business applications that represent a much better opportunity in a larger market. These would include systems involving more loosely coupled architectures for information analysis and cloud computing. It would likely be difficult and expensive to use these business systems to meet the more demanding requirements associated with the DOE missions.

Since much of the talent required to advance high-performance computing software already exists in the DOE laboratories, it is likely they will have to play a major role in any exascale development program. The DOE will have to bear a significant portion of the responsibility for the leadership and cost that will be required to achieve exascale capability to accomplish its missions.

The benefits to the nation of developing an exascale computing capability will extend well beyond the DOE's programs. Exascale computing capability will benefit other government agencies and a broad range of industries in this country, including energy, pharmaceutical, aircraft, automobile, entertainment, and others. More powerful computing capability will allow these diverse industries to more quickly engineer superior new products that could improve the nation's competitiveness. In addition, there are considerable flow-down benefits that will result from meeting both the hardware and software exascale challenges. These would include enhancements to smaller computer systems and many types of consumer electronics, from smartphones to cameras. Other countries recognize these benefits. Some have launched extreme-computing programs.

The key finding of the Panel is that there are compelling needs for exascale computing capability to support the DOE's missions in energy, national security, fundamental sciences, and the environment. The DOE has the necessary assets to initiate a program that would accelerate the development of such capability to meet its own needs and by so doing benefit other national interests. Failure to initiate an exascale program could lead to a loss of U.S. competitiveness in several critical technologies.

TABLE OF CONTENTS

FOREWORD V
EXECUTIVE SUMMARY VII
TABLE OF CONTENTS IX
INTRODUCTION 1
IMPORTANCE OF COMPUTATIONAL SCIENCE 1
ENERGY AND ENVIRONMENT PROGRAMS 2
- Combustion 2
- Nuclear Energy 3
- Biology and Biofuels 3
- Fusion 3
- Materials 4
- Climate Modeling 4

NATIONAL SECURITY PROGRAMS 5
HIGH-ENERGY PHYSICS, NUCLEAR PHYSICS, AND ASTROPHYSICS 6
BEYOND U.S. DEPARTMENT OF ENERGY MISSIONS 7
- Flow-Down Benefits to Smaller Computer Systems 7
- Hardware Flow-Down Benefits 7
- Software Flow-Down Benefits 8
- Flow-Down Benefits to Industrial Competitiveness 8
- Broader Benefits for Science and Engineering 9

EXASCALE COMPUTING CHALLENGES 9
- Hardware 9
- Software 10

AN EXASCALE INITIATIVE 10
U.S. DEPARTMENT OF ENERGY EXASCALE INITIATIVE 12
OTHER NATIONAL PROGRAMS 14
- China 14
- Europe 14
- Japan 14

CONCLUSIONS 14
REFERENCES 15
APPENDICES 17
APPENDIX 1: EXASCALE REVIEW PANEL MEMBERS .... APPENDIX 1-1
APPENDIX 2: LIST OF WORKSHOP AND TOWN MEETING REPORTS PROVIDED TO PANEL PRIOR TO REVIEW MEETING .... APPENDIX 2-1
APPENDIX 3: EXASCALE WORKSHOP PANEL REPORT AGENDA .... APPENDIX 3-1
APPENDIX 4: PRESENTERS AT PROGRAM BRIEFING .... APPENDIX 4-1
APPENDIX 5: ACRONYMS AND ABBREVIATIONS .... APPENDIX 5-1
APPENDIX 6: ACKNOWLEDGEMENTS .... APPENDIX 6-1

INTRODUCTION

The Exascale Review Panel (hereafter "Panel") consists of 12 scientists and engineers with experience in various aspects of high-performance computing and its application, development, and management. The Panel heard presentations by several representatives of the workshops and town meetings convened over the past few years to examine the need for exascale computation capability and the justification for a U.S. Department of Energy (DOE) program to develop such capability.

This report summarizes information provided by the presenters and the substantial written reports provided to the Panel in advance of the meeting in Washington, D.C., on January 19-20, 2010. The report also summarizes the Panel's conclusions with regard to the justification of a DOE-led exascale initiative.

- "Importance of Computational Science" outlines the history and importance of computational science as it has evolved since the 1940's Manhattan Project.
- "Energy and Environment Programs" describes how some of DOE's energy and environmental programs would benefit from exascale-computing capability.
- "National Security Programs" outlines how DOE national security programs have critically depended on the use of high-performance computing and their requirements for exascale-computing capability.
- "High-energy Physics, Nuclear Physics and Astrophysics" outlines some of the computing requirements of DOE programs such as high-energy physics, nuclear physics, and astrophysics.
- "Beyond DOE Missions" identifies some additional benefits to the nation in developing an exascale capability.
- "Exascale Computing Challenges" identifies some of the technical challenges associated with developing an exascale computing capability.
- "An Exascale Initiative" provides the rationale for government investment in an exascale initiative.
- "DOE Exascale Initiative" describes the rationale for an exascale initiative led by DOE.
- "Other Nations' Programs" identifies some of the major high-performance computing activities of other nations.
- "Conclusions" presents the Panel's principal conclusions.

Note: References have been left out of this text since they are in the written reports provided to the Panel and listed in Appendix 2 for those who wish to follow up on any details.

IMPORTANCE OF COMPUTATIONAL SCIENCE

Computational science and technology, begun during World War II, was a key element in the Manhattan Project for its development of nuclear weapons. Since then, through major advances in computers and computational techniques, computation has become an essential tool on par with experiments and analytical theory in solving many of today's highly complex scientific and engineering problems. In some cases, experiments are too costly or impossible to conduct, and computation is the only tool available for prediction and analysis. Computational science and engineering have advanced to the stage that they are used not only for advancing science and for solving important national security problems, but also by industry to design products such as automobiles, aircraft, and drugs, thereby reducing the time and cost of bringing new products to market. Despite the advances in scientific computing over the last 60-plus years, many important, complex problems cannot be realistically addressed using today's most powerful computers. Exascale computing could provide a problem-solving capability that is at least a thousand times more powerful than today's highest performance computers. This advance could facilitate more realistic and accurate simulations, including major advances from descriptive to predictive simulations.

ENERGY AND ENVIRONMENT PROGRAMS

The DOE faces significant challenges in its research and development programs that seek to develop technologies that could reduce dependence on foreign oil. The goal of reducing the emission of greenhouse gases is also important. The following are some examples of how development of energy technologies and research in some areas of basic science could be substantially advanced using application of exascale computing.

COMBUSTION

Burning coal and natural gas produces about 70% of our nation's electricity. Transportation accounts for 60% of national oil consumption. It may be possible to improve the efficiency of burning fossil fuels and reduce environmental impacts through a more detailed understanding of combustion processes that could enable the design of improved automobile engines and electricity-generating power plants.

For example, achieving advances in efficiencies and fuel options for both internal combustion engines and gas turbines will require new computational tools that can model complex chemical and turbulent fluid flows with much greater fidelity than can be obtained today. Promising new energy-efficient internal combustion engines provide opportunities for 25% to 50% improvements in efficiency using both advanced fuels and new low-temperature concepts. Approaches, such as low-temperature compression ignition with homogeneous-charge compression ignition, offer diesel efficiencies with environmental emissions similar to today's gasoline-powered vehicles. Compared with today's engines, these new approaches operate under extreme conditions for combustion involving high-pressure, low-temperature, fuel-lean, oxygen-rich conditions. These types of concepts will necessarily rely heavily on control systems based on a fundamental understanding of combustion science. However, these concepts also require operation with unexplored thermodynamic environments and new chemical fuel properties, resulting in new regimes of complex interactions among multiphase fluid mechanics, thermodynamic properties, and chemical kinetics. Effective design of engines operating in these new regimes will require new computational tools that provide a much greater level of fidelity for these highly complex chemical and fluid behaviors. Producing a validated predictive simulation capability that can be used by industry in the design cycle requires exascale capability to incorporate the combined complexities of chemistry and fluid flow. Having that design tool will allow industry to reduce the development time for significantly more efficient engines.

NUCLEAR ENERGY

Nuclear fission provides about 20% of electricity generation on a world scale without direct emission of greenhouse gases. Nuclear fission will be an important future contributor to the nation's goals of energy security and reduced greenhouse gas emission. Exascale computing could accelerate the design and deployment of new reactor systems that are cost effective to construct and operate, safe and secure, and proliferation resistant. An exascale simulation tool for designing new reactors would provide high fidelity, coupled simulations of neutronics, structure, and thermal hydraulics. In addition, an exascale capability would enable predictive materials analysis tools for fuels, cladding, and reactor vessel welds, all of which are needed for increased fuel utilization, power upgrades, and reactor life extensions.

BIOLOGY AND BIOFUELS

The genomics revolution, driven primarily by tremendous advances in the rate and quality of DNA (deoxyribonucleic acid) sequencing, is flooding public and private databases with genome sequence data. Besides sustaining the exceptional progress in understanding basic biology, rapid and efficient access to these data will enable significant contributions to society in medicine, agriculture, bioremediation, biofuels, and other energy applications. The ability to compare DNA sequences—the process referred to as "comparative genomics"—is at the center of contemporary biological discovery. These comparative studies also build on computational advances for understanding evolutionary relationships among species or phylogenetics. The ability to advance comparative genomics and the underlying phylogenetics is increasingly being compromised by limitations in scientific computing, particularly because of necessary associated data ("annotations"). As a consequence, exascale computing is required to address the challenge of comparative genomics and its extensions. Implementing high-performance computation and advanced information technology (IT) infrastructure to characterize the data flowing from today's experimental technology, such as the data from analyzing proteins and metabolism on a genome scale (proteomics and metabolomics), will provide far richer associated data, provide deeper insight, and establish a research platform for applications of the sequence data to environmental and energy challenges. Using the sequence and associated data for further analysis, such as how biological networks operate (the use of comparative genomics at a comprehensive scale), is opening up simulation and modeling applications that will rapidly enhance in-depth understanding across the biological sciences. This level of understanding underpins the development of synthetic genomics, which is the engineering or design of microbes to carry out novel tasks such as effective bioremediation under tight human control or biological approaches to the challenge of carbon sequestration. Thus, biologists can only specify how to re-engineer biological networks by computing at an exascale level on the emerging phylogenetic and systems biology knowledge base. Doing so will enable the design of next-generation microbes that effectively and economically produce biofuels, implementing the goals of bioenergy in reducing dependence on fossil fuels.

FUSION

Theoretically, fusion energy based on either magnetic or inertial confinement has great promise as an energy source. However, it is one of the most difficult scientific and engineering challenges ever attempted. Producing and maintaining magnetically confined fusion plasmas at temperatures similar to the interior of a star results in a host of complex physical processes that occur on vastly differing length and time scales. Some of these phenomena include turbulent energy transport, magnetohydrodynamic

(MHD) instabilities that can drive major plasma disruptions, local plasma heating with radiofrequency waves for plasma control, and plasma interactions with chamber walls. To obtain a predictive capability for plasma performance, these very different physics regimes must be coupled. Exascale computing is necessary to incorporate and couple all of the complex scientific phenomena for interpreting data from ITER and for eventually designing a cost-competitive fusion reactor system.

Inertial confinement using lasers or Z-pinches to compress a small pellet of fusion fuel also involves many complex phenomena over micro- to meso-time and length scales. Some of these phenomena include nonlinear optics of plasmas, plasma absorption of laser energy at the intensity frontier, and generation of large currents and associated magnetic fields.

A complete modeling of inertial confinement requires coupling of hydrodynamic compression with full kinetic and relativistic models. This complexity of modeling requires exascale capability. If this capability were available, it would allow much more accurate simulations of promising approaches, such as fast ignition, that could significantly advance inertial confinement fusion as an energy source.

MATERIALS

Materials are crucial to all aspects of energy production, distribution, and utilization. It is possible to design many components for energy production and use, but it is currently impossible to design the materials from which those components are made. For example, it is possible to design an airplane wing on a computer with computational fluid dynamics but currently impossible to design the material from which the wing is made. Like so many other important problems, the fundamental equations that govern materials properties have been known for years. The challenge is to bridge the enormous gap in length scale between atoms and the macroscopic length scales of practical materials applications. The general approach to bridge this gap is to connect different computational methods at different length scales and multiscale modeling techniques. Current petascale computers are not powerful enough to accurately predict the properties of materials and provide the desired design tool.

Exascale computing will enable analysis and design of materials at the nanometer scale, improving both capability and reliability of materials used in nuclear reactors and photovoltaic power generators, as well as microcomputers and the devices that will be used in the second generation of exascale computers, to name only a few. Such calculations will enable modeling of environmental degradation under extreme conditions (e.g., in fission and fusion reactors) and over long time scales (i.e., nuclear weapon reliability in the absence of testing). They will also afford analyses of material changes at much shorter time scales; e.g., improving understanding of femtosecond processes in chemistry as well as the design of electronic circuits. Exascale calculations potentially could lead to a thorough understanding and design of high-temperature superconductors for greatly reduced energy losses in electricity transmission.

CLIMATE MODELING

The challenge of understanding the Earth’s climate and its evolution is based partly on computers and codes that seek to model its behavior. Current debate suggests the ability to make regional predictions on the scale of decades with desirable certainty remains a goal. The report prepared by scientists in this field and briefed to the Panel makes a compelling case for exascale computing capability to build on previous work and advance the state of knowledge in this discipline to the level needed to adequately and accurately understand the processes and forces that govern the Earth's climate. Given the scale of investment that the nation and globe will possibly be making in the next few decades to reduce greenhouse gas emissions, it is also imperative that uncertainties are quantified in the complex models of the Earth's climate system.

NATIONAL SECURITY PROGRAMS

Computations at the exascale are required to maintain U.S. national security. In particular, as a result of adherence to the tenets of the Comprehensive Nuclear Test Ban Treaty, the United States is no longer conducting nuclear explosions to help ensure our nuclear stockpile is safe, secure, and reliable. Instead, this assurance now is based in large measure on the current generation of petascale supercomputers used in concert with nuclear test history and the results from laboratory scale facilities. In the absence of nuclear explosive testing, advanced computing serves as the method for assessing the impact of a nuclear weapons surveillance finding, evaluating strategies to address those findings, and designing the life extension program for a warhead in the stockpile.

Following the cessation of nuclear weapons testing in 1992, the primary goal of the required advanced computing program was to establish a high-resolution, three-dimensional weapons simulation capability. Therefore, such capability was needed. Assessing aging and other stockpile issues tend to involve "off-center" defects such as corrosion and cracks. A simple estimate shows that an increase in computer capability of 10 to 100 thousand was needed to perform the job (a factor of 1000 for the third dimension and another factor of 10 to 1000 for the science needed to describe small-scale aging effects). Twenty years later, this increase in computing capability was achieved at the petascale level, and our nuclear weapons stockpile continues to be certified as safe, secure, and reliable.

However, the experience of the past 20 years has also shown the need for additional capability. The events of September 11, 2001, led to a need for higher levels of intrinsic safety and security for nuclear warheads. The requirement to conduct life extensions into the foreseeable future has increased those demands. Specifically, a first-principles simulation capability, rather than one based in part on historical data, is now needed. In practice, this means developing full-scale three-dimensional weapons design codes that can model proposed safety and security concepts and predict the performance of life extensions on a sound scientific basis. To effectively execute full-scale computations, which currently can take months to perform (when possible at all), these analyses will require an increase in speed of factors on the order 1000 (petascale to exascale).

Demands on the reliability of these simulations will become even more severe under any future arms control agreements. As the types and quantities of stockpile weapons decrease, there will be a need for more sophistication and quantification of uncertainty in the models as the relative importance of each element in the stockpile increases. In addition, a large suite of complex simulations is required for nuclear forensics (to support counter-proliferation efforts) and possible eventual Comprehensive Nuclear Test Ban Treaty verifications, making exascale computing an imperative for rapid response to national needs in these areas.

HIGH-ENERGY PHYSICS, NUCLEAR PHYSICS, AND ASTROPHYSICS

Scientists in high-energy physics, nuclear physics, and astrophysics seek to understand the fundamental building blocks of matter and unlock the secrets of how the universe was formed. Although separate workshops were held for these fields, their needs for exascale-computing capability to advance their respective frontiers are similar and compelling. DOE has provided the major source of funding for high energy and nuclear physics dating back to the Atomic Energy Commission. DOE does not have a major responsibility for astrophysics. However, fundamental knowledge that comes from high energy and nuclear physics is essential to understanding the universe. Conversely, observations from astronomical instruments have led to discoveries of the unexplained phenomena of dark matter and dark energy. The physics of high-density, high-temperature materials—such as those found in stars—has many similarities to physics encountered in nuclear weapons design, which is an essential national security mission of DOE.

These fields of science have made extensive use of supercomputers to advance our understanding of the fundamental building blocks of matter at the smallest conceivable scales and to understand the largest imaginable scales that govern universe behavior. The workshops amply demonstrated the requirements for exaflops computing to advance each of these fields, including design of accelerators, analysis of massive amounts of data, developing a predictive capability from theoretical models, and simulation of astrophysical phenomena such as supernova explosions.

Experiments that explore the fundamental constituents of matter and other phenomena at today's very powerful accelerators—such as the Tevatron, the Relativistic Heavy Ion Collider, and the Large Hadron Collider—produce petabytes to exabytes of data. The World Wide Web was developed in part to allow the efficient transfer of data and its analysis from the accelerator sites to physicists around the world. Analysis of such large data sets currently requires powerful computers and eventually will require exascale computers. Similarly, large data sets will be produced from future dark matter/dark energy searches in the universe and will also require exascale computers for analysis.

Extending the energy frontiers of high-energy physics research is crucial for further progress. However, the cost of extending the energy frontier based on current accelerator technology appears daunting. Accelerator simulations play an important role in design optimization for cost and risk reduction. In the long term, new accelerator technologies will likely be necessary. Massive computations are required to model many important accelerator dynamics such as collective beam effects, the properties of complex electromagnetic structures, and plasma and laser wakefield acceleration techniques. Exascale computing will allow designers to decrease the time required to optimize accelerators by decreasing the frequency of trial-and-error procedure for producing prototype components. Accelerator technology has many applications beyond high energy and nuclear physics, and advances in the technology can lead to new applications and reduction in the cost of present applications. For example, accelerators are employed in the production of isotopes used in nuclear medicine such as those used for positron emission tomography (PET). Additionally, accelerators are used for producing radiation beams used in cancer treatment.

Another key tool in developing an understanding of the fundamental building blocks of matter is the use of simulations based on lattice quantum chromodynamics theory. These simulations have the potential of computing fundamental properties of matter, but they require exascale computing to do so.

Simulation is an essential tool for astrophysicists to develop an understanding of the formation and evolution of the universe on all scales. The largest structures in the universe are 22 orders of magnitude larger than the smallest compact objects that are one-hundredth of the size of Earth but more massive than the sun. Understanding the universe at each size scale and tying these size scales together presents an exascale computing challenge.

BEYOND U.S. DEPARTMENT OF ENERGY MISSIONS

While the value proposition to DOE missions alone would favor proceeding with an exascale program, the Panel believes, in addition to meeting DOE programmatic needs, such an effort also would be of substantial value to the science and engineering programs of other government agencies. An accelerated federal-private industry partnership in a sustained research and development (R&D) effort to create working exascale solutions at reasonable power densities will simultaneously lead to significant value across many segments of the nation's economy. United States onshore R&D and manufacture of exascale computing environments—hardware and software—ahead of international competition would help trade imbalance with positive net exports of smaller petascale and terascale computers based on exascale technologies. Even greater benefits are expected beyond electronics and computers. Some other nations have recognized the benefits of access to extreme-scale computing capability and have initiated aggressive programs in high-performance computing (described in "Other Nation's Programs").

FLOW-DOWN BENEFITS TO SMALLER COMPUTER SYSTEMS

It is important not to underestimate the impact of exascale breakthroughs to more capacity-oriented machines, as well as to smaller machines that share the same technologies, architecture, software, and applications. Exascale innovations will be widely used in the more commercially oriented cloud computing systems, whose requirements are quite similar, especially the need for low-power, low-cost components and highly efficient, autonomic systems management. Many of these same innovations should also yield relatively inexpensive petascale-class systems, as well as smaller ones. Expanding access to such families of systems should lead to richer overall ecosystems, including applications, users, and technologies.

HARDWARE FLOW-DOWN BENEFITS

While the development of exascale systems faces many technical challenges, successfully meeting these challenges will have benefits far beyond the development of exascale systems. For example, the need for dramatic reductions in the amount of energy expended per computation required to keep exascale systems at acceptable power levels is a serious matter and is discussed in a subsequent section. However, an exascale solution to this problem will directly impact future electronic systems at the opposite end of the performance scale, where the volume of systems sold is greater by a factor 10^8 or even 10^9 as some have predicted. Handheld computers or "smartphones" and other battery-operated information devices have as great a need for energy-efficient computing advances as the exascale systems. To achieve an order-of-magnitude or more improvement in energy efficiency in the low-volume exascale computers, these improvements will be immediately embraced and adopted in the extremely high-volume systems, such as phones, cameras, media players, and a variety of embedded devices. The "Internet of Things" may depend as much on such energy-efficient technology as do the handful of exascale systems. A 100-Gflops smartphone with an array of near-human perceptual abilities may be commercially available shortly after the first exascale machine is powered up.

SOFTWARE FLOW-DOWN BENEFITS

Exascale software advances also could find wide application outside of the extreme-scale applications they will enable. As more and more computing moves to the cloud, where vast server farms consisting of millions of processors provide an array of sophisticated services to client devices, the software for managing the tens or even hundreds of millions of processor cores in those servers may be direct descendants of the software used to manage the millions or even billions of compute elements to be found in exascale systems. Such software will be capable of providing an apparently reliable computing surface in the face of nearly continuous failures in the underlying hardware. If they are to take full advantage of the extremely low-power circuits developed for the exascale systems, cloud server farms will require similar software. The extremely low operating voltages expected in these circuits exacerbate circuit reliability issues as noise margins drop to near zero. To provide programmers and users with the kind of reliability and availability expected of such systems, a sophisticated combination of hardware and software reliability techniques will be necessary in both classes of extreme-scale systems.

FLOW-DOWN BENEFITS TO INDUSTRIAL COMPETITIVENESS

Two types of flow-down benefits to the nation's industrial competitiveness are expected from exascale development. The first of these comes directly from the family of exascale technologies that must be developed to make such systems a reality. For example, the advent of new generation of dynamic random access memory (DRAM) devices and a new generation of nonvolatile random access memory (NVRAM) devices will be critical to exascale storage. In both areas, most of the manufacturing base for both kinds of devices has left the United States. The leading manufacturer of both kinds of memory chips is now located in South Korea. Exascale provides the United States with the opportunity to bring the memory chip business back onshore if the support and incentives are in place during the exascale development phase. Again, the hardware technologies driven by the extreme needs of exascale will find wide application at the opposite end of the power and performance spectrum.

The other flow-down benefit to industrial competitiveness comes from the availability of lower-cost, high-performance computers based on exascale technologies. As more industries integrate high-performance computing into their research, development, and manufacturing processes, the need for more affordable high-performance computer systems likely will continue to grow. A compelling example, which forms a virtuous cycle with the direct application of exascale technology, is the growth of computational photolithography in the integrated circuit industry. The industry has turned to computational techniques to deal with the errors introduced when the features being exposed are smaller than the wavelength of the light source. The errors can be corrected by computing the amount of distortion needed in the mask patterns to allow the feature to be correctly exposed. While the high-performance computer systems used to compute these mask corrections are not among the TOP500 list of systems, it is assumed that several of them would make the list. As feature sizes continue to shrink, computational needs will continue to grow to exascale levels.

The preceding example provides a glimpse of the importance of high-performance computing to one industry—electronics. Similar compelling cases can be made for other industries, including nuclear energy, aircraft, automobile, pharmaceutical, and entertainment, where enhanced computing capabilities will enable better engineering of new and improved products for a competitive advantage. The Council on Competitiveness has studied the benefits of high-performance computing for a number of U.S.

industries and has produced several reports (see Council on Competitiveness 2010), including *U.S. Manufacturing—Global Leadership Through Modeling and Simulation* (Council on Competitiveness 2009). The conclusion section of the abovementioned report is entitled, "For U.S. Manufacturers to Out-Compete, They Must Out-Compute."

BROADER BENEFITS FOR SCIENCE AND ENGINEERING

Exascale computing will help address a class of highly complex problems that have been beyond solving, not just due to their sheer size, but also because of their inherent uncertainties and unpredictability. The way to deal with such uncertainty is to run multiple ensembles or copies of the same applications simultaneously using many different combinations of parameters to explore the solution space of these otherwise unpredictable problems. This will permit searches for optimal solutions to many problems in science and engineering, as well as enable calculation of the probabilities of extreme events. This new style of predictive modeling is necessary for many of the energy-oriented problems already discussed. However, beyond science and engineering, there are many disciplines that will benefit from such predictive capabilities, from economics and medicine to business and government.

EXASCALE COMPUTING CHALLENGES

HARDWARE

There are numerous challenges in exascale system development. None is greater than the challenge of reducing the amount of electrical power required to operate them, including the refrigeration to cool them. The goal of developing exascale computers will not be accomplished with incremental improvements to current hardware. Exascale computer systems will likely have between 10 million and 100 million processing elements, which is 100 to 1000 times the number in today's petaflops machines. A simple model using a large collection of petaflop computers would lead to an electric power requirement that would likely exceed a gigawatt. To achieve a realistic level of power consumption for an exascale computer, a substantial reduction in power consumption from today's petaflops computers is essential. This would require a reduction in the power consumption of the various individual components that comprise a modern supercomputer. The required reduction in power consumption per operation is estimated to be in the range of 300 times less than that required for today's petascale systems that consume roughly 7 megawatts (MW) of power and exascale data centers that are expected to have a practical power limit of about 20 MW.

Historically, voltage scaling has been the principal means of reducing the power of integrated circuitry because power is directly proportional to the square of the supply voltage. The unreliability of conventional integrated circuits when operated at voltages affecting the threshold (or turn-on voltage) of the transistors also impedes further voltage reductions. New circuit designs are required, which compensate for the reduced operating margins and allow reliable circuit operation. The incentive to operate close to the threshold voltage is great. A nearly ten-fold improvement in energy efficiency over today's integrated circuits is possible.

There are other hardware factors at work that would further complicate the design of exascale systems. With the continued reduction of transistor dimensions expected in the latter half of this decade, transistor characteristics are expected to be extremely variable. The high degree of variability will cause many circuits to fail occasionally or consistently. While circuit design can compensate for some of the device variability, it cannot fully eliminate it. Techniques will be needed to manage the errors that escape the circuit level. It may even be necessary for programming techniques and operating system designs to handle those errors that are behind hardware detection and correction mechanisms.

Memory design is another area replete with engineering challenges. Today's commonly used memory, DRAM, has designs and interfaces that are quite energy inefficient. DRAMs must be redesigned to increase efficiency within the storage arrays and dramatically reduce the amount of energy needed to move data between the memory devices and the processor.

Additional hardware challenges include the move to solid-state disk storage, the use of on-die voltage regulation, and the transition to photonic interconnection networks for moving the massive amount of information between processor-memory elements and bulk storage and networking devices. In total, the hardware challenges of exascale dwarf those faced in moving from gigascale to terascale in the 1990s and from terascale to petascale in the last decade. A large and sustained hardware research effort in the first half of this decade will be required if exascale systems are to become a reality in the second half of the decade.

SOFTWARE

The effective control and operation of exascale systems pose daunting and diverse new software challenges, including resilience and reliability management, energy control, resource scheduling, input/output (I/O), and programming. The potential to successfully leverage mainstream personal computer (PC) and server software to address these challenges is doubtful.

As a motivating example of the need for alternative software approaches, consider systemic reliability. Several DOE and Defense Advanced Research Projects Agency (DARPA) studies have estimated that exascale systems will need billion-way parallelism to deliver exascale performance. Given the large number of components at this scale, permanent and transient hardware failures will be common, and the system and application software will need to detect, recover from, and continue operating across these failures.

Such scaling issues from exascale systems are unlikely to be addressed by the broad open-source software community or by the commercial software industry without government investment. The market is simply too small and the problems too irrelevant to the mainstream.

For the same reasons that incremental extension of PC-class hardware technology will not enable exascale, incremental extensions to existing software approaches are unlikely to yield adequate solutions. Instead, some fundamental software assumptions must be reconsidered and new systems designed as when the transition from vector supercomputers to message-passing systems occurred.

AN EXASCALE INITIATIVE

Historically, federal government programs have driven the development of high-performance computing for science and engineering. Several government agencies have had the mission need (often led by national security requirements, especially those within the nuclear weapons program) and the vision to invest heavily in high-performance computing. Unquestionably, without the federal government's investments, computational science and engineering would not have advanced to its current state, where it is considered an essential tool for advancing many areas of science and technology. The federal government has made investments in the development of advanced computer hardware, often in collaborative and partnership relationships with the computer vendors who carried out much of the hardware development burden. In parallel, the federal government has made large investments over the years in applied mathematics, applications and systems software, and other essential tools for using the new computers.

Throughout the history of high-performance computing, which dates back to the 1940s, the government has been the primary market for the highest-performance computers. Software and hardware developed for government missions have migrated to industry primarily for engineering applications, usually with a lag time of several years. Industry purchases of high-performance computers have focused on smaller machines that require much smaller capital investments. In most cases, these smaller computers were the previous high-end computers developed for government applications. The high-performance computer market for science and engineering remains a relatively small segment when compared with computers designed to handle business applications, personal computers, and embedded computers in consumer and military products. For computer vendors, high-performance computers designed for science and engineering applications do not offer sufficient sales revenue to justify the required large investment in the development of specialized hardware and software. Only a few computer companies have been willing to undertake the risk and long development time inherent in the creation of very advanced computers.

The advance to exascale computing will require overcoming significant technical challenges (as summarized in "Exascale Computer Challenges"). Only the federal government can mobilize the technical talent and make the required investments to achieve an effective exascale computing capability in a timely manner. Moreover, only the federal government's missions can mobilize the nation's capabilities into a program where a broad range of scientific applications would drive development of computer software, architectures, and instruments to produce both the necessary computer hardware and a problem-solving environment that is optimized for computational science and engineering. It is important to recognize the benefits to the federal government and the nation would be significant in achieving this goal of optimized hardware and software.

The effectiveness and efficiency with which scientists and engineers develop applications codes to use high-performance computers is strongly enhanced by having an optimized computational environment. An optimized computing environment that is user friendly can significantly reduce the cost of code development, reduce time-to-solution, and improve interpretation of results. This would also greatly accelerate the rate at which U.S. industry will be able to use the next generation of high-performance computers to improve its competitiveness in global markets. In the absence of a government initiative, the next generation of computers would most likely be driven by business applications and the technical community would be left to adapt whatever computers emerge to solve their problems, which would result in costly inefficiencies to the government and industrial users. This does not take into account that some nations, such as China (see "Other Nation's Programs"), are devoting sufficient effort to develop and market supercomputers that have capabilities projected to equal the capabilities of U.S. supercomputers. While it might be inconvenient, such competition should not be ignored.

U.S. DEPARTMENT OF ENERGY EXASCALE INITIATIVE

DOE and its predecessor agencies, the Atomic Energy Commission and the Energy Research and Development Administration, have a long history of commitment to computational science and engineering to aid in solving its many complex mission problems. DOE's commitment has led to sustained, large investments in all aspects of high-performance computing, including applied mathematics, applications and systems software, computer architecture research, and acquisition of leading-edge supercomputers. DOE has been an early adopter of new supercomputer hardware technologies, buying several of the first in each new generation of supercomputers. The early driver for DOE's investments in high-performance computers was the design and testing of nuclear weapons. These investments in supercomputers started with the purchase of the Univac LARC (Livermore Advanced Research Computer) in 1960 by Lawrence Livermore National Laboratory. The LARC had a peak speed of 250 kiloflops (or 250,000 floating-point operations per second)—an impressive number for its time. Since 1960, DOE has continued to stay at the leading edge of supercomputer technology through its purchases of hardware from IBM®, Control Data Corporation, Cray®, and Intel.[1] As transitions were made from fast scalar computers to vector architectures and from vector computers to today's massively parallel computers based on commodity processors, DOE laboratories developed much of the software required to use each of the new advances in computer hardware.

Over the years, use of DOE supercomputers spread from applications in nuclear weapons to fusion research, chemistry, combustion, materials, nuclear physics, high-energy physics, accelerator design, climate modeling, and biology. In the mid-1970s, the Magnetic Fusion Energy Computer Center (MFECC) was established to provide access via a broadband network to fusion researchers in national laboratories, universities, and industry. This was the first time a supercomputer had been made available to a geographically dispersed research community. Until that time, supercomputers were only available to scientists and engineers at a few government laboratories. The MFECC was later expanded to provide supercomputer access to researchers across the broad set of research fields funded by DOE's then-named Office of Energy Research. MFECC was then renamed the National Energy Research Supercomputer Center (NERSC).

In recent years, DOE's leadership has continued through substantial investments by both the National Nuclear Security Administration (NNSA) and DOE's Office of Science. NNSA's Advanced Strategic Computing Initiative (ASCI) successfully led the development of terascale and petascale computers that are an essential part of NNSA's Stockpile Stewardship Program, which ensures the reliability and safety of the nation's nuclear weapons. Advancements in computer hardware driven by the ASCI program also have benefited many other R&D communities.

The ASCI program successfully expanded many of the elements of the DOE computing efforts from those earlier eras. Beginning in 1995, it set a clear goal of increasing computing power by more than 10,000-fold in the next decade to meet its mission of a safe and reliable nuclear weapons stockpile without nuclear testing. ASCI formed robust partnerships with the computer manufacturers at all levels, so procurements became joint-development projects rather than simply purchases. It established scores of focused academic partnerships, some in the several millions of dollars per year range, to pursue many of

---

[1] IBM is a registered trademark of IBM Corporation, U.S.
Cray is a registered trademark of Cray, Inc.

the software and application challenges and cultivate talented, young scientists in the supercomputing field. In addition, it developed clear leadership at DOE Headquarters in concert with Laboratory leaders who had full responsibility for implementing the program. The result exceeded expectations—petascale computing capability; a revitalized supercomputer industry; and resolution of many nuclear weapons issues that have maintained a safe, secure, and reliable stockpile.

While advances in computer hardware generally receive the most attention, DOE's long-term investments in applied mathematics and computer science provide the underlying computational techniques that have contributed substantially to increased computational capabilities. Developments of new computational techniques have sometimes provided larger advancements in computational speed than those produced by new hardware. In recent years, DOE's Office of Science has formed partnerships between the Advanced Scientific Computing Research program and its other scientific programs to bring together applied mathematicians, computer scientists, and applications scientists into teams to develop efficient applications codes that take full advantage of today's supercomputer capabilities. The program, called Scientific Discovery through Advanced Computing (SciDAC), is often regarded as a model program for scientific computing. It further demonstrates DOE's mission-driven leadership in computational science and its understanding of managing the multidisciplinary approach required for computational science. Within its laboratories, DOE currently has the most powerful collection of supercomputers of any science and technology organization in the world. Furthermore, located primarily in its laboratories but also among its university community, DOE has a large, talented collection of computer scientists, applied mathematicians, applications scientists and engineers, and scientific computer software developers who have decades of experience working together to solve critical problems.

Achieving exascale performance presents a number of difficult technical challenges. New approaches will be required for both hardware and software. While possible avenues to solutions have been identified, research is required to determine the best solutions. Disciplined management will be required to conduct down selections among competing approaches. Critical path items will have to be identified. Resources must be appropriately focused. Schedules will need to be established and maintained within the uncertainties associated with R&D. DOE has a long history of both managing research and large, complex scientific projects that require R&D to extend the state of the art. Two recent examples of successful large science facility construction projects with significant R&D requirements are the National Ignition Facility and the Spallation Neutron Source. The success of these projects is a demonstration of the type of management skills found at both the DOE laboratories and its headquarters that are required for a successful exascale initiative. The DOE capabilities and its longtime commitment and leadership in advancing computation make it the federal government agency to lead and manage an exascale initiative.

Other federal agencies such as DARPA, the National Security Agency (NSA), and the National Science Foundation (NSF) also have capabilities that can make important contributions to the success of an exascale initiative. DARPA and NSA have historically invested in advancing specific hardware architecture, although both have specific areas of application. The NSF has traditionally provided high-performance systems to the computational science and engineering community to advance both system software and applications. In addition, DARPA and the NSF have been instrumental in creating broadly accessible networks that richly connect the research community. Today, DARPA is fielding a program to increase the *productivity* of high-performance computing. This involves measuring performance, programmability, portability, and robustness of high-performance computations. This effort will be quite complementary to a DOE exascale initiative. Of note, the NNSA, DOE Office of Science, and NSA all are partners in the DARPA high-productivity program.

OTHER NATIONAL PROGRAMS

Other governments have clearly recognized the value of developing exascale-computing capabilities (briefly described in the following sections). The primary motivation for these countries is to use supercomputers to drive industrial competitiveness.

CHINA

China is vigorously pursuing supercomputer technology with state investments. The intent is to use supercomputers to drive industrial competitiveness. By November 2010, it is expected that China will have a 3- to 6-petaflops computer largely made by integrating U.S.-developed components. This Chinese-developed computer will likely be classified as the most powerful computer in the world at that time. The fact that in the same timeframe China is likely to have a 1-petaflops computer made entirely from Chinese technology is, perhaps, even more noteworthy.

EUROPE

Europe has initiated three supercomputer activities.

1. The Partnership for Advanced Computing in Europe (PRACE) is a program to install a number of petascale supercomputers with the European Union in the short term. The systems to be deployed later in the cycle likely will be in the range of 10 petaflops.
2. The second European program targets commercial use of supercomputers and is intended to improve the competitiveness of European companies. Centers of applied research have been established in the United Kingdom, France, and Germany to provide support to the automotive, aerospace, and maritime industries.
3. The third is a longer-term effort to advance supercomputer capabilities and is likely to focus on developing exascale systems.

JAPAN

The Japanese government had declared its intent to produce at least a 10-petaflop-computer system by 2011-2012. The government will provide one company with $500 million to complete this effort. Japan's intentions with regard to exascale computing are unknown.

CONCLUSIONS

The key finding of the Exascale Review Panel is that there are compelling needs for exascale computing capability by the DOE's mission programs in energy, national security, fundamental sciences, and the environment. The DOE has the necessary assets to initiate a program to accelerate the development of such capability to meet its own needs and, by so doing, benefit other national interests. Failure to initiate such a program could lead to a loss of U.S. competitiveness in several critical technologies.

The development of exascale computing will provide an essential tool to accomplish a number of DOE programs in energy, national security, and basic science. In addition, there are substantial benefits to be gained for the nation's economy from developing an exascale capability.

Other nations have recognized the economic value to be gained from development of extreme-computing capabilities and have initiated aggressive, government-funded programs.

The development of exascale computing for computational science and engineering research requires overcoming significant technical challenges and will require a partnership between the government and industry.

DOE is the logical agency to take on the challenge of developing an exascale-computing capability. This would be the appropriate way to pursue the challenge and advance all elements at once, thereby managing a balance between hardware, networks, power consumption reduction, operating system and communication control software, disciplined operating environments, and applications across broad areas. The ASCI and SciDAC activities have yielded not just technical results, but mature teams of people with the capability to pursue the multiple elements that must be balanced to achieve useful exascale computation in a timely way. The DOE laboratories have critical expertise in mathematics, computer science, and multiple applications discussed earlier in this document. No other agency has the breadth, critical mass, or recent management experience to accept this challenge. Other agencies will continue on their paths of complementary (but different) programs and likely will contribute on multiple dimensions; define and measured performance; programmability; portability; robustness; and, ultimately, productivity in the high-performance computing domain.

REFERENCES

Council on Competitiveness. 2009. *U.S. Manufacturing—Global Leadership Through Modeling and Simulation.* Last accessed June 18, 2010, at http://www.compete.org/images/uploads/File/PDF%20Files/HPC%20Global%20Leadership%20030509.pdf.

Council on Competitiveness. 2010. *High Performance Computing and U.S. Manufacturing Roundtable*. Washington, D.C. Last accessed June 18, 2010, at http://www.compete.org/publications/detail/1333/high-performance-computing-and-u.s.-manufacturing-roundtable/.

APPENDICES

APPENDIX 1: EXASCALE REVIEW PANEL MEMBERS

APPENDIX 2: LIST OF WORKSHOP AND TOWN MEETING REPORTS PROVIDED TO PANEL PRIOR TO REVIEW MEETING

APPENDIX 3: EXASCALE WORKSHOP PANEL REPORT AGENDA

APPENDIX 4: PRESENTERS AT PROGRAM BRIEFING

APPENDIX 5: ACRONYMS AND ABBREVIATIONS

APPENDIX 6: ACKNOWLEDGEMENTS

APPENDIX 1: EXASCALE REVIEW PANEL MEMBERS

Exascale Review Panel Members

Jim Decker is a Principal and one of the founders of Decker, Garman, Sullivan & Associates, LLC, a management and government relations consulting firm. Prior to founding DGS, he spent 33 years at the Department of Energy (DOE). Among the positions held at DOE, he was the first Director of the Scientific Computing Staff, and he led the first interagency planning efforts in high-performance computing and networking for the Office of Science and Technology Policy in the mid-1980s. For his last 21 years at DOE, he served as Principal Deputy Director of the Office of Science. Dr. Decker was recognized by three Presidents for outstanding performance with two Presidential Meritorious Rank Awards and a Presidential Rank Award for Distinguished Executives. He was also awarded the Secretary of Energy's Gold Award and the National Nuclear Security Administration Gold Medal. Before joining DOE, he was a Member of Technical Staff at Bell Telephone Laboratories. Dr. Decker has a B.S. from Union College and M.S. and Ph.D. degrees from Yale University, all in physics.

Tom Everhart is President Emeritus of the California Institute of Technology. He served as President from 1987-1997. Prior to that, he served as Chancellor of the University of Illinois at Urbana-Champaign (1984-87), Dean of Engineering at Cornell (1979-84) and Professor of EECS at the University of California, Berkeley (1958-78). He has served as a consultant to several companies and on the Boards of Directors of GM (1989-2002), Hewlett-Packard, Agilent Technologies, Hughes Electronics, and Raytheon, to name several. He currently serves on the Caltech Board of Trustees and the Boards of the Keck Foundation and the Kavli Foundation.

Anita Jones is a University Professor Emerita of the University of Virginia and a Professor of Computer Science in the School of Engineering and Applied Science. She served as the Director of Defense Research and Engineering for the U.S. Department of Defense in 1993-1998 and later as vice-chair of the National Science Board. She is a senior fellow of the Defense Science Board and a member of the American Philosophical Society and the National Academy of Engineering. She is a Fellow of the Association for Computing Machinery, the Institute of Electrical and Electronics Engineers, the American Academy of Arts & Sciences, and the American Academy for the Advancement of Science. She has published more than 50 technical articles and two books in the area of computer software and systems, cyber-security, and science and technology policy.

Arthur Kerman is Emeritus Professor of Physics at the Massachusetts Institute of Technology. His research interests in theoretical nuclear physics include nuclear QCD-relativistic heavy-ion physics, nuclear reactions, and laser accelerators. Over the course of 30 years, he has had various long-standing consulting relationships with Argonne, Brookhaven, Lawrence Berkeley, Lawrence Livermore, Los Alamos, and Oak Ridge national laboratories, as well as with several private companies and the National Bureau of Standards (now NIST). He currently has active consulting relationships with LLNL and LANL. He has similarly served on numerous professional committees for DOE, the National Science Foundation, Stanford University, the University of California, and the White House. He currently serves on the Director's Advisory Committee of LLNL, the Physics and Space Technology Advisory Committee at LLNL, the Relativistic Heavy Ion Collider Policy Committee at Brookhaven, the Physics Division Advisory Committee at LANL, and the Theory Advisory Committee at LANL. He was a member of the DOE ICFAC. He is a fellow of the American Physical Society and the American Academy of Arts and Sciences.

Cherry Murray has B.S. and Ph.D. degrees in physics from MIT. Her expertise is in condensed matter and materials physics, phase transitions, light scattering and surface physics, as well as the management of science and technology.

She became Dean of the Harvard School of Engineering and Applied Science and John A. and Elizabeth S. Armstrong Professor of Engineering and Applied Science and Professor of Physics in July 2009. As Dean, she manages new faculty recruitment and faculty relations; directs and leads strategic planning; coordinates fundraising and alumni relations; determines and implements educational, research, and administrative goals for the most recent new School of Harvard. Previous to that, she was Deputy Director and then Principal Associate Director for Science and Technology at Lawrence Livermore National Laboratory from December 2004. She led the Laboratory's science and technology activities including management of 3500 scientists and engineers and the development of the strategic science and technology plan. Formally Senior Vice President for Physical Sciences and Wireless Research at Bell Labs, Lucent Technologies, Dr. Murray first joined Bell Labs in 1978. She is a member of the National Academy of Science, the National Academy of Engineering and the American Academy of Arts and Sciences. Dr. Murray is on the board and fellow of the American Association for the Advancement of Science (AAAS) and past chair of the Physics Section of AAAS, as well as fellow and Past President of the American Physical Society (APS). She serves on the APS and American Institute of Physics Governing Boards and is chair of the National Research Council Division of Engineering and Physical Science. She is a member of the National Academies' Committee on International Security and Arms Control. In 1989, she won the APS Maria Goeppert-Mayer Award, and, in 2005, the APS George E. Pake Prize. She is the author of two patents and more than 75 publications. *Discover* Magazine named her one of the "50 Most Important Women in Science" in 2002.

Aristides A.N. Patrinos is president of Synthetic Genomics, Inc. (SGI), a privately held company founded in 2005 applying genomic-driven commercial solutions that address global energy and environmental challenges.

Prior to joining SGI, Dr. Patrinos was instrumental in advancing the scientific and policy framework underpinning key governmental energy and environmental initiatives while serving as associate director of the Office of Biological and Environmental Research in the U.S. Department of Energy's Office of Science. He oversaw the department's research activities in human and microbial genome research, structural biology, nuclear medicine, and climate change. Dr. Patrinos played a historic role in the Human Genome Project; the founding of the DOE Joint Genome Institute; and the design and launch of the DOE's Genomes to Life Program, a research program dedicated to developing technologies to use microbes for innovative solutions to energy and environmental challenges.

Dr. Patrinos currently serves on the National Academy of Science committee on the Economic and Environmental Impacts of Increasing Biofuels Production. He is a member of the Department of Energy Panel on the Exascale Computer; a fellow of the American Association for the Advancement of Science and of the American Meteorological Society; and a member of the American Geophysical Union, the American Society of Mechanical Engineers, and the Greek Technical Society. Dr. Patrinos is also the recipient of numerous awards and honorary degrees, including three Presidential Rank Awards and two Secretary of Energy Gold Medals, and an honorary doctorate from the National Technical University of Athens. A native of Greece, he received an undergraduate degree from the National Technical University of Athens and a Ph.D. from Northwestern University.

Justin Rattner is vice president and chief technology officer (CTO). He is also an Intel Senior Fellow and head of Intel Labs. In the latter role, he directs Intel's global research efforts in microprocessors, systems, and communications including the company's disruptive research activity.

In 1989, Rattner was named Scientist of the Year by *R&D* Magazine for his leadership in parallel and distributed computer architecture. In December 1996, Rattner was featured as Person of the Week by ABC World News for his visionary work on the Department of Energy ASCI Red System, the first computer to sustain one trillion operations per second (one teraFLOPS) and the fastest computer in the world between 1996 and 2000. In 1997, Rattner was honored as one of the Computing 200, the 200 individuals having the greatest impact on the U.S. computer industry today and subsequently was profiled in the book *Wizards and Their Wonders* from ACM Press.

Rattner has received two Intel Achievement Awards for his work in high-performance computing and advanced cluster communication architecture. He is a member of the executive committee of the Intel's Research Council and serves as the Intel executive sponsor for Cornell University where he is a member of the External Advisory Board for the School of Engineering.

Rattner is also a trustee of the Anita Borg Institute for Women and Technology. Rattner joined Intel in 1973. He was named its first Principal Engineer in 1979 and its fourth Intel Fellow in 1988. Prior to joining Intel, Rattner held positions with Hewlett-Packard Company and Xerox Corporation. He holds B.S. and M.S. degrees from Cornell.

Daniel Reed is Microsoft's Corporate Vice President for Technology Strategy and Policy and Extreme Computing. Previously, he was the Chancellor's Eminent Professor at UNC Chapel Hill, as well as the Director of the Renaissance Computing Institute (RENCI) and the Chancellor's Senior Advisor for Strategy and Innovation for UNC Chapel Hill.

Dr. Reed has served as a member of the U.S. President's Council of Advisors on Science and Technology (PCAST) and as a member of the President's Information Technology Advisory Committee (PITAC). As chair of PITAC's computational science subcommittee, he was lead author of the report "Computational Science: Ensuring America's Competitiveness." On PCAST, he co-chaired the Networking and Information Technology subcommittee (with George Scalise of the Semiconductor Industry Association) and co-authored a report on the National Coordination Office's Networking and Information Technology Research and Development (NITRD) program called "Leadership Under Challenge: Information Technology R&D in a Competitive World."

In June 2009, he completed two terms of service as chair of the board of directors of the Computing Research Association, which represents the research interests of Ph.D.-granting university departments, industrial research groups and national laboratories.

He was previously Head of the Department of Computer Science at the University of Illinois at Urbana-Champaign (UIUC), where the held the Edward William and Jane Marr Gutgsell Professorship. He has also been Director of the National Center for Supercomputing Applications (NCSA) at UIUC, where he also led National Computational Science Alliance, a 50-institution partnership devoted to creating the next generation of computational science tools. He was also one of the principal investigators and chief architect for the NSF TeraGrid. He received his B.S. from Missouri University of Science and Technology and his M.S. and Ph.D. in computer science in 1983 from Purdue University.

Charles V. Shank served as Director of the Lawrence Berkeley National Laboratory from 1989 to 2004. He is now a professor of chemistry, physics, and electrical engineering and computer sciences at UC Berkeley. He has been honored with the R.W. Wood Prize of the Optical Society of America, has received the George E. Pake Prize and the Arthur L. Schawlow Prize of the American Physical Society, and is a member of the National Academy of Sciences and National Academy of Engineering. After graduating with a Ph.D. in electrical engineering from the University of California, Berkeley, in 1969, Charles V. Shank headed to the most successful industrial research institution in the world, the AT&T

Bell Laboratories in New Jersey. He spent 20 years there, as both a researcher and director, and his work changed the way scientists see and understand the most fundamental events that shape our world.

C. Bruce Tarter served as the Director of the Lawrence Livermore National Laboratory from 1994-2002. Prior to becoming Director, he spent most of his career at the Laboratory as a theoretical physicist with responsibilities in weapons research, fusion energy, global climate, and space physics. He has a B.S. from MIT and a Ph.D. from Cornell University, all in physics.

His other activities include service on the Defense Science Board, the California Council on Science and Technology (CCST), and with the American Physical Society (APS). He is a fellow of the APS, CCST, the AAAS, and has received the Roosevelt Gold Medal for Science, National Nuclear Security Administration Gold Medal, and the Secretary of Energy's Gold Award. He recently served as Commissioner on the Strategic Posture Review Commission chartered by the Congress.

Alvin W. Trivelpiece was the Director of Oak Ridge National Laboratory (ORNL) from 1989 to 2000. From 1981 to 1987, he was the director of the Office of Energy Research (now the Office of Science) at the U.S. Department of Energy. He was elected to the National Academy of Engineering in 1993. His research interests has focused on plasma physics, controlled thermonuclear research, and particle accelerators. He was granted several patents on accelerators and microwave devices and is the author or co-author of many papers and two books. He is a fellow of the AAAS, the American Physical Society, and the Institute of Electrical and Electronics Engineers and is a member of the American Nuclear Society, the American Association of University Professors, Tau Beta Pi, and Sigma Xi.

Dr. Irving Wladawsky-Berger retired from IBM in May of 2007 after a 37-year career with the company, where his primary focus was on innovation and technical strategy. He was responsible for identifying emerging technologies and marketplace developments critical to the future of the IT industry, and organizing appropriate activities in and outside of IBM in order to capitalize on them. He led a number of IBM's companywide initiatives including the Internet and e-business, supercomputing, Linux, and Grid computing. He continues to consult for IBM on major new market strategies like Cloud Computing and Smart Planet.

In March of 2008, Dr. Wladawsky-Berger joined Citigroup as Strategic Advisor, helping with innovation and technology initiatives across the company. He is helping to formulate Citigroup initiatives related to the future of global banking, including mobile banking, Internet-based financial services, and financial systems modeling and analysis.

He is Visiting Lecturer at MIT's Sloan School of Management and Engineering Systems Division, Adjunct Professor in the Innovation and Entrepreneurship Group at the Imperial College Business School, and Senior Fellow at the Levin Institute of the State University of New York. He is a member of the InnoCentive Advisory Board, the Spencer Trask Collaborative Innovations Board, the Visiting Committee for the Physical Sciences Division at the University of Chicago, and the Board of Visitors for the Institute for Computational Engineering and Sciences at the University of Texas at Austin.

He was co-chair of the President's Information Technology Advisory Committee, as well as a founding member of the Computer Sciences and Telecommunications Board of the National Research Council. He is a former member of the University of Chicago Board of Governors for Argonne National Laboratories, of the Board of Overseers for Fermilab, and of BP's Technology Advisory Council. He is a Fellow of the American Academy of Arts and Sciences. A native of Cuba, he was named the 2001 Hispanic Engineer of the Year.

Dr. Wladawsky-Berger received an M.S. and a Ph. D. in physics from the University of Chicago.

APPENDIX 2: LIST OF WORKSHOP AND TOWN MEETING REPORTS PROVIDED TO PANEL PRIOR TO REVIEW MEETING

"Scientific Impacts and Opportunities for Computing," January 9-12, 2008, Maui, Hawaii. Available at http://www.science.energy.gov/ascr/ProgramDocuments/Docs/ScientificImpacts&Oppor.pdf.

"Science Based Nuclear Energy Systems Enabled by Advanced Modeling and Simulation at the Extreme Scale," May 11–12, 2009, Washington, D.C. Available at http://www.science.energy.gov/ascr/ProgramDocuments/Docs/SC-NEWorkshopReport.pdf.

"Computational Research Needs in Alternative and Renewable Energy," September 19-20, 2007, Rockville, Maryland. Available at http://www.science.energy.gov/ascr/ProgramDocuments/Docs/CRNAREWorkshopReport.pdf.

"Discovery in Basic Energy Sciences: The Role of Computing at the Extreme Scale," August 13-15, 2009, Bethesda, Maryland.

"Challenges in Climate Changes Science and the Role of Computing at the Extreme Scale," November 6-7, 2008, Washington, D.C. Available at http://www.science.energy.gov/ascr/ProgramDocuments/Docs/ClimateReport.pdf.

"Scientific Grand Challenges in Fusion Energy Sciences and the Role of Computing at the Extreme Scale," March 18-20, 2009, Gaithersburg, Maryland. Available at http://www.science.energy.gov/ascr/ProgramDocuments/Docs/FusionReport.pdf.

"Challenges for Understanding the Quantum Universe and the Role of Computing at the Extreme Scale," December 9-11, 2008, Menlo Park, California. Available at http://www.science.energy.gov/ascr/ProgramDocuments/Docs/HEPReport.pdf.

"Scientific Grand Challenges in National Security: the Role of Computing at the Extreme Scale," October 6-8, 2009, Gaithersburg, Maryland.

"Scientific Grand Challenges: Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale," January 26-28, 2009, Washington, D.C. Available at http://www.science.energy.gov/ascr/ProgramDocuments/Docs/NPReport.pdf.

"Modeling and Simulation at the Exascale for Energy and the Environment - Town Hall Meetings Report," April 17-18, 2007, Lawrence Berkeley National Laboratory, Berkeley, California; May 17-18, 2007, Oak Ridge National Laboratory, Oak Ridge, Tennessee; May 31-June 1, 200, Argonne National Laboratory, Argonne, Illinois. Available at http://www.er.doe.gov/ascr/ProgramDocuments/Docs/TownHall.pdf.

APPENDIX 2: LIST OF WORKSHOP AND TOWN MEETING REPORTS PROVIDED TO PANEL PRIOR TO REVIEW MEETING

"Workshop on Computing Science/Math Institutes and High Risk/High Payoff Technologies for Applications," October 7, 2008, Chicago, Illinois. Available at http://www.er.doe.gov/ascr/ProgramDocuments/Docs/MathCSWorkshopReport.pdf.

"Mathematical Research Challenges in Optimization of Complex Systems," December 7-8, 2006, Bethesda, Maryland. Available at http://www.er.doe.gov/ascr/Research/AM/Docs/ComplexSystemsWorkshopReport.pdf.

APPENDIX 3: EXASCALE WORKSHOP PANEL REPORT AGENDA

Tuesday, January 19, 2010

9:00 a.m.	Al Trivelpiece: Opening Remarks	Monet Suite IV, Second Floor
	Introductions and background: All (~ one minute each)	
	Michael Strayer: Purpose of Exascale Workshops	
9:30 a.m.	Rick Stevens and Andy White: Decadal DOE Plan to Provide Exascale Applications and Technologies for DOE Mission Needs	
10:30 a.m.	Break (15 minutes)	
10:45 a.m.	Executive Session	
11:00 a.m.	Alan Bishop: Scientific Grand Challenges for National Security	
Noon	Lee Hood On Biology And Medicine Via Call In	
1:00 p.m.	Lunch	
2:00 p.m.	Justin Rattner: Review of Limitations and Challenges Related to Stability, Power, and Scalability	
2:30 p.m.	Thomas Zacharia: State Of The Art with Petascale Capability – Successes And Limitations	
3:30 p.m.	Break	
4:00 p.m.	Jim Hack: Challenges in Climate Change Science and the Role of Computing at the Extreme Scale	
5:00 p.m.	Art Kerman: Science Based Nuclear Energy Systems Enabled by Advanced Modeling and Simulation at the Extreme Scale	
5:30 p.m.	Executive Session	
6:00 p.m.	Adjourn	

Wednesday, January 20, 2010

9:00 a.m.	Martin Savage: Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale	Monet Suite IV, Second Floor
10:00 a.m.	Horst Simon: State of the Art with Petascale Capability – Successes and Limitations	
11:00 a.m.	Executive Session and Writing Assignments	
Noon	Lunch	
1:00 p.m.	Executive Session and Writing Assignments	
2:00 p.m.	David Dean: DOE (Senior Advisor to DOE Under Secretary Steve Koonin)	
3:00 p.m.	Adjourn	

APPENDIX 4: PRESENTERS AT PROGRAM BRIEFING

Presenters at Program Briefing on January 19 and 20, 2010 in Washington, D.C.

Alan Bishop is Associate Director for Theory, Simulation and Computation, Los Alamos National Laboratory. His professional training is in theoretical nonlinear and nonequilibrium electronic and structural phenomena in condensed matter, materials science, biophysics, and chemical physics. He received his Ph.D. from the Cavendish Laboratory, University of Cambridge, and, after periods in Oxford, Cornell and London Universities, he became a technical staff member in the Theoretical Division at Los Alamos, where he was a founding member of the Center for Nonlinear Studies. He was Director of the Theoretical Division (1999- 2006), after which he became Laboratory Associate Director, leading a technical emphasis on the integration of interdisciplinary theory, modeling, computer sciences, and high-performance computing. He is a Fellow of the APS, AAAS, and Los Alamos National Laboratory; a Humboldt Foundation Senior Fellow; and he is a recipient of the DOE E.O. Lawrence Award in Materials Science.

Jim Hack is Director, National Center for Computational Sciences, Oak Ridge National Laboratory. An atmospheric scientist, Dr. Hack also directs ORNL's Climate Change Science Institute. Dr. Hack received a Ph.D. in atmospheric dynamics from Colorado State University in 1980. He was a research staff member at the IBM Thomas J. Watson Research Center until 1984. He joined the National Center for Atmospheric Research in 1984, where he held numerous roles including senior scientist, head of the Climate Modeling Section, and deputy director of the Climate and Global Dynamics Division. He was one of the principal developers of the NCAR global climate model that provided simulation data jointly contributed by the Department of Energy and the National Science Foundation to the United Nations' Intergovernmental Panel on Climate Change's 4th Assessment Report. He has been an Editor for the Journal of Climate, has held an adjoint professor position at the University of Colorado at Boulder, and is author or co-author of approximately 100 scientific or technical publications.

Leroy Hood is President, Institute for Systems Biology, Seattle, Washington. He received an M.D. degree from John Hopkins and a Ph.D. from Caltech. Dr. Hood was a professor at Caltech from 1970-1992 and focused on fundamental biology (immunity, evolution, genomics, and neurobiology) and on bringing engineering to biology. In the 1980s, Dr. Hood began thinking about cross-disciplinary biology and systems biology. In 1992, Dr. Hood moved to the University of Washington with support from Bill Gates as founder and Chairman of the first cross-disciplinary biology department (Molecular Biotechnology--MBT). In 2000, he co-founded the independent non-profit Institute for Systems Biology in Seattle, Washington to develop the strategies and technologies for systems approaches to biology and medicine. Dr. Hood is a member of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. Dr. Hood was awarded the Lasker Prize in 1987 for immune diversity, the 2002 Kyoto Prize in Advanced Technology for technology development, the 2003 Lemelson–MIT Prize for the development of the DNA sequencer, and the 2006 Heinz Award in Technology, the Economy and Employment. Dr. Hood is currently pioneering medicine's transition from its current reactive mode to a predictive, preventive, personalized, and participatory mode (P4 medicine).

Martin Savage is a Professor of Physics at the University of Washington. He received his Bachelors and a Masters degree from the University of Auckland and a Ph.D. from the California Institute of Technology.

After postdoctoral positions at Rutgers University and the University of California in San Diego and a faculty position at Carnegie Mellon University, he assumed a faculty position at the University of Washington in 1996. Martin is a theoretical nuclear physicist and a member of the NPLQCD collaboration.

He has received a Superconducting Super Collider Fellowship and a DOE Outstanding Junior Investigator Award, and he is a Fellow of the American Physical Society.

Horst Simon is Associate Laboratory Director for Computing Sciences and the Division Director for the Computational Research Division at Lawrence Berkeley National Laboratory and Adjunct Professor of Computer Science at the University of California, Berkeley. His research interests are in the development of sparse matrix algorithms, algorithms for large-scale eigenvalue problems, and domain decomposition algorithms. His recursive spectral bisection algorithm is a breakthrough in parallel algorithms. He was honored twice with the prestigious Gordon Bell Prize, most recently in 2009 for the development of innovative techniques that produce new levels of performance on a real application (in collaboration with IBM researchers) and in 1988 in recognition of superior effort in parallel processing research (with others from Cray and Boeing). He has served as a senior manager for Silicon Graphics, the Computer Sciences Corporation, Boeing Computer Services, and has been a member of the faculty at the State University of New York. He is currently a member of the advisory boards of more than five research organizations located throughout the world and is a member of many journal editorial boards and one of four editors of the twice-yearly "TOP500" list of the world's most powerful computing systems.

Rick Stevens is a professor of computer science at the University of Chicago and holds senior fellow appointments in the University's Computation Institute (CI) and the Institute for Genomics and Systems Biology, where he teaches and supervises graduate students in the areas of computational biology, collaboration and visualization technology, and computer architecture. He co-founded and co-directed the CI, which provides an intellectual home for large-scale interdisciplinary projects involving computation. Rick is also Associate Laboratory Director responsible for Computing, Environment, and Life Sciences research at Argonne National Laboratory, where he developed Argonne's research program in computational biology.

Recently, Rick has been co-leading the DOE laboratory planning effort for exascale computing research aiming to develop computer systems one thousand times faster than current supercomputers and apply these systems to fundamental problems in science including genomic analysis, whole cell modeling, climate models, and problems in fundamental physics and energy technology development. Rick is a fellow of the American Association for the Advancement of Science, and his research groups have won many national awards over the years, including an R&D 100 award for the Access Grid. He sits on many government, university, and industry advisory boards.

Thomas Zacharia is Deputy for Science and Technology at DOE's Oak Ridge National Laboratory (ORNL), overseeing one of the nation's largest research and development programs. Zacharia is a materials scientist by training and until recently led ORNL's advanced scientific computing efforts. His vision and leadership have transformed ORNL into a world leader in computing, making the Laboratory home to the world's most powerful computer systems and greatly enhancing its capability to support national priorities in areas such as climate change, sustainable energy technologies, nanotechnology, and biotechnology.

Zacharia is a member of the NOAA Science Advisory Board, NSF's Foundation-wide Advisory Committee for Cyberinfrastructure, and the High Performance Computing Advisory Board of the Council on Competitiveness. He is the author or co-author of more than 100 publications and holds two U.S. patents. The recipient of several prestigious scientific and technical awards, Zacharia is also a professor in the Electrical Engineering and Computer Science Department at the University of Tennessee. He holds a Ph.D. from Clarkson University.

Andy White is the Deputy Associate [Laboratory] Director of the Theory, Simulation and Computing Directorate, co-chair with Rick Stevens (ANL) of the DOE Exascale Initiative Steering Committee and the Roadrunner Project Director. From 1989 to 1998, he was founder and Director of the Advanced Computing Laboratory at Los Alamos. He was an Associate Director of the NSF Science and Technology Center for Research on Parallel Computation (CRPC) from 1988–1999 and founded the Delphi Project in 1998. His research interests are in the areas of applied mathematics; high-performance computing; computational simulation; and modeling, resilience, and predictive computational capabilities.

APPENDIX 5: ACRONYMS AND ABBREVIATIONS

ASCI	Advanced Strategic Computing Initiative
DARPA	Defense Advanced Research Projects Agency
DNA	deoxyribonucleic acid
DOE	U.S. Department of Energy
DRAM	dynamic random access memory
I/O	input/output
IT	information technology
LARC	Livermore Advanced Research Computer
MFECC	Magnetic Fusion Energy Computer Center
MHD	magnetohydrodynamic
NERSC	National Energy Research Supercomputer Center
NNSA	National Nuclear Security Administration
NSA	National Security Agency
NSF	National Science Foundation
NVRAM	nonvolatile random access memory
PC	personal computer
PET	positron emission tomography
PRACE	Partnership for Advanced Computing in Europe
R&D	research and development
SciDAC	Scientific Discovery through Advanced Computing

APPENDIX 6: ACKNOWLEDGEMENTS

The Exascale Review Panel greatly appreciates the support provided by Oak Ridge Institute for Science and Education (ORISE). ORISE did an excellent job of making all the arrangements needed for the Panel to meet in Washington, D.C., hear presentations, and learn about the value proposition of a U.S. Department of Energy exascale initiative.

U.S. DEPARTMENT OF ENERGY

Office of Science

Office of Advanced Scientific Computing Research

Prepared for the U.S. Department of Energy under Contract DE-AC05-76RL01830

Production support provided by Pacific Northwest National Laboratory,
Fundamental & Computational Sciences Directorate

PNNL-19515