SUPERCOMPUTING AND AMERICAN TECHNOLOGY LEADERSHIP

HEARING

BEFORE THE

SUBCOMMITTEE ON ENERGY

COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

HOUSE OF REPRESENTATIVES

ONE HUNDRED FOURTEENTH CONGRESS

FIRST SESSION

JANUARY 28, 2015

Serial No. 114–03

Printed for the use of the Committee on Science, Space, and Technology

Available via the World Wide Web: http://science.house.gov

U.S. GOVERNMENT PUBLISHING OFFICE

93–881PDF WASHINGTON : 2015

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

COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

HON. LAMAR S. SMITH, Texas, *Chair*

FRANK D. LUCAS, Oklahoma	EDDIE BERNICE JOHNSON, Texas
F. JAMES SENSENBRENNER, JR.	ZOE LOFGREN, California
DANA ROHRABACHER, California	DANIEL LIPINSKI, Illinois
RANDY NEUGEBAUER, Texas	DONNA F. EDWARDS, Maryland
MICHAEL T. McCAUL	FREDERICA S. WILSON, Florida
STEVEN M. PALAZZO, Mississippi	SUZANNE BONAMICI, Oregon
MO BROOKS, Alabama	ERIC SWALWELL, California
RANDY HULTGREN, Illinois	ALAN GRAYSON, Florida
BILL POSEY, Florida	AMI BERA, California
THOMAS MASSIE, Kentucky	ELIZABETH H. ESTY, Connecticut
JIM BRIDENSTINE, Oklahoma	MARC A. VEASEY, Texas
RANDY K. WEBER, Texas	KATHERINE M. CLARK, Massachusetts
BILL JOHNSON, Ohio	DON S. BEYER, JR., Virginia
JOHN R. MOOLENAAR, Michigan	ED PERLMUTTER, Colorado
STEVE KNIGHT, California	PAUL TONKO, New York
BRIAN BABIN, Texas	MARK TAKANO, California
BRUCE WESTERMAN, Arkansas	BILL FOSTER, Illinois
BARBARA COMSTOCK, Virginia	
DAN NEWHOUSE, Washington	
GARY PALMER, Alabama	
BARRY LOUDERMILK, Georgia	

SUBCOMMITTEE ON ENERGY

HON. RANDY K. WEBER, Texas , *Chair*

DANA ROHRABACHER, California	ALAN GRAYSON, Florida
RANDY NEUGEBAUER, Texas	DANIEL LIPINSKI, Illinois
MO BROOKS, Alabama	ERIC SWALWELL, California
RANDY HULTGREN, Illinois	ELIZABETH H. ESTY, Connecticut
THOMAS MASSIE, Kentucky	MARC A. VEASEY, Texas
BARBARA COMSTOCK, Virginia	KATHERINE M. CLARK, Massachusetts
DAN NEWHOUSE, Washington	EDDIE BERNICE JOHNSON, Texas
BARRY LOUDERMILK, Georgia	
LAMAR S. SMITH, Texas	

CONTENTS

January 28, 2015

SUPERCOMPUTING AND AMERICAN TECHNOLOGY LEADERSHIP

WEDNESDAY, JANUARY 28, 2015

HOUSE OF REPRESENTATIVES,
SUBCOMMITTEE ON ENERGY
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY,
Washington, D.C.

The Subcommittee met, pursuant to call, at 9:08 a.m., in Room 2318 of the Rayburn House Office Building, Hon. Randy Weber [Chairman of the Subcommittee] presiding.

LAMAR S. SMITH, Texas
CHAIRMAN

EDDIE BERNICE JOHNSON, Texas
RANKING MEMBER

Congress of the United States
House of Representatives

COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

2321 RAYBURN HOUSE OFFICE BUILDING

WASHINGTON, DC 20515-6301

(202) 225-6371
www.science.house.gov

Subcommittee on Energy

Supercomputing and American Technology Leadership

Wednesday, January 28, 2015
9:00 a.m. – 11:00 a.m.
2318 Rayburn House Office Building

Witnesses

Mr. Norman Augustine, Board Member, Bipartisan Policy Center

Dr. Roscoe Giles, Chairman, DOE Advanced Scientific Computing Advisory Committee

Mr. David Turek, Vice President, Technical Computing, IBM

Dr. James Crowley, Executive Director, Society for Industrial and Applied Mathematics

U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY
SUBCOMMITTEE ON ENERGY

HEARING CHARTER

Supercomputing and American Technology Leadership

Wednesday, January 28, 2015
9:00 a.m. – 11:00 a.m.
2318 Rayburn House Office Building

Purpose

The Energy Subcommittee will hold a hearing titled *Supercomputing and American Technology Leadership* at 9:00 a.m. on January 28th in room 2318 of the Rayburn House Office Building. This hearing will assess the Advanced Scientific Computing Research (ASCR) program within the U.S. Department of Energy's (DOE) Office of Science as a mechanism to support technological advancement in the United States. This hearing will focus on high performance computing (HPC) facilities' unique ability to accelerate innovation and inform the Committee regarding the applications and benefits from sustained investment in the ASCR program.

Witnesses

- **Mr. Norman Augustine**, Board Member, Bipartisan Policy Center
- **Dr. Roscoe Giles**, Chairman, DOE Advanced Scientific Computing Advisory Committee
- **Mr. Dave Turek**, Vice President, Technical Computing, IBM
- **Dr. James Crowley**, Executive Director, Society for Industrial and Applied Mathematics

Background

The DOE Office of Science is the federal government's largest supporter of basic research in the physical sciences. The Office of Science provides direct funding to researchers and also develops, constructs, and operates large-scale user facilities.[1] Within the Office of Science, the ASCR program develops and maintains world-class computing facilities and provides funding for research in applied mathematics, computer science, and advanced networking.[2] ASCR was funded at $541 million for FY15.

[1] See generally, DOE Office of Science website: http://science.energy.gov/about/
[2] See generally DOE ASCR website: http://science.energy.gov/ascr/about/

ASCR's programs fund discovery-based science in the areas of modeling, analysis, and simulation that may enable breakthroughs in other fields of research and technology development. In May 2013, the Subcommittee held a hearing on one of ASCR's major new initiatives involving exascale computing.[3] The exascale initiative is an effort to develop the next generation of high performance computing systems that are three orders of magnitude faster than current systems, and therefore able to process larger amounts of data more quickly. Scientific discovery in which large volumes of data is gathered and mined to exploit information, sometimes referred to as "big data," has transformed computing technology needs. The greater availability and utilization of these high-speed supercomputers allows increasingly complex scientific research to be achieved. Medical research, energy and environment system simulations, computational chemistry, and innumerable other scientific problems directly benefit from HPC modeling.

Additionally, in a recent report, the Council on Competitiveness, an organization comprised of U.S. corporations, universities, labor organizations, and laboratories, stated the following:

> "High performance computing (HPC) is inextricably linked to innovation, fueling breakthroughs in science, engineering, and business. HPC is a tool used by leaders in diverse fields to help design new products, to improve existing products, and to bring products to market more efficiently. HPC is viewed as a cost-effective tool for speeding up the R&D process, and two-thirds of all U.S.-based companies that use HPC say that 'increasing performance of computational models is a matter of competitive survival.'"[4]

Office of Science and ASCR Budget

ASCR's budget is divided into two main subprograms. The High Performance Computing and Networking Facilities subprogram constructs and maintains cutting-edge facilities supporting researchers across the United States. The Mathematical, Computational, and Computer Sciences Research subprogram supports, among other things, research in mathematics for modeling systems, software, middleware, and new algorithms to efficiently make use of new high performance computing systems. This subprogram develops applied mathematicians and computer scientists to addresses challenges associated with managing large amounts of raw data, a key input for maintaining technical leadership in this field.

[3] See House Committee on Science, Space, and Technology Subcommittee on Energy Hearing- America's Next Generation Supercomputer: The Exascale Challenge *Available at:* http://science.house.gov/hearing/subcommittee-energy-hearing-exascale-computing-challenges-and-opportunities

[4] See "The Exascale Effect: the Benefits of Supercomputing Investment for U.S. Industry", Council on Competitiveness, October 2014. *Available at:* http://www.compete.org/images/uploads/File/PDF%20Files/Solve_Report_Final.pdf

Chairman WEBER. Well, good morning and welcome to today's Energy Subcommittee hearing titled "Supercomputing and American Technology Leadership."

The Committee will come to order.

Without objection, the Chair is authorized to declare recesses of the Subcommittee at any time.

Without objection, the Chair authorizes the participation of Mr. Lipinski, Mr. Swalwell, Mr. Grayson, Ms. Esty, Mr. Veasey, and Ms. Clark for today's hearing. And I understand Ranking Member Johnson will serve as the Ranking Minority Member today and give an opening statement a little later.

In front of you are packets containing the written testimonies, biographies, and truth-in-testimony disclosures for today's witnesses. And I recognize myself for five minutes for an opening statement.

At the outset let me say that this is my first Committee hearing as a Chairman of this Subcommittee and it is truly an honor to be selected to serve in this capacity. And I want to say a personal thanks to Chairman Lamar Smith for his help and his guidance. He has been just a stalwart friend of mine. I really appreciate that.

This Committee will tackle a number of important issues related to America's competitiveness and energy future, and I am excited to be part of these important discussions.

Today, we are going to hear from a distinguished panel of witnesses about the importance of high-performance computing to American technological competitiveness, specifically focusing on the Department of Energy's Advanced Scientific Computing Research program, also known as the ASCR program within the Office of Science.

High-performance computing provides a platform for breakthroughs in all scientific research and accelerates applications of scientific breakthroughs across our economy. Progress in computing has paved the way for breakthroughs in medical imaging, genetics research, manufacturing, engineering, and weapons development. Faster computing speeds have revolutionized the energy sector, improving the efficiency of energy production and aiding in distribution technologies. Advances in modeling and algorithm development offer opportunities for scientific discovery in fields where experiments are too difficult, too costly, or too dangerous to conduct. They are reducing costs and opening the door to more innovative discoveries.

The work underway in the ASCR program drives breakthroughs in high-performance computing. The Department of Energy's national labs host world-class computational science facilities, and the Department funds the applied mathematical and computational science research that will drive the next stage of advancement in this field.

As we face the reality of ongoing budget constraints in Washington, it is our job in Congress to ensure that taxpayer dollars are spent wisely on innovative research that is in the best national interest and provides the best chance for broad impact and long-term success. The basic research conducted within the ASCR program clearly meets this requirement. High-performance computing can lead to scientific discoveries, economic growth, and will help maintain America's leadership in science and technology.

I want to thank the witnesses in advance for participating in today's hearing and look forward to further discussion.

[The prepared statement of Mr. Weber follows:]

PREPARED STATEMENT OF SUBCOMMITTEE CHAIRMAN RANDY WEBER

Good morning and welcome to today's Energy Subcommittee hearing titled "Supercomputing and American Technology Leadership."

Today, we will hear from a distinguished panel of witnesses about the importance of high performance computing to American technological competitiveness, specifically focusing on the Department of Energy's Advanced Scientific Computing Research program, also known as the "ASCR" program within the Office of Science.

High performance computing provides a platform for breakthroughs in all scientific research, and accelerates applications of scientific breakthroughs across our economy. Progress in computing has paved the way for breakthroughs in medical imaging, genetics research, manufacturing, engineering, and weapons development. Faster computing speeds have revolutionized the energy sector, improving the efficiency of energy production and aiding in distribution technologies. Advances in modeling and algorithm development offer opportunities for scientific discovery in fields where experiments are too difficult, costly, or dangerous to conduct, reducing costs and opening the door to more innovative discoveries.

The work underway in the ASCR program drives breakthroughs in high performance computing. The Department of Energy's national labs host world-class computational science facilities, and the department funds the applied mathematical and computational science research that will drive the next stage of advancement in this field.

As we face the reality of ongoing budget constraints in Washington, it is our job in Congress to ensure that taxpayer dollars are spent wisely, on innovative research that is in the national interest, and provides the best chance for broad impact and long-term success. The basic research conducted within the ASCR program clearly meets this requirement. High performance computing can lead to scientific discoveries, economic growth, and will maintain America's leadership in science and technology. I thank the witnesses for participating in today's hearing and look forward to further discussion.

Chairman WEBER. I now recognize Ranking Member Johnson for an opening statement.

Ms. JOHNSON. Thank you very much, Mr. Chairman, and I thank you for holding this hearing. And I want to thank our very excellent panel of witnesses for their testimony and being here today.

America has historically been a leader in advancing new energy technologies, as well as the fundamental sciences of physics, chemistry, engineering, mathematics, and computational science that support energy innovation. But our leadership in technology is challenged by the growing investments of other countries in education and research, investments that are now predicted to quickly outpace our own investments here at home.

High-performance computing or supercomputing is one area that we have led in for decades and the United States currently holds more than 45 percent of the 500 fastest supercomputers in the world. These computers are capable of processing vast amounts of data and mathematical equations at amazing speeds.

In the past, high-performance computers were needed primarily for specialized scientific and engineering applications. Now, as we enter the world of big data where thousands of devices all around us are generating millions of bytes of data to be analyzed, high-performance computing is needed not just by scientists and government researchers but by many civic and commercial enterprises as well.

Public policies play a critical role in supporting the advancement of high-performance computing and in enabling our society and economy to directly benefit from this capability. Our policies allow researchers and private industry to access the Department of Energy's computing systems, which are some of the most powerful in the world. We set policies that support the development of the software necessary to operate and optimize the use of high-performance systems, software that is unlikely to be developed by private industry because the potential sales market is too small to support the initial research and development costs. And our policies ensure that our investments in new computer architectures are diverse and flexible enough to meet our national security needs, in addition to our research and private industry needs. Federal investments in high-performance computing open this technology up to the future development of proprietary products. They grow our technology economy and they advance our technological leadership internationally.

Now, while every witness on this panel is extremely distinguished and I am grateful that each of you could be here today, I hope you won't mind if I thank Dr. Augustine in particular for taking time to speak with us this morning as he has been a great friend to this Committee for well over a decade. As a former Chairman of Lockheed Martin and the Chair of the National Academy of Sciences Committee that produced the seminal Rising above the Gathering Storm report in 2005, he has a broad and deep perspective on the challenges facing our Nation in research and technological innovation. That report laid the foundation for one of our Committee's landmark bipartisan achievements, the *America COMPETES Act* of 2007, which we reauthorized in 2010 and I hope the next reauthorization is a top priority for the Committee and this Congress.

I look forward to hearing Mr. Augustine's thoughts and indeed those of all of our witnesses on where we need to go in scientific research and innovation to grow our economy and to improve the quality of life for all Americans. Working together, our Committee has the opportunity to renew our commitment to scientific and technological leadership by our actions, and I look forward to any input our panelists have toward that goal.

With that, I thank you for coming and I yield back the balance of my time.

[The prepared statement of Ms. Johnson follows:]

PREPARED STATEMENT OF COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY RANKING MEMBER EDDIE BERNICE JOHNSON

Thank you Chairman Weber for holding this hearing, and I also want to thank this excellent panel of witnesses for their testimony and for being here today.

America has historically been a leader in advancing new energy technologies, as well as the foundational sciences of physics, chemistry, engineering, mathematics, and computational science that support energy innovation. But our leadership in technology is challenged by the growing investments of other countries in education and research; investments that are now projected to quickly outpace our own investments here at home.

High performance computing, or supercomputing, is one area we have led in for decades, and the U.S. currently hosts more than 45% of the 500 fastest supercomputers in the world. These computers are capable of processing vast amounts of data and mathematical equations at amazing speeds. In the past, high performance computers were needed primarily for specialized scientific and engineering applications. Now, as we enter the world of 'big data', where thousands of devices all around us are generating millions of bytes of data to be analyzed, high performance computing is needed not just by scientists and government researchers, but by many civic and commercial enterprises as well.

Public policies play a critical role in supporting the advancement of high performance computing, and in enabling our society and economy to directly benefit from this capability. Our policies allow researchers and private industry to access the Department of Energy's computing systems, which are some of the most powerful in the world. We set policies that support the development of the software necessary to operate and optimize the use of high performance systems—software that is unlikely to be developed by private industry because the potential sales market is too small to support the initial research and development costs. And our policies ensure that our investments in new computer architectures are diverse and flexible enough to meet our national security needs, in addition to our research and private industry needs. Federal investments in high performance computing open this technology up for future development of proprietary products, they grow our technology economy, and they advance our technological leadership internationally.

Now, while every witness on this panel is extremely distinguished and I am grateful that each of you could be here today, I hope you won't mind if I thank Dr. Augustine in particular for taking time to speak with us this morning, as he has been a great friend to the Committee for well over a decade. As the former Chairman of Lockheed Martin and the Chair of the National Academy of Sciences Committee that produced the seminal Rising Above the Gathering Storm report in 2005, he has a broad and deep perspective on the challenges facing our nation in research and technological innovation. That report laid the foundation for one of our Committee's landmark bipartisan achievements, the America COMPETES Act of 2007, which we reauthorized in 2010, and I hope the next reauthorization is a top priority for the Committee in this new Congress.

I look forward to hearing Mr. Augustine's thoughts—and indeed those of all of our witnesses - on where we need to go in scientific research and innovation to grow our economy and to improve the quality of life of all Americans. Working together, our Committee has the opportunity to renew our commitment to scientific and technological leadership by our actions, and I look forward to any input our panelists have towards that goal.

With that, I thank you all for coming, and I yield back the balance of my time.

Chairman WEBER. I thank the lady, and if there are Members who wish to submit additional opening statements, your statements will be added to the record at this point.

Chairman WEBER. At this time I would like to introduce our witnesses. Our first witness, who comes with high commendations, is Mr. Norman Augustine, Board Member of the Bipartisan Policy Center. Mr. Augustine served as the Undersecretary of the Army and later as acting Secretary of the Army from 1975 to 1977. Mr. Augustine also served as the President and CEO of Lockheed Martin until he retired in 1997. He has been a member of advisory boards to the Department of Homeland Security, Energy, Defense, Commerce, Transportation, and Health and Human Services, as well as NASA, Congress, and the White House.

Is there any other—are there boards that you weren't a member of, Mr. Augustine?

Our second witness today who is actually joining us by video is Dr. Roscoe Giles, Chairman of the Advanced Scientific Computing Advisory Committee at the Department Of Energy and a Professor at Boston University. Dr. Giles has served in a number of leadership roles in the community, including Member of the Board of Associated Universities Incorporated, Chair of the Boston University Faculty Council, and General Chair of the SC conference in 2002. Welcome, Dr. Giles.

Dr. GILES. Thank you.

Chairman WEBER. Our next witness today is Mr. David Turek, Vice President of Technical Computing at IBM. Previously Mr. Turek—am I saying that name correctly? Okay. Previously, Mr. Turek helped launch IBM's grid computing business and ran IBM's Linux cluster business. He also helped lead IBM's initiative in support of the U.S. Accelerated Strategic Computing Initiative at Lawrence Livermore National Laboratory, which I believe is in Mr. Swalwell's district.

Mr. SWALWELL. That is right.

Chairman WEBER. Yes. So welcome.

Our final witness today is Dr. James Crowley, Executive Director at the Society for Industrial and Applied Mathematics. Dr. Crowley has held this position since 1995. Prior to this, he served in the Air Force for 22 years retiring as Lieutenant Colonel. Dr. Crowley is a fellow of the American Mathematical Society and a fellow of the American Association for the Advancement of Science.

In order to allow time for discussion, please limit your testimony to five minutes, we ask the witnesses, and your entire statement will be made part of the written record.

I now recognize Mr. Augustine for five minutes to present his testimony.

TESTIMONY OF MR. NORMAN AUGUSTINE, BOARD MEMBER, BIPARTISAN POLICY CENTER

Mr. AUGUSTINE. Well, thank you very much, Chairman Weber, Ranking Member Johnson, and Members of the Subcommittee, and thank you, Ranking Member Johnson, for all those kind words.

I am particularly appreciative that this Committee is going to devote some time to the topic at hand and certainly high-performance computing is a key element of research.

I will submit a statement for the record.

I would like to begin by offering a few words about the basic nature of research. It is through research that new knowledge is created that permits engineers like myself to translate that research, knowledge into products and services that, working with entrepreneurs, can go into the marketplace and improve people's lives. We often think of Apple, the great things it has done, deservedly. Think of the iPod, iPads, and so on. But it wasn't Apple that made those things possible; it was researchers working decades ago on such things as quantum mechanics and material sciences, solid-state physics, and so on.

One of the things about basic research in particular is that you can't know or priority what will be the outcome of it and that sure makes it particularly difficult in your roles, to build support for it, yet there are so many examples of where basic research that was curiosity-driven led to greater improvements in people's lives. Three things that come to my mind, one is research on seals in Antarctica that led to a surgical procedure that saved the lives of many children undergoing lung surgery. Another was study of the chemistry of butterfly wings of that led to an ingredient that is used in chemotherapy. Still another of course would be the accidental discovery of penicillin when someone was studying research on bacteria many, many decades ago, Sir Alexander Fleming.

I would like to quickly touch on the importance of research and I will cite three areas where I think it has particularly had an impact. One is on the creation of jobs and there is evidence that if you want to one percentage point to the average number of jobs in America, you have to add about 1.7 percentage points to the GDP of America. There have been a number of studies, one of which was the basis of a Nobel Prize and it has shown that between 50 and 85 percent of the growth of GDP in our country during the last half-century is directly attributable to advancements in two fields: science and technology. And of course those advancements are entirely dependent upon research.

Health is an example. In the last century life expectancy in the United States grew from 47 to 79 years. In fact, I am 79 years old so this is really important to me. The life expectancy gain that came about was in considerable part attributable to advancements in biomedical research.

A third example is things that we take for granted in our everyday life, be they television, electric cars, DVDs, GPSs, CAT scans, or what have you, are dependent upon the knowledge that came through basic research.

Touching briefly on high-performance computing, it impacts field across the entire technological spectrum. My own field of aerodynamics is an example, another would be genomics, high-energy physics. It truly is of broad importance.

The Department of Energy, as you know, operates 17 laboratories. They are able to do things that industry really can't do under the pressures of today's marketplace for quick returns, financial returns. The examples, things that they could do so well are high-risk, high-return payoff research or long-term research, large research projects. They are particularly well suited to that. And work in the past, for example, sponsored by the Department of Energy on hydraulic fracturing, as you know, has had an enormous impact today in the political world, as well as the economic world.

How are we doing in the United States in research? The answer has to be not very well. Research funding as a percentage of GDP of the United States has dropped from 1st place to 7th place in the last decade or so. The fraction of research in a country that is sponsored by the government, United States is down in 29th place. As a fraction of GDP—R&D to GDP we are in 10th place now. In five years China is very likely to pass us in research in the absolute sense and as a fraction of GDP.

Finally, I would note that H.R. 5120 that was introduced last year contributes in a major way to solving what I think are some of the problems we have at translating the research that goes on in the DOE laboratories to the commercial sector, and I would be happy to address that further should the Committee wish. Thank you very much.

[The prepared statement of Mr. Augustine follows:]

TESTIMONY
OF
NORMAN R. AUGUSTINE
BEFORE THE
SUBCOMMITTEE ON ENERGY
(COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY)
OF THE
UNITED STATES HOUSE OF REPRESENTATIVES

JANUARY 28, 2015
WASHINGTON, DC

Good morning, Mr. Chairman and Members of the sub-committee. Thank you for inviting me to speak with you today on a subject that I consider to be of great importance to the nation: scientific research. This of course includes research related to high performance computing, the subject that I understand to be the principal focus of this hearing.

In the spirit of disclosure, I am a member of the board of directors of the Bipartisan Policy Council and co-chair of the American Energy Innovation Council, the latter is a group created by seven of us who are, or were, CEO's of non-energy companies. Our objective is to seek greater support for energy research because of its important potential impact upon the nation. I am an engineer and businessperson, not a researcher, and the views I express will be my own.

In my remarks today I would like to indicate why I consider research to be of such importance; very briefly address high performance computing; speak to the role of the National Laboratories; and conclude with an assessment of where America stands in the global, increasingly competitive research enterprise.

But first let me offer a few words about the nature of research. It is through research that scientists discover knowledge that enables engineers to create new products and services that through the efforts of entrepreneurs are introduced into the market to better people's lives. It has been pointed out, for example, that while Apple deservedly receives a great deal of credit for the iPhone, iPad and iPod, it was not Apple that made those devices possible. Rather, that distinction goes to scientists who conducted research decades ago in such fields as solid state physics and quantum mechanics. It is highly likely that as they conducted their research they had no inkling of the impact their efforts would have on products that would change people's lives; they were simply seeking to unlock secrets of nature.

I find it convenient to conceptualize research into three related categories. The first of these is purely curiosity-driven basic research—for example, studies in astrophysics of black holes. The second category is directed basic research which, for example, includes fundamental studies of the human genome with the intent of learning something useful in the prevention of cancer. The third is applied research, an example of which would be studying various materials in order to create faster electronic components to be used in high-performance computers.

Inherent to *basic* research is that one cannot know *a priori* what may be its ultimate application or impact. There is an abundance of examples of this, including research on seals in Antarctica that unexpectedly led to saving the lives of many children undergoing lung surgery; research on the chemistry of butterfly wings that led to an ingredient used in chemotherapy; or research on bacteria that cause abscesses which led to the accidental discovery of penicillin.

As to the overall importance of research, I will cite just three categories of examples. The first of these concerns the creation of jobs. In the aggregate it can be shown that adding jobs generally requires growing the nation's Gross Domestic Product. Numerous studies, one of which formed the basis of a Nobel Prize, demonstrated that 50 to 85 percent of the increase in America's GDP is attributable to advancements in just two fields: science and technology.

Further, the very foundation of advancements in science and technology is knowledge—and knowledge is the product of research.

Turning next to the subject of the health of the nation's citizenry, during the past century life expectancy in America grew from 47 years to 79 years—with much of the gain attributable to advancements in the field of medicine that were made possible through research in such widespread areas as vaccines, artificial joints, antibiotics, stents and advanced surgical procedures. Recently, when the Ebola outbreak suddenly burst forth, it was to research laboratories that the world's citizens looked for a solution—as had been the case with smallpox, polio, yellow fever, malaria, cholera and many other diseases.

Thirdly, virtually every product that we now take for granted in enhancing our everyday lives had its origin in research, whether it be television, video games, DVD's, electric lights, microwave ovens, GPS, electric cars, e-books, and much more.

Turning to the topic of research in high-performance computing, although perhaps an abstract sounding endeavor, past work related to this subject has produced extremely constructive impacts on our economy and our personal lives. It was earlier generations of such work that helped enable the human genome to be deciphered, criminals to be apprehended through computer matching of fingerprints, and medical imaging devices that "see" in three-dimensions. In my own early career as an aeronautical engineer we used giant wind tunnels requiring enormous amounts of power to determine the aerodynamic characteristics of potential aircraft designs. Today, this is accomplished in microseconds using computational methods. Given the massive databases that exist today, it is only through high-performance computing that we will be able to fully realize the benefits they potentially offer. Clean energy research is another example of a field that is highly dependent upon high-performance computing technologies.

Within the United States the Department of Energy operates 17 laboratories located throughout the country, the efforts of which are principally focused on energy research and the provision of weapons that underpin the nation's nuclear deterrent. Because the laboratories enjoy *relatively* stable funding they are ideally suited to fill a role that would otherwise be largely neglected; namely long-term, high-risk/high-payoff, often-large projects with applicability that may not be evident at their outset. Support of research in commercial nuclear fusion and hydraulic fracturing to produce shale gas would be but two examples of such endeavors.

Next, I would like, particularly given its importance, to address the health of America's research enterprise. In the past century the federal government financially supported two-thirds of the nation's research and development activity but that has gradually declined to one-third. Industry, on the other hand, has increased its share from one-third to about two-thirds. The problem is that, because of financial market pressure for rapid returns, industry focuses largely on "D," not "R." The result has been that in terms of arguably the most significant measure of national research investment, research funding as a fraction of GDP, the United States has recently dropped from first to seventh place in the world. Even funding of biomedical research, generally strongly supported by the public, has been cut by 22 percent in real dollars over the past decade. The extent of America's disinvestment in research is such that America now ranks 29th among developed nations in the fraction of research that is governmentally funded. It is projected that within about five years China will surpass the U.S. in both research funding as a fraction of GDP and absolute funding. This does not portend well for national security, jobs, the economy or the physical health of the citizenry.

Finally, returning to the subject of the National Laboratories, I would call attention to the large body of research that is conducted therein that may have potential applications in industry in fields well beyond energy and national security. Unfortunately, in my view relatively little of this potential is being realized by American industry as it seeks to compete in the global marketplace. Among the many reasons for this, one is that industry, especially small firms, has little idea what research is being conducted at the national laboratories. A second reason is that well-intended conflict of interest rules make it difficult for the laboratories to work closely with industry and also discourage the best means of technology transfer, the movement of people between government and industry. Other nations seem to have found solutions to these problems, albeit not without accepting certain risks. It is my view that the national laboratories are generally well run and are a national treasure that could make an even greater contribution than is the case today.

I will conclude my remarks by addressing the question that often seems to be on people's minds when they observe my commitment to strengthening research in America. Why, they ask, would a fellow creeping up on 80-years of age, a non-researcher, view this as such a critical issue. The reason is that everything I have observed in my roles in industry, government and academia suggests that other than our freedom and Free Enterprise system, education and research are the nation's most fundamental assets.

Thank you.

NORMAN R. AUGUSTINE was raised in Colorado and attended Princeton University where he graduated with a BSE in Aeronautical Engineering, magna cum laude, and an MSE. He was elected to Phi Beta Kappa, Tau Beta Pi and Sigma Xi.

In 1958 he joined the Douglas Aircraft Company in California where he worked as a Research Engineer, Program Manager and Chief Engineer. Beginning in 1965, he served in the Office of the Secretary of Defense as Assistant Director of Defense Research and Engineering. He joined LTV Missiles and Space Company in 1970, serving as Vice President, Advanced Programs and Marketing. In 1973 he returned to the government as Assistant Secretary of the Army and in 1975 became Under Secretary of the Army, and later Acting Secretary of the Army. Joining Martin Marietta Corporation in 1977 as Vice President of Technical Operations, he was elected as CEO in 1987 and chairman in 1988, having previously been President and COO. He served as president of Lockheed Martin Corporation upon the formation of that company in 1995, and became CEO later that year. He retired as chairman and CEO of Lockheed Martin in 1997, at which time he became a Lecturer with the Rank of Professor on the faculty of Princeton University where he served until 1999.

Mr. Augustine was Chairman and Principal Officer of the American Red Cross for nine years, Chairman of the Council of the National Academy of Engineering, President and Chairman of the Association of the United States Army, Chairman of the Aerospace Industries Association, and Chairman of the Defense Science Board. He is a former President of the American Institute of Aeronautics and Astronautics and the Boy Scouts of America. He serves on the Board of Trustees of the National World War II Museum and is a former member of the Board of Directors of ConocoPhillips, Black & Decker, Proctor & Gamble and Lockheed Martin, and was a member of the Board of Trustees of Colonial Williamsburg. He is a Regent of the University System of Maryland (12 institutions), Trustee Emeritus of Johns Hopkins and a former member of the Board of Trustees of Princeton and MIT. He has been a member of advisory boards to the Departments of Homeland Security, Energy, Defense, Commerce, Transportation, and Health and Human Services, as well as NASA, Congress and the White House. He was a member of the Hart/Rudman Commission on National Security, and served for 16 years on the President's Council of Advisors on Science and Technology under both Republican and Democratic presidents. He is a member of the American Philosophical Society, the National Academy of Sciences and the Council on Foreign Relations, and is a Fellow of the National Academy of Arts and Sciences and the Explorers Club.

Mr. Augustine has been presented the National Medal of Technology by the President of the United States and received the Joint Chiefs of Staff Distinguished Public Service Award. He has five times received the Department of Defense's highest civilian decoration, the Distinguished Service Medal. He is co-author of *The Defense Revolution* and *Shakespeare In Charge* and author of *Augustine's Laws* and *Augustine's Travels*. He holds 33 honorary degrees and was selected by Who's Who in America and the Library of Congress as one of "Fifty Great Americans" on the occasion of Who's Who's fiftieth anniversary. He has traveled in 112 countries and stood on both the North and South Poles of the earth.

07/2014

Chairman WEBER. Thank you, Mr. Augustine.
And now, we recognize Dr. Giles.

TESTIMONY OF DR. ROSCOE GILES, CHAIRMAN, DOE ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE

Dr. GILES. Thank you, Chairman Weber, and Ranking Member Johnson, and Members of the Committee. Thank you for inviting me to testify today and thanks for your support of the outstanding scientific and technical activities we are here to discuss.

The Advanced Scientific Computing Advisory Committee, ASCAC, which I chair, is a panel of experts that advises DOE under FACA rules about activities of the Office of Advanced Scientific Computing Research, ASCR. My testimony is largely based on ASCAC reports. I will address the value of research supported directly and indirectly by ASCR and also the technological challenges and rewards represented by U.S. leadership in this field.

The computing needs of science have grown exponentially, paralleling the exponential increases in computer power we have seen in recent decades sometimes pushing the computer industry for new capabilities and sometimes finding novel ways to exploit existing technology. The combination of computing power and the ability to transport, store, and learn from vast amounts of data is critical to U.S. leadership in a wide range of scientific and technical fields.

ASCR has enabled DOE scientists to harness unprecedented computing power to better understand the physical world, design new materials and devices, and engineer new and improved methods for energy production, utilization, and distribution. Recent examples include microscopic modeling of nuclear reactor core startup that can improve reactor efficiency and safety; simulations of complex combustion making the chemistry and physics of fluids and gases to the observed behavior of engines and reactor; predictive modeling of materials for lithium air batteries systems potentially able to store 10 times as much energy as lithium ion batteries; wheat genome sequencing previously impossible to do is now possible in under 32 seconds using new programming methods developed by ASCR; and modeling the surface of human skin to understand its properties and how chemicals might affect it. My written testimony includes many additional examples.

ASCR enables such outcomes by designing and deploying an effective system of world-class facilities for computing, data science, and networking in DOE labs making available expert staff to work with scientists to push the envelope of applications and supporting research in computer science in applied mathematics leading to key advances in software, hardware, algorithms, and applications.

Success also depends on a knowledgeable workforce and an educational pipeline to create that workforce. ASCR supports both training programs for scientists and the renowned Computational Science Graduate Fellowship program, CSGF. ASCR nurtures all elements of the ecosystem for scientific computing.

What about the future? ASCR has consistently provided leadership to DOE, the Nation, and the world by accelerating the development of new computing capabilities that can transform science.

When I last appeared before this Subcommittee in May of 2013, we testified about the importance of funding the development of exascale computing and the dangers to U.S. leadership in computational science if we fail to move expeditiously. Since that time, the urgency has increased, as has our knowledge of how to proceed.

In February 2014, ASCAC reported to DOE on the top 10 exascale research challenges. This report reflected the progress since our earlier 2010 exascale report. In addition to identifying the 10 challenges, our expert panel emphasized both that the United States has the technical foundation to address and overcome them and that it is critical that we do so.

In August 2014 the Secretary of Energy Advisory Board Task Force on Next-Generation Computing, of which I was a participant, made public in its draft report, which included the recommendation that DOE move forward with next-generation computing at the exascale level. The report also endorsed continued use of the co-design process and of government-industry-academic partnering mechanisms. ASCR, in collaboration with the National Nuclear Security Administration, has developed the preliminary plan for such an exascale computing initiative. This plan was provided to ASCAC for review last November. This review is actively in process with the resulting report due in September 2015 and an interim report at the end of March.

I think it is more important than ever for the United States to maintain and extend its leadership in scientific computing. I hope that our presence here today will help to that end. Thank you very much.

[The prepared statement of Dr. Giles follows:]

United States House of Representatives
Committee on Science, Space and Technology
Subcommittee on Energy

Supercomputing and American Technology Leadership
January 28, 2015

Testimony

Roscoe C. Giles, Ph.D.
Chair, Advanced Scientific Computing Advisory Committee (ASCAC)
Office of Science, U.S. Department of Energy

Professor, Department of Electrical and Computer Engineering,
Boston University
roscoe@bu.edu

OPENING

Good morning Mr. Chairman and members of the Subcommittee.
Thank you for inviting me to testify today and thanks for your support for the outstanding scientific and technical activities we are here to discuss.

I am testifying from the perspective of the chair of ASCAC (the Advanced Scientific Computing Advisory Committee), which reports to the DOE Office of Science under FACA rules. I will refer to several published ASCAC reports as part of my testimony. I will address the value of research supported directly and indirectly by ASCR and also the technological challenges and rewards represented by U.S. leadership in this field.

Science's computing needs have grown exponentially – paralleling the exponential increases in computer power we have seen in recent decades – sometimes pushing the computer industry for new capabilities and sometimes finding novel ways to exploit existing technology. Leadership in scientific and technical computing has been critical to U.S. Leadership in science and technology.

ASCR has enabled DOE scientists to harness unprecedented computing power applied to increasing our understanding of the physical world, designing new materials and devices, and engineering new and improved methods for energy production, utilization, and distribution.

Some recent examples described below include simulations that give us new insight into the behavior of Nuclear Reactors, complex burning flames, materials for Li-Air battery systems, the surface of human skin, and the fate of a Type IIb supernova. Overall, ASCAC has been very pleased with the depth, breadth, and significance of ASCR enabled research.

ASCR achieves these outcomes by: designing and deploying an effective system of world class facilities for computing, data science, and networking in DOE labs; making available expert staff to work with scientists to push the envelope of applications; and supporting research in computer science and applied mathematics leading to key advances in software, hardware, algorithms, and applications.

In addition, ASCR has consistently provided leadership to DOE, the nation and the world by accelerating the development of new kinds of computing systems with transformational impact on DOE science and science and engineering more broadly. A productive computing environment requires not only the most advanced hardware and software, but also depends on underlying mathematics and algorithms, and a knowledgeable workforce and educational pipeline to create that workforce. ASCR is active in nurturing all elements of the "ecosystem" for scientific computing.

When I last appeared before this subcommittee , others and I testified about the importance of funding the development of "exascale" computing and the dangers to US leadership in computational science if we failed to move expeditiously.

In February 2014, ASCAC reported to DOE on the "Top 10 Exascale Research Challenges." I will summarize the conclusions of this report the in testimony below. The report identifies significant challenges which U.S. technology has the ability to address and overcome, and which will contribute to our technological leadership more broadly.

DOE and ASCR in collaboration with NNSA have planned an exascale computing initiative (ECI) which was provided to ASCAC for review last November. This review is in process with the resulting report due in September 2015.

I think it is critical for the U.S. to maintain and extend its leadership in scientific computing as evidenced in by DOE and ASCR's activities being discussed today.

Background

My name is Roscoe C. Giles and I am a Professor in the Department of Electrical and Computer Engineering at Boston University. I have been involved for many years in leadership roles in computational science and high performance computing and in computational science education.

In particular, I have been a long time member and am currently chair of the DOE Office of Science Advanced Scientific Computing Advisory Committee (ASCAC). I have also been a participant on the recent Secretary of Energy Advisory Board's (SEAB) Task Force on Next Generation High Performance Computing (NG-HPC).

About ASCAC

ASCAC was first constituted in 1999 and is chartered under the Federal Advisory Committee Act (FACA). ASCAC members are appointed by the Undersecretary for Science and are experts in their fields. We report to the Director of the Office of Science in response to formal charges. We are not paid for our work on ASCAC. We hope to provide a useful external, community perspective on the impact, significance, and directions of ASCR efforts. Our committee meetings and reports are public.

Charges to ASCAC range from reviews of program management and effectiveness - for example we supervise regular Committees of Visitors for ASCR research program areas – to major reviews of strategic areas of emphasis and plans, such as the reports on the Exascale Computing and on Data Intensive Science. Charges are generally handled by subcommittees consisting of a few ASCAC members together with selected external participants chosen for their expertise in the specific area of the report.

Additional Materials

The bulk of the written materials I wish to draw to the attention of the Subcommittee are in the form of published ASCAC reports including the latest reports on the Top 10 Challenges for Exascale Research and the Workforce needs and also the Draft Report of the SEAB Next Generation High Performance Computing task force. These are listed on the accompanying citation page. There is also additional information of interest on the DOE ASCR website.

ASCR Overview

ASCR's mission is "*...to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the Department of Energy (DOE).*"

In pursuit of this mission, ASCR has programs and investments that include:

Research:

- Applied Mathematics Research whose fruits are essential for current applications and which provide the algorithmic framework for future applications and systems.
- Computer Science system and software research whose results both enable applications of current systems and chart the direction for future systems.

Facilities:

- Computer, Networking and Data Management facilities to meet the needs of DOE Science programs and their thousands of users.
- Leadership Computing Facilities with unique high-end capabilities made available both to DOE and to the entire nation, including industry (using DOE's long-established user-facility mission).
- Networking and data management facilities increasingly critical to scientific computing and to all large multi-user scientific instruments.

Human Infrastructure & Computational Ecosystem:

- Programs that enable applications scientists from DOE, Industry, and Academia to use effectively use advanced computing for discovery and problem solving.
- Developing the scientific computing and data science workforce

Advancing the Frontiers of Computing:

- Accelerating the development of new computing paradigms, capabilities, and facilities.
- Accelerating the adoption of advanced computing at DOE and in the nation.

In my last appearance before the subcommittee in May of 2013, my testimony focused mainly on ASCR facilities and the need for future exascale class systems. Today, I would like to touch more lightly on the facilities characterization and discuss in more depth some of the impacts of ASCR enabled science and the human infrastructure programs that have made them possible. I will also discuss ASCR's leadership in advancing the future of scientific computing.

Facilities Overview

The **Energy Sciences Network**, or ESnet, is the Department of Energy's high-speed network that provides the high-bandwidth, reliable connections that link scientists at national laboratories, universities and other research institutions, enabling them to collaborate on some of the world's most important scientific challenges. Managed and operated by the ESnet staff at Lawrence Berkeley National, ESnet provides direct connections to more than 40 DOE sites at speeds up to 100 gigabits per second, allowing scientists to manage, share and analyze massive datasets that are the hallmark of 21st century science. Most recently ESnet completed an expansion to Europe that will provide an overall capacity of 340 gigabits per second to support Office of Science experiments and partnerships there. ESnet derives its effectiveness from the extensive cooperation it enjoys with its user community. It is one of the most widely based and successful cooperative efforts within the Department of Energy.

The **National Energy Research Scientific Computing Center** (NERSC) at Lawrence Berkeley Lab is the primary scientific computing facility for the Department of Energy's Office of Science. As one of the largest facilities in the world devoted to providing computational resources and expertise for basic scientific research, NERSC is a world leader in accelerating scientific discovery through computation. More than 5,000 DOE scientists use NERSC systems annually to perform basic scientific research on more than 600 projects spanning a wide range of disciplines. NERSC users consistently publish more than 1,500 peer-reviewed publications each year.

The **Oak Ridge National Laboratory Leadership Computing Facility** (OLCF) was established in 2004 and was charged with developing an unclassified computing resource 100 times more powerful than the systems of the day. Today the OLCF is home to Titan, the United States' fastest and most powerful supercomputer dedicated to open scientific research. Titan ranked as 2nd most powerful computer in the world according to the November 2014 "Top500" list. In 2014, nearly two billion processor hours on Titan were awarded to projects from universities, private industry, and government research laboratories, representing a wide array of scientific and engineering research, from climate science to critical materials discovery and to nuclear physics. The OLCF operates a Liaisons program to assist INCITE science teams is porting and optimizing software on the OLCF machines.

The **Argonne Leadership Computing Facility** (ALCF) houses world-class supercomputing resources for open science. It supports a wide range of science and engineering research and serves users from academia, industry, and the national laboratories. ALCF's Mira supercomputer ranked as the 5th most powerful in the world. The ALCF provided over 3 billion processor hours on Mira to researchers in 2014 as well as comprehensive services from training to performance engineering to data analysis, and operates a unique catalyst program to assist the individual science teams to achieve optimal performance and results on ALCF systems from day one.

ASCR Impacts

ASCR's success is ultimately reflected in the scientific productivity, deepening insights, results, and technologies of DOE Science.

There exists a growing stream of scientific successes resulting from the computing capabilities enabled by ASCR. The foundation for this is that ASCR has created an ecosystem that enables and encourages application scientists, computer scientists, and mathematicians to work together on world-class DOE lab computing facilities, in order to solve problems that were considered intractable in the past.

Programs leading to significant science impacts

ASCR's "Scientific Discovery through Advanced Computing" Program (SciDAC) has had three incarnations over the years. In each case the focus has been precisely on enabling applications scientists to use the most advanced ASCR facilities to solve their problems. Successive incarnations of SciDAC have refined how directly DOE Science program offices were involved in the program administration and how the requisite interactions between applications, computer scientists, and mathematicians were managed.

A recent ASCAC Committee of Visitors (2) reaffirmed that "SciDAC remains the gold standard nationally and internationally for fostering interaction between disciplinary scientists and HPC." The current focus of SciDAC is insuring that DOE application scientists are able to effectively execute DOE mission science on the current mid-Petascale generation of supercomputers and the program officers from ASCR and other Science offices work together to fund and manage the program. This is an exemplar of cooperation and collaboration between program offices.

INCITE ("Innovative & Novel Computational Impact on Theory and Experiment") has made resources at leadership computing facilities available competitively to DOE and external scientists and engineers, including industry(9). The fundamental criteria for INCITE awards is the potential to perform transformational computing on the leadership computing facilities at OLCF and ALCF. INCITE awards are open to DOE scientists and also scientist from around the nation and the world.

The leadership computing facilities offer considerable staff support for the implementation and development of projects through the previously mentioned Catalyst and Liaisons programs at ALCF and OLCF. In this sense, INCITE projects are all partnerships between ASCR and the application research groups.

The ASCR Leadership Computing Challenge (ALCC) programs at the LCF's offer similar access to leading edge facilities and support in a more flexible way which allows additional opportunities for access by DOE applications and discretionary opportunities for access to LCF's by applications, particularly from industry, which are being developed outside the usual DOE and academic scientific circles.

Finally, ASCR SBIR-STTR (Small Business Innovation and Technology Transfer) activities target high-end computing and data science as an enabler of new commercial enterprises.

Some Recent Examples of ASCR Impact

- **2014:** Impossible to compute wheat sequencing now possible in seconds: ASCR research in programming environments, specifically X-Stack[1] DEGAS technologies, has enabled human genome sequencing within 20 seconds, reducing sequencing time by 2X from previous approaches. Wheat genome sequencing, which has been impossible to do, is now possible under 32 seconds using DEGAS.

- **2014: Confirmed: Stellar Behemoth Self-Destructs in a Type IIb Supernova:** For the first time ever, astronomers have direct confirmation that a Wolf-Rayet star died in a violent explosion known as a Type IIb supernova. Using the intermediate Palomar Transient Factory pipeline, supported by resources at NERSC and ESnet, researchers caught supernova SN 2013cu within hours of its explosion. These stars are interesting because they enrich galaxies with the heavy chemical elements that eventually become the building blocks for planets and life.

- **2014: Simulations Shed Light on Pine Island Glacier's Stability.** The rapid retreat of Antarctica's Pine Island Glacier has perhaps reached a point of no return, say three international modeling teams who ran a number of simulations to model the glacier's behavior. To do this work, they relied on three different ice-flow models including BISICLES[2], a collaborative software package developed in part by ASCR researchers.

Industry, Academia, and the National Labs use ASCR HPC resources to advance basic science and applied research topics in a broad-spectrum of technology areas including: nuclear energy; biofuel production; materials science; enzyme design; photovoltaics; engine combustion; electronics and superconductivity; turbomachinery for wind, carbon -sequestration and gas-turbines; turbulence modeling for noise reduction in wind-turbines and jet engines.

- **Procter & Gamble and Temple University scientists model skin's makeup.** Researchers at Procter & Gamble and Temple University used DOE's Titan to better understand the three-dimensional structure of skin's outermost barrier, the stratum corneum. Access through the INCITE program enabled some major achievements. One is the simulation of 1 million atoms of skin lipid matrix—four separate bilayers, each hundreds of square

---

[1] (exascale systems software stack)

[2] Berkeley-ISICLES

nanometers and made of ceramides (waxy compounds), fatty acids, and cholesterol in water.

This skin simulation, extending over 2 microseconds, was far too large and complex to have been carried out without OLCF resources. A related achievement is the validation of a state-of-the-art empirical model, the Kasting model, for the permeation of compounds through the stratum corneum. The project's third major accomplishment is the modeling of contact between skin and large concentrations of selected chemical compounds to uncover the mechanism by which compounds disrupt the stratum corneum barrier and are transported through it.

- **Related Publication**: M. Paloncyova, R.H. DeVane, B.P. Murch, K. Berka, M. Otyepka, "Rationalization of Reduced Penetration of Drugs through Ceramide Gel Phase Membrane," *Langmuir*, **30** (46) 13942-13948 (2014); DOI: 10.1021/la503289v. Published: NOV 25 2014
- https://www.olcf.ornl.gov/2014/11/14/procter-gamble-and-temple-university-scientists-model-skins-makeup/

- **"The Complexities of Combustion" (Jackie Chen, SNL).** Researchers from Sandia National Laboratories, including PI Jackie Chen and team member Ankit Bhagatwala, employed the direct numerical simulation code (DNS) known as S3D on Titan to simulate a jet flame burning dimethyl ether, an oxygenated biofuel, in an attempt to match the conditions of a companion experiment at The Ohio State University (OSU).

  The jet flame configuration is used to probe fundamental turbulent flame physics associated with local extinction, where parts of the flame extinguish, and are subsequently re-ignited through turbulent mixing, a finite-rate chemical process that may occur in practical combustors including diesel jet flames. If researchers can identify strategies to minimize flame extinction, this will greatly enhance efficiency and minimize undesired emissions in combustion devices such as engines. While Chen's team has simulated jet flames in the past, the latest simulations on Titan were a breakthrough for two reasons: the inclusion, for the first time, of dimethyl ether (DME), a more complex and oxygenated fuel, and the highest Reynolds number ever achieved by the team, 13,050, which is a measure of the turbulent mixing intensity.

  The increased Reynolds number allows the team to resolve a wider range in turbulence scales with detailed chemistry, a major breakthrough when trying to match experimental conditions and also for evaluating turbulent mixing and combustion models. "These simulations represent the first time we've incorporated DME and the highest Reynolds number ever achieved in a fully resolved reacting direct numerical simulation," said Ramanan Sankaran, an OLCF scientific liaison who assists Chen's team with their simulations on Titan.

- Related Publication: A. Bhagatwala, Z. Luo, H. Shen, J. Sutton, T. Lu and J. H. Chen, "Numerical and experimental investigation of turbulent DME jet flames," *Proc. Combust. Inst.* **35**, 1157-1166 (2015).
- https://www.olcf.ornl.gov/2014/11/11/the-complexities-of-combustion/

- **Westinghouse-CASL team wins major computing award for reactor core simulations on Titan.** The Consortium for Advanced Simulation of Light Water Reactors (CASL) team, led by Fausto Franceschini of Westinghouse and Andrew Godfrey of ORNL, performed core physics simulations of the Westinghouse AP1000 pressurized water reactor (PWR) core using CASL's Virtual Environment for Reactor Application (VERA).Westinghouse is deploying the AP1000 worldwide with eight plants currently under construction in China and the United States.

- The simulations, performed on Titan's Cray XK7 system, produced 3D, high-fidelity power distributions representing conditions expected to occur during the AP1000 core startup and used up to 240,000 computational units in parallel. One of the neutron transport components of VERA, the Exnihilo code suite developed at ORNL, was used for the simulations. The code includes deterministic transport solvers such as Denovo, which can take advantage of Titan's NVIDIA graphics processing unit accelerators, as well as a new stochastic transport module known as Shift.

  The results included as many as one trillion particle histories per simulation to reduce statistical errors and provide insights that improve understanding of core conditions, helping to ensure safe startup of the AP1000 PWR core.
  - https://www.olcf.ornl.gov/2014/07/02/westinghouse-casl-team-wins-major-computing-award-for-reactor-core-simulations-on-titan/
  - Thomas M. Evans, Wayne Joubert, Steven P. Hamilton, Seth R. Johnson, John A. Turner, Gregory G. Davidson, and Tara M. Pandya, "Three-Dimensional Discrete Ordinates Reactor Assembly Calculations on GPUs", ANS MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method, Nashville, TN, April 19-23, 2015, American Nuclear Society, LaGrange Park, IL (2015).
  - F. Franceschini, B. Oelrich Jr., J. Gehin,, "Simulation of the AP100 First Core with VERA", Nuc. Eng. Inter. 59, (718), 33-35 (2014):

- **Converting Greenhouse Gas CO_2 Into Fuel and Useful Chemicals** *(Victor Batista, Yale University)***.** Scientists using supercomputers at the National Energy Research Scientific Computing Center (NERSC) have discovered a mechanism that drives the conversion of the greenhouse gas carbon dioxide (CO_2) into usable fuels and chemicals. This mechanism holds the promise of being able to provide carbon-neutral fuels and furthering the goal of U.S. energy independence. The simulations run at NERSC reveal that, unexpectedly, the CO_2 conversion process is initiated by a reaction with hydrogen atoms bound to a platinum metal surface. Other mechanisms had been proposed before these simulations settled the issue. The findings will be useful in the design and development of new technologies that can generate fuels that are consumed without producing CO_2.
    - http://www.nersc.gov/news-publications/news/science-news/2013/turning-greenhouse-gases-into-gold/
    - http://pubs.acs.org/doi/abs/10.1021/jz400183z

- **More Efficient Thin Film Solar Cells** *(Yanfan Yan, U. Toledo).* Researchers have made a discovery that could lead to less expensive, more easily fabricated thin-film solar cells using supercomputer simulations performed at the National Energy Research Scientific Computing Center (NERSC). Scientists had known that treating cadmium-telluride (CdTe) solar cells with cadmium-chloride improves efficiency, but the detailed mechanism had been unknown until a team of researchers combined simulation and experimental analysis to show that the enhanced efficiency occurs when chlorine atoms substitute for many of the tellurium atoms near grain boundaries. Thin-film CdTe solar cells are considered a rival to currently used silicon-based photovoltaic systems because of their theoretically low cost per power output and ease of fabrication, but their efficiency has lagged that of traditional materials. This new understanding could be used to guide the engineering of high-efficiency CdTe materials.
    - http://www.nersc.gov/news-publications/news/science-news/2014/atomic-switcheroo-explains-origins-of-thin-film-solar-cell-mystery/
    - http://journals.aps.org/prl/abstract/10.1103/PhysRevLett.112.156103

- **A Previously Undiscovered Molecular Turn Leads Down the Path to Type 2 Diabetes.** Computing resources at the ALCF have helped researchers determine how proteins misfold to create the tissue-damaging structures that lead to type 2 diabetes. The researchers combined experiments and computation to understand the chemical pathway. The simulations located an entire step that had been missing, an intermediate step in which transient rigid fibrils form, then morph into floppy protein loops, which finally take the form of more tough fibrils and stack up to form the damaging amyloid fibril.

With the new understanding and access to Mira, future work could target a possible treatment, such as designing an inhibitor to interfere with the harmful pathway. In addition, the research collaboration can apply the method to determine the intermediate steps in similar diseases such as Alzheimer's that are linked to the formation of amyloid fibrils.

- **2014: Decreasing time to solution for supersonic turbomachinery** *(R. Srinivansan, Ramgen Power Systems LLC)* Cost-effective methods to capture and store power-plant carbon emissions are a principle barrier to widespread application of CCS. Simulations at OLCF transformed Ramgen's process for designing cost effective carbon capture and sequestration turbomachinery, cutting the projected time from concept to commercial testing by at least two years and the cost by over $4 million. Based off previous HPC simulations at OLCF, Ramgen has been able to complete initial testing of a 13,000-horsepower CO2 compressor, with an additional stage of ongoing technology development utilizing HPC resources. This compressor is projected to reduce the capital costs of CO2 compression by 50 percent and produce a minimum of 25 percent savings in operating costs. Applying these cost savings to a new 400-megawatt clean coal plant would result in significant annual operating cost savings.
    - https://www.olcf.ornl.gov/2012/08/14/ramgen-simulates-shock-waves-makes-shock-waves-across-energy-spectrum/
    - https://www.olcf.ornl.gov/2014/06/13/ramgen-takes-turbomachine-designs-for-a-supersonic-spin-on-titan/
- **2014 Materials Science for Wind Energy –1 Million Molecule Freezing Water Droplet Simulation for Non-Icing Surfaces** *(M. Yamada (GE) at OLCF).* One of the factors restricting the growth of wind energy is that many of the world's windiest places are cold as well. Just as ice growth on aircraft wings limits aircraft operation, ice growth on wind turbine blades will reduce their efficiency, or require that the turbine be shut down. Re-engineering the surface of the turbine blades may solve this problem, but requires fundamental understanding of how ice forms on the blades. While previous computer simulations of freezing had used around 1,000 molecules, GE understood that much larger simulations were needed to study icing on wind turbines. Using Titan, GE simulated hundreds of water droplets, each including one million molecules for 6 different surfaces under a range of temperatures. These represent the most comprehensive simulation of water freezing on a surface ever performed The results will guide future testing in developing new anti-icing materials.
    - https://www.olcf.ornl.gov/2013/10/25/titan-propels-ge-wind-turbine-research-into-new-territory/

- **Predictive Materials Modeling for Li-Air Battery Systems.** Lithium-air (Li-air) batteries are viewed as a possible game changer for electric vehicles, but realizing their enormous potential is a very challenging scientific problem that requires the development of new materials for electrodes and electrolytes. With this INCITE project, scientists from Argonne and IBM Research are teaming up to use Mira to better understand the physical and chemical mechanisms needed to make Li-air batteries a reality.

  With the potential to store up to 10 times the energy of a Li-ion battery of the same weight, Li-air batteries are particularly appealing to researchers because of their theoretical energy density. But developing a viable Li-air battery is a long-range effort that requires scientific breakthroughs in materials design, chemistry, and engineering.

  One of the most significant hurdles is finding suitable materials for Li-ion-conducting electrolytes, which enable the transport of ions between the anode and the cathode and promote the diffusion of oxygen from the environment into the electrochemical cell.

  In a recent study, IBM researchers focused on the zirconium-containing, garnet-like lithium-lanthanum-oxide, known as LLZO (Li7La3Zr2O12), a promising material for solid-state electrolytes. To observe Li-ion migration, the team needed to simulate time scales in nanoseconds rather than picoseconds. They were able to achieve this factor-of-a-thousand improvement by implementing a highly efficient parallel version of metadynamics (a tool for accelerating rare events, such as mapping the conductivity of a material regardless of its complexity) into their simulations and by taking advantage of Mira's substantial power with full machine runs. This enabled the researchers to obtain, for the first time, the free-energy profile for Li-ion conductivity in LLZO. One of their key findings was that the presence of vacancies in cubic LLZO is crucial to lowering its activation energy and enhancing its Li-ion conductivity.

  IMPACT
  This project is providing insight into the complexities of the Li-air battery at the molecular level, including an understanding of the microscopic mechanism for high Li-ion conductivity. The results will help to inform the design of new materials for Li-air electrolytes and electrodes. If realized, Li-air batteries could enable widespread deployment of electric vehicles, greatly reducing U.S. dependence on foreign oil.

- **Improving Aircraft Engine Combustor Simulations.** A jet engine combustor combines air flowing faster than a hurricane with swirling fuel to generate the extraordinary release of heat that ultimately powers the aircraft. Understanding these complex physical and chemical interactions is critical to fuel efficiency and emissions performance, but physical testing can be difficult and time consuming.

Computer simulation of the complex physics of a combustor creates a "virtual test," thus reducing the need for physical testing. Pratt & Whitney explored leading-edge combustor design methods using the Argonne Leadership Computing Facility as part of DOE's INCITE program.

IMPACT

This INCITE project led to improved capabilities and reduced solution times for 3-D combustor simulations. The work has been a key enabler for the depth of understanding needed to meet emissions goals. INCITE simulation technologies were applied to Pratt & Whitney's next-generation, low-emission Geared Turbofan™ engine. This groundbreaking engine delivers unprecedented reductions in emissions, noise, and cost of ownership compared to engines of the previous state of the art.

- **Simulations of Electrochemical Oxide Interfaces at Mesoscale:** Electrochemical interfaces are responsible for the success of electrochemical energy conversion/storage systems, such as fuel cells and batteries, as well as the failure of materials caused by corrosion. This project aims to elucidate nanoscale and mesoscale mechanisms associated with electrochemical processes occurring at reactive material interfaces. Breakthroughs in the fundamental understanding of these interfaces are necessary in the design and development of novel oxide materials for energy applications.

  Dynamical electrochemical processes combine the remarkable complexity of many interfacial reactions, transport phenomena, and microstructural evolution with the formidable subtleties of material defect chemistry at the interface. Focusing on material synthesis and electrochemical interfaces, project researchers set out to understand these behaviors at the atomistic and molecular levels during an electric field-assisted oxide-sintering process.

  Sintering—the process of welding together two metal powders at a diminished liquid phase—is well established, though the processes governing electric field sintering phenomena at the nanoscale were largely unexplored. Using representative oxide models of zirconia and ceria variants, researchers have shown that electric field-assisted sintering can be used to design oxide materials with modifications that would affect their functional properties.

  Using an INCITE award, Argonne researchers conducted calculations of the atomistic scale and interfacial properties of nanoscale oxides and oxide heterostructures. Results have helped explain experimentally observed improvements in oxide quality resulting from the electric field-assisted sintering process.

- **Intensity-Dependent Dynamics in Fermilab and CERN Accelerators** *(James Amundson, Fermilab):*
  Particle accelerators are an enabling technology for both basic research (e.g., studying the fundamental constituents of matter and the structure of nuclei) and the applied sciences (e.g., probing the structure of materials). Accelerator technologies also have broad economic and societal impact. DOE-originated accelerator technologies are behind the tens of thousands of accelerators that are at work every day producing particle beams in hospitals and clinics, in manufacturing plants and industrial laboratories, in ports and printing plants and, literally, on the ships at sea. Adding them all up, some 30,000 particle accelerators operate in the world today in medicine, industry, security and defense and basic science. The market for medical and industrial accelerators exceeded $3.5 billion dollars a year in 2009, and it is growing at more than ten percent annually. All digital electronics now depend on particle beams for ion implantation, creating a $1.5 billion annual market for ion-beam accelerators. All the commercial products that are processed, treated or inspected by particle beams have a collective annual value of more than $500 billion.

  Fermilab researchers are using ALCF resources to perform complex accelerator simulations aimed at reducing the risks and costs involved in developing the world's highest-intensity particle beams.

  The future of high-energy physics requires running today's accelerators at intensities that are higher than ever. Both Fermilab and CERN strive to accurately understand intensity-dependent effects in their accelerator complexes. Such understanding requires detailed numerical modeling that goes beyond the capabilities of desktop machines and simple clusters. This project is working to simulate these accelerators with unprecedented fidelity.

  The Fermilab Recycler and Main Injector form the final high-energy stage of the Fermilab accelerator complex. During each acceleration cycle, the Recycler receives protons in six batches from the Booster. The simulations carried out by this INCITE project have captured theoretically predicted instabilities in the high-intensity neutrino beams that will be produced, including demonstrating the existence of an instability in the Fermilab Booster, a kind of instability generated by bunch-to-bunch effects that could potentially shut down the whole beam. These kinds of insights will help in the planning of a high-intensity beam project called the Fermilab Proton Improvement Plan II, which will create neutrino beams for the Long Baseline Neutrino Facility, and help considerably in the next phase of research at CERN's LHC.

  IMPACT
  While it is generally understood that particle accelerators are used to discover the fundamental matter of the universe, high-intensity beams in accelerators around the world allow scientists to explore the structure and

function of matter and materials more tangible to our everyday lives. The results can advance research in fields as diverse as materials science and renewable energy, as well as drive the long-term development of ideas once thought unfeasible, such as the transmutation of nuclear waste.

Presentations:

- "Accelerator Beam Dynamics on Multicore and GPU and MIC Systems," SIAM Conference on Parallel Processing for Scientific Computing, February 18–21, 2014, Portland, Oregon.
- James Amundson, "Transverse Instabilities in the Fermilab Booster," 16th Advanced Accelerator Concepts Workshop (AAC 2014), July 2014, San Jose, California.
- Eric Stern, "Developments in Synergia Space-Charge-Induced Resonance Trapping," 16th Advanced Accelerator Concepts Workshop (AAC 2014), July 2014, San Jose, California.

- **Increasing efficiency for automobiles through improved under-hood packaging.**

  Improving the air flow through the engine compartment of an automobile can reduce the drag and improve the efficiency of the vehicle in the same way improving the air flow around the automobile can. However, the engine compartment must also be arranged in a way that all engine components receive adequate cooling. The complexity of the problem is so large that Ford used the Jaguar and Titan supercomputers to perform accurate simulations of air flow through more than 50 million individual computational cells. More than 1,600 individual simulations were used to understand the effects of different packaging arrangements and different driving conditions. The results showed enough potential for improving automobile efficiency that Ford both updated its automobile design methods, and began improving its own high-performance computing capability.

  (Additional note: Ford has published papers based on the methods used in these simulations.)

- **In an EERE funded Project, GM and Ford study improved engine efficiency through High-Performance Computing**

  The auto manufacturers, the national laboratories, and the DOE's Office of Energy Efficiency and Renewable Energy have a long history of working together to develop simulations of engine operations that can be used to improve the efficiency of internal combustion engines without compromising their power or reliability. Both GM and Ford, with support from EERE (DOE Office of Energy Efficiency & Renewable Energy) and Oak Ridge National Laboratory, are currently using Titan's computational capability to test strategies for improving engine efficiency. GM is using OpenFOAM, an open-source computational fluid dynamics (CFD) code, to simulate new concepts for fuel injectors. These simulations will help GM develop "spark-ignited direct injection," or SIDI, which will improve the combustion efficiency of engines. Ford is using the commercial CFD code Converge to study the effects of exhaust gas recirculation (EGR), a process which improves fuel efficiency. Simulating EGR requires simulating the engine over multiple combustion cycles, greatly increasing the complexity of the simulation. In both cases, GM and Ford are not only using their simulations to improve engine operations, but are using their experience to develop software that will let other users perform simulations using Titan more effectively.

Example SBIR Activities

Industries are increasingly using modeling and simulation to improve productivity. Commodity IT is becoming more parallel and Industry can benefit from tools, software, and libraries developed in the HPC community. ASCR SBIR investments target small businesses that harden HPC tools and software, including those developed by ASCR researchers, to improve industry access to advanced HPC methods.

Success Examples:

Optimization Algorithms for Power Applications (*Ziena Optimization LLC*). Ziena, a US small business, used ASCR SBIR funds to integrate HPC decision algorithms for power applications into their optimization software program. Test results showed over 50% decreased solution time. The Ziena product is being tested by the European Commission. If successful, Ziena's product will become the leading optimization solver engine for managing power grids in Europe.

2014 Computational Fluid Dynamics Simulations and Real-Time Analysis (*David Philips, Cascade Technologies Inc*). Cascade Technologies increased usability of high-fidelity HPC CFD codes for non-HPC specialists by integrating the code into standalone web based "apps". The resulting product enabled faster analysis, reduction in wasted compute hours, and increased data provenance for end-users. Since the recent product launch, over 30 commercial, government, and academic users have benefited from the system.

2013 Graphical HPC Application for Product Simulation (*Robert M O'Bara, Kitware Inc*). Kitware developed easy-to-use Computational Model Building framework that connects users to HPC simulation codes. Kitware's product provides the ability to create and run proof-of concept product simulations and allows easily scaling to higher-end systems. The Army Corps of Engineers Engineering Research and Design Center are early adopters of the new product.

2013 Big Data Tools for Energy Materials (*Tal Sholklapper, Voltaiq*). Companies developing batteries and related technologies perform extensive cyclic testing, generating enormous quantities of data. Entrepreneurs at Voltaiq are used ASCR SBIR funds to develop cloud-based software to address a Big Data problem in energy materials sector. The new software product can dramatically increase the pace of battery innovation and development for thousands of organizations working in the battery sector.

ASCR Partnerships with other DOE Offices

The Office of Advanced Scientific Computing Research is currently working directly with DOE programs in Nuclear Energy and the Office of Electricity, which maintains resources that are used by EERE programs in Vehicles Technology and Wind Energy. ASCR has partnered with the Office of Electricity and Argonne National Laboratory to perform fundamental research on simulating the electrical grid that will allow development of strategies for making the grid both flexible enough to integrate renewables, and secure enough to provide energy to the nation. In Vehicles Technology, ASCR super-computers are used as part of industrial partnerships designed to increase fuel economy without compromising engine performance. ASCR computing capabilities have also been used by companies in wind energy that are bringing technology to market.

As computing power increases to the exascale, the potential gains from applying advanced simulation to applied technologies and key scientific problems also increases. Because of the challenges of preparing simulations to operate at the exascale, the DOE has already funded three exascale co-design centers: the Center for Exascale Simulation of Advanced Reactors (CESAR), the Center for Exascale Simulation of Combustion in Turbulence (ExaCT), and the Exascale Co-design Center for Materials in Extreme Environments (ExMatEx). CESAR's work will allow both simulation of new reactor designs, and create capabilities that will allow existing reactors to be operated safely and reliably. The capabilities ExaCT develops for simulation turbulent combustion using exascale simulation can be used to develop predictive models used to improve the efficiency of all combustion powered systems, including aircraft, automobiles, and coal and natural gas fired power plants. Finally, the discovery and design through simulation of materials that can operate at high temperatures and high pressures will benefit a range of technology areas, including combustion, nuclear energy, fuel cells, and concentrated solar power. These three examples are not the only potential impacts of exascale computing and ASCR is working with programs within the Office of Science and throughout DOE to identify specific challenges and opportunities for a wide array of applications.

Workforce Development

ASCAC, along with the other Office of Science Advisory Committees, was recently asked to identify disciplines in which significantly greater emphasis in workforce training at the graduate or postdoctoral levels is necessary to address workforce gaps in current and future Office of Science mission needs. Professor Barbara Chapman of the University of Houston led our workforce subcommittee whose report (3) concluded that:

> "Simulation and computing are essential to much of the research conducted at the DOE national laboratories. Experts in the ASCR-relevant Computing Sciences, which encompass a range of disciplines including Computer Science, Applied Mathematics, Statistics and domain sciences, are an essential element of the workforce in nearly all of the DOE national laboratories.
>
> Results of data analyzed are that the Computing Sciences workforce recruitment and retention activities are below the level necessary to sustain ASCR facilities and maintain DOE's high standards of excellence for innovative research and development. In particular, the findings reveal that:
>
> - All large DOE national laboratories face workforce recruitment and retention challenges in the fields within Computing Sciences that are relevant to their mission...
> - Insufficient educational opportunities are available at academic institutions in the ASCR-related Computing Sciences that are most relevant to the DOE mission.
> - There is a growing national demand for graduates in ASCR-related Computing Sciences that far exceeds the supply from academic institutions. Future projections indicate an increasing workforce gap and a continued underrepresentation of minorities and females in the workforce unless there is an intervention.
> - The exemplary DOE Computational Science Graduate Fellowship (CSGF) program, deemed highly effective in every one of multiple reviews, is uniquely structured and positioned to help provide the future workforce with the interdisciplinary knowledge, motivation, and experiences necessary for contributing to the DOE mission.
> - The DOE laboratories have individually developed measures to help recruitment and retention, yet more can be done at the national level to amplify and extend the effectiveness of their locally developed programs.
>
> The subcommittee recommendations:
>
> - Preserve and increase investment in the DOE CSGF program *(discussed in detail below)* while developing new fellowship programs modeled after the CSGF program to increase opportunities for more high-quality students, particularly students from underrepresented populations and demographics, in the computing sciences.

- Develop a recruiting and retention program that increases DOE's visibility on university and college campuses, establish uniform measures across DOE laboratories to improve the attractiveness of careers in DOE laboratories, and examine the laboratory funding model and its impact on recruiting and retention.
- Establish a DOE-supported computing leadership graduate curriculum advisory group to identify and raise visibility of graduate level curricular competencies specifically required to fulfill DOE's Computing Sciences workforce needs.
- Expand support for local laboratory programs, collect workforce data pertaining to the ASCR-related Computing Sciences, and encourage greater inter-laboratory sharing of information about locally successful programs and workforce related data.
- Working with other agencies, develop a strategic plan with programs and incentives to pro-actively recruit, mentor and increase the involvement of significantly more women, minorities, people with disabilities, and other underrepresented populations into active participation in CS&E careers."

The Computational Science Graduate Program (CSGF)
Since its inception in 1991 the DOE Computational Science Graduate Fellowship (CSGF) has played a vital role in ensuring a sufficient supply of skilled interdisciplinary scientists and engineers who can effectively address the Department's most complex challenges using the most advanced scientific modeling and simulation tools. CSGF has fulfilled not only the growing demand for these skills to address DOE mission needs, but has helped to meet the significant growth in national demand by industry, academia, and other government agencies, as these institutions increasingly seek highly trained, computationally skilled scientists and engineers. In recent years, CSGF also has begun to produce trained scientists who can deal with the challenges "big data."

The CSGF program, which is funded jointly by the DOE Office of Science's Office of Advanced Scientific Computing (ASCR) and the National Nuclear Security Administration (NNSA), is viewed as an exemplar in providing doctorate-level training of computational scientists and engineers who have the interdisciplinary knowledge and skills most relevant to the DOE mission areas. CSGF participants are selected competitively through vigorous review and the program provides for four years of support to graduate students and requires an approved interdisciplinary graduate program of study and research practicum with computational scientists at a DOE laboratory. In short, CSGF aims to create a new kind of scientist that is unobtainable presently from traditional academic programs that do not expose students to real-world computational-science applications.

As mentioned above, the Office of Science, in the past year, tasked its Advisory Committees, including ASCAC (7), to provide their respective expert assessments of the disciplines in which significantly greater emphasis in workforce training at the graduate level is necessary to address gaps in the current and future Office of Science workforce needs. All six committees, including ASCAC and representing disciplines from chemistry, materials, biological systems science, and climate science to high energy physics, nuclear physics, fusion and plasma sciences, identified computational sciences as in need of greater emphasis in workforce training.

The resounding multigenerational success of the CSGF program gives evidence to the impact the ASCR can have on the workforce. (Indeed, we have just had our first CSGF Alum join ASCAC this year).

I respectfully urge the Science Committee to put its weight behind this important workforce development endeavor.

Maintaining World Leadership in Computing

ASCR has consistently provided leadership to DOE, the nation and the world by accelerating the development of new kinds of computer systems with transformational impact on DOE science and science and engineering more broadly.

DOE has been a leader

Since the 1950s, DOE and its predecessor organizations have been at the forefront of scientific computing. Although this grew out of the national imperatives associated with nuclear weapons, it was recognized early on that high-performance computing had an essential role in the physical sciences and civilian research in applied mathematics and scientific computing was started in the organizations that have evolved into the Office of Science today. In the 1990s, as massively parallel computing established new frontiers in scientific computing, ASCR brought these tools to the U.S. scientific community, culminating in the petascale computers in use today. During this same time, ASCR supported development of cluster-based computing tools that are in widespread use today in research laboratories, universities and industries.

Emergence of Data Science

During the last decade, there has been increased recognition of the importance of data, often referred to as "big data" across both the research and commercial domains. In the commercial domain, this is apparent though the rise of Google, Facebook, and Twitter, to name just a few. In the scientific domain, DOE Office of Science has been at the forefront of applications of large data in the physical sciences, including data taken at large scientific facilities, such as the Large Hadron Collider, and observational data, such as taken by terrestrial-based telescopes. ASCR has actively supported these developments through applied mathematics and computer science research in the management and manipulation of large scientific datasets. The ESnet system, operated by ASCR, provides the backbone capability for the manipulation of these large datasets across U.S. scientific institutions.

ASCR High End Computing Development Activities

After an extraordinary series of community workshops, engaging DOE applications scientists and engineers, computer scientists, mathematicians, industry representatives and academics, ASCR developed the foundations for the "exascale" initiative: to build the technology – hardware, software, applications frameworks—that would allow for a machine to deliver a computational capability of 10^{18} operations per second and, along the way, enable remarkable advances at all intermediate scales of computing of relevance to science and industry. Key requirements of such computers is that they be energy efficient (goal is 20 megawatts), readily programmable, and reliable over the long run times of DOE's applications.

The term "extreme-scale computing" is used more broadly to refer to leadership systems across these scales, ranging from embedded processors to leadership facilities that will host exascale computers. A recent report from the National Research Council titled "The Future of Computing Performance: Game Over or Next Level?" (11) highlighted the importance of leadership in extreme-scale computing for US competitiveness.

Next Generation Computing Reviews and Reports

ASCAC: Exascale Computing

In 2009, ASCAC was charged with reviewing ASCR's body of work on exascale computing. Dr. Robert Rosner of the University of Chicago led our subcommittee on this charge. We delivered our review report – "The Opportunities and Challenges of Exascale Computing" (8) in Fall 2010. We found the case for exascale compelling and recommended that ***"DoE should proceed expeditiously with an exascale initiative so that it continues to lead in using extreme scale computing to meet important national needs."***

ASCAC: Data Intensive Science

In 2012 ASCAC was asked to consider the synergies between "big data" and the development of exascale computing. Professor Vivek Sarkar of Rice University chaired a subcommittee to address this issue. In 2013, ASCAC delivered the report "Synergistic Challenges in Data-Intensive Science and Exascale Computing." (5)

Its findings included:

1. There are opportunities for investments that can benefit both data-intensive science and exascale computing.
2. Integration of data analytics with exascale simulations represents a new kind of workflow that will impact both data-intensive science and exascale computing.

3. There is an urgent need to simplify the workflow for data-intensive science.
4. There is a need to increase the pool of computer and computational scientists trained in both exascale and data-intensive computing.

Recommendations included:

1. The DOE Office of Science should give high priority to investments that can benefit both data-intensive science and exascale computing so as to leverage their synergies.
2. DOE ASCR should give high priority to research and other investments that simplify the science workflow and improve the productivity of scientists involved in exascale and data-intensive computing.
3. DOE ASCR should adjust investments in programs such as fellowships, career awards, and funding grants, to increase the pool of computer and computational scientists trained in both exascale and data-intensive computing.

ASCAC: Top 10 Exascale Research Challenges
In July 2013, ASCAC was charged to identify "a list of no more than 10 technical approaches (hardware and software) that will enable the development of a system that achieves the Department's exascale goals, particularly the usability goals for the Department's mission critical applications."

Professor Robert Lucas of the University of Southern California Information Sciences Institute chaired a subcommittee of experts to respond to this charge. The report entitled "Top 10 Exascale Research Challenges" (4) was completed on February 10, 2014.

The Top Ten Exascale System Research Challenges
(not in a particular order):

1. Energy efficiency: Creating more energy-efficient circuit, power, and cooling technologies.
2. Interconnect technology: Increasing the performance and energy efficiency of data movement.
3. Memory Technology: Integrating advanced memory technologies to improve both capacity and bandwidth.
4. Scalable System Software: Developing scalable system software that is power-aware and resilience-aware.
5. Programming systems: Inventing new programming environments that express massive parallelism, data locality, and resilience
6. Data management: Creating data management software that can handle the volume, velocity and diversity of data that is anticipated.
7. Exascale Algorithms: Reformulating science problems and redesigning, or reinventing, their solution algorithms for exascale systems.
8. Algorithms for discovery, design, and decision: Facilitating mathematical optimization and uncertainty quantification for exascale discovery, design, and decision making.
9. Resilience and correctness: Ensuring correct scientific computation in face of faults, reproducibility, and algorithm verification challenges.
10. Scientific productivity: Increasing the productivity of computational scientists with new software engineering tools and environments.

In addition to the list the subcommittee strongly felt that it should give voice to the community's assessment of the importance of next generation computing to DOE and the nation. ASCAC concurred and provided the additional findings and recommendations below:

Findings:

1. <u>Exascale computing is critical for executing the DOE mission.</u>
   ASCAC reaffirms its findings from previous reports that leadership in high performance computing (HPC) is critical to achieving the DOE mission of ensuring U.S. leadership in science, engineering, and national security. In the

last six years, this has been documented in many exascale reports from Office of Science programs, the National Nuclear Security Administration, and other U.S. government agencies.

2. U.S. national leadership is at risk.
   Without aggressive investment and technical innovation in HPC, the U.S. risks falling behind rapidly emerging international competitors, not all of whom are friendly to U.S. interests. This in turn threatens to undermine the nation's intellectual leadership in a broad range of science, its economic position, and its security.

3. The U.S. has the technical foundation to create exascale systems.
   The U.S. semiconductor and HPC systems industries are capable of developing the necessary technologies for an exascale computing capability by the early part of the next decade, based largely on evolving commercially driven component fabrication, systems integration, and software engineering capabilities. However, for a truly effective and productive exascale computing capability, the U.S. government will need to focus investments on the research, development, and integration of HPC technologies that will otherwise not be created solely by commercial drivers.

4. An evolutionary approach to achieving exascale will not be adequate.
   The dramatic improvements essential to achieving effective exascale computing will not be satisfied by incremental extensions to today's conventional practices. Commercial market drivers do not provide a viable general path to delivering necessary scalability, time and energy efficiency, and user productivity including performance portability to future generation exascale class computers.

5. The U.S. government's continuous leadership and investment are required to create exascale systems.
   The U.S. computing industry is unlikely to develop effective exascale computer systems without U.S. government investment and focused mission goals. Innovation, sometimes of an incremental nature, and in other areas revolutionary, will be required under DOE direction to enable U.S. leadership in advanced HPC.

Recommendations:

1. DOE should invest in a program of continuous advancement in HPC.
   Exascale is only the next milepost in a half-century of continuous progress towards increasing capability in computational science. The U.S. government requires a stable, long-term investment strategy to ensure continuous U.S. leadership in HPC beyond today's petascale performance regime, extending to exascale and beyond. In the immediate future, much of that research investment should be focused on the top ten challenges identified within this report.

2. DOE should invest in the U.S. industrial base to catalyze the foundation for exascale systems.
   DOE should invest in extending commercial semiconductor, communications, systems integration, and software technologies to prepare the U.S. industrial base for its role and contributions in future HPC scientific, engineering and national security missions. All of these exascale components must be developed by and be available from U.S. sources, otherwise the supply chain is vulnerable to interdiction by foreign powers, which in turn could threaten the nation's security.
3. DOE should invest in exascale mathematics and system software responsive to DOE missions and other U.S. government requirements.
   The mathematical algorithms needed for many DOE missions are unique, and must be reinvented to function at exascale. As with today's Leadership-class systems, much of the software infrastructure of an exascale system will be unique to its scale and the missions for which DOE will deploy it. Therefore, DOE must invest in robust and scalable mathematical algorithms, operating systems, runtime systems, and tools for the management of the data that will be generated and/or processed.

SEAB: Next Generation Computing Task Force

The DOE Secretary of Energy Advisory Board (SEAB) was charged in December 2013 by Energy Secretary Moniz to "review the mission and national capabilities related to next generation high performance computing," including exascale computing. Specifically, SEAB was requested to examine and report on the justification for an exascale computing initiative, related basic research necessary to enable next generation high performance computing, the current state of technology and plans for an exascale program in DOE, the role of DOE in leading the development of exascale computing, and the implications of data centric computing for exascale computing. I was a participant in this task force.

SEAB provided its draft report (10) on August 10, 2014 (http://energy.gov/seab/downloads/report-task-force-next-generation-high-performance-computing) and its final report is expected to be released at the March 2015 SEAB meeting.

SEAB concluded in its preliminary report that investable needs exist for an exascale class machine; significant, but projectable technology development can enable one last "current" generation machine; "classical" high end simulation machines are already significantly impacted by many of the data volume and architecture issues; data-centric at the exascale is already important for DOE missions; common challenges and under-girding technologies span computational needs; the DOE National Labs are an important and unique resource for the development of next generation high performance computing and beyond; a broad and healthy ecosystem is critical to the development of exascale and beyond systems; and it is timely to invest in science, technology and human investments for "Beyond Next".

SEAB's recommendations in its preliminary report to Secretary Moniz include the following:

1. DOE, through a program jointly established and managed by the NNSA and the Office of Science, should lead the program and investment to deliver the next class of leading edge machines by the middle of the next decade. These machines should be developed through a co-design process that balances classical computational speed and data-centric memory and communications architectures to deliver performance at the 1-10 exaflop level, with addressable memory in the exabyte range.
2. This program should be executed using the partnering mechanism with industry and academia that have proven effective for the last several generations of leadership computing programs. The approximate incremental investment required is $3B over 10 years.
3. DOE should lead, within the framework of the National Strategic Computing Initiative (NSCI), a co-design process that jointly matures the technology base for complex modeling and simulation and data centric computing. This should be part of a jointly tasked effort among the agencies with the biggest stake in a balanced ecosystem.
4. DOE should lead a cross-agency U. S. Government (USG) investment in "over-the-horizon" future high performance computing technology.
5. DOE should lead the USG efforts to invest in maintaining the health of the underlying balanced ecosystem in mathematics, computer science, new algorithm development, physics, chemistry, etc.

ASCR path toward exascale

And, indeed, ASCR has been working in partnership with industry, lab personnel, and the community to move us along the path to exascale. Some program elements have included:

- Establishment of Co-Design centers to exploit a key element of effective extreme computing applications – the guided interplay of application/hardware/software in the design of systems.
- Computer Science Research: X-Stack software to develop tools for extreme scale systems, Advanced Architectures
- Applied Mathematics Research: Uncertainty Quantification, Extreme Scale Algorithms
- Prototypes: (joint with NNSA) FastForward, Design Forward
- Community: Exascale Research Conferences

The Exascale Computing Initiative (ECI)

DOE's plan for its Exascale Computing Initiative (ECI) was provided to ASCAC for review(1) on November 21, 2014 (http://science.energy.gov/~/media/ascr/ascac/pdf/meetings/20141121/Exascale_Preliminary_Plan_V11_sb03c.pdf). ASCAC's review is scheduled to be complete by September 2015.

A key strategy in the ECI is to work jointly with U.S. computer companies on the path to exascale computing. This effort is embodied in the FastForward program, which presently seeks to develop critical technologies needed to deliver next-generation affordable and energy-efficient technologies for extreme scale computing for the next decade. FastForward is joint between DOE Office of Science and National Nuclear Security Administration (NNSA) and involves participation by computing industry leaders, including AMD, Cray, IBM, Intel and NVIDIA. FastForward will be followed by successive R&D efforts that integrate these technologies into first single node prototypes and eventually into full exascale systems. Deployment is planned for the 2023 time frame.

Concurrently, DOE and NNSA are, through the ECI, developing the systems and applications software required to effectively use the planned exascale systems. Additionally, research is underway to develop new approaches for realizing the full potential of the expected increase in parallelism in exascale computers.

The next steps in facilities

In the interim until an ECI facility is available to do science, ASCR continues to deploy more powerful super computer facilities, reflecting the state of the art on the path to exascale.

The Collaboration of Oak Ridge, Argonne, and Lawrence Livermore (CORAL) was established by DOE to leverage supercomputing investments, streamline procurement processes and reduce costs to develop supercomputers that will be five to seven times more powerful when fully deployed than today's fastest systems in the US. DOE announced in November 2014 (http://energy.gov/articles/department-energy-awards-425-million-next-generation-supercomputing-technologies) its plans million to build two state-of-the-art supercomputers at the Department of Energy's Oak Ridge and Lawrence Livermore National Laboratories, respectively. Both CORAL awards leverage IBM's Power Architecture, NVIDIA's Volta GPU and Mellanox's Interconnect technologies to advance DOE's priority research initiatives for national nuclear deterrence, technology advancement, and scientific discovery. Oak Ridge National Laboratory's (ORNL's) new system, Summit, is expected to provide at least five times the performance of ORNL's current leadership system, Titan. Lawrence Livermore National Laboratory's (LLNL's) new supercomputer, Sierra, is expected to be at least seven times more powerful than LLNL's current machine, Sequoia. Argonne National Laboratory will announce its CORAL award at a later time.

Conclusions

- DOE through ASCR has enabled truly great computational science serving national needs modeling for the nation transformational impacts of modeling, simulation, and data science on our knowledge and well-being.
- ASCR has pioneered the development of next generation computing for the use of science. Confronting and overcoming the challenges of exascale technology will be a benefit to DOE science and the nation.
- U.S. Leadership in next generation computing should not be squandered.
- Thanks for your continuing support to grow the US capabilities and capacity and expertise in high end computing.

Citation List for Roscoe Giles' Testimony

ASCAC Reports

(All available from links on the report Section of the Advanced Scientific Computing Advisory Committee website: http://science.energy.gov/ascr/ascac/reports/)

(1) Charge letter to ASCAC to review the draft preliminary conceptual design for the Exascale Computing Initiative, November 19, 2014: http://science.energy.gov/~/media/ascr/ascac/pdf/meetings/20141121/Binkley_Exascale_charge.pdf

(2) Report of Committee of Visitors (COV) to review the management processes for the SciDAC portfolio, March 28, 2014: http://science.energy.gov/~/media/sc-2/pdf/cov-ascr/2014/ASCR_COV_2014_SciDAC_Report.pdf

(3) Assessment of workforce development needs in Office of Science research disciplines, February 18, 2014: http://science.energy.gov/~/media/ascr/ascac/pdf/charges/ASCAC_Workforce_Letter_Report.pdf

(4) Ten technical approaches to address the challenges of Exascale computing, July 19, 2013: http://science.energy.gov/~/media/ascr/ascac/pdf/meetings/20140210/Top10reportFEB14.pdf

(5) "Synergistic Challenges in Data-Intensive Science and Exascale Computing", March 2013: http://science.energy.gov/~/media/ascr/ascac/pdf/reports/2013/ASCAC_Data_Intensive_Computing_report_final.pdf

(6) "The ASCAC Facilities Statement," March 21st, 2013: http://science.energy.gov/~/media/ascr/ascac/pdf/reports/2013/ASCAC_facilities_statement_final.pdf

(7) DOE Computational Science Graduate Fellowship Report, November 23, 2011: http://science.energy.gov/~/media/ascr/ascac/pdf/reports/ASCAC_CSGF_Report_2011-Final.pdf

(8) "The Opportunities and Challenges of Exascale Computing", Fall 2010: http://science.energy.gov/~/media/ascr/ascac/pdf/reports/Exascale_subcommittee_report.pdf

(9) INCITE Committee of Visitors Report, August, 2008: http://science.energy.gov/~/media/ascr/ascac/pdf/reports/Incite_cov_report_aug_08.pdf

Other Reports and Documents

(10) Report of the Secretary of Energy Advisory Board (SEAB) Task Force on Next Generation High Performance Computing, August 2014: http://energy.gov/seab/downloads/report-task-force-next-generation-high-performance-computing

(11) "The Future of Computing Performance: Game Over or Next Level?," Samuel H. Fuller and Lynette I. Millett, Editors; Committee on Sustaining Growth in Computing Performance; National Research Council (2011).

Roscoe C. Giles Short Biographical Statement

Roscoe Giles is a professor in the Department of Electrical and Computer Engineering at Boston University. Giles' research focuses on the application of high performance and parallel computing to physics and materials problems.

Giles has served in a number of leadership roles in the scientific community, including as a chair of the DOE Advanced Scientific Computing Advisory Committee (ASCAC), the board of Associated Universities Inc (chair-elect), chair of the Boston University Faculty Council, and General Chair of the SC Supercomputing Conference in 2002.

Professor Giles has also worked to prototype and build computational and educational infrastructure that will enable broad participation of scholars and students in high-performance computing. For his work in increasing the participation of minorities in computer and computational science, in 2000 Giles received the A. Nico Habermann award from the Computing Research Association. Giles was profiled in the History Makers Archive (http://www.thehistorymakers.com/biography/roscoe-giles)

Professor Giles received his PhD in Physics from Stanford University in 1975 - the first in a long line of African-American PhD's in Physics and Applied Physics at Stanford. He received a Bachelor of Arts in Physics from the University of Chicago in 1970.

Chairman WEBER. Thank you, Dr. Giles.

And, Mr. Turek, you are now recognized for five minutes.

TESTIMONY OF MR. DAVID TUREK, VICE PRESIDENT, TECHNICAL COMPUTING, IBM

Mr. TUREK. Good morning, Chairman Weber, Ranking Member Johnson, and Members of the Subcommittee. Thank you for the opportunity to speak with you about the Office of Science ASCR program, supercomputing, and American technology leadership.

I have been involved in many of IBM's activities and supercomputing over the last 25 years. During that time, I have worked closely on supercomputing projects with both the Office of Science and NNSA such as ASCI White, Blue, and Purple systems at Lawrence Livermore; the Blue Gene systems, Mira, and Sequoia at Argonne and Livermore respectively; the Roadrunner system at Los Alamos; and as well as key software projects at Pacific Northwest National Lab. I have witnessed firsthand the magnitude of innovation possible courtesy of the collaboration between private industry and the national labs.

I would like to pose today three questions with respect to the linkage between supercomputing and technological leadership. First, why be concerned about supercomputing leadership? The Council on Competitiveness has stated that to out-compete you must out-compute. I believe this to be true. Supercomputers, as the other panelists have said, are tools for inside strategic advantage with broad and diverse application in areas such as oil discovery, fraud detection, efficient automobile and aerospace design, and even many areas of basic science. It is nearly axiomatic that better supercomputers give one a chance for more insight and greater advantage than those with lesser supercomputers. That is why you see the Europeans, the Chinese, the Japanese, and others making a concerted push through public funding of major supercomputer projects. They want to out-compete us.

But there is a fundamental understanding we must also have. Supercomputers are nothing without the software programs and applications that run on them and software engineers only want to produce software for the best machine, not the second, third, or fourth best. Without the best supercomputers available in the United States, software developers will migrate to develop their innovations elsewhere. Once that trend starts, it is very hard to stop or reverse. It is much more costly to catch up than it is to stay ahead.

The second question is what technology problems are in the way of maintaining leadership? The first problem is the need to make supercomputers more energy efficient. The fastest Western economy-based supercomputers in the world today consume about 10 megawatts of energy or $10 million a year. As supercomputers get bigger and more powerful, without some real breakthroughs, by the beginning of the next decade the energy bill could easily be 100 megawatts or $100 million to run. This means the cost of energy will begin to overtake the cost of the computer itself, that becoming a limiting factor in supercomputer usage. A slowdown in usage will ultimately correlate with a slowdown in innovation and impact economic competitiveness.

The second problem is how to handle huge amounts of data. It is clear that the explosive growth of data is challenging some of the fundamental design principles of supercomputers. For example, 500 e-books is about a billion bytes of data. With today's technology, that amount of data can be moved through a computer network in a matter of minutes or less. But suppose we multiplied that amount of data by a million? That would represent the amount of data many supercomputers are working on today and in short order there will be problems a thousand times beyond that.

Old design principles don't solve this problem. We cannot simply do what we did in the past at greater scale to fix this. The temptation, therefore, would be to ignore portions of data to make the problem more tractable, but data left unanalyzed is insight undiscovered, so we have to find ways to make future supercomputers more accommodating to the vast amounts of data they will be asked to explore. New innovations are requiring networking, memory design, storage innovation, and data management software to remedy this circumstance.

The third problem is application software. Most application software running on supercomputers today are based on mathematical approaches more than 40 years old, which is the last time there was a major systematic government investment in new algorithms. The software is now horribly mismatched to modern supercomputers simply because 40 years ago no one could have guessed what today's supercomputers would look like. Access to modern software and new algorithms will have a dramatic impact on the utility of modern supercomputers. There must be a plan to modernize application software. There is no silver bullet to solve these problems. Inventions required to maximize impact, all the problems must be addressed in concert.

The third question is what needs to happen to maintain leadership? From my experience, collaboration with the national labs has been a proven means to stimulate innovation in supercomputers. The labs work on problems of such complexity they always stretch the limits of computing technology. In fact, a crude rule of thumb is the computing requirements of the national labs are about five to seven years advanced over the rest of the market. Finding the ASCR program will present the opportunity to address the problems I described and contribute to maintain the pace of innovation competitiveness demands. If this commitment is made, U.S. leadership in supercomputing should be preserved for years to come.

Thank you very much and I would be happy to answer your questions.

[The prepared statement of Mr. Turek follows:]

David W. Turek
Vice President, Technical Computing OpenPOWER
IBM Systems Group
House Committee on Science, Space and Technology
Subcommittee on Energy
Supercomputing and American Technology Leadership
January 28, 2015

Good morning, Chairman Weber, Ranking Member, and members of the subcommittee. Thank you for the opportunity to speak with you about The Office of Science ASCR program, supercomputing and American Technology Leadership.

Supercomputing and National Competitiveness

Supercomputers have become critically important tools for driving innovation and discovery across an array of industries and scientific domains. By virtue of digital simulation and modeling, supercomputers have helped displace expensive, time consuming and sometimes dangerous physical models and experiments. This, in turn, has led to the development of dramatic insights into a host of very complex domains which have yielded profound commercial and scientific results. For example, supercomputers have been used to accelerate the pace of oil discovery, design new materials, predict the evolution of disease, discover fraud in consumer transactions, design fuel efficient cars and planes, provide insight into the operation of the brain, model the national nuclear stockpile, and contribute to a host of basic science discoveries. In fact, in a report from the US Council for Competitiveness, it has been noted that "modeling, simulation and massive data analysis are the ... game changing drivers for innovation." The Council put an exclamation point on this by succinctly declaring that "to out compete, you must out compute"!

Of course, the problems of tomorrow are not the same as the problems of today, and there is a growing chorus of concern that without adequate investment in supercomputing, the ability of the US to continue to out-compete its sovereign competitors will diminish. It is our contention that investment in the Advanced Computing Science Research Programs in the DOE's Office of Science are instrumental to ameliorating this concern: maintaining innovative leadership in supercomputers goes hand in hand with maintaining scientific and economic leadership.

International Supercomputing Activities

The relationship between advanced computing and economic competitiveness is well understood internationally. In Europe, under the banner of Horizon 2020 and PRACE (Partnership for Advanced Computing Europe), there are a myriad of advanced computing efforts covering basic research (e.g. The Human Brain Project), support for university access, and the development of new computing approaches for complex scientific and engineering problems. In China, there has been a well-funded set of government programs focused on things like microprocessor design and software which are expected to contribute to the performance of domestic industries as a prelude to the establishment of a supercomputing export industry. Japan, this past fall, launched a government funded program for advanced supercomputing to achieve worldwide leadership (in a rudimentary way) within the next six years. There are even nascent plans for sovereign funded efforts to drive advanced supercomputing being contemplated in Russia, India, and Korea. Nothing will stop or otherwise impede these international efforts so competition with the US is here to stay. However, to maintain economic advantage it is very important for the US to be the leader in advanced supercomputing: there is a big difference between being present and being the best. It has long been the case that software engineers gravitate to do their work on the best and fastest supercomputers. In the absence of true supercomputing leadership the software engineers we depend upon to build the applications solving some of most vexing business and scientific problems the US faces will migrate to wherever leadership exists. What we lose in terms of early access to these skills, someone else in the world will recognize as a gain at our expense.

State of Supercomputing in the United States

The problems for which supercomputers are designed are of such extreme complexity that conventional computer technologies are often stretched to the breaking point. However, it is when computer designers are working at the extremes of technology that innovation and invention is oftentimes spurred. For example, in the early nineteen nineties when commercial computing was struggling with data files measured in billions of bytes, supercomputers were being designed to work on files 10,000 times larger. Many of the technologies pioneered in supercomputers for big data have, over the course of time, been adopted in conventional computers helping usher in a new era of advanced analytics widely used in companies of all sizes to gain strategic business insight

Another critical issue facing supercomputer users is the amount of energy required for operation. In economic terms, the rule of thumb is that the fastest supercomputer in the world at any point in the last 50 years has cost between $100 million and $200 million dollars. We forecast that absent serious and sustained innovation in energy efficiency the cost for electricity for the fastest computer at the end of this decade or early the next could approach $100 million per year! Clearly, the growing operating costs for supercomputers, left unchecked, could put a material damper on adoption and propagation throughout the economy and have a serious downside effect on the nature and rate of economic and scientific innovation. Conversely, if major advances in energy efficiency were realized one could envision a future where a single researcher could expect to have a desk-side computer equivalent in capability to the fastest computers in the world today. The ubiquity of this caliber tool in the hands of large numbers of people would be a tremendous catalyst to accelerating the pace of innovation in the economy.

There is also the need for dramatic innovation in network technology, computer memory designs and storage technology. Today, one can move a file containing a billion bytes of data across the country in a few minutes under ideal conditions. Supercomputer applications often deal with a million billion bytes of data! This requires transcendent technologies and architectural designs, especially as supercomputer problems demand ever greater amounts of data. We need innovations to make memory dramatically more affordable to contain these volumes of data; it is commonplace today that much more of the acquisition cost of a supercomputer is tied to the cost and speed of memory than it is to the microprocessor. And we need innovations in storage technologies that affordably house these volumes of data but also are capable of making the data readily available to the supercomputer. Rotational speeds of spinning disks are relatively stalled with respect to the growth rate of data and promising new technologies are marginally affordable at best given the scale of data being entertained. Solving these types of problems holds huge promise even for a single computer user: with supercomputer inspired innovations one could imagine a handheld device holding the contents of the Library of Congress , continuously updated, with near instantaneous intelligent cognitive capability in any language, allowing the user to get answers to almost any question imaginable in real time.

Software is perhaps one of the greatest challenges facing supercomputer use today. The fastest supercomputers currently employ millions of computer cores orchestrated and managed by software to work in concert on solving very complex problems. However, much of the software in use is based on mathematical and algorithmic approaches that are more than 40 years old, and created at a time when there was no concept of current supercomputer design. It is imperative that software innovation be driven in lockstep with the innovations coming in computer hardware.

Addressing the challenges I've outlined will involve a host of players including supercomputer companies, technologies companies, software companies, universities, and the national laboratories. Much of the work is necessarily collaborative: a hardware company must work closely with software companies to produce the best solutions for example. All of the work is necessary to drive the pace of innovation to the requisite competitive level.

Supercomputing Collaborations with the Department of Energy

For more than twenty years, IBM has been engaged in a partnership with the Department of Energy to produce advanced supercomputing technologies in support of critical problems facing the National Nuclear Security Agency and the Office of Science. In the 1990s our collaboration with the NNSA resulted in advanced networking technologies enabling successful deployment of systems in support of the Accelerated Strategic Computing Initiative; later in the decade we extended the partnership through co-design of the Blue Gene system utilizing innovations enabling massive scale and dramatic improvements in system energy efficiency; and in the last decade we collaborated on the development of the Roadrunner system demonstrating the effectiveness of accelerators in the solution of highly complex science problems. In each case the innovations and ideas derived from our collaborations were adopted, copied, or extended by other computer companies helping to usher in an era of dramatic growth in the use of supercomputing. In a recent National Academy study, it has been noted that "virtually every sector of society-manufacturing, financial services, education, government, the military, entertainment, and so on-has become dependent on continued growth in computing performance to drive industrial productivity, increase efficiency, and enable innovation". In essence, an advanced economy needs advances in supercomputing to remain competitive.

Pathway to Leadership for the US

The longstanding collaboration with the Department of Energy is entering a new phase with the launch of the CORAL program targeted to produce systems for Lawrence Livermore and Oakridge National Labs beginning in 2017. These systems will be more than 100 times more powerful than the Roadrunner system and many times faster than the current fastest system in the world in China. These systems are also a prelude to the design and development of exascale systems in the early part of the next decade. (An exascale system will perform about one quintillion (1 followed by 18 zeros) operations per second.) The path we are on with CORAL will stimulate a new wave of innovation in supercomputing giving rise to new technologies and designs that we would expect to be broadly beneficial across the entire computing industry. We have to invent new designs for managing the massive amounts of data that supercomputers are asked to analyze, amounts of data so large that simply ingesting them into the supercomputer can take weeks. We have to invent new software techniques to program systems with millions of microprocessors working in concert. We have to find ways to make these systems reliable when the aggregate mean time to failure of all the parts implies continuous failure of the system. And we have to find a way to make these systems energy efficient so that they can become affordable.

We endorse funding for the Advanced Computing Science Research Program of the DOE because our historical experience working collaboratively with the DOE has proven to be of tremendous value in terms of stimulating the kinds of innovation we expect the CORAL and follow-on programs to require. Through the kind of public-private partnership that has characterized much of what we have done with the DOE over the last few decades we expect to see the US continue to lead the international supercomputing community for years to come.

Thank you for this opportunity to speak with you today. I would be happy to answer any questions.

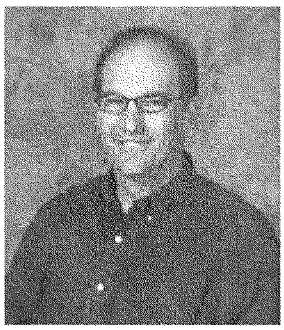

David William Turek

Dave is currently the Vice President, Technical Computing OpenPOWER at IBM with responsibility for IBM's overall high performance computing strategy. In previous executive positions Dave helped launch IBM's Grid Computing business, and started and ran IBM's Linux Cluster business. As a development executive he had responsibility for IBM's SP supercomputer program as well as the mainframe version of AIX and other Unix software. In that capacity he orchestrated the initial IBM development effort in support of the US Accelerated Strategic Computing Initiative at Lawrence Livermore National Laboratory. He has been recognized for his contributions to the Roadrunner program at Los Alamos National Laboratory in the Congressional Record and sits on the Advisory Committee to the National Simulation Center at the University of Tennessee at Chattanooga. Dave has degrees in Philosophy and Mathematics from the University of Rochester, a Masters Degree from Trinity College, and advanced study at the University of Pennsylvania in Operations Research.

Chairman WEBER. Thank you, Mr. Turek.
And now, Dr. Crowley, you are recognized for five minutes.

TESTIMONY OF DR. JAMES CROWLEY, EXECUTIVE DIRECTOR, SOCIETY FOR INDUSTRIAL AND APPLIED MATHEMATICS

Dr. CROWLEY. Good morning, Chairman Weber, Ranking Member Johnson, and Members of the Committee.

As noted in my introduction, I am Executive Director of the Society for Industrial and Applied Mathematics, or SIAM. SIAM comprises over 14,000 members who work in industry, government and national labs, and in academia. They represent over 500 universities, corporations, and research organizations from around the world. SIAM is dedicated to solving real-world problems through applied mathematics and computational science.

Thank you very much for allowing me to testify and for highlighting the critical work of the Department of Energy's Office of Science and its Advanced Scientific Computing Research program. SIAM greatly appreciates your Committee's continued leadership on, and the recognition of, the critical roles of the Office of Science and ASCR in enabling a strong U.S. economy, workforce, and society through mathematical, scientific, and engineering research relevant to the DOE mission.

The Office of Science supports basic research to address pressing challenges in energy, computing, physical sciences, and biology and this support has been critical to the applied mathematics and computational science community.

I wish to focus on three topics: ASCR support for mathematical and computational science research, the potential benefits of exascale and the technological challenges to reach it, and finally workforce and training needs. First, the role of ASCR in supporting key mathematical and computational research.

ASCR supports the development of new modeling simulation and data tools to help researchers solve scientific and energy challenges. Modern life as we know it, from search engines like Google to the design of modern aircraft, would not be possible without the unique contributions of mathematicians and computational scientists. Likewise, DOE depends on mathematical and computational techniques to make predictions, model and simulate systems that would be costly or impossible to experiment on, and manage and make sense of ever-growing data that is produced by scientific experiments such as DOE's particle accelerators and light source facilities.

The Nation faces critical challenges in energy efficiency, renewable energy, future energy sources, and environmental impacts of energy production and use. These challenges all involve complex systems such as the power grid or the U.S. nuclear stockpile. Mathematical and computational tools help us model and understand these systems, design new solutions to problems, and predict the impact of new technologies. ASCR programs not only support new mathematical tools but also develop software so that DOE, industry, and the academic community can use these tools. And I note that the PETSc team at Argonne just was awarded the ACM SIAM prize in computational science and engineering and that shows the power of the people working at DOE.

Second, I would like to address the possibilities and challenges of exascale. For all the advances that ASCR has already enabled, today, there are still challenges that are too complex for current computers to model. Exascale computing has the potential to spur revolutionary advances in modeling and simulation, expand our capacity to analyze complex systems in great detail, and capture more complexity with better predictive abilities than ever before.

I will note that the investments in modeling, algorithm research, and software development are essential to realizing the full benefits of exascale computers so that we can use these machines to solve pressing scientific and energy challenges. It is not just the hardware; the computer science and the math are essential.

Finally, I would like to discuss an important workforce development program within ASCR. Researchers trained to use high-performance computers to solve key scientific challenges are central to DOE's mission. The Computational Sciences Graduate Fellowship program is a critical program that maintains the pipeline of this workforce by supporting the training of new scientists and engineers with strong computational research experience and close ongoing ties to DOE and the national labs. The CSGF has a long history of success at DOE and SIAM strongly supports its continuation.

I thank you again for the opportunity to provide this testimony today and I am happy to answer any questions. I have provided additional details in my written testimony. Thank you.

[The prepared statement of Dr. Crowley follows:]

SIAM®
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Phone +1-215-382-9800
Fax +1-215-386-7999
www.siam.org · siam@siam.org

Testimony by Dr. James Crowley
Executive Director, Society for Industrial and Applied Mathematics (SIAM)

Testimony on Supercomputing and American Technology Leadership
before the
Subcommittee on Energy
Committee on Science, Space, and Technology
United States House of Representatives, Washington, DC

January 28, 2015

My name is Jim Crowley, and I am the Executive Director of the Society for Industrial and Applied Mathematics (SIAM). I have served as Executive Director of SIAM since 1995.

SIAM has over 14,000 members, including applied and computational mathematicians, computer scientists, numerical analysts, engineers, statisticians, and mathematics educators. They work in industrial and service organizations, universities, colleges, and government agencies and laboratories all over the world. In addition, SIAM has over 500 institutional members—colleges, universities, corporations, and research organizations. SIAM members come from many different disciplines, but have a common interest in applying mathematics in partnership with computational science towards solving real-world problems.

Thank you very much for the opportunity to testify today and for highlighting the critical work of the Department of Energy (DOE) Office of Science and the Office of Advanced Scientific Computing Research (ASCR). I would like to emphasize how much SIAM appreciates your Committee's continued leadership on and recognition of the critical role of the Office of Science and its support for mathematics, science, and engineering in enabling a strong U.S. economy, workforce, and society. The DOE Office of Science supports basic research to address pressing challenges in energy, computing, physical sciences, and biology. This support has been critical to the applied mathematics and computational science community. DOE was one of the first federal agencies to champion computational science as one of the three pillars of science, along with theory and experiment, and SIAM deeply appreciates and values DOE activities.

The Role of Mathematics and ASCR in Meeting Energy Challenges

ASCR supports research to develop new modeling, simulation, and data tools and to connect those tools to researchers across the Office of Science and DOE for use in solving scientific and energy challenges. Modern life as we know it – from search engines like Google to the design of modern aircraft, from financial markets to medical imaging – would not be possible without the techniques developed by mathematicians and computational scientists. Likewise, the Department of Energy depends on mathematical and computational techniques to advance science and engineer new energy solutions.

SIAM®
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Phone +1-215-382-9800
Fax +1-215-386-7999
www.siam.org · siam@siam.org

The nation faces critical challenges in energy, including in energy efficiency, renewable energy, improved use of fossil fuels and nuclear energy, future energy sources, and reduced environmental impacts of energy production and use. As DOE and the research community design a long-term strategy to tackle these issues, the tools of mathematics and computational science (theory, modeling, and simulation) have emerged as a central element in designing new materials, predicting the impact of new systems and technologies, and better managing existing resources.

To tackle many of these challenges, DOE must be able to understand complex systems such as the US power grid, the behavior of nanomaterials relevant to energy, and the dispersion of nuclear radiation after a disaster. These and other complex systems have high levels of uncertainty, lack master plans, and are susceptible to breakdowns that could have catastrophic consequences. Understanding complex systems helps mitigate these risks and facilitate the development of controls and strategies to make systems more efficient.

Applied mathematics and computational science play a key role in predictive modeling and analysis to understand complex systems. Already, mathematical and computing researchers using these tools have made substantial progress improving our understanding across fields such as genomics, biofuels, materials fabrication, and nuclear security.

Activities within ASCR play a key role in supporting research that begins to fulfill these needs through programs such as the Applied Mathematics program, the Scientific Discovery through Advanced Computing (SciDAC) program, and programs to maintain the pipeline of the computational workforce. These programs have a long history of not only developing novel mathematical methods and algorithms but also taking these the next step and developing robust software that is used by DOE, the academic community, and industry.

Mathematics and the Development of Exascale

SIAM supports ASCR's new all-in approach on research to develop exascale computing, noting that investments in mathematical modeling, algorithm research, and software development are essential to realizing the full benefits of this next generation of high performance computers and to transferring their capabilities to industry for broad economic benefit.

While achieving exascale computing has the potential to allow for revolutionary advances in many fields critical to solving our energy challenges, getting to exascale and realizing its benefits requires overcoming significant computing challenges, including in applied mathematics. ASCR is currently supporting research to address these challenges, which are described below.

At the hardware level, computer chips must change in radical ways to achieve exascale speeds, such as containing 100 times more processing elements than today. This will require new breakthroughs in algorithm development to take advantage of parallelism on the chip as well as

siam.
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Phone +1-215-382-9800
Fax +1-215-386-7999
www.siam.org • siam@siam.org

parallelism between computational nodes comprised of multiple chips. In order to achieve high rates of performance, algorithms that minimize data movement, possibly by increasing the relative number of computations, will be the most efficient. Algorithm developers will need to take these facts into account as they develop multi-scale, multi-physics algorithms.

Utilizing exascale to study complex systems will require the construction of new predictive mathematical models of large, nonlinear dynamic networks that span several spatial and temporal scales. Understanding and manipulating these systems will require large, multi-scale, nonlinear, and hybrid models, which either do not exist or are in their infancy today. Simulations carried out at finer scales often will require the modeling of physical effects that are ignored in current simulations, which requires both the development of new mathematical models as well as new algorithms to run them. DOE is additionally supporting research in uncertainty quantification, which will allow the confidence in computational forecasts and predictions to be measured.

In addition to allowing revolutionary advances in modeling and simulation, exascale can potentially be utilized to manage ever-growing data volumes in science, from DNA sequence data to information on materials from DOE's light source facilities. These data need to be stored in databases that are easily accessible and searchable, requiring increasingly sophisticated and scalable data mining algorithms. In addition, the data from heterogeneous sources need to be integrated, within databases as well as within models. Once accessible in databases, the typically high dimensional data sets need to be analyzed using advanced mathematical and statistical methods.

Co-design of hardware, software, and applications are critical for successful exascale development so that each component of an exascale computer is optimized in an integrated fashion. ASCR Co-Design Centers are addressing the challenge of designing exascale systems to enable a variety of predictive mission and science outcomes. To achieve co-design requires frequent interactions among hardware architects, systems software experts, designers of programming models, and implementers of the science applications that provide the rationale for building extreme-scale systems.

Mathematical and computational scientists with appropriate awareness and ability to tackle interdisciplinary exascale challenges will be essential to realizing the benefits of exascale. Programs should be implemented to ensure this pipeline of mathematicians and computational scientists at the undergraduate, graduate, post-doctoral, and early career levels. In addition, programs that encourage interdisciplinary collaboration between computational scientists and domain scientists are essential to incentivizing these collaborations, which are crucial for the realization of exascale's scientific benefits. Programs that support partnerships, workshops, travel to technical conferences, and summer internships also help to create a new community of researchers more alert to and equipped to conduct interdisciplinary research.

siam.
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Phone +1-215-382-9800
Fax +1-215-386-7999
www.siam.org · siam@siam.org

Supporting the Pipeline of Mathematicians and Computational Scientists

Before I close my testimony, I would like to discuss an important workforce development program within ASCR. SIAM is grateful to Congress for its support of the Computational Sciences Graduate Fellowships (CSGF), a critical program that maintains the pipeline of the computational science and engineering workforce. Researchers trained in computational science and working in universities, national laboratories, and industry are central to DOE's mission and essential to propel advances in DOE mission-critical fields such as nanotechnology, biofuels, genomics, and materials fabrication. CSGF helps ensure the existence of an adequate supply of scientists and engineers with strong computational research experience and close ongoing ties to DOE to meet future national workforce needs.

The CSGF has a long history of success at DOE. Connections to the national labs are integral to CSGF's success, as fellows train at DOE national labs and program requirements are closely tied to DOE mission needs. A 2011 Committee of Visitors (COV) report[1] evaluating CSGF found that the program has been highly successful at producing alumni with strong computational research experience and close ongoing ties to DOE. CSGF is a valuable program and a unique source of talent in the area of computational science where high performance computing is applied to challenging and important science and engineering problems.

Conclusion

The programs in the Office of Science, particularly those discussed above, are important elements of DOE's efforts to fulfill its mission. They contribute to the goals of dramatically transforming our current capabilities to develop new energy supplies and improve energy efficiency to ensure energy independence and facilitate DOE's effort to increase U.S. competitiveness by training and attracting the best scientific talent into DOE laboratories and the American research enterprise.

I would like to conclude by thanking you again for your ongoing support of the DOE Office of Science and the actions you have already taken to enable DOE and the research and education communities it supports, including thousands of SIAM members, to undertake the activities that contribute to the health, security, and economic strength of the U.S. The DOE Office of Science is critical to maintaining our competitive edge in science and technology.

I appreciate the opportunity to provide testimony to the Committee and am happy to answer any questions.

[1] http://science.energy.gov/~/media/ascr/ascac/pdf/reports/ASCAC_CSGF_Report_2011-Final.pdf

Dr. James Crowley Biography

Dr. James Crowley is the Executive Director of the Society for Industrial and Applied Mathematics (SIAM), a position he has held since 1995. Dr. Crowley came to SIAM following a 22-year career in the Air Force where he retired as a Lieutenant Colonel. During his Air Force career Dr. Crowley served in a number of positions, including Assistant Chief Scientist for Air Force Systems Command, Director of Math and Information Sciences at the Air Force Office of Scientific Research, and Tenured Associate Professor at the U.S. Air Force Academy. Dr. Crowley also served as program manager at the Defense Advanced Research Projects Agency (DARPA) for two years. Dr. Crowley received a PhD in Applied Mathematics from Brown University, an M.S. in Mathematics from Virginia Polytechnic and State University, and a Bachelor's degree in Mathematics from the College of the Holy Cross. His scientific interests are in control theory and large scale computing. Dr. Crowley is a Fellow of the American Mathematical Society (AMS), and a Fellow of the American Association for the Advancement of Science (AAAS).

Chairman WEBER. Thank you. I thank the witnesses for their testimony. Members are reminded that the Committee rules limit the questioning to five minutes and the Chair recognizes himself for five minutes.

Good grief, where do we start? You all have just raised a whole bunch of questions. Dr. Crowley, can you provide an overview in plain English so our constituents can understand? You kind of went through it there toward the end of what ASCR program does but why is it important to the U.S. economy?

Dr. CROWLEY. The tools that are provided for modeling and simulation are used across—I gave you the example of the award—the prize that went to the PETSc team at Argonne National Lab. PETSc is a team that has developed computational tools for high-performance computers. These tools are used by industry that do modeling and simulation for advancing materials, some of the things that Roscoe Giles mentioned, and without those tools, one can't use the computers efficiently to do that. And so it is the research into not only developing the tools that the people can use the computers but also the models that run on them.

I mean to take an example that is not necessarily a DOE model but just looking at one that came home to me recently because of the weather prediction in Philadelphia that almost kept me from coming here, this latest snowstorm missed Philadelphia but it was predicted to dump more than a foot of snow on us. Improved modeling, better tools, and higher performance computing would have made that ability to make those predictions much better. And that same thing applies for any other kind of thing that is modeled across science and engineering, that with better tools for modeling, better computational tools, we can advance our ability to produce better materials, to simulate anything that we need—and understand better scientific things in fusion or in any other area.

Chairman WEBER. Okay. I think it was you that said that the algorithms or the math that was used 40 years ago—it was Mr. Turek so this question is probably for you—couldn't predict or you couldn't see what computers look like today. What did you say about that?

Mr. TUREK. Yeah, what I meant by that was that at that time frame, the nature of what was considered to be a supercomputer bears no resemblance to the kinds of computers that exist today, so people designed the algorithms and the corresponding software to map to that kind of computer. Those approaches don't translate well over four decades to the kinds of things we are doing today.

Chairman WEBER. So here is my question. Do we have the capability today to look out 40 years in the future and predict how effective those algorithms will be or will there be new techniques? With advanced computing today are we able to look out 40 years in the future?

Mr. TUREK. Nobody can look out 40 years correctly. However, what I would say is that we know that many of the algorithms and the software implementations today are obsolete for what we are trying to do. The way to characterize it would be the following: Today's modern supercomputers typically use order of millions of microprocessors. Many of the algorithms and the software implemented only scale to maybe a handful of hundred of microprocessors not because it can be done; it is because it is a byproduct of the fact that that invention is 40 years old. A reinvestment in algorithmic development, the fundamental mathematics and the associated software, has been demonstrably proven in places like Argonne, Lawrence Livermore, and Oak Ridge that these approaches to common problems can be modified to accommodate this nature of supercomputing we have today. You would have a material effect on dramatically improving the insight that people gain from the application of the supercomputing tool.

Chairman WEBER. Is part of the aim of ASCR, for example—because we hear a lot in today's society about hacking and so we invest the money and I am a great believer that we need to be on the cutting edge because it helps national security, for example, but are we at risk with supercomputing of investing money, time, and resources, and then having that technology stolen from us by other countries

Mr. TUREK. So there is this notion of internationalism if you will, but I would characterize it this way: The Chinese program is very parochial to China. The European program is very parochial to Europe and they are making investments that are very much wedded to the parochial interests of companies and institutions in those geographies. There is always the chance that through regular commerce or more nefarious means technology can escape geographic boundaries, but I think the deployment of technology in the economy is what really makes a difference, so the more supercomputing that can be made available, the more and diverse kinds of people who can get access to it and use it is what really spurs the economic kind of innovation we have all alluded to here today.

Chairman WEBER. Yeah. Well, I appreciate that. And I am out of time so the Chair will now recognize Ranking Member Johnson.

Ms. JOHNSON. Thank you very much, Mr. Chairman.

I am so delighted we have such able witnesses today and I know that this hearing is focused on our investments in supercomputing research in particular, but I would like to take advantage of your presence, Mr. Augustine, to ask a few broad questions to help us guide the future in how we are able to continue research, whether or not we are producing the researchers. In 2005 the National Academies' Gathering Storm Panel, which you chaired, recommended increasing science agency budgets by ten percent annually.

The 2007 COMPETES bill, which was very graciously accepted and supported by President Bush, had bipartisan support for a positive growth trajectory of R&D, and unfortunately, appropriations for the last eight years have not come close to keeping up with what was projected. It was changed to a more conservative recommendation to at least four percent annually in 2014.

In the current budgetary and political environment, how would you continue to make the case for increased funding for R&D to politicians across the political spectrum? And what do you believe are the consequences if we do not even achieve this modest four percent annual growth target for federal investment in basic research and development? And, finally, do you believe that a robust reauthorization of America COMPETES should be a top priority for this Committee this year?

Mr. AUGUSTINE. Well, thank you for that question.

Chairman WEBER. Mr. Augustine, turn your mike on, please.

Mr. AUGUSTINE. I thought it was on.

Chairman WEBER. Oh, there you go.

Mr. AUGUSTINE. Sorry. To deal with the last part of your question first, I think America COMPETES is perhaps the most important thing that this Committee could take on. It drew more attention to the problems we face in this area and took further steps to improving the situation than anything else I am aware of that we have done. So I would strongly urge that.

With regard to the status of the research and where we have come since the various reports that you allude to, the bad news is that we are declining in our investment in research as a percentage of GDP. Other countries are growing. Even at NIH, which is—research there is strongly favored by the American public—we have seen a 22 percent cut in the last decade in real dollars and it is continuing to decline. This of course discourages young people from going into research and basically it means that we are going to have a lower quality of life, impact on our health will be very real, and the economy today is so heavily dependent on technology that without doubt we will be hurt economically seriously.

I would cite an example from my own field of the impact of research and particularly high-performance computing. I am an aerodynamicist, design airplanes, among other things. The way we used to design airplanes when I was early in my career was built giant wind tunnels. We built them when they were plugged into the Tennessee Valley Authority by and large because that was the only place we could get enough power. We ran them at night we didn't shut down the lights in the southern part of the country.

Today, we don't use wind tunnels. We put the airplane and a high-performance computer if you will, use a mathematical model and within a nanosecond have the answers that we are researching, just one example of the enormous impact that investment in technology can have and also the negative impact of not investing in science, research, and technology.

Ms. JOHNSON. Well, thank you very much.

The National Research Council report entitled "Rising to the Challenge: U.S. Innovation Policy for the Global Economy," states the assumption that the output of the U.S. innovation process will be captured by U.S.-based industry has been rendered obsolete by globalization, and that knowledge created through federally funded research at universities and national laboratories can be commercialized and industrialized virtually anywhere. The report goes on to say that a more comprehensive innovation policy is needed to anchor new and existing companies here in the United States.

The American Academy of Arts and Sciences panel that you recently chaired addressed some of this issue in a report released this fall. What recommendations do you have for what federal policies are necessary to ensure that U.S. companies benefit from U.S. innovation?

Mr. AUGUSTINE. Well, thank you for that question. And as you point out, research is a global commodity or global asset, and it raises a question why not just let others do the research and then apply their research? The answer, I would cite Craig Barrett, who ran Intel some years ago. Craig says that on the last day of any calendar year 90 percent of the revenues that Intel receives are for products that didn't exist on the first day of the calendar year, and so the only answer to your question that I can see is that we just have to be faster than others in applying the results of research. We have got to be fast.

And your question what do we do about it and the answer is remove every bureaucratic obstacle, every obstacle we can think of, particularly in technology transfer from the labs, that causes time delay because time is everything.

Ms. JOHNSON. Thank you very much. My time has expired. Thank you.

Chairman WEBER. Thank you.

And the Chair now recognizes the Vice Chairman of this Committee, Congressman Newhouse.

Mr. NEWHOUSE. Thank you very much, Mr. Chairman. I appreciate that and appreciate you gentlemen being here this morning and talking about this very important subject. It is certainly enlightening me as to the nature of our responsibility here.

Not to let you dominate the program this morning, Mr. Augustine, but a question that arose in my mind after reading through your testimony that a lot of the body of research at our national laboratories is maybe not being utilized as much as it could be so to speak, not to put words in your mouth, but there are certain obstacles that stand in the way of getting that research to industries. So could you talk a little bit about maybe what you see as solutions to that issue that we have? Is it communication, some of the conflict-of-interest issues that you mentioned, and those kinds of things?

Mr. AUGUSTINE. Well, thank you, Congressman Newhouse, for that question. And I do believe that the Nation doesn't begin to benefit from the asset that our national labs represent. It certainly benefits importantly but it could be so much more, and the reason for that is that we need to do a better job of getting knowledge out of the laboratories and into industry so that we can commercialize and distribute the results.

And as to impediments, there are many. One that certainly stands in my mind is that firms simply don't know what is going on in the national laboratories. They tend to be rather isolated. And we could do a much better job of letting people, industry, know what is happening at the laboratories.

Secondly, the best way to transfer technology that I have ever been able to find is by transferring people. You move the knowledge that is in their minds. And today, well-meaning conflict-of-interest laws make it very difficult to transfer people among industry, government, and academia. In my career I had the opportunity to put in two tours in government and today I doubt that I could do that under the conflict-of-interest laws that exist.

A third one that I would cite is that we are very concerned, properly so, about favoring one firm over another. What do we do about it? Without taking a great deal of time, one is for the labs to do a better job of letting the world know what they are working on, the industrial world if you will. Other things that are cited in H.R. 5120, for example, giving the labs more latitude to create industry partnerships, give the labs more latitude to negotiate technology transfer agreements. These are a few of the things that could be done but I don't have answers to the conflict-of-interest one because obviously we don't want conflicts of interest. On the other hand, the inability to move people and to move ideas in and out of the labs is a huge burden on our country.

Mr. NEWHOUSE. Thank you. I appreciate that.

Mr. AUGUSTINE. Thank you.

Mr. NEWHOUSE. Quickly, a question then perhaps for Mr. Turek and perhaps Dr. Giles as well. It is—my limited understanding is that the largest supercomputers are rarely able to operate at full capacity due to their complexity, some components almost always in need of attention or repair. If that is a true statement, could you tell me what is being done to improve the reliability of these systems and are we devoting enough resources to this aspect of advancement?

Mr. TUREK. I will take the first shot at it. We are doing a lot for that. A lot of that is actually handled by software so soft recoveries of problems. What you see with supercomputing are problems of scale. If you have a million parts of anything, the likelihood is you are going to see something failing pretty regularly, even if it is integrated circuits. It is a problem that has been understood for quite some time and principally is handled by software techniques to overcome it. So in the vast majority of cases you actually can get to full capacity if you have the software capability on the application level to utilize it. That is the bigger impediment right now. Again, most people who gain access to commercial software are gaining access to software that is archaically designed relative to the scale of the kinds of computers being built today and that is the limiting factor.

Dr. GILES. Can I add something?

Mr. NEWHOUSE. Absolutely, Dr. Giles.

Dr. GILES. I think that—yes, I think that also our sense of what the capacity of a system is reflects some of the archaic history in the sense that we often measure or think of a capacity is how much data can you sort of crunch, transform from one form to another, which is an artifact of the time when the critical component of a computer was the processor that made that transformation. Now, people are looking at systems with millions of processors and redundancy in processors is not a negative to have multiple processors comparing results one to another. So, as Mr. Turek said, there are lots of opportunities for new ways of ensuring the reliability of the final answers we get.

And if we get discouraged about thinking about that problem, I would remind us all that our brains, with millions and millions of—and billions of neurons and interconnections have faults on the neuron level all the time and they don't materially affect the ultimate outcome, and I think we are in the process of building computers that can function more like that.

Mr. NEWHOUSE. Thank you very much.

Thank you, Mr. Chairman.

Chairman WEBER. Thank you.

And the Chair now recognizes Congressman Hultgren from Illinois for five minutes.

Mr. HULTGREN. Thank you all so much for being here. Thank you, Chairman. I especially want to thank the Chairman for working out a way for Dr. Giles to be with us remotely.

I am very fortunate to represent Fermilab and I have Argonne right down the road from me. Because of this, I have been able to see the fruits that grow out of our Nation's commitment to basic curiosity-driven scientific research. The impacts of this research I believe are limitless. Just as we didn't go to the moon to invent Velcro, we didn't build particle colliders so that we could invent the magnet for our MRI machines.

This topic, supercomputing, is close to home for me because physics is where big data began. Besides the maintenance of our nuclear stockpile, it is either astrophysics or high-energy physics that is driving the research necessary to build the most sophisticated computer networks we have today. Because of this, it was largely DOE that began the genome project before NIH realized it was a feasible endeavor. As interested as I am in technology transfer and local economic development, if our research enterprise is focused on the short-term photo op and press release-style research, which it appears the Administration is more prone to advance, we will lose out on the long-term benefits we all say we should be focused on. If we are going to stay at the forefront of technology or technological development, we must reaffirm our commitment to basic scientific research.

Dr. Giles, in our previous hearing, you had a chance to review a draft copy of my legislation, which in the 113th House eventually passed, H.R. 2495, the American Supercomputing Leadership Act. My bill called for a lab-industry-university partnership to develop two different exascale machines. I wondered if you would be willing to describe what industry's role should be in such a partnership and then describe the benefits of having a university as part of this partnership?

Dr. GILES. Yes, I would be happy to address that and some of my written testimony does get to that point. I think that ASCR's work has helped to start a virtuous cycle with industry, academia, and the labs in developing and looking forward to the path for exascale so that in collaboration with industry we are able to have government funds help to stimulate research and investigation in areas that are important for building the next generation scale of computers before that is actually competitive or something that is in the competitive spirit of the industry, but then industries impact is to help define what is sufficiently along the lines of work that they can build and build on into something that they would be interested in from their perspective, that we find an accommodation.

In the co-design methodology that I mentioned represents the pattern of developing new software and algorithms as—in the context of hardware that is evolving and to help use those needs from the scientific community, from the universities and the labs to help define what kind of hardware makes sense so that the—this goes back to the idea of building an ecosystem that supports rapid advances in scientific computing that links together all those elements.

I do want to thank you so much for the legislation you propose that we discussed last time and which made it out of the House, as I understand it, but not all the way through the end of the process. You know, I think it is a really important step that we explicitly fund the development of that next generation better systems.

Mr. HULTGREN. Thanks, Dr. Giles.

Quickly, Mr. Augustine, I would first like to thank you for all of your work. You have been a leader in this and in so many other spaces, it is amazing. Thank you.

I had the pleasure of sitting down with your colleague Dr. Neal Lane to discuss local economic development potential for the national labs in reference to the Restoring the Foundation report. Many of the recommendations from this discussion echoed my previous passed legislation, the DOE Labs Modernization and Technology Transfer Act, which the Bipartisan Policy Center listed in their doable items, which there aren't too many of, for the 114th Congress. I wonder if you could make a comment more generally on this bill and the needs and benefits for making the labs more nimble and open to the public?

Mr. AUGUSTINE. Well, yes. One of the things that certainly relates to what you raise is that the labs are able to build major facilities that individual firms can't afford to build. Fermilab is a classic example. And if they are not available to the public or industry by and large, then we don't begin to get the value from them that we could get. Some of the legislation that you describe takes important steps in this regard.

I guess I would say in terms of a broad answer—and I realize that we are running out of your time—that the bad news is that we spend, as I said, a 10th of a percent of the GDP on research. The good news is you could double that and only have to allocate a 10th of a percent of the GDP. And so the opportunity is probably there to make major changes.

I go back to one of the studies that you refer to. We discovered that we spend more on potato chips in this country than we spend on research on clean energy. That just doesn't make sense.

Mr. HULTGREN. Well, again, I want to thank you all for being here.

Thank you, Chairman.

And real quick, just thank you, Dr. Crowley, too, for the shout-out to Argonne and the recent recognition there. That is fantastic. So thank you so much.

Chairman, I yield back.

Chairman WEBER. Thank you, sir.

The Chair now recognizes Mr. Massie from Kentucky.

Mr. MASSIE. Thank you, Mr. Chairman.

My question is really for anybody up there that cares to comment, but it seems like 20 years ago there was the apocryphal prediction that we would run out of available computing power with silicon, yet here we are still on silicon. What is the next step after silicon? And since we didn't run out of power with silicon how much further can we go on silicon?

Mr. Turek, it looks like you are interested in answering that.

Mr. TUREK. I will take the first shot at least.

We are at an apocryphal time and to a certain extent you could characterize the industry as putting a Band-Aid over this problem. So the limitations of silicon are embedded in physics. We are at those limits today. I think the last time I saw an advertisement on TV about buy a computer because the processor is faster was January 2001. You don't have a 10 gigahertz processor. You are never going to see one either because the physics are limiting.

So instead what the industry has done is it has spewed out massive amounts of cores, lower-power compute elements that are ganged together to work in concert on the problems at hand. The problem is you don't get a linear scalability of the compute effect. So in other words, if I have four cores, I don't get four times the compute capability of one core. Maybe I get 2.5. And as I scale up to a million, I am not getting a million times; I am getting something far less than that.

So we are Band-Aiding our way through this limitation at the physics level. There are more materials and so on that are coming forth and whether it is carbon nano tubes or something else, but physics is a limiting factor here.

The way you deal with this ultimately is you look at the architecture of how these systems are put together and the composite set of technologies that let you deal with the problem. Advances in networking technology, memory systems, all these things need to be looked at in total to begin to push the ball forward but it is the real slog now. Believe me, in 1996 I knew how to build a Roadrunner system, not a problem; it was just a matter of hard work. That was the first petascale system on the planet. In 2005 I didn't know how to get to exascale and still struggle today. We are up against real limits.

Mr. MASSIE. So does anybody else care to talk about that?

Dr. GILES. Yes, just to add one quick observation. The Secretary of Energy Advisory Board Task Force considered very seriously this question about the relationship of what we are doing now to—for the future, and one of the things that became very clear is that because the limitations and the possibilities and opportunities are physics-based and the DOE labs are the premier research set of facilities for the physical sciences, that in some ways DOE with its computing interest and capability and the labs is in an excellent position to do the research needed to move beyond silicon and CMOS and what we are doing now to the next generation, whether that involves, as David said, superconducting technology or quantum technology, the labs are in a really good position to investigate.

Mr. MASSIE. That was going to be my next question. So obviously we have already hit the physical limits of silicon and the speed of light and energy density and all that stuff, and we have Band-Aided that with architecture or maybe that is the way around it, but we have diminishing returns to putting more cores in there. What are the next promising platforms and what role will our research that we are paying for here in Congress play? What is the next transistor? What is going to be the next paradigm shift and what role does our research play in that?

Mr. TUREK. Well, I will make a brief comment. There is no silver bullet. There is nothing I can point to that says the problems of the future are done; we can simply move along as systematically as we have over the last 50 years or so. When I talk about architecture I mean different approaches to solve the problem.

Today, one of the techniques that is being explored and reflected in the CORAL program at the DOE is the employment of accelerators, specialized processors attached to conventional processors to give an overall speed-up in compute capability. We pioneered this, by the way, with a cell processor at Los Alamos ten years ago, which was an accelerator-based kind of technology. That is a new idea. Accelerators have been thought of over many years but never gained acceptance because we could leverage the evolution of silicon to overcome the limits. No longer possible, now there is an embrace of accelerators. So you see a lot of different kinds of accelerators come into play and applied in very unique and interesting kinds of ways.

Mr. MASSIE. Thank you very much. I am excited to see what the next breakthrough is. I realize there is no silver bullet and we have got to use a shotgun, but I trust that we will come up with something. Thank you.

Chairman WEBER. Thank you. And I thank the witnesses for their valuable testimony and the Members for their questions.

The record will remain open for two weeks for additional comments and written questions from the Members.

So thank you, gentlemen. Thank you, Dr. Giles. The witnesses are excused and the hearing is adjourned.

[Whereupon, at 10:07 a.m., the Subcommittee was adjourned.]

Appendix I

Answers to Post-Hearing Questions

ANSWERS TO POST-HEARING QUESTIONS

Responses by Mr. Norman Augustine

N. Augustine Responses to Questions from Cong. Lipinski

1. This is a very important but very difficult question to answer. While one does not want to tolerate individuals taking advantage of conflicts; on the other hand, if an individual has absolutely no conflicts it is also likely that they have no hands-on experience in the subject at hand. I would recommend seeing what the current Congressionally-directed commission investigating this and other topics reports and what are its specific recommendations on how to improve the flow of technology from the National Laboratories to Industry. Among the matters to be considered are making it easier for individuals to move between industry and government and to serve on government advisory boards; making it easier for industry personnel to work in the laboratories on a temporary basis; and making IP ownership by both academia and government less restrictive. In the end, the best way to transfer technology is to transfer people.

2. I regret that I have nothing to add to the testimony of the other panelists on this particular topic.

3. The problem the question addresses begins during, or even before, grade school. Only 15 percent of U.S. high school graduates (excluding the 1/3 that drop out of high school) are prepared to even begin studying engineering. We must drastically improve the quality of STEM education in pre-K–12 and must reverse the states' disinvestment in higher education. We should also make it easier for highly qualified engineers and scientists who come to America for their education to remain here and work here.

4. The energy demand of high performance computers is becoming a serious problem not only in terms of the cost of that energy but also in the restrictive design implications of eliminating waste heat. Research on such fundamental issues as reducing energy demands of high-performance computing systems will impact not only those computers *per se*, but to other applications of advanced microelectronics.

5. An excellent question. It is essential to (1) select trusted, reliable partners, and (2) assure *quid pro quo* agreements. Given the speed at which knowledge moves around the world today, particularly in the computational sciences, one cannot long maintain a lead through secrecy. The answer is to invest so as to advance the state of the art at a pace that keeps the nation at the leading-edge of new knowledge.

Responses by Dr. Roscoe Giles

Draft Additional Question Responses – R. Giles

1. *Innovation in high performance computing is critical to keeping our nation competitive. An important part of innovation is transferring technologies into the private sector so that they can help drive our economy. Mr. Augustine described in his testimony that it can be difficult for the National Labs to work closely with industry, for example due to well-intended conflict of interest rules. How can we improve the transfer of technology between National Labs and industry?*

In the area of high performance computing (HPC) there are several avenues for transfer of knowledge and for harnessing capabilities and facilities of the DOE labs to advance innovation. These include: (a) programs to offer access to and support for industry to use HPC facilities for their own research; (b) direct interaction with private sector startups through SBIR (Small Business Innovation Research) and STTR (Small Business Technology Transfer) programs; and (c), for the longer term, insuring dissemination of research results more broadly to society and particularly to the next generation through the university engagement.

High performance computing advances science and engineering across many fields by facilitating the discovery of new knowledge and enabling sophisticated engineering design that underlies innovation. For all recent generations of DOE HPC facilities, programs have been developed to give industry and researchers around the country access to top DOE computing capabilities. The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) and the ASCR Leadership Computing Challenge (ALCC) allocation are programs open to all researchers, including those from industry. Additionally the DOE Labs have industrial partnership programs that promote technology transfer.

SBIR activities in ASCR have emphasized software based in ASCR R&D, in particular trying to enable independent software vendors to package and support software based on open source software developed at DOE.

Another mode of "tech transfer" that I would emphasize is engagement of young professionals and students at all stages in DOE research. The knowledge and skills gained by them can significantly impact future innovation in the labs, universities, and private sector.

2. *I am glad to see DOE continuing to push the frontiers of our nation's supercomputing capacity with the CORAL initiative. I am especially glad to see Argonne, who is in my district, as one of the three National Labs in CORAL. In both Dr. Giles and Dr. Crowley's testimony, you mention the importance of an "all-in" approach. In other words, hardware alone will not develop the next generation of super computers. Argonne will play a critical role in this approach by developing a system architecture that is different than the one used by Oak Ridge and Livermore. Can you describe what are the key elements to an "all-in" approach, and how ASCR [pronounced "Oscar"] will address those elements?*

Indeed, Argonne has played a critical role in the DOE scientific computing complex and continues to do so as we move aggressively into the exascale era. Argonne's contributions in pioneering parallel systems and software over the years have been key to creating, deploying, and evaluating supercomputer systems, ranging from parallel clusters that have become the defacto standard for most deployed HPC systems to the current generation of petascale systems.

On a personal note, I recall that when we got our first generation of massively parallel supercomputers at Boston University, I went to Argonne workshops to learn how to approach programming in such a novel environment.

As a leading edge site, Argonne continues to play an essential role in designing and developing the next generation supercomputers.

Key elements of the "all-in" approach to creating the ecosystem for next generation HPC are the balanced development of hardware, software, applications, and their mathematical and computational foundations. Since we cannot explore all possible future paths in such a wide variety of directions, it continues to be important that we coordinate development in these areas and invest artfully so as to maximize the impact on science.

3. *It is our responsibility to train a workforce that can meet our future needs, especially in such a critical area as high performance computing. I was at the launch of the Blue Waters computer at the University of Illinois, which is a powerful tool for training future generations. However, not every university or federal lab has access to such powerful tools. What can we do to improve STEM education and attract people to this important field?*

There are a number of things that DOE might do directly and indirectly to improve STEM education and attract people to HPC fields.

Direct action includes programs like the Computational Science Graduate Fellowships (CSGF) that help train the next generation of computational science leaders, including giving them experience with DOE facilities and at DOE Laboratories. This model could easily be expanded to other STEM fields where direct experience with the best scientific tools can inspire students and transform their understanding of science.

Research experiences, laboratory internships, and cooperative arrangements with universities can impact a range of people along the standard "pipeline." Our efforts to diversify the workforce to include underrepresented minorities, women, and people with disabilities has taught us much more about how to bring people in outside of the conventional pipeline and how to connect people with STEM in ways that respect their interests and goals.

Finally, current and emerging computing technologies (especially high performance, low power exascale technology) can transform what is available in educational and personal settings. Combined with our increasing network capabilities, we can do more distributed science analysis and computing that supports STEM observational science, projects, and modeling.

4. *What kinds of research need to be done to bring down power requirements for exascale computers? In Mr. Turek's testimony, he states that the cost for electricity for the fastest computer in the world could reach as high as $100 million by the end of the decade. If left unchecked, this will seriously impact the adoption of supercomputer technology.*

Ongoing work on delivering HPC at much lower power cost has been a key issue to address in order to enable exascale computing. As you mentioned, straightforward extrapolation of current technologies to exascale levels would lead to totally unacceptable power costs and cooling challenges. The assessment of how to produce systems at a much higher compute power to electrical power ratio has been a research and development focus.

ASCAC identified this as one of the ten top technology challenges for exascale computing in our 2014 report on this subject. The challenges in this report are all challenges that we believe can be addressed and which will be addressed in the exascale development plans.

With regard to the power challenge, ASCR has been working with both researchers and industry to accelerate the development of high performance power efficient systems that incorporate hardware advances and system software. The "fast forward" and "path forward" programs represent the industry side of this collaboration.

On the path to exascale, we are seeing success in the goal of more performance and reduced power. For example, the current Summit system (part of the CORAL acquisition) is expected to deliver 5 times the computing capability of currently deployed petascale systems at only a 10% increase in power.

5. *In the area of high performance computing, there are international partnerships, such as DOE's Energy Sciences Network that recently expanded to include European support. How can these efforts reduce the cost of reaching Exascale capabilities without damaging economic competitiveness or national security? Because on the one hand, it makes sense to combine our efforts with those of other countries, but on the other hand we don't want to give up our nation's lead in high-end computing.*

This is a very good point – finding ways to work together internationally without compromising essential national competitive advantages.

My thoughts on the matter are to focus first on areas where the threat to competitiveness is relatively small or where the goals to be achieved are shared across national boundaries.

There is a significant opportunity for an international approach to creating exascale systems and application software. In this case, open source software developed with DOE support is already visible to the world. Thus there is little to lose and much to gain from collaboration with international partners in this area because so much software innovation is needed for the new exascale systems.

Another area of natural collaboration is for applications of next generation computing, and particularly data science and networking, in service of international, collaborative, large scale scientific instruments. The LHC and ITER are examples of this as are US facilities with significant international participation. Developing applications to effectively analyze data from such instruments is a shared international goal and therefore a natural area for sharing of development responsibilities.

Finally, the applied mathematics that underlies scientific software, like mathematics in general, is an international enterprise.

Responses by Mr. David Turek

Question 1 - Innovation in high performance computing is critical to keeping our nation competitive. An important part of innovation is transferring technologies into the private sector so that they can help drive our economy. Mr. Augustine described in his testimony that it can be difficult for the National Labs to work closely with industry, for example due to well-intended conflict of interest rules. How can we improve the transfer of technology between National Labs and industry?

Improving transfer of technology between the National Labs and industry is an important goal given the amount of creative innovation occurring within all of the Labs. To improve transfer I would recommend a "marketing" effort to more broadly inform industry and the American public more generally about the capabilities embodied in the Labs. I believe that only a very small minority are aware of what is possible in terms of collaboration or partnership with the Labs. Second, I believe that a career development program be instituted within the Labs to encourage and enable personnel to take a 1-2 year assignment in private industry on a periodic basis. The establishment of personal ties goes a long way towards fostering closer collaboration between industry and the Labs. Third, and most important, I would recommend a consideration to streamline the processes governing the sharing and creation of intellectual property. In my experience, agreement with the Labs on collaborative efforts are easily struck at a conceptual level but the negotiations over IP and risk mitigation are long, involved, and costly with terms especially onerous to firms bringing background IP into the discussion. There are times when firms are placed in a position where failure or abandonment of a project puts the firm's own background IP at risk. In any case, the complexity of negotiation over IP can be, by itself, a deterrent to developing industry-Laboratory collaborations.

Question 2 - I am glad to see DOE continuing to push the frontiers of our nation's supercomputing capacity with the CORAL initiative. I am especially glad to see Argonne, who is in my district, as one of the three National Labs in CORAL. In both Dr. Giles and Dr. Crowley's testimony, you mention the importance of an "all-in" approach. In other words, hardware alone will not develop the next generation of super computers. Argonne will play a critical role in this approach by developing a system architecture that is different than the one used by Oak Ridge and Livermore. Can you describe what are the key elements to an "all-in" approach, and how ASCR y will address those elements?

The notion of an "all-in" approach is meant to convey the need to systematically determine the necessary and sufficient conditions that contribute to success:

hardware is necessary but it is not sufficient. To create a comprehensive capability software must be emphasized. By software I mean the base system software that affects how easy it is for the system to be used as well as the application software that enables scientific inquiry. All too often (and this is an international maxim as well) government funds hardware and leaves the development of software as an after-thought. I would argue that the state of much of the application software in use today on supercomputers is archaic and falling behind the evolutionary hardware architectures. This is hardly surprising given the imbalance of emphasis on hardware and software. ASCR can play a role in remedying this situation by funding and orchestrating programs designed to modernize supercomputing software. By careful management of the process, ASCR can also insure that the software can be substantially portable so that software running on the Argonne architecture could move with relative ease to the architecture of the other CORAL systems. Having a coherent, well designed program on software will minimize costs but more importantly create focus so that the best possible software can be created.

Question 3 - It is our responsibility to train a workforce that can meet our future needs, especially in such a critical area as high performance computing. I was at the launch of the Blue Waters computer at the University of Illinois, which is a powerful tool for training future generations. However, not every university or federal lab has access to such powerful tools. What can we do to improve STEM education and attract people to this important field?

Improving STEM education is a necessity for enhancing US capability in computational science. I think that the attractiveness of the field is a consequence of the problems it attempts to solve, and it is the nature of the problems themselves which can motivate interest in STEM education. In the 1990s much conversation around Supercomputing was embodied in discussions of so-called "Grand Challenge Problems"—the kind of science problems previously considered unfathomable that could be pursued, studied, and maybe solved with supercomputers. Today, it would be interesting to revisit the notion of Grand Challenges but perhaps cast them in an industrial context. For example, if senior representatives of many industries came together to define their most vexing engineering and technology problems, perhaps a new class of computational science problems would emerge that would fuel the imagination of younger people and stimulate interest in STEM education. The linkage of problem to training seems to me to always motivate a better outcome than the pursuit of training without a well-defined goal.

Question 4 - What kinds of research needs to be done to bring down power requirements for exascale computers? In Mr. Turek's testimony, he states that the cost for electricity for the fastest computer in the world could reach as high as $100 million by the end of the decade. If left unchecked, this will seriously impact the adoption of supercomputer technology. Also, what kinds of benefits could research into energy efficient high performance computing bring to other computing systems?

The power requirements for exascale class supercomputers are indeed a challenge. There is much to be done in terms of fundamental research in materials and nanotechnology. But there is also much to be done in both systems architecture and software design. Systems architecture specifies the way in which a computer is constructed and operated. Today, for example, the computer industry builds supercomputers out of large numbers of multicore processors and operates the system in a massively parallel fashion. This helps to constrain power usage but also puts a substantial burden on software developers to write software utilizing this kind of architecture. Also, we are beginning to see the emergence of software that pays close attention to which parts of the system are consuming power and making utilization adjustments accordingly. While it is well understood generally that microprocessors are very energy consumptive, the energy profile of memory, networks, and data storage devices are also becoming critical elements in energy efficiency of supercomputers. Research into energy efficiency for supercomputers must, therefore, be comprehensive with respect to all the component technologies. Cloud based computer deployments will clearly benefit as will computers of any scale that utilize these component technologies.

Question 5 - In the area of high performance computing, there are international partnerships, such as DOE's Energy Sciences Network that recently expanded to include European support. How can these efforts reduce the cost of reaching Exascale capabilities without damaging economic competitiveness or national security? Because on the one hand, it makes sense to combine our efforts with those of other countries, but on the other hand we don't want to give up our nation's lead in high-end computing.

International partnerships on high-end computing can be a double edged sword but, on balance, I think the benefits out-weigh the risks. One key benefit of international partnership is the participation and creation of standards. Standards have always helped to drive innovation in the computer industry and a "go-it-alone" approach always carries the risk of becoming viewed as non-standard. Standards also help to reduce costs and speed technology deployment. I expect that there will be a tremendous amount of open source software created in support of exascale computing and the more there are standards in place, the more this body of open source software can be utilized across US systems. This is likely to be both a significant cost saving over trying to build all the software in the US as well as an acceleration of innovation courtesy of the larger software community working in concert. Leadership can be preserved and extended by developing home-grown skills in using supercomputers. The number of people who can effectively use large supercomputers is relatively small but the majority reside in the US; training to expand the size of that community will have the greatest impact on the ability of the US to maintain leadership in deriving value from supercomputing.

Responses by Dr. James Crowley

siam®
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS

3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Phone +1-215-382-9800
Fax +1-215-386-7999
www.siam.org · siam@siam.org

Jim Crowley Responses to Hearing Questions for the Record
House Science, Space, and Technology Committee Subcommittee on Energy
Supercomputing and American Technology Leadership

Question 1
Innovation in high performance computing is critical to keeping our nation competitive. An important part of innovation is transferring technologies into the private sector so that they can help drive our economy. Mr. Augustine described in his testimony that it can be difficult for the National Labs to work closely with industry, for example due to well-intended conflict of interest rules. How can we improve the transfer of technology between National Labs and industry?

Jim Crowley Response: Technology transfer and collaboration between the national labs and industry are critical for US competitiveness, and many innovations developed by ASCR have enabled advances in commercially-available computing that help industry and academia. SIAM supports efforts to enhance collaboration and technology transfer, but notes that any changes to regulations should be considered with an eye to national security and conflict of interest concerns.

As a society with members from national labs, academia, and industry, SIAM is well-suited to help seed communication and collaboration between these sectors, and is happy to help with any efforts to bring these groups together. For example, SIAM is deeply involved in the European Service Network of Mathematics for Industry and Innovation, which connects mathematics experts with industry to advance innovation (http://www.eu-maths-in.eu/).

Many companies both big and small cannot employ computational scientists and engineers at the critical mass needed to tackle computational research problems. However, industry domain expertise coupled with the National Labs computational science and engineering expertise can help industry and the National Labs succeed and advance US competitiveness.

Question 2
I am glad to see DOE continuing to push the frontiers of our nation's supercomputing capacity with the CORAL initiative. I am especially glad to see Argonne, who is in my district, as one of the three National Labs in CORAL. In both Dr. Giles and Dr. Crowley's testimony, you mention the importance of an "all-in" approach. In other words, hardware alone will not develop the next generation of super computers. Argonne will play a critical role in this approach by developing a system architecture that is different than the one used by Oak Ridge and Livermore. Can you describe what are the key elements to an "all-in" approach, and how ASCR *[pronounced "Oscar"]* will address those elements?

Jim Crowley Response: The "all-in" approach as described by ASCR head Steve Binkley, is to consider hardware, software, applications, large data, and underpinning applied math and computer science "all in" the exascale program. The "all-in" approach is critical to the success of exascale computing, because almost all technical challenges require an integrated solution approach that takes software, hardware, computer science, applied mathematics, and science into account.

siam®
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Phone +1-215-382-9800
Fax +1-215-386-7999
www.siam.org · siam@siam.org

History has shown that increased computational capabilities have grown in equal measure by improvement in hardware, algorithms, and modeling advances --- enabled by new software. The most expensive of these is the hardware advances. But to maximize the return on investments in hardware, we must also invest in algorithms and modeling. This can be because new hardware architecture can change which algorithm is most efficient for running an existing model or for analyzing the increased data stream.

As the number of nodes increases, transferring data across nodes can become a bigger computational bottleneck. This fact means that new algorithms must be designed for optimum performance.

It also can be because the increased spatial resolution enabled by new hardware can introduce spatial scales at which the models used by current codes are inadequate.

For example, when the finest resolution in a weather simulation is six miles then clouds must be (and are) modeled rather simply. As the resolution enabled by new machines drops below .25 miles then the structure of clouds comes into focus. Current models of cloud physics need to be improved to gain full advantage of this increased resolution. These new models also require new algorithms. In either case, new software is needed.

As another example, the anticipated billion-way concurrency (exascale machines will be able to perform a billion operations at the same time) fundamentally changes the characteristics of high-performance computing. Whereas traditionally, parallel computing was limited by the number of parallel operations that could be performed, future systems will be limited by the amount of data that can be transferred between compute nodes. Consequently, in order to achieve the fastest speeds and realize the full capability of exascale machines, applied mathematicians must re-design the underlying algorithms to minimize communication between nodes.

In addition, exascale computing opens the possibility to new capabilities such as the design of new materials through optimization. Problems such as these require an "all-in" approach, where application scientists formulate new models, applied mathematicians develop new algorithms to solve these models, and computer scientists ensure that the algorithms run efficiently on emerging architectures.

It will be important that we not wait for the final exascale system to address these algorithm, software, data, and application challenges but begin immediately to work across industry, academia and national labs to address these problems on all fronts allowing for an evolution rather than a revolution in solutions.

SIAM expects ASCR to develop the exascale program going forward to include all of these efforts. SIAM particularly supports, and underscores the importance of underpinning applied mathematics and computer science research aimed at developing more effective mathematical models, algorithms, and scientific codes to assure success in delivering scientific results to the country.

siam.
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS

3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Phone +1-215-382-9800
Fax +1-215-386-7999
www.siam.org · siam@siam.org

Question 3

It is our responsibility to train a workforce that can meet our future needs, especially in such a critical area as high performance computing. I was at the launch of the Blue Waters computer at the University of Illinois, which is a powerful tool for training future generations. However, not every university or federal lab has access to such powerful tools. What can we do to improve STEM education and attract people to this important field?

Jim Crowley Response: Researchers trained in computational science and working in universities, national laboratories, and industry are central to DOE's mission and essential to propel advances in DOE mission-critical fields such as nanotechnology, biofuels, genomics, and materials fabrication. The DOE Computational Science Graduate Fellowship program plays an important role in developing the future workforce that will take advantage of exascale computing in support of broad scientific advancement for the nation. SIAM encourages that expansion of this program, and the development of similar programs targeted at developing a well-trained applied mathematics and computer science workforce that will be able to support the DOE exascale program. It is important that these students have access to high performance computers during their training, such as those at the DOE Labs, and so any such program should include such access.

Every STEM field is currently evolving its curriculum to meet the challenges brought on by data science. Computational science is a central part of these developments. The transformation of data to information to knowledge to action is largely mediated by algorithms and models. SIAM and other professional societies have been working to nurture changes in the STEM curriculum that will help train a workforce capable of meeting these challenges. Important to these efforts are programs at the National Science Foundation that fund research and implementation activities to improve STEM education and workforce development across STEM disciplines.

Question 4

What kinds of research needs to be done to bring down power requirements for exascale computers? In Mr. Turek's testimony, he states that the cost for electricity for the fastest computer in the world could reach as high as $100 million by the end of the decade. If left unchecked, this will seriously impact the adoption of supercomputer technology.

Also, what kinds of benefits could research into energy efficient high performance computing bring to other computing systems?

Jim Crowley Response: SIAM agrees that reducing power requirements for exascale computers is a key research challenge, but we are not experts in this area and thus leave it to other witnesses to respond. I would note that we expect any improvements leading to exascale technology will "trickle down" into smaller computers that will be available for small industry, and academic departments, bringing broad benefit to US competitiveness.

siam®
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Phone +1-215-382-9800
Fax +1-215-386-7999
www.siam.org · siam@siam.org

Question 5

In the area of high performance computing, there are international partnerships, such as DOE's Energy Sciences Network that recently expanded to include European support. How can these efforts reduce the cost of reaching Exascale capabilities without damaging economic competitiveness or national security? Because on the one hand, it makes sense to combine our efforts with those of other countries, but on the other hand we don't want to give up our nation's lead in high-end computing.

Jim Crowley Response: At pre-competitive research stages international collaborative research can be very helpful. For example, the development of fundamental applied mathematics and computer science tools that will support exascale should leverage international collaborations to catalyze advances and reduce costs. These collaborations would not damage economic competitiveness or affect national security issues as it is already regular practice for applied mathematicians to collaborate internationally.