

Statement of Rick Stevens
Associate Laboratory Director
for Computing, Environment and Life Science
Argonne National Laboratory

Before the
Subcommittee on Energy
of the Committee on Science, Space and Technology
U.S. House of Representatives
May 22, 2013

Thank you Chairman Lummis, Ranking Member Swalwell, and Members of the Subcommittee. I appreciate this opportunity to talk to you about the future of high performance computing research and development, and about the importance of continued U.S. leadership in the development of Exascale computing.

I am Rick Stevens, the Associate Laboratory Director responsible for Computing, Environment, and Life Sciences research at Argonne National Laboratory, one of America's first and largest multipurpose science and engineering laboratories. My research focuses on finding new ways to increase the impact of computation on science – from advancing new architectures to developing large-scale applications for computational genomics targeting infectious disease research, energy and the environment. I also am a Professor at the University of Chicago in the Department of Computer Science, where I hold senior fellow appointments in the University's Computation Institute and the Institute for Genomics and Systems Biology. In addition, I am the co-founder of the University of Chicago/Argonne Computation Institute, which advances interdisciplinary projects involving computation.

I believe that advancing American leadership in high-performance computing (HPC) is vital to our national interest. High-performance computing is a critical technology for the nation, it is needed by all branches of science and engineering, and it is a critical policy tool for government leaders. More importantly, its availability is a pacing item for much of science and for many technological developments on the horizon. Today the United States is the undisputed leader in the development of high-performance computing technologies both hardware and software. However, the nation that succeeds in leading in high-performance computing and large-scale data analysis for the long term will gain an insuperable competitive advantage in a wide array of sectors, including advanced manufacturing, energy, health care, space, defense, transportation, education, basic science and information technology.

The next stage of international leadership in computing is the development of systems capable of 100x to 1000x times more performance on both modeling and simulation tasks, as well as the analysis of ever-increasing large and complex datasets from commerce, the internet, healthcare and science. This next stage of development is known as Exascale (i.e. 10^{18} operations per second or 10^{18} bytes of storage). The availability of Exascale computing capabilities will improve our economic competitiveness, giving companies the capability to use modeling and simulation at unprecedented speed and fidelity to spur creativity and speed development of innovative new products. Just as importantly, Exascale computing will be an extraordinarily powerful scientific tool, enabling more accurate predictive models and facilitating analysis of massive quantities of data, making it possible for researchers to tackle and solve problems where experiments are dangerous, impossible, or inordinately costly. For example, it has been remarked that practical high-throughput genome sequencing would have been impossible without high-performance computing. The rise of 3D printing and digital fabrication, which have the potential to transform our economy, likewise relies on computing in fundamental new ways.

Today, high-performance computing plays a central role in our nation's scientific, technical and economic enterprise, and the global community looks to our fiercely competitive HPC industry for cues on where this vital technology is heading next.

Visitors from around the world visit the DOE National Laboratories to learn about how we develop, deploy and use large-scale computers to attack the hardest problems in science and engineering. The DOE National Laboratories are recognized worldwide as thought leaders in the development and use of the largest systems, they are developers of the most advanced systems software and tools, and they pioneered the idea of open source math libraries long before open source was recognized as a key to rapid progress in software.

In the DOE Office of Science, the Advanced Scientific Computing Research program supports the operation of four national scientific user facilities:

- Energy Sciences Network (ESnet), a high-speed network serving thousands of Department of Energy researchers and collaborators worldwide. Managed and operated by the ESnet staff at Lawrence Berkeley National Laboratory, ESnet is one of the DOE's most widely based and successful cooperative efforts, providing direct connections to more than 30 DOE sites at speeds up to 10 gigabits per second. ESnet allows scientists to access unique DOE

research facilities and computing resources independent of time and location with state-of-the-art performance levels.

- Oak Ridge National Laboratory Leadership Computing Facility (OLCF) serves all areas of the research community through programs such as the Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. Last year, INCITE awarded nearly a billion processor hours on OLCF's Jaguar system to 35 projects from universities, private industry, and government research laboratories.
- Argonne Leadership Computing Facility (ALCF) enables research spanning a diverse range of scientific areas, providing expertise and assistance in support of user's projects to achieve top performance of applications and maximize resources. The ALCF has allocated over 2.1 billion processor hours via INCITE to over 37 projects for the current year.
- National Energy Research Scientific Computing (NERSC) Center, a division of Lawrence Berkeley National Laboratory located at the University of California, Oakland Scientific Facility, serves over 3,000 scientists throughout the United States each year. These researchers work at DOE laboratories, universities, industrial laboratories and other Federal agencies. Computational science conducted at NERSC covers the entire range of scientific disciplines, but is focused on research that supports DOE's missions and scientific goals. The NERSC staff delivers critical computing resources, applications and information enabling users to optimize the use of their computer time allocation.

Over the years, the National Laboratories have formed important collaborations and partnerships that have strengthened our nation's computing leadership. Argonne and Lawrence Livermore National Laboratory (LLNL), along with industry partner IBM, pioneered a highly successful co-design approach to develop Petascale computers for grand challenge science resulting in the BG/L, BG/P and BG/Q systems. Additionally, Oak Ridge National Laboratory (ORNL), Argonne and LLNL have developed a partnership, known as CORAL (Collaboration of Oak Ridge Argonne Livermore), to develop the next round of pre-Exascale machines. Our success has been made possible by strong public/private partnerships, supported by thoughtful ongoing investments by our leaders in Washington.

In recent years, other nations have challenged our present dominance and future information superiority by making substantial, sustained investment in high performance computing research and systems deployments. Today, the Department of Energy's Titan supercomputing system at ORNL tops the current list of the world's most powerful supercomputers, and U.S. machines hold five of the top 10

spots. (Of those, three are sited at DOE National Laboratories: Titan at Oak Ridge, Sequoia at Lawrence Livermore, and Mira at Argonne.) But other nations are increasingly aware of the strategic importance of HPC and are creating supercomputing research programs that rival our own. (Most notably, China has announced plans to build more than a dozen supercomputing centers, with an announced goal of reaching Exascale capability by 2018.) Japan is planning a next-generation supercomputing project with an estimated budget of 110 trillion yen, and Europe has established PRACE (Partnership for Advanced Computing in Europe) to advance high-performance computing and to re-establish an HPC industry in Europe.

Right now, our competitors are relying primarily on American technology to create these powerful machines. But increasingly, other nations are developing the expertise and technology to build supercomputers that could rival or even surpass American-made high-performance computing systems.

I have been working over the past five years with my colleagues in the National Laboratory system, academia and private industry to develop an integrated, ambitious plan to keep the United States at the forefront of high-performance computing, a plan that will enable us to reach the Exascale performance level on critical DOE applications by the end of this decade. Our plan is based on a continuing, long-term partnership between computer hardware manufacturers and laboratory based teams of scientists, mathematicians and engineers who develop mathematical models, write parallel computer programs that implement those models on hundreds of thousands of processors, and develop the software that enables the computers operate efficiently. These teams are supported in turn by hundreds of scientists and engineers who develop and operate the machines.

Already, our joint efforts have led to major advances in nearly every area of computational science. For example:

- Researchers at the University of Chicago are using the Argonne Leadership Computing Facility's resources to pursue transformative breakthroughs in materials for batteries and fuel cells. At present, much of material science discovery still relies on the traditional "Edisonian", intuition-based "trial and error" approaches. The ALCF enables multi-scale modeling of charge transport processes in materials relevant to fuel cell and battery technologies – an effort that could have a significant impact on chemistry and material science while improving our nation's energy security. This research will have important implications for the work of the Joint Center for Energy

Storage Research, the new Battery and Energy Storage Hub headquartered at Argonne and funded through the DOE's Office of Basic Energy Sciences.

- Researchers at LLNL, working with colleagues at the IBM T. J. Watson Research Center in New York, used Sequoia to develop a highly scalable code, called Cardioid, that replicates the electrophysiology of the human heart. By accurately simulating the activation of each heart muscle cell and the cell-to-cell electric coupling, Cardioid will give researchers groundbreaking new insights into serious arrhythmia, a heart malfunction that causes 325,000 deaths in the United States each year.
- Pratt & Whitney is working with the Argonne Leadership Computing Facility (ALCF) to perform "virtual testing" of jet engine designs. Engine improvements based on these computer simulations have contributed to 15% improved fuel burn, saving \$1.5 million per airplane each year while reducing carbon emissions by 3,000 tons. These design upgrades were made possible by ALCF researchers who were able to implement code improvements that resulted in a 10-fold performance advance in the Pratt & Whitney simulations.
- The Consortium for Advanced Simulation of Light Water Reactors (CASL), a group of national laboratories, universities and industry partners led by ORNL, is using the Titan system to develop a virtual nuclear reactor simulation toolkit that will model the interior of a nuclear reactor and gain information to aid design of next-generation reactors, and to improve the safety and performance of reactors currently in use.
- Procter & Gamble researchers used the ALCF to investigate the molecular mechanisms of bubble formation in foams, performing unprecedented computer simulations of the dissolving of soap and foaming of suds. By better understanding the process of "sudsing," Procter & Gamble can evaluate new materials more quickly and create better, less expensive consumer products, foods, and materials for fire control.
- A group at the National Center for Atmospheric Research is collecting weather data from the past 150 years to achieve a more detailed understanding of climate, with a goal of predicting severe events such as Superstorm Sandy. Using NERSC, the team is analyzing massive datasets that use the history of global weather patterns to validate climate and weather models and enable more accurate forecasting.
- Researchers from the University of California, San Diego are using the ALCF to target Parkinson's disease, a progressive and devastating disorder of the nervous system that affects more than 2 million Americans and costs an estimated \$25 billion annually in medical treatments and lost productivity.

Supercomputers accelerate drug development research by making it possible for medical researchers to model the proteins that cause Parkinson's and rapidly identify possible therapeutic interventions.

The HPC systems now being used by these researchers, and by thousands more nationwide, offer speeds that are measured in Petaflops – a quadrillion sustained floating-point operations per second. These systems took roughly five years from design to deployment, and have productive life spans of about five years. At present, the HPC industry in the United States is capable of delivering new frontier-level systems every three to four years, through the efforts of a complex and interwoven “ecosystem” that brings together industry, national laboratories and universities. As we look to the future, however, it is clear that the domestic demand for HPC access is outpacing supply, and the gap between our HPC capabilities and our national needs will widen dramatically as more and more leaders in private industry come to understand how high-performance computing can accelerate and support their R&D programs.

To better understand our nation's HPC needs for the future, DOE held twelve discipline-oriented workshops over the past few years to determine the scientific and engineering opportunities and priorities for HPC for this decade and beyond. These workshops brought together more than 1,200 scientists and engineers from around the country and around the world to discuss the central problems in each field, to identify important questions that might be addressed through advances in computation, and to determine the performance requirements for those computations.

The DOE workshops led to identification of pressing questions and important problems that will require Exascale computing capability to solve. For example, we want to use first principles to design new materials that will enable a 500-mile electric car battery pack. We want to build end-to-end simulations of advanced nuclear reactors that are modular, safe and affordable. We want to add full atmospheric chemistry and microbial processes to climate models, and to increase the resolution of climate models to provide detailed regional impacts. We want to model controls for an electric grid that has 30 percent renewable generation. We want to create personalized medicines that will incorporate an individual's genetic information into a specific, customized plan for prevention or treatment. In basic science, we would like to study dark matter and dark energy by building high-resolution cosmological simulations to interpret next generation observations. All of these challenges require machines that have more than 100x-1000x the processing power of current supercomputers.

The challenges of creating Exascale systems will require significant changes both in the underlying hardware architecture and in the many layers of software above it. To address those challenges, researchers at the National Laboratories and throughout the HPC community are working together to design and develop architectures, operating systems, runtime, storage, languages and libraries, and application codes. We are thinking about new algorithms that will be “Exascale ready,” and we are building co-design teams that bring together computer scientists, mathematicians and scientific domain experts to work on solving these problems. We also are working with existing applications communities to guide their decisions about rewriting codes for near-term opportunities so their work will transition smoothly to Exascale systems. As we move ahead, this process will require increasing communication and continued refinement of ideas among a larger-than-normal group of stakeholders. In the past, architects could use rules of thumb from broad classes of applications or benchmarks to resolve design choices. But Exascale will require an aggressive co-design process that makes visible to the whole team the costs and benefits of each set of decisions on the architecture, software stack, and algorithms.

As we transition to Exascale, the hierarchy of systems will largely remain intact. That means the advances needed for Exascale will be hugely beneficial to Petascale computing and so on down the computing space. So, for example, if improved energy efficiency and better software solutions for resilience are developed as part of Exascale research, then it becomes possible to build Petascale computers out of less expensive components.

We have identified five major hurdles that must be overcome if we are to achieve our goal of pushing the computing performance frontier to the Exascale by the end of the decade:

- We must reduce power consumption by at least a factor of 50.
- We must increase the parallelism of our applications software and operating systems by at least a factor of 1,000.
- We must develop new programming methods to increase dramatically the number of programmers that can develop parallel programs.
- We must improve memory performance and cost by a factor of 100.
- We must improve systems reliability by at least a factor of 10.

Let me now address each of these major challenges in a bit more detail.

Power Consumption

The majority of performance improvements are achieved by increasing the number of processors, which requires an increase in power consumption. Today's most powerful supercomputer systems require a few megawatts of electricity; with current power prices at roughly \$1 million per megawatt year, it is clear that we must find ways to reduce power consumption of both processors and memories by 50x over the next decade if we are to create Exascale systems that are affordable to operate. At present, industry is on track to lower power consumption by approximately 5x, but additional research is needed to lower it by another factor of 10, resulting in Exascale systems that will consume approximately 20 megawatts.

To reach our power consumption goals, we must redesign the processors and the memories that feed the processors, each of which contribute about equally to the power used. We will also need to replace copper wiring that supports communications between the processors with lower-power optics. Preliminary research work has indicated that major changes in processor design could save up to 20x on power consumption, so there is reason to be optimistic that we can achieve the necessary energy efficiency. These improvements will not only make it feasible to build and operate Exascale supercomputers for science; they also will have a positive cascading impact on energy consumption across the entire information technology (IT) sector. Incorporating these advances in products in the broad market would improve power consumption in large-data centers and extend battery life in laptops and handheld devices, making computer systems more power-efficient and therefore more affordable, while reducing their environmental impact. Given that an estimated 5 percent of global energy consumption, and of global carbon dioxide emissions, is attributed to computing services, energy efficiency improvements of this magnitude could have a significant impact on the environment.

Increasing Parallelism

In the past, supercomputer performance gains have been achieved by improving the speed of individual processors. Now, however, we have reached practical limits in both features sizes and power consumption, which means nearly all future performance improvements will come through increasing the number of CPU "cores," or processing units, that can be applied to a single problem. Today, your laptop has a few cores; the biggest systems have approximately 1,600,000 cores. In the future, we expect this number to grow linearly with overall performance. So to reach our goal of

100-fold to 1,000-fold improvement in performance, we will need to increase the number of cores by approximately 100x -1,000x, resulting in systems of 100,000,000 cores or more. This transition, from faster single processors to increased numbers of processors, will require an equal increase in the parallelism, or concurrency, of our applications and systems software. In response to this challenge, the community has been working hard to develop parallel programming technology and tools, parallel programming languages and new parallel algorithms all aimed at enabling this transition. This shift to increased concurrency as a means to improve performance will impact all sectors of IT, from business servers, to desktops and laptops, to cellphones and personal electronics. HPC will lead the way, but the transformation will be ubiquitous, impacting all forms of computing.

New Programming Models

Today, programmers who develop codes for scientific applications must indicate to the computer precisely how information is divided among the processors in the system, and how the different parts of the problem are communicated between processors to enable all the processors to work together to solve the problem – a challenge roughly equivalent to writing out detailed instructions for managing all the traffic in New York City. The tool most commonly used for this part of the programming task is a language extension called the Message Passing Interface (MPI). It works well; nearly all-scientific groups use MPI to enable their programs to run on today's massively parallel systems. In the future, however, we would like to make it easier for programmers to develop parallel software, eliminating the need to explicitly manage millions of processors by developing new parallel programming languages and tools that will allow programs to be written at a higher level, making programmers more productive while increasing the scope of applications for highly parallel systems. This is an active area of computer science research and one that will impact industry broadly; indeed, Microsoft, Intel, IBM and others are now working closely with U.S. universities and national laboratories to address this challenge.

Memory

To solve larger and more complex problems, computational scientists need computers that offer increased memory as well as increased speed. In fact, memory capacity is often the limiting factor in determining whether it is possible to solve a particular problem. For example, in climate modeling we might want to increase the resolution (by using a finer mesh) of the

simulation to resolve details relevant to regional impacts – a challenge that requires more grid points, and therefore more memory.

Ideally scientists would like memory capacity to grow proportionally with the number of processors. However, that goal is not feasible in the near future, given that increasing the number of processor cores is a much simpler task than increasing memory capacity. We also need memory that is faster, to communicate with the processor at a higher throughput. Because increasing memory bandwidth consumes even more power, new ideas are needed to simultaneously improve memory performance and reduce memory power consumption. With balanced investments, it appears that we can increase memory capacity by approximately 100x as we move toward Exascale. However, it should be noted that the United States faces serious international competition in this challenge; of the top 10 global suppliers of memory, only one – Micron – is U.S.-based.

Reliability

As we build ever-larger systems, it becomes necessary to improve the reliability of every component; otherwise, the risk of diminished overall system reliability increases with the addition of each new component. While all modern large-scale computers have sophisticated mechanisms in place to identify and manage failures, these systems must be improved to ensure that our future systems will stay up long enough to do useful work. Overall reliability can be improved through new hardware designs that make fewer assumptions about individual components being failure-free, and through new ideas in software that can identify and isolate failed components, enabling the system to stay up while users' jobs are restarted on different parts of the machine.

In summary, Exascale computing represents a critical technological and economic opportunity for the United States. Right now, the HPC global market is estimated at \$10 billion, and that market is expected to grow to \$40 billion over the next decade. At present, we lead the world both in the development of HPC and in the use of HPC for advancing science and engineering, and we are working hard to achieve the next great milestone.

I believe that we can – and must – continue American leadership in HPC, to Exascale and beyond. But to fulfill the promise of Exascale, we must make sure that our efforts to develop the next-generation supercomputer are matched by increased outreach to American industry – to identify new industrial partners, to show them

how HPC can support their work, and to address any gaps in industry-specific technologies. Exascale computing will create enormous opportunities in modeling and simulation, but it may have even more impact on the large-scale data problems at the heart of many enterprises. So as we go forward, we must continue to work with our partners in industry to identify sectors where big data will enable smarter, faster decisions and outcomes, and where highly accurate modeling and simulations could lead to better results – from choosing the most effective treatment for an individual patient with breast cancer, to improving car engine combustion efficiency, to creating new energy technologies that will protect our environment and our national security.

All of us who are working in this community believe that Exascale supercomputing will be a reality by the end of this decade, and that the next-generation machines will provide tremendous benefits in terms of scientific impacts, national security and international economic competitiveness. We also understand that reaching Exascale will require many breakthroughs in science and engineering, supported by a strong public/private sector partnership.

Ultimately, however, this is a race, not against our international competitors, but against ourselves. Exascale computing is necessary to the achievement of our most urgent goals in energy, in medicine, in science and in the environment. I believe we have a duty to move as swiftly as we can, and I sincerely hope that we will seize this opportunity to maintain and extend our record of success in HPC over the next decade by making a national commitment to Exascale computing.

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