THE VITAL IMPORTANCE OF HIGH-PERFORMANCE COMPUTING TO U.S. COMPETITIVENESS

STEPHEN J. EZELL
ROBERT D. ATKINSON

APRIL 2016

ITIF
INFORMATION TECHNOLOGY
& INNOVATION FOUNDATION

The Vital Importance of High-Performance Computing to U.S. Competitiveness

BY STEPHEN J. EZELL AND ROBERT D. ATKINSON | APRIL 2016

Leadership in high-performance computing remains indispensable to a country's industrial competitiveness, national security, and potential for scientific discovery, yet heightened global competition places America's HPC leadership under increasing threat.

High-performance computing (HPC) refers to systems that, through a combination of processing capability and storage capacity, can rapidly solve difficult computational problems across a diverse range of scientific, engineering, and business fields. HPC represents a strategic, game-changing technology with tremendous economic competitiveness, science leadership, and national security implications. Because HPC stands at the forefront of scientific discovery and commercial innovation, it is positioned at the frontier of competition—for nations and their enterprises alike—making U.S. strength in producing and adopting HPC central to its competitiveness. But as competitor nations rapidly scale up their investments in and applications of high-performance computing, America will need concerted public and private collaboration and investment to maintain its leading position in both HPC production and application.

TABLE OF CONTENTS

INTRODUCTION

High-performance computing has become indispensable to the ability of enterprises, scientific researchers, and government agencies to generate new discoveries and to innovate breakthrough products and services. Indeed, high-performance computers are contributing significantly to scientific progress, industrial competitiveness, national security, and quality of life. Accordingly, many leading nations are engaged in an intensifying contest to develop the most sophisticated high-performance computing systems—and to broadly deploy them throughout their academic, industry, and government institutions—in order to advance their industrial competitiveness and scientific leadership, ensure their national security, and help address social challenges such as health, public safety, weather forecasting, climate change, and environmental protection. These nations recognize that more sophisticated and faster computers can give their countries a comparative advantage.

This report explains what high-performance computing is and why both HPC production and use matters; articulates how industry, academia, and governments leverage HPC to solve frontier challenges; details the contours of the intensifying competition for global HPC leadership; and assesses U.S. policy toward high-performance computing. The report finds that robust levels of public investment—and effective public-private partnerships to diffuse the availability and accessibility of HPC systems—has been foundational to America's leadership in high-performance computing. The report makes the following policy recommendations to ensure America's continuing HPC leadership into the future.

Congress should:

- Hold hearings on the National Strategic Computing Initiative (NSCI) and the intensifying race for global HPC leadership.
- Authorize and appropriate funding levels for the National Strategic Computing Initiative as requested in the Obama administration's FY 2017 budget for FY 2017 and for future years.
- Reform export control regulations to match the reality of current high-performance computing systems.

The administration, or its agencies and departments therein, should:

- Continue to make technology transfer and commercialization activities a priority focus of America's network of national laboratories.
- Emphasize HPC in federal worker training and retraining programs.
- Emphasize HPC in relevant Manufacturing Extension Partnership (MEP) engagements, helping facilitate small- to medium-sized enterprises' (SME) access to high-performance computing.

WHAT IS HPC, AND WHY DOES ITS PRODUCTION AND USE MATTER?

This section defines and describes high-performance computing and then assesses specific reasons why leadership in both HPC adoption and production matters to nations.

What is HPC?

High-performance computing entails the use of "supercomputers" and massively parallel processing techniques to solve complex computational problems through computer modeling, simulation, and data analysis.[1] High-performance computing brings together several technologies, including computer architecture, programs and electronics, algorithms, and application software under a single system to solve advanced problems quickly and effectively. Whereas a desktop computer or workstation generally contains a single processing chip (a central processing unit, or "CPU"), an HPC system essentially represents a network of CPUs (e.g., microprocessors), each of which contains multiple computational cores as well as its own local memory to execute a wide range of software programs.[2] The software programs that coders write to run on supercomputers are divided up into many smaller independent computational tasks, called "threads," that can be executed simultaneously on these cores. For supercomputers to operate effectively, their cores must be well-designed to communicate data efficiently, for modern supercomputers can consist of over 100,000 "cores" or more. (For example, America's *Titan*, currently the world's second-fastest supercomputer, contains just under 300,000 cores, which are capable of operating more than 6,000,000 concurrent threads.)[3]

The use of high-performance computing has become globally widespread across all branches of government and academia and virtually all fields of industry and commerce.

In essence, a supercomputer can be likened to tens of thousands of workstations performing together like a symphony orchestra to process billions and trillions of bits of data every second, sometimes for hundreds of users simultaneously. (Large programs can actually take weeks or more to complete even on the largest HPCs.) Some supercomputers are general or multipurpose machines that perform diverse tasks such as modeling and simulation or advanced business data analytics; others may be dedicated to specific tasks, such as operating cloud-based services, such as music streaming or managing telecommunications infrastructure.[4]

The use of high-performance computing has become globally widespread across all branches of government and academia and virtually all fields of industry and commerce. The impact of HPC touches almost every aspect of daily life—energy, transportation, communications, medicine, infrastructure, finance, business management, and the manufacture of both new and traditional consumer products.[5] HPCs are particularly well suited to tasks that are either computationally, numerically, or data intensive as well as tasks that require a large number of complex computations to be executed on vast data sets rapidly. This makes high-performance computers useful in all computation-intensive research areas, including physics, earth sciences, national security, biology, engineering, climate modeling, aerospace, and energy.[6] (The terms "supercomputers" and "high-performance computers" are often used interchangeably, but "supercomputers" generally refers to the most-sophisticated of high-performance computers. For example, the research firm IDC defines supercomputers as those costing $500,000 or more.)

HPCs are particularly indispensable in modeling the behavior of complex, iterative, multivariate physical systems—such as weather patterns, the dynamics of living cells and complex organs, or the movement of air and spacecraft—for they help reveal interactions and processes (many invisible to observational research) governing the behavior of the various components of these adaptive, dynamic systems. Accordingly, HPCs touch every facet of automotive and aerospace product development, oil and gas exploration, drug discovery, weather prediction and climate modeling, complex financial modeling, consumer product design and optimization, 3-D animation, and advanced business analytics, among a litany of additional applications (as the following section on HPC applications elaborates).[7] HPCs also play essential national security roles in communications, cryptography, signals processing, weapons design and testing (particularly for nuclear weapons), and war gaming. In short, advanced computing increasingly determines a nation's economic as well as defense security.[8]

Advanced, high-performance computing increasingly determines a nation's economic as well as defense security.

Analysts measure the speed of computers by their ability to calculate floating-point operations per second (or "flops"). As Figure 1 shows, growth in supercomputer operating speeds has increased exponentially over the past two decades.

Figure 1: Speed of World's Fastest and 500th-Fastest Supercomputer, by Year, 1993-2015[9]

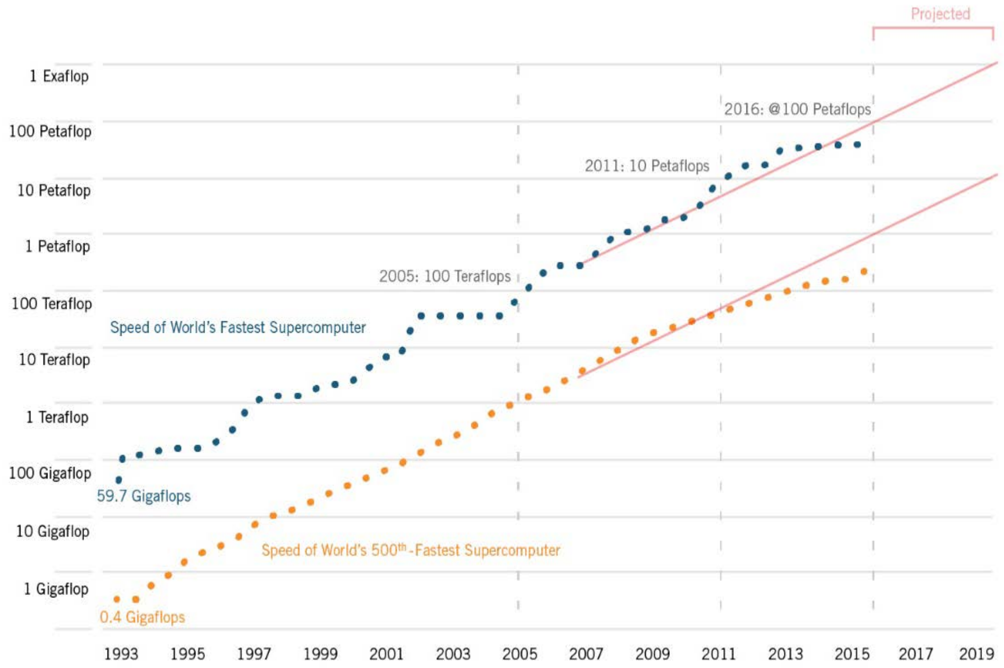

Only in the late 1990s did the fastest supercomputers break the teraflop barrier—performing a trillion arithmetic floating point operations per second (10^{12} floating point operations per second). But by the end of 2005, the fastest supercomputers had reached the petaflop range, performing one quadrillion floating operations per second (10^{15} floating point operations per second). To put this growth into context, the speed of the world's fastest supercomputer in 1993 reached only 59.7 gigaflops (one billion floating point operations per second), meaning that the speed of the world's fastest supercomputers has increased by a factor of roughly half a million over the past 23 years.[10] This represents a truly rapid transformation without comparison in any other industry.

As of November 2015, China's *Tianhe-2*, shown in Figure 2, rates as the world's fastest high-performance computer, with a peak theoretical performance speed of 54.9 petaflops, double the speed of the world's second-fastest computer, America's *Titan*, which operates at a maximum speed of 27.1 petaflops at the Oak Ridge National Laboratory in Tennessee. Surprisingly, the ranking of the world's top-five fastest supercomputers has remained unchanged from June 2013 to November 2015, but that reality is soon to change as countries significantly ramp up their investments in HPC and new machines currently in development come online. For example, later in 2016, analysts expect China to bring two 100-petaflop-capable supercomputers online.[11] In 2015, the U.S. Department of Energy (DOE) contracted with IBM and NVIDIA to launch two 150-petaflop supercomputers in the 2017-2018 timeframe: one, *Summit*, to be focused on open science and to be built at the Oak Ridge National Laboratory, and the other, *Sierra*, to be built at the Lawrence Livermore National Laboratory (LLNL). Analysts expect another DOE supercomputer, *Aurora*, due in 2019, to deliver 180 petaflops.[12]

The speed of the world's fastest supercomputers has increased by a factor of roughly a half million over the past 23 years, an extremely rapid transformation that no other industry has experienced.

Figure 2: China's *Tianhe-2* Supercomputer[13]

But the future of high-performance computing—and a race that commands the attentions of no less than a half-dozen nations—centers on which country will be the first to develop "exascale" high-performance computing: that is, a supercomputer operating at one thousand petaflops or greater. The speed of exascale computers will be measured in exaflops, or the ability to perform one quintillion (i.e., million trillion) calculations per second (that's 10^{18} floating point operations per second). China, the European Union, Japan, and the United States are vying to be the first to create an exascale supercomputer, with aspirations to do so by 2020.[14] However, while the peak operating performance speeds of supercomputers are important, it is also critical that their architectures are designed so that coders can write effective programs to run on them. Indeed, the scalability of software—that is, the ability to use a large portion of the computational capability of an HPC on a single program—is viewed as the leading barrier to getting to 10 times greater HPC systems scalability.[15] And that explains why "Adequate investments in software are one of the most important determinants [of a country's] future HPC leadership."[16] Another barrier toward realizing the next-generation of HPC systems is developing much more energy-efficient systems. China's *Tianhe-2*, for example, consumes $100,000 worth of electricity per hour. Exponentially increasing FLOP speeds while not exponentially increasing power consumption is a key challenge that next-generation HPC designers will need to overcome. But the point here is that real national leadership in HPC comes from the combination of superfast systems, designed in a functionally operational, system-efficient, and cost-efficient manner, something at which the United States has long excelled compared with peer nations.

Why HPC Adoption and Use Matters

While HPC adoption and use matters for many reasons, at least five are conceptually distinctive: 1) Each step-change in HPC represents an order of magnitude change that unlocks potential new applications, or the better use of existing ones; 2) HPC is transforming the scientific method itself with the introduction of computational simulation (or "theoretical") science; 3) HPC will be needed as a platform of innovation to handle the tremendous growth of data; 4) HPC represents one avenue to address the erosion of Moore's Law, at least for high-performance systems; and 5) Declining prices (and increasing capabilities) are democratizing HPC systems and making them accessible to a greater diversity of institutional and commercial users, including small- to-medium-sized enterprises. (The distinct economic and commercial impact of HPCs will be addressed in an ensuing section.)

Figure 3: Conceptualizing the Growth in Supercomputer Processing Speeds

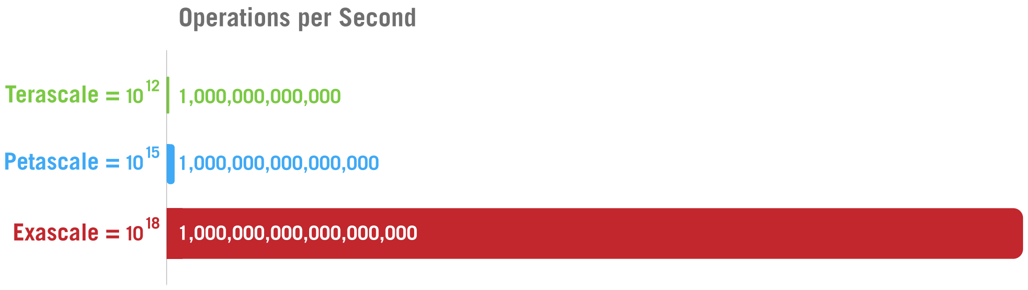

Each progression is a thousand-fold increase.

First, it is important to remember, as Figure 3 shows, that each step-change in computer processing speeds—from gigaflops, to teraflops, to petaflops, to exaflops—represents a 1,000-fold increase in peak computing speeds: that is, an increase in three "orders of magnitude" (an "order of magnitude" generally being understood as an increase in something by a factor of 10). And, as the Dutch computer scientist Edsger Dijkstra noted, "A quantitative difference is also a qualitative difference, if the quantitative difference is greater than an order of magnitude."[17] Thus, one key reason why the push to exascale matters is because for every order of magnitude increase in computing capability, one enjoys a qualitative increase in what one achieve with that computing power. The types of applications one can run on exascale platforms—such as for 3-D modeling and simulation—are fundamentally different from the types of applications one can run on petascale platforms.[18] As Bert Still of the Lawrence Livermore National Laboratory explains, "We have large problems that need to be solved, and they require models and computing that are beyond today's reach, which is what's driving our interest in exascale."[19] As following sections elaborate, exascale-level computing will be needed to solve challenges in numerous fields, ranging from life sciences to defense to energy-efficient manufacturing.

Simulations can be used to answer basic scientific questions, making simulation science an equal partner with traditional "theory and experiment" methods as an avenue of scientific discovery.

Another conceptual reason why high-performance computing not only matters—but will be disruptively transformative—is that it is helping unlock a new paradigm of scientific discovery that holds the potential to fundamentally transform the scientific method itself. Heretofore, the fundamental steps in the scientific method have been: 1) research, 2) form a hypothesis, 3) conduct an experiment, and 4) analyze the data and draw a conclusion. But HPC enables the introduction of an entirely new step through its simulation and prediction capabilities. That is, the model of "theory/experiment/analysis" in sciences or "theory/build a physical prototype/experiment/analyze" in product development is changing to one of "theory/predictive simulation/experiment/analyze."

Today, theory is being inserted into computer models in a new way—simulation science (also described as "theoretical science"). HPCs allow researchers to develop highly precise simulations on many types of phenomena that could not be analyzed before because it was not possible or was too expensive to gather sufficient data or to understand component interactions quickly enough to achieve scientific discovery. But now simulations can be used to answer even basic scientific questions, making simulation science an equal partner with traditional "theory and experiment" methods as an avenue of discovery. Moreover, many scientific experiments are now being designed up front with simulation tools in mind. For example, when scientists began to design a new tokomak thermonuclear fission reactor, the design was composed in simulation before it was built in reality, a good example of simulation science informing experimental science. Likewise, when Goodyear sets out to design new tires today, it does not start by building physical prototypes and testing them on roads; rather, it builds virtual designs and tests them in a variety of simulated conditions on high-performance computers before laying rubber down for a physical prototype that will be much closer to the end product. Many believe this advent of simulation science—with its implications for scientific discovery and commercial innovation alike—portends the most significant change to the scientific method in over half a millennium, since Galileo pioneered observational science with the telescope and gave humankind the opportunity to perceive beyond the unaided human eye.[20]

Another reason high-performance computing matters—both for a nation's science leadership and industrial competitiveness—is the continuing rapid growth of data and the need to be able to extract actionable insight from massive amounts of information rapidly. But with 2.5 quintillion bytes of data—that is, 2.5 exabytes of data—being generated daily (although granted some of this is from videos and the like), the world will need computers that can apply exaflops of computing power to exabytes of data.[21]

Indeed, scientists rely on HPC to solve problems across a wide variety of data-intensive disciplines. For example, the European Organization for Nuclear Research's (CERN's) Large Hadron Collider collides millions of particles a second, each of which produces a megabyte of data, leaving scientists with 30 petabytes (30 million gigabytes) of data to analyze even after they extract the 1 percent of data they need.[22] To help make sense of such a large amount of information, researchers at DOE's Argonne National Laboratory use an HPC system dubbed *Mira* to simulate and analyze these collisions, leveraging *Mira's* enormous 10-petaflops computing power, which performs 10^{16} calculations per second. For perspective, *Mira* can compute as much data in a day as an average personal computer could in 20 years.[23]

A key reason HPC matters—both for a nation's science leadership and industrial competitiveness—is the continuing explosion of data and the need to be able to extract actionable insight from information rapidly.

Likewise, the rapid growth in genomic data offers enormous opportunities for the application of high-performance computing to help researchers develop new insights into genetic diseases, tailor treatments to individual patients, and potentially even help cure cancer. Ongoing advances in sequencing technology have made it possible to sequence an entire human genome for as little as $1,000 and as quickly as 26 hours.[24] Considering that a comprehensive sequence of a whole human genome amounts to approximately 200 gigabytes of information, working with this data also requires massive amounts of computing power.[25] Now, researchers at the University of Chicago using an Argonne National Laboratory supercomputer named *Beagle* can analyze 240 full human genomes in just 50 hours.[26] While the Large Hadron Collider and gene-sequencing cases are just two examples coming from different scientific fields, a later section on commercial applications provides examples of companies leveraging HPC to tackle big data analytics challenges.

Another reason the ability to develop massively parallel computing systems is growing in importance is that the ability to pack more transistors onto a single processor is beginning to reach its physical limits and experience diminishing returns. For over half a century, the information technology industry has been driven by the dynamics of Moore's Law, Gordon Moore's revolutionary observation/prediction made in 1965 that the number of transistors on a chip would double every 12 to 18 months, and thus, roughly, so would computer processing speeds. Moore's Law has proven stunningly prescient, as over the past 50 years computer processing speeds have increased over one million-fold, unleashing a wave of innovation across industries ranging from aerospace to life sciences that have played a transformational role in driving the global economy and improving quality of life for citizens throughout the world.[27]

Yet, possibly as soon as 2020, the dominant silicon-based complementary metal–oxide–semiconductor (CMOS) architecture could start to hit physical limits that threaten to compromise Moore's Law unless a leap can be made to radically new semiconductor chip architectures or radically new systems architectures. (In fact, the next generation of semiconductors will be designed at the 3-nanometer scale—12 atoms across—and it will soon become physically impossible to build semiconductors at a smaller scale.)[28] This is one of the most critical technology issues the world faces today, because without significant investment in research into new semiconductors and new computer architectures, it's likely that Moore's Law will falter before many, much-needed next-generation technologies are available at commercial scale. If so, the negative consequences of a slowdown in transistor chip technology would be enormous, for new innovations in robotics, intelligent machines, data analytics, defense technology, and many other domains all require continued progress. (In fact, one reason there has been such little turnover in the world's top 10 fastest supercomputers in recent years is that Moore's Law is already starting to slow down, meaning it's getting harder and harder to get speed performance gains from individual microchips.) Accordingly, foundational innovation in semiconductor electronics and systems architecture is needed from both the public and private sectors to ensure computing power continues to advance and improve our future digital economy.

Moore's Law is already starting to slow down, meaning it's getting harder and harder to get speed performance gains from individual microchips.

However, one way to deal with the increasing challenge of making individual chips more powerful is to effectively link those chips together in massively parallel computer systems, so more chips can work together in tandem to solve complex computational problems. In other words, using existing CMOS architectures, engineers are not likely going to be able to get individual microprocessors to go much faster, but they may be able to get more chips to work together simultaneously and also to position computing functions in different parts of the system itself. This so-called "distributed computing" approach pushes processing capability out to other system components, even to the storage platforms, instead of having it all centrally focused on the CPU. In this new distributed architecture approach, instead of just driving information in and out of the CPU itself, engineers consider which tasks they are asking the machine to perform and where can they be more efficiently performed within the system. HPC systems are breaking new ground in distributed computing architectures, which may provide an avenue to sustaining the continued increases in computer processing speeds users have come to expect thanks to Moore's Law. However, while this may be true for large, massively scalable computing systems, the challenge of developing faster chips for use in robotics or cell phones, for example, will continue.

A final conceptual reason why HPC matters is because it has become increasingly accessible and cost-feasible for a growing number of commercial and institutional (e.g., academic and government) customers. A decade ago, supercomputers were the province of only the most deep-pocketed government agencies or the largest multinational corporations, as their cost often exceeded $10 million, with the global market for supercomputers confined to only a few hundred per year. But as the cost of computing power has dropped precipitously (driven largely by Moore's Law), the price of entry-level supercomputers has declined to the $500,000 to $1 million range, opening up new markets, including a range of academic, commercial, and institutional customers from around the world. This democratization of supercomputing capacity has also been furthered by the ability to access shared supercomputing resources through cloud computing-enabled platform-as-a-service type business models powered by increasingly faster broadband connections. This increasing accessibility has propelled the number of HPC systems sold annually into the thousands, expanding the global commercial market for high-performance computers into the $15 to $20 billion range.[29] Moreover, as leading-edge research pushes supercomputer speeds into the exaflop range, the expensive high-teraflop (or low-petaflop) HPC machines of today will become the cheaper, commodity machines of tomorrow, repeating a dynamic recognizable to all consumers of information technology products, from desktop computers to mobile phones.

Indeed, as supercomputing becomes cheaper, more powerful, more standardized, and more capable of embracing a blend of computational and data-driven approaches to problem solving, all sectors and tiers of the economy stand to benefit. As Bright Computing's Kristin Hansen asserts, "This convergence of HPC and big data will bring supercomputing to the masses, enabling more and more of us to participate in solving the world's biggest challenges."[30]

Why HPC Production Matters

The United States remains the leader in both developing HPC systems and in deploying them, although that lead is shrinking. Some might ask why it matters that the United States launched a National Strategic Computing Initiative in July 2015, with a goal of ensuring continued U.S. leadership in the development and application of HPC systems. Likewise, others might argue that so long as HPC users in the United States—whether enterprises, academic researchers, or government agencies—can get access to the HPC systems they need, it does not matter which enterprises in the world manufacturer those machines (so long as U.S. entities have access to them), and so policymakers should be agnostic to the issue. However, such contentions are misguided for a number of reasons.

The United States remains the leader in both developing HPC systems and in deploying them, although that lead is shrinking.

First, a key objective of the National Strategic Computing Initiative is to ensure that the United States retains leadership of the world's most high-powered supercomputers because these play vital mission-oriented roles, particularly with regard to defense and national security, as they long have. For instance, America's ability to leverage supercomputers to model the effectiveness and reliability of its nuclear stockpile was instrumental to its signing the comprehensive nuclear test-ban treaty (although the U.S. Senate has never ratified this). As *The Washington Post* writes, "U.S. weapons laboratories, armed with some of the fastest computers on the planet, are peering ever deeper into the mystery of how thermonuclear explosions occur, gaining an understanding that in some ways goes beyond what was learned from explosive tests, which ended in 1992."[31] Supercomputers allow scientists "to attempt to create a realistic model of what happens inside a nuclear explosion," with one study modeling the behavior of 9 billion individual atoms in an atomic explosion in an analysis that took over a week and used 212,000 microprocessors.[32] Furthermore, in 2011, supercomputers at the Lawrence Livermore National Laboratory revealed a weakness in America's process for storing and maintaining nuclear weapons that could have led many of them to "fail catastrophically" if ever needed for use.[33]

Put simply, supercomputers are a vital enabler of the U.S. nuclear defense posture. In fact, one could substitute nuclear weapons themselves for high-performance computers and ask if it would be troubling if the United States depended on China or the European Union for its nuclear weapons systems. And if the United States' relying on other nations to supply its nuclear arsenal sounds like an untenable proposition, then so is the notion of it relying on other nations for the most-sophisticated HPC systems. From a national security perspective, the United States needs assurance of access to the best high-performance computers in the world, simply because it gives U.S. defense planners a competitive edge and allows the U.S. defense industrial system to design leading-edge weapons systems and national defense applications faster than anyone else can.

Second, the notion that U.S. enterprises would certainly enjoy ready access to the most sophisticated high-performance computing systems for commercial purposes should they be predominantly produced by foreign vendors constitutes an uncertain assumption. If Chinese vendors, for example, dominated globally in the production of next-generation HPC systems, it's conceivable that the Chinese government could exert pressure on its enterprises to supply those systems first to their own country's aerospace, automotive, or life-sciences enterprises and industries in order to assist them in gaining competitive advantage in global markets. The notion that U.S. enterprises can rely risk-free on access to the world's leading HPC systems if they are no longer being developed in the United States amounts to a tenuous expectation that could place broad swaths of downstream HPC-consuming industries in the United States at risk if that situation ever came to fruition.

Third, and perhaps the most compelling reason why U.S. leadership in HPC matters, is that HPC systems are not developed in a vacuum: HPC vendors don't go off into a room and draw up designs and prototypes for new HPC systems by themselves, hoping someone will purchase them later. Rather, HPC vendors often have strong relationships with their customers, who co-design next-generation HPC systems in partnership with them. So-called "lighthouse [or 'lead'] users"—which, in fact, are government agencies such as the Department of Energy or Department of Defense, equally as often as leading-edge corporate users—define the types of complex problems they want to leverage HPC systems to solve, and then the architecture of the system (e.g., how the cores will be designed to handle the threads calculating the solutions) is co-created. This ecosystem exists between the HPC vendors and some of the more advanced users in both the commercial and government sector, and this symbiotic relationship pushes the frontier of HPC systems forward. So when a country has a leadership position in HPC, this enables close collaboration with the end users who buy the machines, and that creates a supply and demand for systems that are best for U.S. domestic competitiveness.

Further implicit in this relationship is the notion of *first-mover advantage.* As this report has shown, HPC is not just about "the machines," but also how HPC systems need to be optimally designed so that the software (i.e., code) running on them can best solve complex problems, something that requires close interaction between system users and designers—a relationship that benefits from geographical proximity. Moreover, both large companies such as Boeing, General Motors, and Procter & Gamble (P&G), as well as SMEs such as L&L Products and Zipp, are leveraging HPCs to solve frontier technical challenges in their industries, and using those insights to develop innovative products that give them first-mover advantage in their markets. Thus, America's HPC-consuming enterprises are best positioned when they have ready access to the leading developers of HPC systems; so both sides of the equation can stand at the leading edge in defining and solving frontier challenges using HPC. Moreover, these ecosystems go well beyond HPC vendors and commercial or government HPC-using entities; they extend to universities, private research institutions and investigators, and the talent base on both the HPC hardware and HPC software side whose competencies are likewise optimized by working on the most sophisticated problems that HPC systems can be employed to solve.

Finally, the HPC sector generates direct economic benefits for the U.S. economy. First, it appears that the industry generates a net trade surplus, particularly given that most of the microprocessors in the world's leading HPC machines—regardless of which final vendors assembled the HPC system—were of U.S. origin. HPC is also an important component of the broader computer manufacturing subsector that in the United States employs approximately 1 million individuals, 600,000 in production roles.[34] In 2015, this employment included 28,370 computer hardware engineers; 22,570 semiconductor processors; 38,010 electrical and electronic engineering technicians; and 97,200 electronic and electronic equipment assemblers.[35] While the data does not break out employment in the production of high-performance computers specifically (as opposed to broader computer and electronic product manufacturing), HPC constitutes an important source of this employment. Moreover, this employment supports substantial downstream, or indirect employment. In fact, electronic computer manufacturing generates one of the highest employment multipliers of any industry: a multiplier effect of 16 jobs, meaning 15 other jobs are dependent on one job created in the industry.[36] Further, high-performance computing supports high-wage U.S. jobs. For example, in 2015, the average computer hardware engineer earned an annual wage of $115,050, nearly 2.5 times the national mean of $48,320 earned by U.S. workers in 2015.[37]

If the United States were to cede global competitive advantage in yet another technology industry (i.e., HPC), it would mean stiffer economic headwinds for the U.S. economy and slower per-capita income growth.

Put simply, the production of high-performance computers is a robust source of employment, exports, and economic growth for the United States. If the United States cedes our leadership and global competitive advantage in this sector, it will represent yet another technology industry ceded by the United States, which will mean stiffer economic headwinds for the U.S. economy and slower per-capita income growth.

COMMERCIAL APPLICATIONS AND BENEFITS OF HPC USE

The ability to leverage high-performance computing has become indispensable not only to virtually all advanced manufacturing industries, but also to a wide range of commercial sectors. Indeed, the use of high-performance computers has led to breakthroughs in a wide range of commercial applications, including discrete parts manufacturing, pharmacology, chemical engineering, electronics design, content management and delivery, and the optimal development of energy sources, among many others.[38]

HPC provides manifold benefits and value when used for commercial and industrial applications. In particular, HPC enables advanced modeling, simulation, and data analytics that can help address manufacturing challenges and aid in decision-making, optimize processes and design, improve quality, predict performance and failure, and accelerate or even eliminate prototyping and testing.[39] In particular, HPC enables the breakthrough discoveries that fuel innovation. HPC provides a cost-effective tool for accelerating the research and development (R&D) process, particularly by helping to design new products, to improve existing products, and to bring products to market much more efficiently and quickly. Further, as noted, future products, business models, industrial processes, and companies are being built on the ability to collect, analyze, and leverage data, making supercomputers a necessity in analyzing massive amounts of data in industries such as biotechnology, finance, manufacturing, and oil and gas exploration. [40]

The economic benefits of investments in high-performance computing are significant. As of August 2015, the research firm IDC found that enterprises' investments in high-performance computing systems generate $515 in revenue and $43 in profits and/or cost savings per dollar of HPC investment.[41] IDC's August 2015 estimate of revenues generated per dollar of investment in HPC represented a 44 percent increase over its 2013 estimate that each dollar invested in HPC generated $357 in revenues. IDC further estimates that enterprises' average HPC investment per innovation project is $3 million. The average number of years before enterprises realize a return on their HPC investments is approximately three years. A study of HPC in the European Union (EU) found even higher returns, concluding that (for projects that generated financial returns), each euro invested in HPC on average returned €867 in increased revenue and €69 in profits.[42] Total increased revenue for the 59 HPC-enabled, quantifiable projects in Europe reached €133.1 billion, or about €230 million per project on average.[43]

The economic benefits of investments in HPC are significant; one study found that each $1 invested in HPC generates $515 in revenue and $43 of profits and/or cost savings.

So pervasive has the use of HPC become that another IDC study found that 97 percent of companies that had adopted HPC said they could no longer compete or survive without it.[44] That finding aligns with those from an October 2014 Council on Competitiveness study, *Solve*, which reported that 72 percent of enterprises believe that HPC is a cost-effective tool for R&D and that 76 percent believed that "increasing performance of computational models is a matter of competitive survival."[45] That study further reported that 86 percent of respondents agreed with the statement that "HPC is critical to the future direction of our business."[46] The study also reported that responding companies believed they could consume up to a 1,000-fold increase in computing capability and capacity in a relatively short period of time.[47]

The following section examines applications of high-performance computing in the manufacturing, life sciences, energy, and other industries.

Manufacturing and Industrial Applications of HPC

Ever-fiercer global competition continuously amplifies the need for manufacturing innovation, explaining why even long-standing manufacturing companies—such as Kennametal, Parker Hannifin, Timken, and United Technologies—strive to have 20 percent or more of their products be new or at least substantially revamped each year.[48] High-performance computing has become a critical enabler of innovation, new product design and development, and product testing and validation in virtually all manufacturing companies, meaning HPC is helping manufacturers both cut costs and create new revenues. The following section examines applications of HPC in manufacturing industries, including aerospace, automotive, and consumer packaged goods. It further examines the unique challenges of bringing HPC to a country's SME manufacturing base.

Aerospace

HPC has helped aircraft manufacturers, such as Boeing, significantly reduce the design-to-production timeline for new aircraft, which has already saved the aerospace industry multiple tens of billions of dollars. Large passenger jets contain well over 2 million individual parts that need to be simulated both individually and as part of a larger system.[49] Those millions of parts (and their interactions) must withstand varied pressures and strains over the 30 or more years that mark a typical jet's lifetime, while at the same time maintaining operational reliability in excess of 99 percent.[50]

HPC simulation allows aircraft developers to improve the design and to simulate the functional operation of many critical aircraft components—such as wing and fuselage design—before a physical prototype is ever tested in a wind tunnel. In other words, aircraft designers can use computational modeling and simulation to explore various design options before building physical prototypes and conducting experimental testing. In fact, computational modeling of different aircraft components in combination with supercomputers has enabled a 50 percent reduction in wind tunnel testing needed for a new aircraft's development.[51] Supercomputing is now used to model a significant portion of a new commercial jet aircraft, from cabin design, cabin noise, interior air quality, high-speed wing design, wing tip design, exhaust system design, vertical tail and aft body design, and much more.[52]

As a specific example, Boeing physically tested 77 prototype wing designs for its 767 aircraft (which was designed in the 1980s), but for its new 787 Dreamliner, only 11 wing designs were physically tested (a 7-fold reduction in the needed amount of prototyping), primarily because over 800,000 hours of supercomputer simulations had drastically reduced the need for physical prototyping.[53] In this case, Boeing used supercomputers located at the Oak Ridge National Laboratory to assess aeroelasticity—that is, the effect of aerodynamic loads on airplane structures. As Doug Ball, Boeing's Enterprise Director for Computational Fluid Dynamics, notes, HPC "lets engineers design better airplanes with fewer resources, in less time, with far less physical simulation based on wind tunnel testing."[54] In essence, supercomputers enable Boeing to bring safer, more efficient aircraft to market sooner and cheaper, which both lowers Boeing's costs, speeds its innovation cycles, and increases its competitiveness.

But it's not just U.S. firms using HPC; their global competitors are as well. Boeing rival Airbus operates three supercomputers rated in the global Top 500, with a combined 80,000 cores and approximate combined total peak performance around 1.5 petaflops.[55] Airbus's supercomputing power gives it the ability to tackle "computational fluid dynamics multiphysics problems at scale," for it likewise recognizes that HPC enables "the virtualization of aircraft development" with "numerical simulation [being] one of the most important means to realize this objective."[56] Like Boeing, Airbus virtually models new aircraft features early in the design phase, assessing both physical flight behavior as well as the behavior of different aircrafts structures (such as ailerons or flaps) and systems. Airbus design teams then apply different improvement plans and "what if" studies to this model, which allows for quick analysis of the consequences of any modification or optimization suggested. To get a sense of just how much calculating and data creation occurs, the definition of the overall aircraft model ultimately translates into between 500 million and 1 billion nonlinear equations that must be resolved during each reiteration in a real-time manner.[57] Given this, Airbus notes that one of the key challenges it faces lies at the software level, particularly in software scalability, so that it can take full advantage of the computing architectures it possesses.[58]

Supercomputers have transformed not only the design of aircraft, but also the engines that power them. As one designer of jet engines commented, "We're removing design cycles from jet engine component technology, [and] doing full modeling of individual components of an engine: compressors, combustors, turbines, rotating elements, etc."[59] In one example, General Electric (GE) Global Research, in collaboration with Cornell University, used supercomputers at the Lawrence Livermore and Oak Ridge national laboratories to improve jet engine efficiency through simulation.[60] In this case, GE used a supercomputer to reveal a new aspect of turbine behavior that is already providing the company with a competitive advantage in fuel efficiency. GE estimates that every 1 percent reduction in fuel consumption saves users of these products approximately $2 billion per year.[61] Supercomputers have also helped GE to reduce new jet engine development timelines by at least half a year.

So pervasive has the use of HPC become that one IDC study found that 97 percent of companies that had adopted HPC said they could no longer compete or survive without it.

Again, U.S. competitors are using HPC as well. Rolls-Royce, for example, uses HPC systems to model and run engine design tests. At Rolls-Royce, single components and subsystem design, analysis, and optimization through to whole engine modeling all rely heavily on HPC.[62] The company's requirements for high-fidelity modeling of complex geometries with multiphysics and multidisciplinary approaches demands extreme computational power. Within a design cycle, this modeling and analysis has to be accomplished within an acceptable and challenging time scale and accuracy, and the platforms on which these simulations are performed have to be robust, stable, scalable, and reliable in their availability and usability. Rolls-Royce notes that HPC is the only computational resource that can meet such high demands.[63] Indeed, HPC has become central to engineering activities across all of Rolls-Royce's global business sectors, from aerospace to marine and nuclear.[64] To ensure it has sufficient access to HPC resources, Rolls-Royce was the first company to sign on to a new supercomputing brokerage scheme launched by the United Kingdom's Engineering and Physical Sciences Research Council's HPC Midlands Center, giving it ready access to £60 million worth of on-demand high-performance computing resources.[65]

Automotive

Just as with aircraft, HPC has transformed how vehicles and their components are designed, modeled and simulated, safety tested, and ultimately manufactured, playing a key role in reducing vehicle design costs, introducing innovative new features, and improving the fuel efficiency and safety of vehicles. As Nand Kochhar, a chief engineer at Ford, explains, "The combination of HPC and computer-aided engineering (CAE) simulation technology is a key enabler of our product development process. We provide advanced computational capabilities for Ford not just as a service, but as an integrated enabler of company business strategy."[66]

While General Motors has used high-performance CAE tools such as design, modeling, and simulation software since 1998, GM Engineering General Manager of Global CAE Strategy Martin Isaac notes that, over the past decade, GM's "use of simulation has grown exponentially."[67] Simulation has greatly reduced both the number of prototypes that GM needs to create for each new vehicle design as well as the number of physical models required for testing. The tools allow GM's engineers to simulate crash tests from every angle, testing restraint and airbag performance and even running digitalized pedestrian impact scenarios to improve pedestrian safety. The tools enable engineers to run aerodynamic and airflow models simulating air-conditioning, heating, and electrical systems, and the interactions among all of them. So effective is the simulation approach that many of the tests with physical prototypes "are simply final checks after long, iterative tests in GM's computer farm."[68] As Isaac notes, "The technology has enabled us not only to reduce development time but to get a much better engineering solution."[69] All told, GM has invested close to two decades of work in accelerating the company's product development cycle to make it more like software development, which has significantly accelerated GM's time to market for new vehicles while boosting the productivity of its vast engineering workforce.[70]

HPC has helped aerospace and automotive manufacturers accelerate design-to-production timelines by as much as a factor of three, significantly reducing costs while facilitating design of superior products.

The impacts of high-performance computing have left their mark on Europe's automotive industry as well. In fact, one IDC study estimated that HPC has helped European automakers reduce the time needed for developing new vehicle platforms from an average of 60 months to 24 months, while greatly improving crashworthiness, environmental friendliness, and passenger comfort.[71]

HPC has impacted truck design as well. BMI Transportation used supercomputers to design more aerodynamically efficient long-haul trucks that can achieve fuel savings of between 7 and 12 percent, a design which, if applied to all 1.3 million Class 8 big-rig trucks in the United States, could save 1.5 billion gallons of diesel fuel and $5 billion in fuel costs in the U.S. trucking industry annually.[72] Moreover, BMI's use of HPCs significantly accelerated the company's design process and time to market. Running design simulation models on supercomputers allowed BMI to move from concept to a design that could be turned over to a manufacturer in 18 months instead of the usual 42 months.[73]

As noted earlier, Goodyear has leveraged HPC-enabled predictive modeling and simulation techniques to test virtual tire models and significantly reduce time to market, in the process transforming how tires are designed, developed, and tested.[74] As Joseph Gingo, then Goodyear's senior vice president of technology and global products planning, explained the development of the breakthrough all-weather *Assurance*® tire in the early 2000s: Back then the company only "used HPC modeling to augment the conventional tire development procedure: building physical prototypes and then subjecting them to extensive environmental testing."[75] But knowing that the accelerated product development schedule anticipated for the *Assurance* tire did not provide enough time for "the conventional design, build, and test cycle" (especially when real-world tests like tread-wear can take four to six months to get useful results), Gingo decided, "Let's flip our procedure—let's use modeling and simulation from the very beginning."[76]

However, Goodyear lacked adequate (hardware and software) HPC resources to meet the challenge at the time, so it partnered with Sandia National Laboratories in New Mexico, with Goodyear engineers working with Sandia's supercomputer experts to develop complex, state-of-the-art software to run Goodyear's HPC clusters.[77] The jointly developed software allowed Goodyear to run more advanced simulations and to maximize the performance of its computers, helping Goodyear to reduce its R&D costs and cut time to market for its *Assurance* tire. In total, HPC modeling and simulation enabled Goodyear to reduce its product design time from three years to less than one year and to decrease tire building and testing costs from 40 percent of the company's R&D budget to 15 percent.[78] As Gingo concludes, "Computational analysis tools have completely changed the way we develop tires. They have created a distinct competitive advantage for Goodyear, as we can deliver more innovative new tires to market in a shorter time frame."[79] Or, as Loren Miller, Goodyear's IT director, frames it, "high-performance computing is more than just a very sophisticated tool … it's a strategic asset that … makes Goodyear a formidable competitive force in today's global market."[80]

But high-performance computing is not just for large companies in the automotive industry; HPC-enabled modeling and simulation software tools are vital for the competitiveness of automotive-industry SME manufacturers as well. For instance, L&L Products, a Romeo, Michigan-based automotive supplier that makes high-strength adhesives and structural composites for strengthening vehicles, was able to develop a new structural composite line for automakers that doubled the size of its business.[81] L&L leveraged HPC by first creating a virtual model of parts to be added to a vehicle design and then running simulations to understand how the products would behave in vehicle crashes. As Steven Reagan, L&L's computational modeling manager noted, "There is no way to compete in this market without that [HPC] tool."[82]

Likewise, Zipp Speed Weaponry, a small, Indiana-based specialty manufacturer of performance-biking gear such as wheels and tires (and which is the only remaining U.S. manufacturer of advanced high-performance cycling components) leveraged HPC to conceive of innovative racing tires for bicycles. Specifically, Zipp used HPC-enabled virtual simulations to better understand computational fluid dynamics (CFD) and to resolve turbulence challenges it was unable to solve with traditional wind tunnel experiments, allowing Zipp to jump ahead of the global competition in its unique market niche.[83] The results helped Zipp introduce its aerodynamically revolutionary Firecast wheels in 2010, enabling it to double its product category revenues in just two years. Moreover, given the increased demand, Zipp was able to add120 new manufacturing jobs. As Zipp Technical Director Josh Poertner explained, "We continue to spend the same on prototypes and development, but are finding that our prototypes are significantly more successful as we are able to cull hundreds of ideas into dozens, whereas before we would have to guess at the dozens of prototypes to produce. We are now optimizing for twice the variables, but in the same amount of time."[84] In this case, Zipp collaborated with Intelligent Light, a consulting firm specializing in computational fluid dynamics analysis, thus leveraging Intelligent Light's HPC resources.

Steel and Welding

HPC has also transformed manufacturing processes in key sectors supplying the automotive industry, such as steel and welding. For example, HPC has facilitated development of a cloud-based tool that simulates welding processes used in metallic product manufacturing. The application, being developed by the Ohio Supercomputer Center (OSC) and the Engineering Mechanics Corporation of Columbus, in part through a Small Business Innovation Research (SBIR) grant awarded by DOE, is a welding design software package called Virtual Fabrication Technology that enables SMEs to tap into HPC resources, so they can validate the integrity of welds in assembled components.[85]

Likewise, researchers use simulations to decrease the materials scrapped during the continuous casting process. Decreasing the material scrapped due to defects such as cracks, even by a small percentage, results in a large net savings to steel manufacturers and customers. Based on the roughly 100 million tons of steel produced each year in the United States and the approximately $400 per ton net cost of scrapping, the 1 percent reduction in yield loss enabled by HPC is helping save the U.S. forging industry about $400 million per year.[86]

America's network of national laboratories constitutes a national treasure that has proven indispensable to facilitating American industries' ability to access, understand, and leverage HPC systems to bolster their competitiveness.

Consumer Packaged Goods Manufacturing

Pioneered largely by Proctor and Gamble, high-performance computing has also revolutionized the consumer packaged goods (CPG) industry. As the representative of one CPG manufacturer elucidated, with HPC:

> I can figure out whether a bottle will break when it drops. I can figure out how the handle will fit small hands and big hands. I can figure out whether a diaper will leak. I can figure out whether the closure on a diaper will mark a baby's leg because the elastics are too tight. Whether a formula will remove a stain and still protect a new fabric. How many washes will it take for jeans to fade? Can we smell a new perfume on laundry after it's been washed? ... All of those things we now do with high-performance computing.[87]

P&G leverages supercomputing to understand formulations down to the molecular level in a wide range of products such as shampoos, soaps, and diapers, thereby improving product performance, in part because HPC helps P&G to identify molecular characteristics that are not observable experimentally. In total, P&G's use of simulation and modeling has allowed it to reduce the number of steps involved in process design by over 50 percent.[88] According to a representative of one CPG company, HPC-powered "modeling and simulation has accounted for hundreds of millions of dollars of value over the last decade, and I can point to several products in the marketplace that would not have been there had it not been for modeling."[89]

Like Goodyear, P&G has also tapped into HPC hardware and programming resources at U.S. national laboratories to turbocharge its innovation efforts. In this case, P&G partnered with the Los Alamos and Sandia National Laboratories to "tap the labs' supercomputers and brain trusts to create new eco-friendly materials for consumer products." As P&G's Thomas Lange, the company's legendary director of modeling and simulation, noted, the national laboratories "are the only places I can go in the world that have such a range of world-class physicists, chemists, biologists, production engineers, and computational scientists. These labs are national treasures."[90] P&G's collaboration with Los Alamos in computer simulation has saved the company over $1 billion while helping it to develop more environmentally friendly products.[91]

Energy Consumption and Production

HPC is transforming how energy is both consumed and produced, and it will play an important role in advancing innovative clean energy technologies, improving energy efficiency, and reducing energy and resource consumption. For example, supercomputer models are playing a key role in the design of more energy-efficient buildings. Buildings consume approximately 40 percent of the energy used in the United States and are responsible for nearly 40 percent of greenhouse gas emissions.[92] A 50-percent reduction in buildings' energy usage would be equivalent to taking every passenger vehicle and small truck in the United States off the road. The Greater Philadelphia Innovation Cluster serves as a national center for energy-efficient building research, education, policy, and technology commercialization. It is leveraging HPC to develop integrated end-to-end code for simulating building fluid/thermal flows.[93]

The oil and gas industry likewise makes extensive use of HPC, particularly in 3-D seismic modeling to identify oil and gas deposits. For example, the French company Total recently tripled the power of one of its supercomputers to develop more complete visualizations of seismic landscapes and run simulations at 10 times the resolution of existing oil and gas reservoir models. This new capability will enable more efficient upstream oil and gas exploration, as well as the discovery of reserves under more challenging geological conditions.[94]

SMEs in the energy sector also leverage HPC for innovation and cost savings. For example, the Seattle-based start-up Ramgen has used HPC-enabled computer simulations to design highly efficient gas compressors that can potentially reduce the capital costs of CO_2 compression by 50 percent and produce a minimum 25 percent savings in operating costs. Ramgen's compressors, based on shock-wave technology used in supersonic flight applications, will have significant impact on the broader turbomachinery industry. When its compressors are used in traditional 400 megawatt (MW) clean coal plants, Ramgen anticipates capital cost savings of approximately $22 million and an annual operating cost savings of approximately $5 million.[95]

Bringing HPC to America's SME Manufacturers

While large manufacturers have made great progress in leveraging high-performance computing systems for innovation, the penetration of HPCs into America's SME manufacturing base has come much more slowly and sporadically. SME manufacturers (those with fewer than 500 employees) account for 248,155 of the 251,857—or 98.5 percent—of U.S. manufacturing companies as of year-end 2013.[96] Moreover, 94 percent of all U.S. manufacturers employ 100 or fewer workers.[97] This vast number of SMEs constitute the so-called "missing middle" of HPC adoption in U.S. industry. Here, the term "missing middle" refers not directly to company size but rather to a company's computing capacity; the term specifically refers to the group of HPC users between low-end, mostly workstation-bound HPC users, and the kind of high-end HPC uses typically performed at national labs and some universities.[98] Nevertheless, in industry parlance, the term has come to refer to the wide swath of small- and mid-sized manufacturers who could be leveraging HPC in their product development or manufacturing processes, but are not.

For example, one 2013 study estimated that only 8 percent of U.S. manufacturers with fewer than 100 employees are using HPC.[99] Earl Joseph, an HPC analyst at IDC, estimates that at least 25,000 U.S. manufacturers, the vast majority of them SMEs, would benefit from having access to HPC-empowered modeling and simulation systems.[100] Others peg that figure even higher, estimating that as many as half of all U.S. manufacturing SMEs could leverage HPC-enabled modeling and simulation tools in design, prototyping, and testing of their parts, components, and finished products.[101] In fact, some estimate that if America can increase HPC application and use among this missing middle of small- and medium-sized enterprises, "this represents a [potential] market nearly as large as the entire global HPC segment today."[102]

However, at least three major barriers have prevented America's SME manufacturers from leveraging HPC solutions in greater numbers. First, there exists a general lack of knowledge about how to apply HPC tools to solve engineering challenges, an especially acute problem because many of the engineers working at America's SMEs simply were not exposed to computational sciences in their electrical or mechanical engineering training. Second, taking those engineers "off the line" to train them in modern modeling and simulation tools takes them away from the urgent needs of the business and represents an expense many SMEs cannot incur. Third, and more subtly, many existing modeling and simulation packages (e.g., designed to model aircraft and engines) are often too complex or overdesigned for the needs of smaller manufacturers.

Several initiatives have been launched to help remedy the lack of availability, accessibility, or approachability to HPC tools for SME manufacturers. For example, the National Center for Manufacturing Sciences (NCMS) has created a dozen centers throughout the United States (located near universities and national labs to tap into local expertise) to connect manufacturing firms with HPC resources. NCMS's network of "Predictive Innovation Centers" represents public-private collaborations (the public component thanks mostly to state-level matching investments) providing U.S. manufacturers with high-performance computing tools aimed at increasing product design cycles, improving manufacturing processes, and reducing the need and costs of laboratory testing of new products.[103]

Likewise, the Ohio Supercomputer Center's (OSC's) AweSim program is a partnership among OSC, simulation and engineering experts, and industry to assist SME manufacturers with simulation-driven design to enhance innovation and strengthen economic competitiveness. As AweSim Director Alan Chalker explains, "Simulation-driven design replaces physical product prototyping with less expensive computer simulations, reducing the time to take products to market, while improving quality and cutting costs. Smaller manufacturers largely are missing out on this advantage."[104] AweSim levels the playing field by giving smaller companies equal access to HPC technologies. AweSim invites SMEs to bring in their technical challenges and then work with experts to understand how HPC-enabled modeling and simulation tools can help solve their problems. Launched in December 2013, in its first year AweSim served over 100 SMEs, including one that reported that the use of virtual prototyping reduced the number of physical prototypes needed to develop a new product (at a cost of $25,000 each) from 100 to 1, saving the company over $2 million.[105] The National Center for Supercomputing Applications, a hub of transdisciplinary research and digital scholarship led by Merle Giles at the University of Illinois at Urbana-Champaign, has also played a pivotal role in helping U.S. enterprises, large and small alike, understand how they can leverage HPC tools to bolster their competitiveness.[106]

Likewise, the Chicago-based Digital Manufacturing and Design Innovation Institute (DMDII), one of the institutes within the National Network for Manufacturing Innovation (NNMI), is developing a new cloud-based system to democratize SME manufacturers' access to HPC resources. Expected to launch in April 2017, DMDII envisions its Digital Manufacturing Commons (DMC) as a free, open-source software project to develop a collaboration and engineering platform that will serve as an online gateway for digital manufacturing.[107] Akin to an "app store for manufacturing," the DMC will be a digital services marketplace with a software development kit and collaboration platform at its core, essentially equipping SME manufacturers with the modeling and simulation tools they need to address technical design challenges as well as access to shared HPC resources.[108]

Solving the democratization of HPC for a nation's industrial base is one of the most significant challenges for countries seeking HPC leadership.

Making HPC accessible to all manufacturers in a country can be a tremendous differentiator, and no nation has cracked the puzzle yet. Ensuring that many more companies, including SMEs, can effectively use HPC is a critical challenge; the country that solves it first will gain a considerable competitive advantage. Furthermore, one reason that so-called Tier 1 OEMs (original equipment manufacturers, the firms such as Boeing or GM at the top of their industrial value chains) care that SMEs have access to and facility with high-performance computing is so they can be certain their key supplier base can interface with their product development systems and also produce the most innovative and cost-effective parts and components of their own.

What's Next for HPC in Manufacturing?

Despite HPC's already tremendous impact on manufacturing, as computer processing speeds accelerate and the technology matures, new opportunities for industrial applications will open up. As Vice President, Manufacturing Technology Tim Shinbara of the Association for Manufacturing Technology (AMT) explains, "Only in the last five years have manufacturers really started to understand how to move HPC applications from the modeling and simulation realm down to assisting in real-time production on the shop floor."[109] For example, in additive manufacturing, successive layers of material are built up to synthesize a three-dimensional solid object composed in a digital file, with each layer a thinly sliced horizontal cross-section of the eventual object.[110] Heretofore with additive manufacturing, if a problem is discovered as layers are being printed (e.g., there are undesired thermal effects), computers were not sophisticated or fast enough to detect (and even potentially solve) the problem in real-time and adjust so that the next layer printed does not have the defect. In other words, a key choke point has been that computers have historically not been fast enough to collect, assimilate, and assess information in real-time so intelligence can be instantaneously injected back into manufacturing processes, whether the machine is a mill, a lathe, or something more sophisticated such as a 3-D printer. In other words, HPC has traditionally been used from a pre-manufacturing simulation or a post-manufacturing quality verification standpoint; the frontier is leveraging HPC *in situ* by embedding it into manufacturing control systems. Moreover, the ability to identify and resolve potential defects in real-time in manufacturing processes should tremendously reduce defects, allowing quality to be "built into" the system.[111] As Shinbara notes, "HPC has been great for solving equations in mass volume; the next step is to inject intelligence so that HPC systems generate context, nuance, and genuine intelligence we can use in real-time in the manufacturing process domain."[112]

HPC has traditionally been used from a pre-manufacturing simulation or a post-manufacturing quality verification standpoint; the frontier is leveraging HPC in situ by embedding it into manufacturing control systems.

Another frontier use of HPC in manufacturing processes pertains to the testing and certification of exotic materials, such as metals, polymers, and hybrid composites (particularly relevant in defense domains) to understand how their physical properties perform and react under extreme conditions. HPC allows 1,000-fold iterations of virtually testing these materials under extreme conditions, significantly accelerating the speed at which the safety and sustainability of new materials can be validated, with significant implications not just for defense but also for much faster time-to-market introduction of new materials technologies.

Finally, in February 2016, as part of its HPC4Mfg (HPC for Manufacturing) challenge, the Department of Energy announced $3 million in funding for 10 projects that will allow manufacturers to tap into the power of HPC systems at DOE-managed national laboratories.[113] Each of the projects is designed to leverage HPC to improve efficiency, enhance product development, or reduce energy consumption. For example, one initiative will help Global Foundries optimize semiconductor transistor design, and in another GE will leverage advanced HPC particle physics simulations to improve the efficiency and lifespan of its aircraft engines.[114] The vision is to grow this concept from just HPC4Mfg into an HPC4X template where the same process can be applied to HPC4transportation, HPC4life sciences, etc.

As noted, previous examples have explained how Boeing, Goodyear, and P&G, among many other enterprises (large and small alike), have been able to partner with various U.S. national laboratories to bring their technical challenges to the table and collaborate with the national laboratories, leveraging both their extensive high-performance computing resources (often beyond the reach of even the largest companies) and technical experts across a range of scientific fields—from computational fluid dynamics, to chemistry, to biology, and beyond—to jointly solve engineering challenges and bring new or improved products to market. Far from such collaborations representing so-called "corporate welfare," they effectively leverage public-private resources to collaboratively solve frontier engineering challenges, thus enabling U.S. enterprises to create innovative new products (or improve existing ones), which bolsters those enterprises' competitive position in fiercely contested global markets. That produces a range of spillovers that benefit the American public, from new (or retained) jobs to new revenues, taxes from which fill public coffers.

Elsewhere on the HPC frontier, the Network for Computational Nanotechnologies (NCN) has launched a virtual laboratory that develops modeling and simulation tools to better predict behavior at the device, circuit, and system level for nanoelectronics, nanomechanics, and nanobio systems. Serving over 180,000 users, the NCN mounts over 10,000 simulations each year, providing users access to supercomputers as needed.[115] Put simply, the application of HPC to modern manufacturing challenges is only just beginning.

Health-Care-Related Applications of HPC

Biotechnology researchers and companies alike are leveraging the power of HPC to understand fundamental biological processes, to develop new drug therapies, and to improve the delivery of health care through personalized medicines. As Earl Joseph of IDC explains, "Biology is fast becoming a digital science, and HPC is increasingly important for advanced medical research, biomedicine, bioinformatics, epidemiology, and personalized medicine—including 'Big Data' aspects."[116]

Biotechnology researchers and companies alike are intensively leveraging the power of HPC to understand fundamental biological processes, to develop new drug therapies, and to improve the delivery of health care through personalized medicines.

However, the application of high-performance computing to biotechnology challenges is not new. Indeed, supercomputing proved key to the initial Human Genome Project, and played a pivotal role in completing the mapping of the human genome three years ahead of the originally projected timeline.[117] However, what has changed is the advent of the field of computational biophysics, supported by HPC-powered next-generation sequencing (NGS) techniques, that has reduced the time and cost of sequencing a complete human genome—some three billion nucleotides—to a few days for just $1,000 (an effort that required over $1 billion and many person-years of effort just 15 years ago).[118] Moreover, HPC has proven instrumental in helping researchers improve their understanding of and treatments for a wide range of diseases and conditions.

The Center for Pediatric Genomic Medicine at Children's Mercy Hospital in Kansas City, Missouri, for example, has been using HPC to help save the lives of critically ill children. In 2010, the center's work was identified as one of *Time*'s top 10 medical breakthroughs of the year. Roughly 4,100 genetic diseases affect humans, and they are the main causes of infant deaths. But identifying which genetic disease is affecting a critically ill child is extremely difficult. For one infant suffering from liver failure, the center used 25 hours of supercomputer time to analyze 120 billion nucleotide sequences and narrowed the cause of the illness down to two possible genetic variants. Thanks to this highly accurate diagnosis, the baby is alive and well today. For 48 percent of the cases the center works on, HPC-powered genetic diagnosis points the way toward a more effective treatment.[119]

Heart disease remains the leading cause of death in the United States, with 611,000 fatalities from heart disease in 2014 making it the cause of one in four deaths in America.[120] Scientists from IBM Research, in collaboration with the Department of Energy's Lawrence Livermore National Laboratory, have now been able to model the human heart in much greater detail than ever before, using one of the world's most powerful supercomputers in doing so.[121] For the first time, researchers developed a computer model that comprehensively simulates the human heart down to the cellular level, as depicted in Figure 4. To use the laboratory's powerful supercomputer most efficiently, the researchers created new ways to divide heart tissue into small pieces of equal work units that could be evenly distributed across the supercomputer's 1,572,864 elements.[122] This innovation has significant potential for saving health-care costs by reducing heart disease and improving heart health and could lead to breakthroughs in identifying, preventing, or remedying coronary heart disease, which costs the United States over $100 billion annually.[123]

Figure 4: Computer Visualization of the CARDOID Heart Model[124]

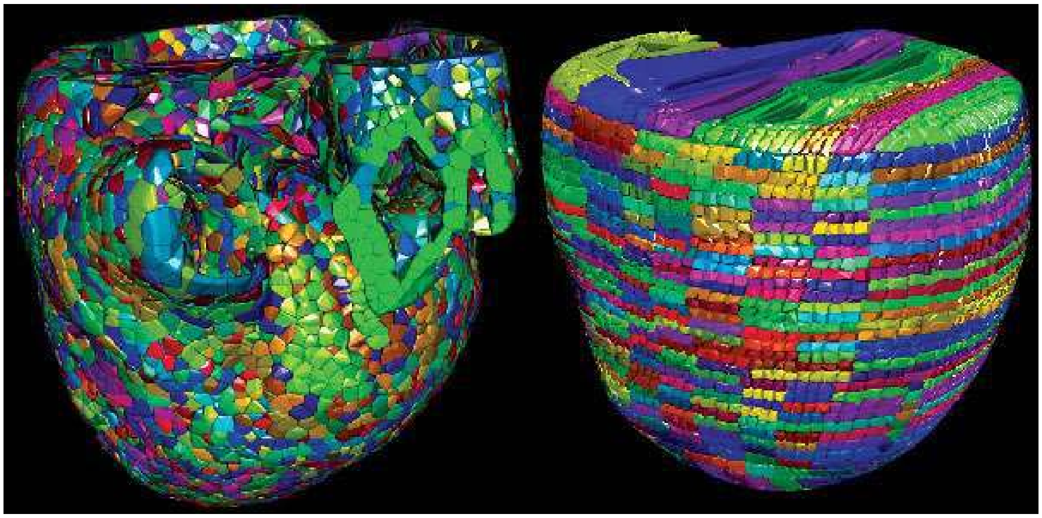

Supercomputers are being used not only to unravel the morphology of cancer cells, but also to diagnose and treat cancers and improve the safety of cancer treatments.

Just as supercomputers are reshaping how researchers understand the heart, so with the head. For example, researchers at San Diego's Salk Institute are using supercomputers at the nearby National Science Foundation (NSF)-funded San Diego Supercomputer Center to investigate how the synapses of the brain work.[125] The researchers have made considerable progress in modeling chemical transmission across nerve synapses, which involves an array of complex electrochemical processes. Supercomputers have enabled a 150-fold increase in simulation speed as well as an increase in simulation complexity. The research has the potential to help people suffering from mental disorders such as Alzheimer's, schizophrenia, and manic-depression.[126]

Supercomputers are being used not only to unravel the morphology of cancer cells, but also to diagnose and treat cancers and improve the safety of cancer treatments. For example, researchers at the Harvard Medical School have developed practical strategies for reducing radiation doses associated with computerized tomography (CT) and Positron Emission Tomography (PET) scans.[127] Currently, a patient receiving a single CT scan can receive as much as 3 millisieverts of radiation, an amount equal to the average person's annual exposure. Traditional scans work by injecting patients with radiated glucose and having the scan identify tissues consuming the most energy (potentially signaling the presence of cancer). Supercomputers help reduce radiation exposure by permitting the use of a less radioactive form of glucose, which produces a grainier image on the scan; HPC-powered software can extract the low signal from the high noise, thus enhancing resolution and detecting the cancer.[128] Likewise, the Mary Bird Perkins Cancer Center in Baton Rouge, Louisiana, has made important advances that can lower the incidence of secondary cancers caused by radiation in children receiving radiation therapy. The researchers, who saved more than $12 million by using high-performance computing, estimate the achievement has accelerated radiation toxicity research by more than a decade.[129]

Supercomputers have also contributed to breakthroughs in treating hepatitis C and AIDS. Researchers from the Centers for Disease Control (CDC) and the Cornell University Center for Advanced Computing created a detailed model of the hepatitis C virus, a major cause of liver disease.[130] Through faster computations (more than 175 times sped up), a better understanding of networks of coordinated amino-acid variation opened the door for the discovery of new therapeutic targets for the hepatitis C virus (HCV). Over 500,000 jobs ran on Cornell's computers over two years, generating new scientific insights and publications in condensed matter physics, gravitational wave detection, biomedical imaging, orthopedics, neuroscience, and optics. The team's detailed modeling of HCV was rated among the top 50 innovations in the last decade.[131] Similarly, supercomputers recently helped unlock the structure of the human immunodeficiency virus (HIV), thus making significant progress toward more effective treatments.[132]

While there are many specific examples, broadly, as Klaus Schulten, a professor (and pioneer in the field) of computational biophysics at the University of Illinois at Urbana-Champaign, explains, supercomputers are "an extremely useful instrument for pharmacological research because the targets of most of our medical treatments are specific molecules in the body. That's what the computer can help with—describing the target molecules and suggesting pharmacological treatments against what the computer sees."[133] And HPC can help in testing the candidates identified. For example, the biotech firms Novartis and Schrödinger recently teamed to accelerate the testing of drug candidates by using HPC. The companies tested 21 million drug candidate molecules on the Amazon public cloud, using a new technical computing (HPC-based) algorithm. (Their successful use of HPC for the test run cost only €10,000.)[134]

But returning to the concept of and principles behind computational biophysics, the field has the potential to transform the process of biomedical discovery by leveraging simulation and modeling to detect and ferret out processes occurring at the subcellular level that are unidentifiable by traditional observational or experimental science. In essence, computational biophysics can be conceived of as a "computational microscope" for studying living systems.[135] Schulten elaborates on how computer simulations can deliver advantages over real-world observations:

> Physical measurements can only be taken under certain conditions. For example, the light microscope is really a versatile instrument. But it can only resolve things down to a certain size, and that size is limited by the wavelength of light—a wavelength much too big to see details like molecules in living cells. In an electron microscope, you have a much higher resolution, but you have to use a vacuum environment and dry conditions to examine what you're looking at. So, most of these experimental instruments are very limited. That's where the computer comes in, to be a microscope where real microscopes don't work. Just as Boeing uses a computer to simulate airplanes, we simulate what we know is in the cell.[136]

In other words, HPC-enabled modeling and simulation can divine underlying biological processes in some cases where experiments can only see static processes or structures. Thus, computational biophysicists are becoming competitors to their experimental colleagues in the sense of making discoveries. Supercomputers are giving researchers the ability to look into living cells and resolve them in mechanical detail to understand the processes that occur there, at resolutions in the low nanometers.[137] As Schulten noted in his Biophysical Society 2015 National Lecture, "Without computing, there would be no discovery."[138]

He continues, scientists "are now studying the macromolecules of thousands of proteins working together. That is a big step that was completely impossible to study before. Today, that work is being done on the biggest petascale computers. But with the future exascale one [e.g., an exascale supercomputer], we will be able to do even more: to chemically resolve the details of the cell."[139] As Schulten concludes in explaining why getting to exascale matters for a country's leadership in biomedical science:

> The goal of modern life science is to characterize biological systems from the atom to the cell. We are now somewhere in the middle. A human cell is around 10 micrometers long, and we can simulate it at a scale of about one-hundredth to one-thousandth of that. To reduce it by a factor of 10—a factor of 1,000 by volume—we will need a computer 1,000 times as powerful.[140]

Computational biophysics has the potential to transform the process of biomedical discovery by leveraging simulation and modeling to detect processes occurring at the subcellular level that are unidentifiable by traditional observational or experimental science.

While HPC unlocks tremendous opportunity to understand fundamental biological processes with a degree of fidelity never before imagined, another critical challenge is making the transition from the R&D domain to the clinical application of these technologies. Here, HPC-enabled modeling and simulation tools are benefitting researchers in bioengineering, or how bioinformatics uses HPC for data analytics and insight.[141]

For example, in the field of precision genomic medicine, the New York Genome Center (NYGC) uses IBM's Watson technologies as a medical research application to help physicians choose the best treatments for patients who have unusual conditions. Using the most advanced gene-sequencing machines, this year NYGC will sequence 65 million base pairs every second (4 billion every minute).[142] But to make this technology truly relevant to human disease requires a matchup between the vast amount of data generated and powerful analytical tools capable of making sense of that amount of genetic information—the big data of human biology.[143]

Thus, researchers at the Genome Center are collaborating with IBM scientists to feed Watson vast quantities of disease, treatment, and outcome data to find hidden patterns and correlations. This "learning systems" approach can help answer questions such as: Do people with a combination of genetic dispositions and health problems react better to a certain treatment? Would a drug that's used now for one cancer be useful for another? Would a combination of drugs be better? The hope is that the system will harness the power of innovation and discoveries to improve people's lives by giving physicians a head start in identifying effective treatments, testing them in clinical studies, publishing the results, and quickly getting solutions into the hands of physicians all over the world.[144] While these are just a few examples, they show that the application of HPC to improve patient outcomes in the life-sciences realm is only beginning to take off.

Other Commercial Sectors Leveraging HPC

HPC applications are not limited to the industrial or life-sciences sector, but are also widely used across the finance, entertainment and media, and even sports sectors. A few examples:

Finance

After the May 6, 2010, mega-glitch, aka the Flash Crash, caused the Dow Jones Industrial Average to plummet by about 10 percent, only to bounce right back, a researcher from the University of Illinois, Urbana-Champaign collected two years of data and put it on two supercomputers, one at Pittsburgh and one at DOE. The work uncovered a source of market manipulation that prompted the U.S. Securities and Exchange Commission (SEC) to enact more transparent reporting requirements.[145]

JP Morgan uses a supercomputer "to measure risk in its fixed-income operations by assessing tens of thousands of possible market scenarios, allowing the calculation of complex scenarios in just minutes that previously took hours."[146]

Sports and Entertainment

At least one Major League Baseball team uses a Cray supercomputer to evaluate how batters fare against different types of pitchers.[147]

Scientific Research Applications of HPCs

Beyond commercial applications, high-performance computing has transformed scientific research across numerous disciplines and made tremendous contributions in a number of fields, including in meteorological forecasting and space research. The following section provides several illustrative examples.

Weather Forecasting

Weather—and, by extension, climate—represent two of the most complex physical systems humans encounter on earth, making the role of supercomputers indispensable in modeling and simulating the behavior of these multivariate systems. Accurately anticipating, predicting, and tracking the movements of dangerous weather systems such as hurricanes, cyclones, and tornadoes can save lives and potentially prevent or mitigate damage that can run into the tens of billions of dollars.

When hurricanes appear in the Atlantic Basin each late summer and early fall, hurricane forecasting models from dozens of universities and research institutions vie to divine the most accurate track using thousands of input variables. In recent years, the European Center for Medium-Range Weather Forecasting (ECMWF) forecast model has more often accurately predicted the movements of hurricanes and other high-impact storms than the U.S. Global Forecast System (GFS) model.[148] While several factors explain this, a key contributor has been Europe's application of more powerful supercomputers, better data, and more effective software programs to assess these weather events.

As Richard Rood, a professor at the University of Michigan, elaborates in *The Washington Post*, "In the United States, we [in the weather forecasting community] remain largely reactionary to the evolution of high-performance computing systems. Therefore, each shift in computing technology is a moment in time that the forecast gap is increased."[149] As

Rood notes, ECMWF's "attention to the fundamental issues of software and systems allows not only better operations, but supports the scientific method of investigation."[150] Moreover, it is not only the speeds of the machines that matter, but the quality of the data they are analyzing, and a key advantage for Europe is that the "ECMWF pioneered the examination of failed forecasts and discovered that the forecast busts often hinged on the inclusion or exclusion of a small number of specific observations."[151] As Rood concludes, ECMWF appears to do a better job of "identify[ing] the observations that would most improve the forecast," and its "attention to the entire weather forecasting system and the infrastructure that supports its operations" enables it to produce superior, "science-based" forecasts.[152] As noted before, success in supercomputing is not just about the speeds of machines, but also the quality of the software they run and the data inputs the software uses. In this case at least, Europe's use of supercomputers appears to be delivering stronger results than America's.

Success in supercomputing is not just about the speeds of machines, but also the quality of the software they run and the data inputs the software uses.

But a new U.S. research initiative may help close the gap. The U.S. Naval Research Laboratory (NRL) in Monterey, California, leverages high-performance computing to develop software that more accurately models tropical cyclone forecasting and tracking. Effectively identifying the formation of hurricanes or cyclones and predicting their tracks depends on understanding the complex relationships among a number of physical systems, including ocean circulation, temperature, and salinity; ocean surface waves; and their interactions with the atmosphere and its temperature, moisture, and winds.[153] The Coupled Ocean/Atmosphere Mesoscale Prediction System for Tropical Cyclones (COAMPS-TC™) model recently developed by NRL has given forecasters more accuracy in predicting the track, intensity, and size of tropical cyclones with a nearly 120-hour lead time. The longer warning times give communities more time to evacuate, prepare to protect buildings and other physical property, and to implement contingency plans, potentially saving hundreds of millions of dollars a year in property damage and evacuation costs while mitigating storm-related injuries or fatalities.[154]

Likewise, supercomputers are being used to provide more detailed predictions of coastal storm flooding impacts in vulnerable, low-lying areas. The North Carolina-based Renaissance Computing Institute (RENCI) provided HPC resources to enable the development of a robust forecast system, which for the last four years has been an important tool during East Coast tropical and extra-tropical storms, providing high-resolution predictions of storm surge and waves for vulnerable parts of the U.S. Atlantic and Gulf of Mexico coasts.[155]

Space Research

Understanding the origins of the cosmos depends on the ability to mathematically model the conditions that unfolded in the microseconds following the (apparent) Big Bang. The purpose of the Large Hadron Collider is to reveal particles that were generated in those critical seconds after the Big Bang. As noted previously, supercomputers being used in research institutes such as the Large Hadron Collider play a pivotal role in modeling and simulation of those critical moments after the Big Bang.

Another project, the Square Kilometer Array (SKA), represents one of the most ambitious scientific projects ever undertaken. Its designers envision SKA as a massive radio telescope comprised of more than half a million antennas spread out across vast swaths of Australia and southern Africa. Expected to be completed by 2024, the SKA, which endeavors to collect radio signals from outer space for a multitude of purposes, including better understanding the origins of the universe, monitoring for intelligent life, and understanding the properties of celestial systems, will collect 14 exabytes of digital data per day. As IBM's John Kelly notes, processing all those signals "is the ultimate cognitive computing and big data challenge" and will fundamentally rely on high-performance computing to extract insight from the deluge of signal.[156]

THE GLOBAL HPC MARKET

U.S. vendors remain the most globally competitive in the HPC servers and systems market. The latest full-quarterly period, publicly available data regarding the global HPC marketplace comes from the first quarter of 2015. In that quarter, HP Enterprise, Dell, and Lenovo accounted for 67 percent of HPC systems revenues, with the companies holding 36.1 percent, 16.9 percent, and 15 percent of the market respectively, as Figure 5 shows.

Figure 6 shows Q1 2015 revenues from high-performance computing by the country in which the enterprises are headquartered, showing that U.S.-headquartered companies (notably HP Enterprise, Dell, IBM, SGI, and Cray) accounted for 62.2 percent of global total revenues, followed by China (Lenovo and Sugon) at 17.1 percent, Japan at 3 percent, France at 0.9 percent, and others at 16.9 percent. The U.S. share in Q1 2015 actually declined by 17.7 percentage points from FY 2014, primarily due to Lenovo's acquisition of IBM's x86 line.[157]

Figure 5: Company Share of Global HPC Revenues, Q1 2015[158]

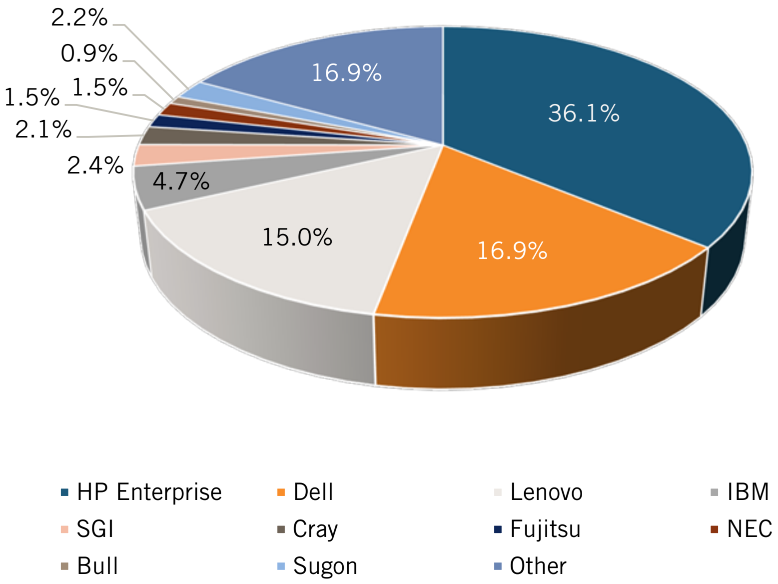

Figure 6: Country Share of Global HPC Revenues (by Headquartered Company), Q1 2015[159]

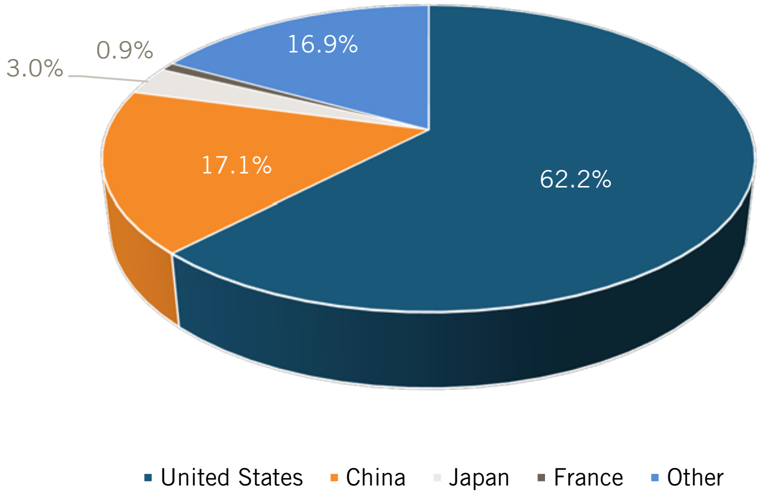

Apart from revenues, calendar year (CY) 2015 data is available for “processor package volumes”; in effect, how many processors HPC players shipped in 2015. As Figure 7 shows, HP Enterprise, Dell, and Lenovo unsurprisingly again rank in the top three, capturing 30.2 percent, 25.7 percent, and 17.2 percent, respectively, of processor package volumes, for a combined share of close to three-quarters.

When assessed by headquartered country of the HPC vendor, U.S.-headquartered HPC vendors accounted for 68.8 percent of processor package volumes for CY 2015, Chinese players (Lenovo and Sugon) for 20.9 percent, Japanese companies (Fujitsu and NEC) for 2.2 percent, France (Bull Atos) for 0.8 percent, and miscellaneous others for 11 percent, as Figure 8 shows. It is worth noting that almost all of the processors counted here—regardless of which vendor used them in their systems—were of U.S. origin.

Of the 100-fastest high-performance computers in the world as of November 2015, the United States leads as the country whose companies were principal manufacturers of the chipsets in those machines (although local players are more often involved in design of the interconnects and software running on those machines). Nonetheless, U.S. HPC vendors were the principal manufacturers of 69 of the world’s 100-fastest HPCs as of November 2015, followed by China with 11, Japan with 10, France with 7, and Russia with 3, as Figure 9 illustrates.

Figure 7: Processor Package Volume CY 2015, by HPC Company[160]

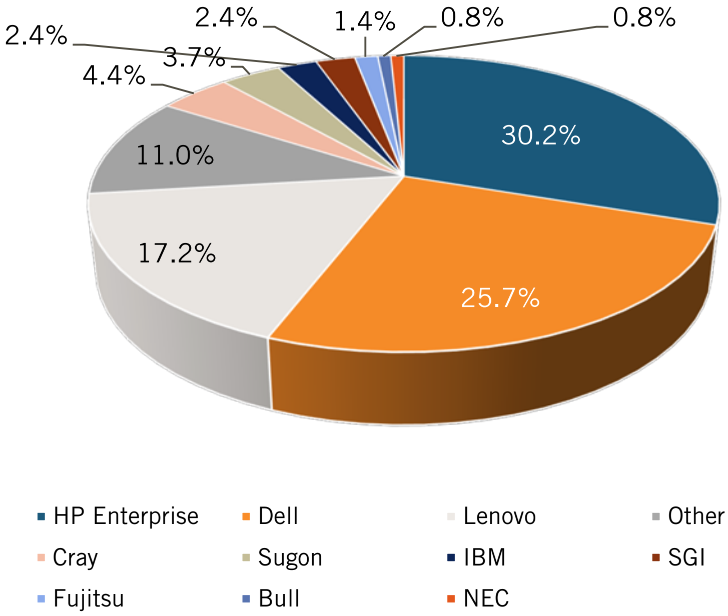

Figure 8: Processor Package Volume CY 2015, by Country Headquarters of Company[161]

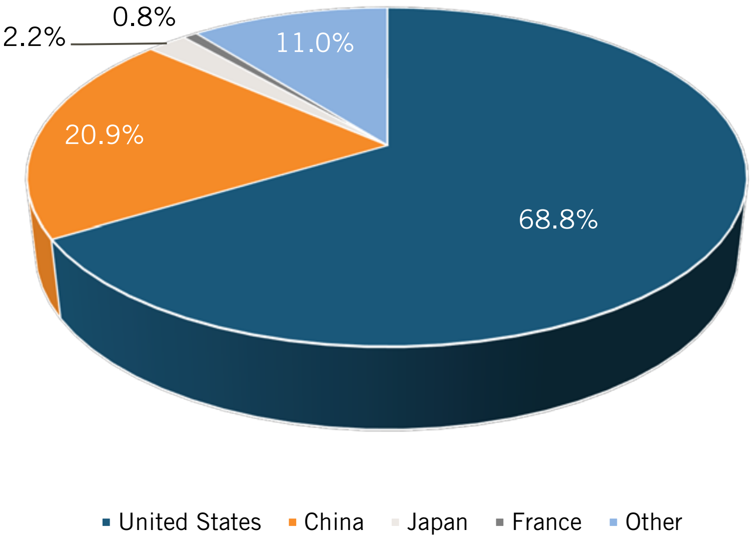

Figure 9: Headquarters Country of Principal HPC Vendor of World's 100-Fastest HPCs[162]

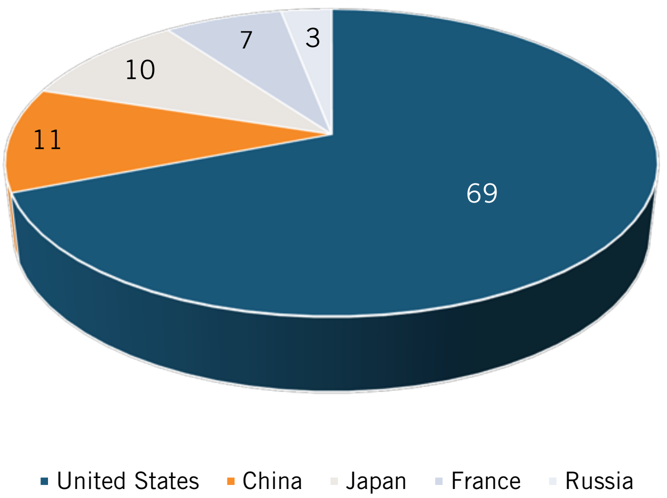

In terms of where the HPC systems sold in CY 2015 were purchased, by geographical region, the United States deployed 40.9 percent; Europe, the Middle East, and Africa (EMEA) deployed 33.9 percent; the (non-Japan) Asia-Pacific region deployed 18.2 percent; Japan deployed 6 percent; and the rest of the world deployed 1.1 percent, as Figure 10 shows.

Figure 10: HPC Systems Deployments, by Region, CY 2015[163]

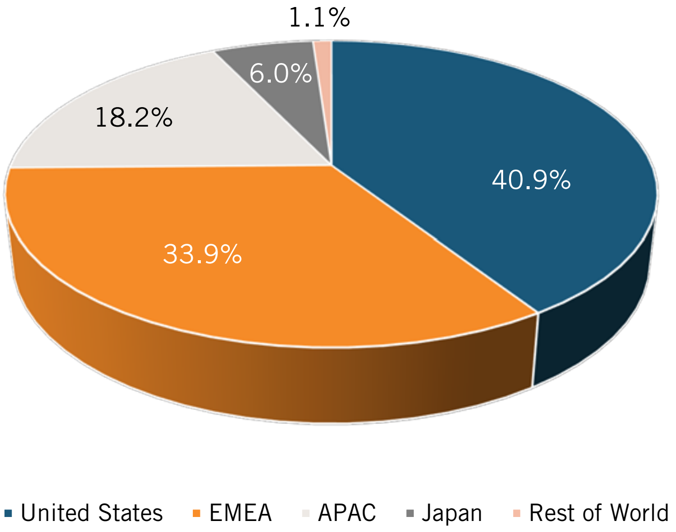

Unfortunately, international trade balance data specifically for high-performance computers is unavailable, the United Nations Comtrade Database only carrying data for "computer products" trade, for which it reported a U.S. deficit of $56 billion in 2015.[164] However, the U.S. trade balance in the high-performance end of this market is likely positive.

INTERNATIONAL COMPETITION FOR HPC LEADERSHIP

As the preceding sections elucidated, leadership in both the development and use of high-performance computing is vital for countries' economic competitiveness and innovation potential. Accordingly, many countries have made significant investments and implemented holistic strategies to position themselves at the forefront of the competition for global HPC leadership. As IDC explains, "Other nations and global regions including China, the United States, Japan, and Russia, are racing ahead and have created national programs that are investing large sums of money to develop exascale supercomputers. What this global race is really about is supremacy in supercomputing and in all the disciplines and markets that depend heavily on this game-changing technology."[165]

Perhaps the most well-known measure of how countries are faring in the competition for HPC leadership is the Top 500 list, which every six months ranks the world's 500 fastest supercomputers. As Figure 11 shows, the United States continues to lead these rankings, with 199 supercomputers making the most recent, November 2015, list.

Compared to a year earlier, in 2015 the number of U.S.-based Top 500 supercomputers dropped 14 percent.

Figure 11: Number of Supercomputers in Top 500, by Country, November 2015[166]

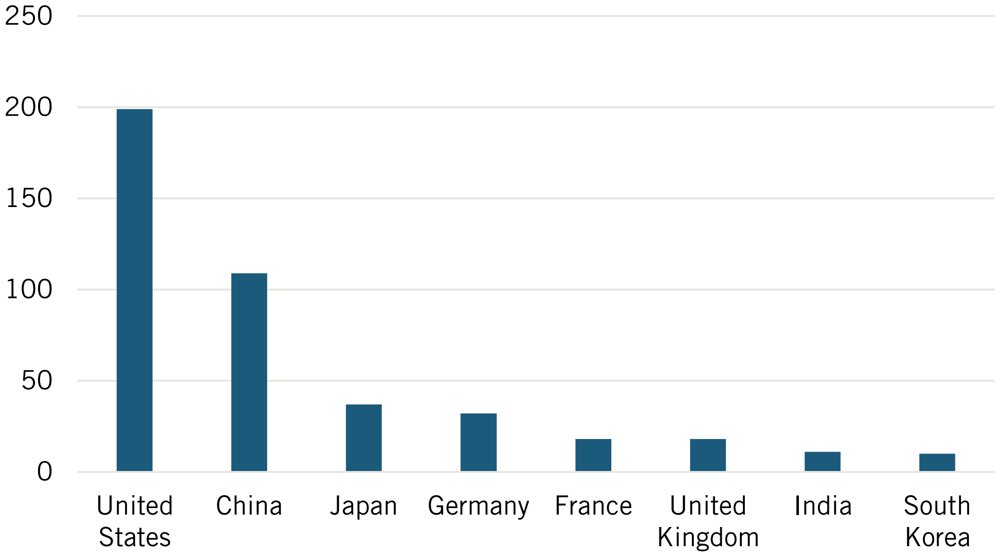

Yet this graph does not tell the full story. The United States' 199 supercomputers on the Top 500 list in November 2015 represent the fewest number of supercomputers the United States has placed in the Top 500 since the list's inception in June 1993. Compared with a year earlier, the number of U.S.-based Top 500 supercomputers dropped 14 percent from the 231 the United States placed on the Top 500 list in November 2014.[167] In short, the global competition for HPC leadership has been intensifying, as the following section elaborates.

China

Following the United States, China boasts one of the world's most well-developed HPC ecosystems, having ramped up HPC investment faster than any other nation or region. Since June 2013, China's *Tianhe-2* supercomputer has claimed the mantle of the world's fastest, delivering a peak performance of 54.9 petaflops. China's National University of Defense Technology built *Tianhe-2* as part of a government-academic project at the National Super Computer Center in Guangzhou at a cost of almost $400 million.[168] The *Tianhe-2*'s top performance speed nearly doubles that of the United States' fastest supercomputer, the *Titan*. China stands second to the United States' 199 supercomputers, with 109 Chinese supercomputers in the world's Top 500 fastest.[169] Just as important as the performance of its most advanced HPC system is how the *Tianhe-2* fits into China's broader HPC ecosystem and how rapidly this is changing in China. From June to November 2015, China tripled the number of supercomputers it placed in the Top 500, increasing its number of Top 500 supercomputers in just six months by 196 percent, from 37 to 109 machines.[170]

Moreover, in 2016, China is expected to be the first country to bring online not one, but two, 100+ petaflop supercomputers, with China's machines expected to come online about one year before the United States' 100+ petaflop-capable supercomputers being developed as part of the Collaboration of Oak Ridge, Argonne, and Lawrence Livermore (CORAL) initiative (i.e., the *Aurora, Sierra,* and *Summit* supercomputers).[171] The first of these two 100+ petaflop Chinese machines is being developed at the Jiangnan Institute of Computer Technology and will use a next-generation, Chinese-designed and manufactured chip, the ShenWei chip.[172] The second domestically designed chip will be the "China Accelerator" that the National University of Defense Technology is developing for the *Tianhe-2* supercomputer (in part as a result of U.S. export control regulations).[173] China's government is encouraging Chinese HPC systems, interconnect, and component makers to aggressively export their products on international markets, part of China's "go outside" strategy.[174]

Clearly, China has made HPC leadership a national priority. A key reason for this is that, for China, leadership in high-performance computing is central to the country's goal of transitioning away from reliance on foreign technology to using homegrown technology. As Li Na, a spokesperson for the *Tianhe-2* project, explains, "We are producing supercomputers with a fundamental purpose of providing a driving force for the construction of an innovation-oriented country."[175] As IDC's Rajnish Arora explains, "The Chinese government and companies want to become the creators and not just producer of products that are being designed elsewhere."[176] Or, as Chinese President Xi Jinping himself puts it, China has built its HPC capabilities in part to demonstrate that the country has become a "cyber power."[177]

The Chinese government's push to enhance the country's HPC capabilities began in 1989 with the National Computing and Networking Facilities initiative. Over the last decade, China's National High-Tech Research and Development (863) Program, which promotes innovation through public investment, took over, spawning the National High Performance Computing Environment (NHPCE) initiative, which supported development of the *Tianhe-2* supercomputer and others such as the *Nebulae* at China's National Supercomputing Centre (NSCS) in Shenzhen.[178] China's NHPCE now focuses on providing HPC services to nearly 3,000 Chinese government, industrial, and academic clients. Yet China's 25-year push to build supercomputing capability has been as much about gaining competencies in integrated circuits and other supercomputing hardware components as it has been about producing the end machines. As Alspector et al. write in *China—A New Power in Supercomputing Hardware*:

> Supercomputer development in China, pushed largely to promote industrial development, has contributed to developing domestic companies capable of producing highly competitive servers, handheld mobile appliances, integrated circuit designs, multicore processors, digital signal processors, secure cryptographic chips, secure operating systems, and HPC software.[179]

As Alspector et al. continue, "China's achievements in HPC have been remarkable.…There is reasonable agreement on the part of respected computer scientists that the Chinese HPC community is rapidly catching up with the West in the hardware arena."[180] Nevertheless, while China can fairly claim that 109 of the world's 500 fastest supercomputers are located in China, more than 85 percent of the components in China's supercomputers have heretofore come from foreign vendors, and that has made those imports a target of China's efforts to supplant foreign-enterprise-developed information technology component and end-product imports with domestically produced equipment.

Following the United States, China boasts the world's most well-developed HPC ecosystem, having ramped up HPC investment faster than any other nation or region.

Indeed, the ability to supply indigenously produced semiconductor chips to power China's HPC aspirations has been a key animating factor behind China's decision to announce its "National Guidelines for Development and Promotion of the Integrated Circuit (IC) Industry."[181] The guidelines—essentially China's national semiconductor industry development strategy—articulate a goal of creating a completely closed-loop semiconductor ecosystem, from design and prototyping to manufacturing, assembly, testing, and packaging.[182] The strategy unabashedly calls for China to reduce imports of U.S. semiconductors by half in 10 years and to eliminate them entirely within 20 years. It also establishes specific revenue targets of 20 percent compound annual growth and sets a goal of increasing the industry's size to $140 billion by 2020.[183] To achieve these goals, China has launched a National Integrated Circuit Industry Investment Corporation that intends to invest more than $100 billion in China's semiconductor industry over the next decade. It's also gone on an acquisitions spree, including Lenovo's 2014 acquisition of IBM's x86 server business, which has made Lenovo one of the world's top four HPC server system vendors. But whether through acquisition or development, there is little doubt that a core goal of China's National IC Strategy is to supplant foreign-designed semiconductor chips with domestic suppliers as China looks to develop its own HPC systems going forward.

However, while China has shown it can build massively parallel, fast supercomputers, it lags at developing innovative software applications that can leverage these supercomputers to generate new insights and discoveries across a wide range of fields. As *HPCWire's* Tiffany Trader puts it, "China's challenge has been a dearth of application software experience."[184] For example, the *Tianhe-2* "is reportedly difficult to use due to anemic software and high operating costs [including] electricity consumption that runs up to $100,000 per day."[185] In fact, because of a lack of software support from the *Tianhe-2* project's backers, users have been forced to write application programs themselves, making the expensive machine far less useful than it otherwise could be. That's why one official described it as "a giant with a super body but without the software to support its thinking soul," and, as of June 2014, it had only 120 clients using only 34 percent of its capacity.[186] In short, China's HPC approach thus far appears to have emphasized performance speeds over practical applications, meaning the functionality of its machines lag those in Europe and the United States (although system efficiency is a challenge everywhere). Nevertheless, China's HPC capabilities grow daily. China has clearly become America's leading competitor in HPC systems production, and China's HPC vendors hope to leverage that into strong exports on the international market.

The European Union

In February 2012, the European Union released a report, *High Performance Computing: Europe's Place in the Global Race*, which articulated a goal (and set of action plans) "to achieve HPC leadership" including by "acquiring at least one exascale computer in the same timeframe as the U.S., Japan, and China."[187] Europe's HPC strategy recognizes that "HPC is a strategic resource for Europe's future" with "computational science already the 'third pillar' of science" and "industry relying more and more on HPC to innovate in products and services."[188] In accordance with this initiative, the European Union has significantly increased funding for HPC systems research and development, including launching the European Technology Platform on High-Performance Computing (ETP4HPC) in 2012, part of a €700 million public-private investment in HPC through the Contractual Public Private Partnership (cPPP) initiative.[189] That investment has been complemented by €400 million in commitments (mostly in-kind contributions) made through the Partnership for Advanced Computing in Europe (PRACE). However, IDC notes that Europe will need to amass €1 billion more in funding if it is to meet a stated European Commission goal of "acquiring two pre-exascale supercomputers in 2019-2020 and two additional exascale supercomputers in 2022."[190]

Reviewing Europe's efforts since 2012 to bolster its HPC capabilities, market research firm IDC wrote in late 2015 that, "The European Commission, PRACE, and ETP4HPC have done an admirable job of advancing Europe's position in the few years since the 2012 Communication."[191] However, in the same report, IDC commented that, despite Europe's ambition to compete with China, Japan, and the United States, "No clear funding path exists to acquire/operate exascale systems in competitive timeframes."[192] Perhaps responding to that feedback, on February 25, 2016, under the auspices of the European Union's Horizon 2020 R&D funding program, the European Union announced the launch of ExaNeSt, a consortium of 12 HPC hardware and software developers and suppliers (all European) that will "seek to build a straw man exascale HPC prototype in 2016 and a full prototype version by 2018 as a means to drive indigenous HPC technology development across the EU."[193] As IDC's Robert Sorensen commented, "The ExaNeSt project is an ambitious project from a technological [and organizational] perspective.... However, if these hurdles can be managed, this could be a most promising effort to develop an indigenous exascale infrastructure that can compete with counterpart efforts in China, Japan, the United States, and other countries."[194]

Though it is narrowing the gap, Europe has had some ground to make up in the HPC race. For instance, Europe's five leading HPC-using nations placed a combined 74 supercomputers on the November 2015 Top 500 list, including Germany with 32, France and the United Kingdom with 18 each, and Switzerland and Poland with 6 each. Combined, the European Union placed fewer than 100 systems on the November 2015 Top 500, half the United States' 199 and fewer than China's 109 systems.[195] Moreover, Europe has not produced a substantial supercomputer vendor, and the European marketplace represents only one-quarter of global spending on HPC systems.[196] Europe is, however, much more competitive in HPC software applications. This disparity is shown by the fact that, while 83 percent of HPC application software used in Europe was created there, Europe's share of HPC system vendors stands at less than 5 percent.[197]

IDC notes that, historically, "Europe's HPC community has been more oriented toward science than industry."[198] Further, IDC points out that, "Industrial access to Europe's leading supercomputers has grown but remains limited."[199] However, Europe is working to address this, with the PRACE, SME HPC Adoption Programme in Europe (SHAPE), and Fortissimo initiatives focused on expanding industrial use. Thus, while Europe does "have some world-leading HPC centers for collaborating with industrial users, including SMEs (such as HLRS, Teratec, SURFsara, CINECA, and LRZ)" democratizing access to HPC resources for industrial purposes remains as much a challenge in Europe as it is in the United States.[200] In terms of sectoral investment, IDC reports that Europe's manufacturing sector invested approximately €638 million in HPC systems in 2015, while Europe's "bio-life sciences" sector invested €416 million in HPC systems in 2013, with that figure expected to grow to €510 million by 2018.[201]

As IDC concludes, "Today, there is a European [HPC] consciousness, a European-wide scheme for categorizing HPC centers, more of the world's top 50 supercomputers [than ever before], and improved access for industry of all sizes (including SMEs)" to go along with the new ExaNeSt exascale initiative. IDC also notes that in "Europe it is inherently more challenging to organize and collaboratively advance HPC than [in] its single-nation competitors for leadership: the U.S., Japan, and China" and that Europe needs "a visible person/evangelist in charge" who can communicate Europe's HPC vision. However, while ExaNeSt represents a European push to get into the exascale game, the major thrust of the European HPC strategy appears to be positioning its HPC investments so it can become a leader in HPC software and applications development. In other words, Europe appears to be trying to position itself to lead in HPC adoption broadly, and in terms of production, to focus on the software rather the hardware side.

Japan

Japan ranks third on the November 2015 Top 500 list, with 37 of the world's 500 fastest supercomputers, including the *K* computer, Japan's fastest computer and the world's fourth-fastest overall, which operates at 10.5 petaflops. An earlier version of the *K* computer had been the world's fastest, until 2011. But Japan has now launched an initiative to recover its world-leading HPC position. Japan's national *Flagship2020* program calls for investment of over $1 billion, with the goal of delivering by 2020 a supercomputer with 100 times more application performance than the current *K* computer.[202] The *Flagship2020* program clearly puts Japan in the middle of the race to reach exascale by 2020. Japan's new supercomputer intends to use a six-dimensional design that can facilitate connections for more simultaneous CPUs, memory, and storage compared to traditional system configurations.

While supercomputers in Japan certainly support industrial applications, Japan's *Flagship2020* program has actually prioritized development of a supercomputer particularly attuned for social and scientific purposes. That thinking has been driven by realization among Japanese leaders that health care and aging, natural challenges including earthquakes and climate change, and the need to sustainably produce clean energy are among the nation's most significant challenges. Interestingly, as Japan designs its new supercomputer, it is already thinking about how to best tailor it to suit the types of analytical challenges that software programs will face in dealing with such social and scientific issues.[203] For example, understanding DNA at the molecular level and achieving an understanding of full-body medicine (in other words, modeling how a change to one organ or system in the body affects others) are some of the most difficult challenges in modern biology. Japan literally intends to build its computer architecture in a way that can help solve those questions; in essence, Japan is following a "design thinking" approach toward the development of the most functionally effective and useful HPC system possible.[204] And this is important because, as with China, in Japan to date most application programs have failed to use the full capacity of existing HPC systems. However, Japan has been a world leader in applying HPC to large scale metagenomics challenges, leading in developing the first ultra-high-sensitive metagenome sequence analysis of the human oral microbiome.[205]

Real national leadership in HPC comes from the combination of superfast systems, designed in a functionally operational and system-efficient way.

India

In October 2014, India announced its National Supercomputing Mission (NSM), which calls for $730 million in investment from 2016 to 2023 to build out a vast supercomputing grid comprised of 73 high-performance computing facilities.[206] It also calls for India to produce three petascale machines operating in the 25-30 petaflop range. The seven-year mission will take place in two phases: the first three years focused on construction of the 73-networked systems and the remainder on application development.[207] Professor Rajat Moona of India's Centre for Development of Advanced Computing noted that a key success factor for the mission would be training Indian scientists to develop "home-grown applications" in the fields of agriculture, medicine, space research, and manufacturing technology.[208] Moona further noted the NSM would be designed to tackle, "Engineering problems such as weather and climate modelling, computational fluid dynamics, computational structural mechanisms are other areas of applications."[209] As *HPCWire's* Trader observes, India has chosen to focus on practical applicability with its networked grid HPC approach, putting "usability before [speed] rankings."[210] As of November 2015, India has 11 supercomputers in the Top 500, and one, the 97th-ranked *SERC*, in the Top 100.[211]

South Korea

In June 2011, South Korea passed the National Supercomputing Act, which established a five-year master plan (covering the years 2013 to 2017), which intends "to place Korea in the top seven nations of supercomputing by 2017."[212] The plan focuses on three key areas: applications, infrastructure, and technology, with major emphasis given to building a balanced ecosystem that addresses all key components in the workflow, from systems to solutions. The program will be led by the Korea Institute of Science and Technology Information (KISTI), which has established the National Institute of Supercomputing and Networking (NISN). NISN will manage $100 million in R&D investments as part of the "SuperKorea 2020" initiative, which seeks to indigenously develop a petascale-range "national leadership" supercomputer.[213] KISTI is also developing a National Supercomputing Education and Training Framework that will both train scientists and engineers to work on development of HPC systems and work with enterprises to show them how they can leverage HPC systems for competitiveness. South Korea placed 10 supercomputers on the November 2015 Top 500 list.[214]

Russia

Russia has recently redoubled its efforts to raise its game in the global HPC competition. Then-Russian President Dmitry Medvedev announced the Medvedev Modernisation Programme in 2009, which called for domestically produced high-performance supercomputers "in order to provide means for the complex calculations for nuclear technology centers, aircraft industry and other major clients which need high-performance computing for computer simulation of their projects."[215] Medvedev explained that use of advanced supercomputers would be vital for Russia to remain globally competitive in its aerospace and other manufacturing sectors, stating that, "Any sort of airframe or engine that is not produced with the aid of supercomputers is unlikely to trigger interest among buyers in a few years, because even now there are standards already set and so far we are doing practically nothing to meet them."[216] Beyond aerospace, Russia's most significant use of HPCs for industrial purposes include advanced material creation, manufacturing and process modeling, and 3-D seismic modeling of oil and gas fields and reserves. Sergei Abramov, head of the Russian Academy of Sciences Program Systems Institute, observed, "A supercomputer is the only instrument to beat a competitor."[217] As of November 2015, Russia's fastest supercomputer, the *Lomonosov 2*, located at Moscow State University, ranked as the world's 37th fastest, with another six supercomputers among the Top 500. While Russia is hurriedly trying to make up ground, analysts estimate that Russia still lags some five-and-a-half years behind the United States in supercomputer technology.[218]

Country Summary

International competition for HPC leadership is clearly intensifying, as summarized in Table 1. That is what makes U.S. strategy toward high-performance computing, as described in the following section, so important.

Table 1: Summary of National HPC Strategies, by Country[219]

Country	HPC Strategy/Program and Description	Investment Level
United States	National Strategic Computing Initiative (NSCI)	@$320 million/year
China	13th Five-Year Development Plan (Develop Multiple Exascale Systems)	$200 million/year (for next five years)
Japan	Flagship2020 Program	@$200 million/year (for next five years)
European Union	ExaNeSt; PRACE; ETP4HPC	@$1.1 in billion total allocated through 2020 (annual allocations N/A)
India	National Supercomputing Mission	$140 million/year (for five years from 2016-2020)
South Korea	National Supercomputing Act	$20 million/year (for five years from 2016-2020)
Russia	HPC Focus of Medvedev Modernisation Programme	N/A

WHY A U.S. HPC POLICY IS NEEDED

Some argue that the United States does not need robust investments in the next generation of high-performance computing (as envisioned by the National Strategic Computing Initiative) because U.S. HPC vendors appear to be globally competitive. Indeed, if U.S. companies are already leaders in the HPC space, some might contend that federal investment in HPC is unnecessary. But this contention suffers from several problems. First, it misses that the goal of the NSCI is not to somehow advantage or finance U.S. HPC vendors, but rather to strengthen the entire U.S. HPC ecosystem, from broadening knowledge and use of HPCs, to developing talent and skills, to facilitating their application by government agencies for specific mission-oriented purposes, from defense to renewable energy to weather forecasting.

Second, and more importantly, it misses that federal R&D investment in frontier technologies has in fact been a foundational catalyst in the very development of U.S. technology-based industries (particularly in the information and communications technology sector) that are world-leading. Moreover, as the preceding section documented, other nations are aggressively supporting the development of home-grown high-performance computing competitors. Absent policy steps (as outlined below), it’s conceivable that HPC could follow the path of the personal computer in terms of eroded national competitive advantage.

The United States’ global leadership in developing breakthrough information and communications technologies and systems over the past half-century has been the unique product of both robust public and private sector investment and innovation. In particular, the role of the U.S. federal government as an R&D funder—and early procurer—of ICTs has been indispensable in underpinning the development of a number of core information technologies, from the transistor and integrated circuit to relational databases, graphical user interfaces, the global positioning system (GPS), and search engines.[220] As Singer notes, "One economic sector where federal research funding has worked synergistically with industry is information technology."[221] In fact, the National Research Council points to at least eight IT sectors, all but one of which now constitute a more than $10 billion global industry, that have their roots in federally funded (often academically conducted) scientific research.[222] As Rob Leland of the Sandia National Laboratory explains, "Each new major area in computing has been preceded by five to seven years of a forward-looking investment by the government R&D push … You can trace that back at least five cycles."[223] Further, in many cases, the government's role as an early adopter and procurer of nascent information technologies drove their price down to a point that made their application by industry feasible.

America's National Strategic Computing Initiative defines a multiagency framework for furthering U.S. economic competitiveness and scientific discovery through orchestrated advances in high-performance computing.

The history of supercomputers has been no different. As Signer notes, "Driven by the demands of nuclear research, the U.S. national laboratories worked with private companies to develop new supercomputers and to provide the requirements that shaped the field."[224] And by pushing the leading edge of computation, "new technologies and capabilities first funded by government ultimately become available for business—in aeronautics, pharmaceuticals, finance, energy, automotive, and many other sectors."[225] Indeed, as the Council on Competitiveness's *Solve* report noted, 62 percent of study respondents agreed (or strongly agreed) with the statement, "Past government investments in new generations of supercomputing have had a benefit on your company/industry."[226] Moreover, 56 percent of respondents agreed that work done by national government research organizations "act[s] as a major driver for advancing HPC technology, leading to products and software that we will use in the future."[227] In short, the federal government's role in investing in HPC systems' R&D and acquisition; in deploying HPC systems throughout America's network of national laboratories and facilitating their use by industry; and in convening and coordinating the activities of academic, government, and commercial actors in America's HPC ecosystem has been instrumental to America's historical leadership in high-performance computing.

The National Strategic Computing Initiative

Recognizing the heightening global competition for high-performance computing leadership, in July 2015 President Barack Obama, by Executive Order, announced the U.S. National Strategic Computing Initiative. The NSCI seeks to create a coordinated federal strategy for HPC research, development, and deployment and defines a multiagency framework for furthering U.S. economic competitiveness and scientific discovery through orchestrated HPC advances.[228] The NSCI represents a whole-of-government effort designed to create a cohesive, multiagency strategic vision and federal investment strategy, executed in collaboration with industry and academia, to maximize the benefits of HPC (in terms of both production and adoption) for the United States. Tim Polk, assistant director of cybersecurity with the White House Office of Science and Technology Policy (OSTP), explained the importance of the NSCI and of exascale computing for the maintenance of

U.S. leadership over the coming decades: “The United States must make strategic investments in high-performance computing to meet increasing computing demands and emerging technological challenges.”[229] Indeed, the NSCI is positioned to “transform the world’s capacity to calculate, analyze, and ultimately address some of the most pressing challenges global society faces.”[230] Before President Obama issued the NSCI Executive Order, there was no coordinated federal activity for dealing with the nation’s HPC needs.[231]

The NSCI articulates five strategic objectives:

1. Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10- petaflop systems across a range of applications representing government needs;
2. Increasing coherence between the technology bases used for modeling and simulation and those used for data analysis in supercomputing;
3. Establishing a viable path forward for future HPC systems, even after the limits of current semiconductor technology are reached (the “post- Moore’s Law era”);
4. Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development; and
5. Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the U.S. government and industrial and academic sectors.[232]

The NSCI represents a serious effort to sustain U.S. leadership in HPC for the foreseeable future.

The Department of Energy (DOE), the Department of Defense (DOD), and the National Science Foundation (NSF) will serve as the three coordinating lead agencies for the NSCI. (DOE and DOD are also the leading government-agency consumers of HPC systems.) In terms of conducting foundational HPC R&D activities, the Intelligence Advanced Research Projects Activity (IARPA) and the National Institute of Standards and Technology (NIST) will play important roles. IARPA will focus on future computing paradigms offering an alternative to standard semiconductor computing technologies. NIST will focus on measurement science to support future computing technologies. In terms of funding, the Obama administration’s FY 2017 budget proposal calls for $285 million for the National Strategic Computing Initiative through the Department of Energy and a further $33 million for activities being pursued by the National Science Foundation.[233] If such funding levels are authorized and appropriated by Congress—and sustained at commensurate levels in coming years—U.S. funding for high-performance computing would approach $1 billion over the ensuing three-year period. (This follows on DOE’s November 2014 announcement that it has committed to investing $325 million to research extreme-scale computing and build two new supercomputers, the previously mentioned *Aurora* and *Sierra* supercomputers.)[234]

The U.S. Department of Energy will play the leading role in the NSCI, through two distinct efforts: the Exascale Computing Initiative (ECI) and the Exascale Computing Project (ECP) (one roughly corresponding to hardware and the other software). As DOE explains, "We are once again at a critical turning point in high-performance computing (HPC) technology, with industry innovations in hardware and software architectures driving advances in computing performance, but where the performance of application codes is suffering because the technology advances are not optimized for memory intensive, floating point HPC use."[235] In other words, NSCI recognizes that it's the combination of well-integrated HPC hardware and software that truly matters for national HPC leadership. Indeed, the NSCI cites as a core goal: "To revolutionize our problem-solving capabilities by combining the best attributes of today's 'computing intensive' and 'data intensive' architectures." As Bright Computing's Hansen comments, NSCI is about "Systems that can perform complex modeling and simulation to derive insightful theoretical outcomes but that are also fast and nimble enough to process and respond to massive volumes of real—rather than theoretical—information."[236]

The NSCI represents a serious effort to sustain U.S. leadership in HPC for the foreseeable future. As the Department of Defense's Will Koella states, "The NSCI is one of the more elegantly put together and well-founded initiatives I've seen come out of government." As he continues, "Our hope is that this will be an 'Apollo' project for our nation. We need to ignite the public's imagination around computers."[237]

However, IDC's Robert Sorensen notes that at least five critical challenges confront America's NSCI:

1. Fostering a robust commercial HPC sector that can supply systems to critical U.S. government missions;
2. Keeping the United States as the leading supplier nation in an increasingly competitive global HPC sector;
3. Building up an HPC workforce that ensures an adequate number of qualified job applicants and workers for the HPC R&D and deployment disciplines;
4. Training a wide range of non-HPC scientists and engineers across a broad range of technical areas to introduce or improve their use of HPC in their overall business processes; and
5. Managing coordination between agencies that have not worked together before as part of this whole-of-government effort. [238]

In other words, articulation of the NSCI is a great start, but more needs to be done, as the following section elaborates.

POLICY RECOMMENDATIONS

To ensure America's continuing leadership in high-performance computing, the Information Technology and Innovation Foundation offers the following policy recommendations:

Congress should hold hearings on the National Strategic Computing Initiative and the intensifying race for global HPC leadership. For the NSCI to succeed, there must be a broad-based consensus on the technical potential of this program and of the importance of U.S. leadership in high-performance computing. There has to be a national commitment, and it will need to be multi-administration because the investments will take that long to make and payoff. Congressional hearings can highlight the importance of high-performance computing to America's industrial competitiveness, scientific leadership, and national security.

Congress should authorize and appropriate NSCI funding levels as requested in the administration's FY 2017 budget for FY 2017 and future years. Leadership in high-performance computing will require a steady, stable, robust, and predictable stream of funding. Congress should fund NSCI and related high-performance computing initiatives at a level of at least $325 million per year over the next five years.

While some will assert that it is industry's job alone to make these types of investments, the reality is that public-private partnerships and investment will be needed to ensure the United States remains at HPC's cutting edge. As noted, forward-looking federal investment has been foundational to U.S. leadership across a range of IT sectors historically, and supercomputers are no different. Industry alone (though it will make significant investments in next-generation HPC) cannot make the kinds of dramatic investments necessary to reach exascale, in part because many of the benefits of reaching exascale will inure to public missions (i.e., defense, science, health, etc.) and in part because industry has limited capacity to invest in technology systems exhibiting extreme risk. As ITIF wrote in *Innovation Economics: The Race for Global Advantage*, "Even 'rational' companies are reluctant to invest in next-generation technologies, especially when it involves high levels of risk and exceedingly lengthy R&D time frames."[239] That explains why Defense Advanced Research Projects Agency (DARPA) investment was so vital in supporting the initial development of the Internet, and why it remains so vital in the race to exascale. Moreover, as noted previously, robust investment in the NSCI is warranted because it would be imprudent to rely on foreign-produced HPC systems both because of the national security implications and because a disruption to the supply of those systems could imperil the competitiveness of U.S. industries that consume HPC systems.

The administration should make technology transfer and commercialization activities a priority focus of America's network of national laboratories. As this report has shown, a number of U.S. national laboratories, including the Lawrence Livermore National Laboratory and the Sandia National Laboratories, particularly, have played crucial roles in assisting U.S. industry in leveraging HPC resources and know-how for innovation and industrial competitiveness. It is good to see that role recognized and reaffirmed by the National Strategic Computing Initiative. But Congress should likewise continue to affirm technology transfer and commercialization as a core mission of the national laboratories, including by increasing the weighting attached to technology transfer and commercialization activities as part of the labs' Performance Evaluation Management Plan, or PEMP, process.[240] As the previous examples of productive outcomes from U.S. enterprise and U.S. national laboratory collaborations on HPC-related projects illustrate, U.S. companies benefit immensely by tapping into the national laboratories' latent expertise; increasing the expectation within the national laboratories that translating technology and insights to the private sector is a priority will only amplify these types of collaborations.

Congress should expand funding for the National Network for Manufacturing Innovation. The NNMI and its institutes, notably the Digital Manufacturing and Design Innovation Institute, are playing key roles in helping America's industrial base leverage HPC. Other IMIs, such as the Youngstown, Ohio-based America Makes, which focuses on additive manufacturing, are also assisting in helping America's SMEs leverage high-performance computing to innovate, to eliminate product development costs, and to speed time to market. Maintaining committed funding for these institutes—and providing funding for the five additional manufacturing institutes requested in the administration's FY 2017 budget—would bring the number of IMIs in the NNMI to at least 15 and demonstrate America's commitment to lead the world in manufacturing product and process innovation (and applying HPC thereto).[241]

The Department of State and the Department of Commerce should continue to pursue export control reform to match the reality of current high-performance computing systems. The U.S. export control regime governing exports of high-performance computers has failed to keep up with the pace of innovation in the field. As Figure 1 demonstrated, HPC systems that were cutting-edge just 10 years ago are run-of-the-mill today, yet may still be treated the same for export control purposes.

U.S. export control rules should be updated so that only the newest and most-sensitive HPC systems and technologies are subject to export control rules.

Accordingly, the system should be updated so that only the newest and most-sensitive HPC systems and technologies are subject to export control rules. HPCs fall under export rules for "digital computer systems," meaning that when U.S. vendors wish to export an HPC system or component, they must undertake an exhaustive analysis of what they are shipping, who the system is going to, what the device is going to be used for, and determine if an export license will be required or not. Such determinations entail an arduous, time-consuming process that on some occasions has cost U.S. vendors sales.

Requirements under the International Traffic in Arms Regulations (ITAR) should be amended so they don't inadvertently control commercial components and technology that may be used by government and commercial customers. Current ITAR controls often trigger an automatic review for any device (and device is a broad term here, meaning hardware, software, services, etc.) that is being used for a defense application (and that's also a very broad term, because it can apply to any HPC system or component being used by a national laboratory). Accordingly, these controls snare HPC systems and components that are also being sold in other contexts for government, academic, or commercial purposes, meaning that often these machines cannot be sold overseas (or only sold after triggering an extensive and often times costly ITAR compliance review). Put simply, current ITAR controls are at odds with the increasing trend toward the democratization of HPC, and again should be updated to narrowly apply only to the most sophisticated and sensitive cutting-edge HPC systems. These regulations hinder the ability of commercial companies that export products and technology worldwide to collaborate effectively with government agency customers for fear their technology will be inadvertently captured under the ITAR. Going forward, applying export controls to exascale systems may be sensible, but not to petascale or lesser systems.

The intent of this recommendation is to open a dialogue about reevaluating how commercial hardware, technology, and software (much of which is ubiquitous across the HPC ecosystem) is inadvertently captured by the ITAR if and when the commodity in question is being used in a computer system being deployed by a government agency customer. The recommendation to reevaluate U.S. export control laws also recognizes that in some cases overly stringent export control regulations have prevented the sale of noncritical HPC systems to customers in some nations, a policy decision that (as with China) has had the unintended consequence of further spurring these nations to pursue their own HPC development programs. HPC vendors from a number of countries, from China and Japan to Korea and Taiwan, have benefitted by being able to step in and make sales in situations where potential sales of U.S.-made HPC systems have been impeded by export control regulations. When Chinese makers of HPC interconnects and high-speed network interface chips are able to support development of HPC systems nearing speeds of 100 petaflops, as *Scientific Computing World* reports, U.S. export controls preventing exports of similar, U.S.-produced components are unlikely to achieve their intended purpose.[242]

Federal programs involved in supporting technical education programs should emphasize HPC-related skills. Broadly, America's HPC community is having trouble attracting sufficient HPC talent, something that goes both for the talent needed to develop exascale HPC systems and the talent in industry to apply HPC to industrial needs to the maximum extent possible.[243] Worker training programs can play an important role in accelerating the so-called "blue-collar computing movement." As HP Enterprise's Stephen Wheat explains:

> We need to turn skilled workers on the manufacturing floor into innovators. If we can provide them with simulation tools to allow them to conceive of something they would like to build, and be able to model it with confidence, we can then drive that level of innovation. So we can retool an entire set of people (most of whom lack engineering degrees) and turn their manufacturing know-how into innovation, thus helping to deal with the skills mismatch for where the jobs are right now in manufacturing.[244]

The Manufacturing Extension Partnership (MEP) should emphasize HPC in relevant engagements. The U.S. Manufacturing Extension Partnership plays a key role in helping America's SME manufacturing base adopt next-generation manufacturing processes.[245] Each state has its own MEP center (or more than one in some cases), and some centers excel in working with SME manufacturers to leverage high-performance computing as a platform for innovation. This should become a core competency for every MEP center in the United States and a screened-for element in the current round of MEP center recompetes.

Furthermore, at the regional level, communities should encourage and facilitate access to shared HPC resources. For example, the Massachusetts Green High Performance Computing Center in Holyoke, Massachusetts, provides state-of-the-art infrastructure for computationally intensive research, principally supporting thousands of researchers throughout Massachusetts.[246] This model could be followed with other regional high-performance computing centers across the country. For example, a university in West Virginia might not be able to afford its own HPC, but could share it with other universities.

CONCLUSION

HPC is helping humanity solve some of its hardest problems.[247] National leadership—both in HPC adoption and production—will remain vital to countries' industrial competiveness, national security, and ability to lead in frontier science. The global race for HPC leadership is intensifying as China, Japan, and the European Union vie to develop exascale supercomputers by 2020. The United States cannot take its leading position in HPC for granted, but must continue to invest to ensure it leads in producing and deploying next-generation HPC systems. It must also enact proactive policies to ensure that existing and future HPC systems reach "the missing middle" so that firms of all sizes can reap the benefits of using HPC systems. HPC has and will remain critical to U.S. economic and industrial competitiveness.

ENDNOTES

1. “High-Performance Computing (HPC),” Techopedia, accessed [March 10, 2016], https://www.techopedia.com/definition/4595/high-performance-computing-hpc.
2. “What is HPC?” The National Institute for Computational Sciences, accessed [March 10, 2016], https://www.nics.tennessee.edu/computing-resources/what-is-hpc.
3. Anand Lal Shimpi, “Inside the Titan Supercomputer: 299K AMD x86 Cores and 18.6K NVIDIA GPUs,” *Anand Tech*, October 31, 2012, http://www.anandtech.com/show/6421/inside-the-titan-supercomputer-299k-amd-x86-cores-and-186k-nvidia-gpu-cores.
4. Pete Carey, “Supercomputers a Hidden Power Center of Silicon Valley,” *The San Jose Mercury News*, May 7, 2015, http://www.mercurynews.com/business/ci_28071868/supercomputers-hidden-power-center-silicon-valley.
5. Don Johnston, “HPC Matters to Our Quality of Life and Prosperity,” *Scientific Computing*, November 11, 2014, http://www.scientificcomputing.com/articles/2014/11/hpc-matters-our-quality-life-and-prosperity.
6. Senate Energy and Natural Resources Committee, “Next Frontier in High-Performance Computing to Usher in New Chapter in Scientific Discovery,” news release, August 1, 2012, https://www.highbeam.com/doc/1P3-2725052741.html.
7. IDC, *High Performance Computing in the EU: Progress on the Implementation of the European HPC Strategy* (Brussels: European Commission DG Communications Networks, Content & Technology, 2015), 17, http://knowledgebase.e-irg.eu/documents/243153/246094/High+Performance+Computing+in+the+EU+-+Progress+on+the+Implementation+of+the+European+HPC+Strategy.pdf/b0adf617-3f50-4a6f-9217-4e0fbb5edd09.
8. Amber Harmon, “HPC Matters: Funding, Collaboration, Innovation,” *ScienceNode,* November 26, 2014, https://sciencenode.org/feature/hpc-matters-funding-collaboration-innovation.php.
9. Data from Top500.org, “The List: November 2015,” http://www.top500.org/lists/2015/11/.
10. Dr. Mark Seager, “Innovation of HPC and as a Driver for Economic Benefit” (presentation, Washington, DC, February 2010).
11. Brian Wang, “China Will Have Two 100 Petaflop Supercomputers in 2016 and They Will Use Domestic Chips,” *NextBigFuture*, September 18, 2015, http://nextbigfuture.com/2015/09/china-will-have-two-100-petaflop.html.
12. Rich Brueckner, “A Closer Look at Intel’s Coral Supercomputers Coming to Argonne,” *Inside HPC*, April 9, 2015, http://insidehpc.com/2015/04/intel-build-coral-supercomputers-argonne-200-procurement/.
13. Wikipedia Commons, “Tianhe-2 in National Supercomputer Center in Guangzhou,” https://commons.wikimedia.org/wiki/File:Tianhe-2.jpg. CC BY-SA 4.0, https://creativecommons.org/licenses/by-sa/4.0/deed.en.
14. Agam Shah, “Japan Chases Title of World’s Fastest Computer With New System,” *PCWorld*, November 16, 2015, http://www.pcworld.com/article/3004551/servers/japan-chases-title-of-worlds-fastest-computer-with-new-system.html.
15. RichReport, “2015 High Performance Computing Trends,” filmed 2015, Intersect 360 Research, YouTube video, 48:31, posted February 9, 2015, https://www.youtube.com/watch?v=wESqpuFAz4M.
16. IDC, *High Performance Computing EU*, 5.
17. Seager, “Innovation of HPC.”
18. Mark Seager, phone interview by Stephen Ezell, ITIF, March 2, 2016.
19. Robert Sorensen, Earl C. Joseph, and Steve Conway, “The New U.S. National Strategic Computing Initiative” (IDC, September 2015), http://www.idc.com/getdoc.jsp?containerId=259254.
20. Seager, “Innovation of HPC.”
21. Tiffany Trader, “Budget Request Reveals New Elements of Exascale Program,” *HPCWire*, February 12, 2016, http://www.hpcwire.com/2016/02/12/obama-budget-reveals-new-elements-exascale-program/.
22. Jim Collins, “ALCF Helps Tackle the Large Hadron Collider’s Big Data Challenge,” *Argonne Leadership Computing Facility*, November 3, 2015, https://www.alcf.anl.gov/articles/alcf-helps-tackle-large-hadron-collider-s-big-data-challenge.
23. “Computing Resources: Mira,” Argonne Leadership Computing Facility, https://www.alcf.anl.gov/mira.
24. Claire Maldarelli, “We Can Now Sequence a Whole Human Genome in 26 Hours,” *Popular Science*, September 30, 2015, http://www.popsci.com/scientists-can-now-sequence-whole-genome-in-26-hours.

25. Reid J. Robinson, "How Big Is the Human Genome?" *Precision Medicine*, January 5, 2014, https://medium.com/precision-medicine/how-big-is-the-human-genome-e90caa3409b0#.mufn3tb96.
26. Elizabeth Armstrong Moore, "Supercomputer Beagle Can Analyze 240 Whole Genomes in Two Days," *CNET*, February 21, 2014, http://www.cnet.com/news/supercomputer-beagle-can-analyze-240-whole-genomes-in-two-days/.
27. Stephen Ezell, "Robust Federal Investment in Scientific Research Critical to Sustaining Moore's Law," *The Innovation Files*, November 22, 2013, http://www.innovationfiles.org/robust-federal-investment-in-scientific-research-critical-to-sustaining-moores-law/#sthash.NLHHh98T.dpuf.
28. Dr. Mark Johnson, Director of the Advanced Manufacturing Office in the Office of Energy Efficiency and Renewable Energy, U.S. Department of Energy, interview by Stephen Ezell, ITIF, April 7, 2016.
29. Carey, "Supercomputers a Hidden Power."
30. Kristin Hansen, "Why the U.S. National Strategic Computing Initiative (NSCI) Matters," *Bright Computing*, September 15, 2015, http://info.brightcomputing.com/blog/why-the-u.s.-national-strategic-computing-initiative-ncsi-matters.
31. David E. Hoffman, "Supercomputers Offer Tools for Nuclear Testing—and Solving Nuclear Mysteries," *The Washington Post*, November 1, 2011, https://www.washingtonpost.com/national/national-security/supercomputers-offer-tools-for-nuclear-testing--and-solving-nuclear-mysteries/2011/10/03/gIQAjnngdM_story.html?wpisrc=nl_tech.
32. Ibid.
33. Ibid.
34. U.S. Bureau of Labor Statistics, "Industries at a Glance: NAICS 334," (accessed April 25, 2016), http://www.bls.gov/iag/tgs/iag334.htm.
35. U.S. Bureau of Labor Statistics, "Computer and Electronic Product Manufacturing: NAICS 334," (access April 25, 2016), http://www.bls.gov/iag/tgs/iag334.htm#workforce.
36. Ross DeVol et al., "Manufacturing 2.0: A More Prosperous California" (The Milken Institute, June 2009), http://www.milkeninstitute.org/pdf/CAManufacturing_ES.pdf.
37. U.S. Bureau of Labor Statistics, "Computer and Electronic Product Manufacturing: NAICS 334"; U.S. Bureau of Labor Statistics, "Occupational Employment and Wages—May 2015," news release, March 30, 2016, http://www.bls.gov/news.release/pdf/ocwage.pdf.
38. The Council on Competitiveness, *Solve* (The Council on Competitiveness, October 2014), 7, http://www.compete.org/storage/images/uploads/File/PDF%20Files/Solve_Report_Final.pdf.
39. "High Performance Computing for Manufacturing (HPC4Mfg) Accelerating Innovation," Lawrence Livermore National Laboratory, accessed April 1, 2016, https://hpc4mfg.llnl.gov/.
40. The Council on Competitiveness, *Solve*, 4.
41. Earl Joseph, Steve Conway, and Robert Sorensen, "IDC HPC ROI Research Update: Economic Models For Financial ROI and Innovation From HPC Investments" (IDC, August 2015), https://hpcuserforum.com/presentations/colorado-sept2015/PublicIDCDOEROIResearchUpdate8.19.2015.pdf.
42. IDC, *High Performance Computing EU*, 4.
43. Ibid.
44. Earl C. Joseph, Chirag Dekate, and Steve Conway, "Real-World Examples of Supercomputers Used for Economic and Societal Benefits: A Prelude to What the Exascale Era Can Provide" (IDC, May 2014), 1, http://casc.org/wp-content/uploads/2014/07/IDCReportRealWorldExamplesOfBenefitsOfSupercomputers.pdf.
45. The Council on Competitiveness, *Solve*, 18.
46. Ibid., 5.
47. Ibid.
48 Thomas Duesterberg, "Manufacturing's New Innovation Labs," *Harvard Business Review*, November 14, 2011, https://hbr.org/2011/11/manufacturings-new-innovation.
49. Nicole Hemsoth, "The Supercomputing Strategy That Makes Airbus Soar," *The Next Platform*, July 22, 2015, http://www.nextplatform.com/2015/07/22/the-supercomputing-strategy-that-makes-airbus-soar/.
50. Ibid.
51. Joseph, Dekate, and Conway, "Real-World Examples of Supercomputers," 20.
52. Ibid.
53. The Council on Competitiveness, "Case Study: Boeing Catches a Lift with High Performance Computing" (The Council on Competitiveness, 2009), 3,

http://science.energy.gov/~/media/ascr/pdf/benefits/Hpc_boeing_072809_a.pdf; Joseph, Dekate, and Conway, "Real-World Examples of Supercomputers," 20.

54. The Council on Competitiveness, "Case Study: Boeing."
55. Hemsoth, "Supercomputing Strategy Airbus."
56. Ibid.
57. Beth Bacheldor, "Airbus Innovates with High-Performance Computing," *IT World*, September 27, 2011, http://www.itworld.com/article/2737507/data-center/airbus-innovates-with-high-performance-computing-.html.
58. Ibid.
59. The Council on Competitiveness, *Solve*, 20.
60. Johnston, "HPC Matters."
61. IDC, *High Performance Computing EU*, 23.
62. Ibid.
63. "High Performance Computing," Rolls-Royce, accessed March 10, 2016, http://www.rolls-royce.com/about/our-technology/enabling-technologies/high-performance-computing.aspx.
64. Ibid.
65. Cliff Saran, "Rolls-Royce Signs Up to Supercomputing On-Demand from HPC Midlands," *ComputerWeekly.com*, July 8, 2015, http://www.computerweekly.com/news/4500249521/Rolls-Royce-signs-up-to-supercomputing-on-demand-from-HPC-Midlands.
66. Johnston, "HPC Matters."
67. Sean Gallagher, "General Motors Is Literally Tearing Its Competition to Bits," *Ars technica,* September 26, 2013, http://arstechnica.com/information-technology/2013/09/general-motors-is-literally-tearing-its-competition-to-bits/.
68. Ibid.
69. Ibid.
70. Ibid.
71. IDC, *High Performance Computing EU*, 21.
72. Joseph, Dekate, and Conway, "Real-World Examples of Supercomputers," 11.
73. Ibid.
74. The Council on Competitiveness, "Case Study: Goodyear Puts the Rubber to the Road with High-Performance Computing" (The Council on Competitiveness, 2009), http://hpc4energy.org/wp-content/uploads/HPC_Goodyear_032009.pdf.
75. Ibid.
76. Ibid.
77. "Success Stories: Goodyear," HPC For Energy, accessed March 7, 2016, http://www.hpc4energy.org/hpc-road-map/success-stories/goodyear/index.html.
78. The Council on Competitiveness, "Case Study: Goodyear High-Performance Computing"; "Success Stories: Goodyear," HPC For Energy.
79. The Council on Competitiveness, "Case Study: Goodyear High-Performance Computing."
80. Ibid.
81. Patrick Thibodeau, "Automotive Supplier Doubles in Size after Adopting HPC," *ComputerWorld*, November 23, 2011, http://www.computerworld.com/article/2499361/high-performance-computing/automotive-supplier-doubles-in-size-after-adopting-hpc.html.
82. Ibid.
83. Stephen Wheat, Director, HPC Pursuits, Hewlett Packard Enterprise, phone interview by Stephen Ezell, ITIF, March 1, 2016.
84. Intelligent Light, "Intelligent Light Receives Prestigious HPC Innovation Excellence Award from IDC," news release, November 20, 2012, http://www.ilight.com/wp-content/uploads/2012/11/HPC_InnovExcell_Award_Announce.pdf.
85. Johnston, "HPC Matters."
86. Joseph, Dekate, and Conway, "Real-World Examples of Supercomputers," 6.
87. The Council on Competitiveness, *Solve*, 20.
88. Oak Ridge National Laboratory, "Oak Ridge Supercomputer Turns the Tide for Consumer Products Research," news release, August 15, 2014, https://www.olcf.ornl.gov/2014/08/25/oak-ridge-supercomputer-turns-the-tide-for-consumer-products-research/.
89. The Council on Competitiveness, *Solve*, 20.

90. Pete Engardio, “Los Alamos and Sandia: R&D Treasures,” *BloombergBusiness*, September 10, 2008, http://www.bloomberg.com/bw/stories/2008-09-10/los-alamos-and-sandia-r-and-d-treasures.
91. Ibid.
92. Johnston, “HPC Matters.”
93. Ibid.
94. IDC, *High Performance Computing EU*, 23.
95. Joseph, Dekate, and Conway, “Real-World Examples of Supercomputers,” 21.
96. U.S. Census Bureau, Statistics of U.S. Businesses (SUSB) Main, (U.S., all industries; accessed March 1, 2016), https://www.census.gov/econ/susb/.
97. National Association for Manufacturers (NAM), “Facts About Manufacturing” (NAM, November 2012), http://www.themanufacturinginstitute.org/Research/Facts-About-Manufacturing/~/media/A9EEE900EAF04B2892177207D9FF23C9.ashx.
98. Michael Feldman, “Minding the Missing Middle,” *HPC Wire*, March 31, 2011, http://www.hpcwire.com/2011/03/31/minding_the_missing_middle/.
99. Addison Snell, “Modeling and Simulation among U.S. Manufacturers: The Case for Digital Manufacturing” (Intersect 360 Research and National Center for Manufacturing Sciences, September 2010), http://www.intersect360.com/documents/NCMS%20IS360%20exec%20sum.pdf.
100. Thibodeau, “Automotive Supplier Doubles after HPC.”
101. Stephen Wheat phone interview, March 1, 2016.
102. Feldman, “Minding the Missing Middle.”
103. Richard McCormack, “NCMS Creates National Network to Promote Advanced Digital Manufacturing,” *Manufacturing & Technology News*, October 19, 2010, https://www.ncms.org/index.php/2010/10/ncms-innovation-network-cited-in-manufacturing-technology-news-2/.
104. “About,” AweSim, accessed April 11, 2016, https://www.awesim.org/en/about/.
105. Stephen Wheat phone interview, March 1, 2016.
106. “About NCSA,” The National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, April 11, 2016, http://www.ncsa.illinois.edu/about.
107. James Barkley, Digital Manufacturing and Design Innovation Institute, phone interview by Stephen Ezell, ITIF, March 4, 2016.
108. Ibid.
109. Tim Shinbara, VP, Manufacturing Technology, Association for Manufacturing Technology, phone interview by Stephen Ezell, ITIF, March 3, 2016.
110. “What Is 3D Printing?” 3D Printing.com, accessed March 10, 2016, http://3dprinting.com/what-is-3d-printing/.
111. Tim Shinbara phone interview, March 3, 2016.
112. Ibid.
113. Patrick Thibodeau, “U.S. Offers Supercomputing Help to Manufacturers,” *ComputerWorld*, February 17, 2016, http://www.computerworld.com/article/3034309/high-performance-computing/u-s-offers-supercomputing-help-to-manufacturers.html.
114. “HPC4Mfg | Solicitation,” Lawrence Livermore National Laboratory, accessed March 10, 2016, https://hpc4mfg.llnl.gov/proposal-call.php.
115. “Advancing Nanoscience,” The Network for Computational Nanotechnology, https://nanohub.org/groups/ncn.
116. IDC, *High Performance Computing EU*, 21–22.
117. Peter L. Singer, “Federally Supported Innovations: 22 Examples of Major Technology Advances That Stem From Federal Research Support” (Information Technology and Innovation Foundation, February 2014), 12, http://www2.itif.org/2014-federally-supported-innovations.pdf.
118. SGI and Intel, “The Inside HPC Guide to HPC in Life Sciences” (SGI and Intel, 2015), 2, http://insidehpc.com/2015/03/the-insidehpc-guide-to-hpc-in-life-sciences/.
119. Steve Conway, “HPC Speeds Innovation for Suppliers and Other SMBs” (IDC, November 2014), http://www.idc.com/getdoc.jsp?containerId=252680.
120. “Heart Disease Facts,” Centers for Disease Control and Prevention, accessed March 10, 2016, http://www.cdc.gov/heartdisease/facts.htm.
121. Joseph, Dekate, and Conway, “Real-World Examples of Supercomputers,” 15.

122. Ibid.
123. Ibid.
124. Courtesy Lawrence Livermore National Laboratory, 2014; Joseph, Dekate, and Conway, "Real-World Examples of Supercomputers," 15.
125. Joseph, Dekate, and Conway, "Real-World Examples of Supercomputers," 6.
126. Ibid., 6.
127. Tiffany Trader, "Why HPC Matters," *HPCWire*, November 19, 2014, http://www.hpcwire.com/2014/11/19/hpc-matters/.
128. Ibid.
129. Joseph, Dekate, and Conway, "Real-World Examples of Supercomputers," 11.
130. Ibid., 1.
131. Ibid., 2.
132. International Supercomputing Conference (ISC), "Klaus Schulten Talks about the Evolution of Computational Biophysics," *Scientific Computing*, March 14, 2014, http://www.scientificcomputing.com/articles/2014/03/klaus-schulten-talks-about-evolution-computational-biophysics.
133. ISC, "Schulten Evolution of Computational Biophysics."
134. Conway, "HPC Speeds Innovation."
135. ISC, "Schulten Evolution of Computational Biophysics."
136. Ibid.
137. PBS, "Interview Klaus Schulten, 2015 National Lecturer - BPS 2015," YouTube video, 5:34, posted February 10, 2015, https://www.youtube.com/watch?v=8x_WL1ViR_0.
138. ISC, "Schulten Evolution of Computational Biophysics."
139. Ibid.
140. Ibid.
141. Merle Giles, "Why We Care About Industrial HPC," *Top500.org Blog*, April 2015, http://www.top500.org/blog/why-we-care-about-industrial-hpc/.
142. John E. Kelly III, "Leveraging a Universe of Big Data" *Innovation* 12, no. 3 (3rd quarter 2014), http://www.innovation-america.org/leveraging-universe-big-data.
143. Ibid.
144. Ibid.
145. Trader, "Why HPC Matters."
146. Tom Groenfeldt, "Supercomputer Manages Fixed Income Risk At JPMorgan," *Forbes*, March 20, 2012, http://www.forbes.com/sites/tomgroenfeldt/2012/03/20/supercomputer-manages-fixed-income-risk-at-jpmorgan/.
147. Carey, "Supercomputers a Hidden Power."
148. Diana Kwon, "Are Europeans Better Than Americans at Forecasting Storms?" *Scientific American*, October 1, 2015, http://www.scientificamerican.com/article/are-europeans-better-than-americans-at-forecasting-storms1/.
149. Richard B. Rood, "To Be the Best in Weather Forecasting: Why Europe Is Beating the U.S.," *The Washington Post*, March 8, 2013, https://www.washingtonpost.com/blogs/capital-weather-gang/post/to-be-the-best-in-weather-forecasting-why-europe-is-beating-the-us/2013/03/08/429bfcd0-8806-11e2-9d71-f0feafdd1394_blog.html.
150. Ibid.
151. Ibid.
152. Ibid.
153. Joseph, Dekate, and Conway, "Real-World Examples of Supercomputers," 23.
154. Ibid., 24.
155. Ibid., 25.
156. Kelly, "Leveraging Universe Big Data."
157. Robert Sorensen, "Takeaways from the Most Recent HPC QView: Settling Dust But Some Changes in the Air" (IDC, March 2016), http://www.idc.com/getdoc.jsp?containerId=US41140516.
158. John Russell, "IDC: The Changing Face of HPC," *HPCWire*, July 16, 2015, http://www.hpcwire.com/2015/07/16/idc-the-changing-face-of-hpc/.
159. Ibid.
160. IDC, "HPC QView: 4Q 2015" (IDC, March 2016).

161. Ibid.
162. "Current List: November 2015," Top 500.org, accessed February 10, 2016, http://www.top500.org/accounts/login/?next=/lists/2015/11/download/TOP500_201511.xls/.
163. IDC, "HPC QView: 4Q 2015".
164. United Nations Comtrade Database, (USA Trade Flows in HS 8471–Automatic data processing machines (computers); accessed April 1, 2016), http://comtrade.un.org/data/.
165. IDC, *High Performance Computing EU*, 24.
166. "Current List: November 2015," Top 500.org.
167. "November 2014," Top 500.org, accessed March 1, 2016, http://www.top500.org/lists/2014/11/.
168. Shannon Tiezzi, "US to Challenge China for World's Fastest Supercomputer," *The Diplomat*, August 4, 2015, http://thediplomat.com/2015/08/us-to-challenge-china-for-worlds-fastest-supercomputer/.
169. "Current List: November 2015," Top 500.org.
170. Rajeshni Naidu-Ghelani, "China Nearly Triples Number of Supercomputers, Report Says," *BBC News*, November 17, 2015, http://www.bbc.com/news/technology-34840805.
171. "China: Two 100 Petaflop Machines within a Year," *Scientific Computing World*, August 25, 2015, http://www.scientific-computing.com/news/news_story.php?news_id=2711.
172. Ibid.
173. Ibid.
174. Tom Wilkie, "Chinese Government Kicks Commercial Companies Overseas," *Scientific Computing World*, August/September 2015, http://www.scientific-computing.com/features/feature.php?feature_id=468.
175. "China Builds World's Fastest Supercomputer," *Xinhua*, June 17, 2013, http://news.xinhuanet.com/english/china/2013-06/17/c_132461860.htm.
176. Naidu-Ghelani, "China Triples Number of Supercomputers."
177. Ankit Panda, "Xi Jinping: China Should Become a 'Cyber Power'," *The Diplomat*, March 4, 2014, http://thediplomat.com/2014/03/xi-jinping-china-should-become-a-cyber-power/.
178. Thesigers, "High Performance Computing in US-China Relations" *Sovereign Data* 1, no. 3 (September 2015): 1, https://pure.royalholloway.ac.uk/portal/files/25502946/SD_2015_09_01_PUBLIC_.pdf.
179. Joshua Alspector et al., "China—A New Power in Supercomputing Hardware" (Institute for Defense Analysis, 2013), 8, https://www.ida.org/idamedia/Corporate/Files/Publications/IDA_Documents/ITSD/ida-document-ns-d-4857.ashx.
180. Ibid., 11.
181. Bill Whyman and Michael Swierczek, "China Wants Its Own Semiconductor Industry" (presentation, Washington, DC, September 3, 2015).
182. Stephen J. Ezell and Robert D. Atkinson, "False Promises: The Yawning Gap Between China's WTO Commitments and Practices" (Information Technology and Innovation Foundation, September 2015), 7, http://www2.itif.org/2015-false-promises-china.pdf?_ga=1.26342712.409366180.1449181307.
183. Whyman and Swierczek, "China Semiconductor Industry."
184. Tiffany Trader, "China's Supercomputing Strategy Called Out," *HPC Wire*, July 17, 2014, http://www.hpcwire.com/2014/07/17/dd/.
185. Thesigers, "High Performance Computing US-China."
186. Mimi Lau, "China's World-Beating Supercomputer Fails to Impress Some Potential Clients," *South China Morning Post*, June 30, 2014, http://www.scmp.com/news/china/article/1543226/chinas-world-beating-supercomputer-fails-impress-some-potential-clients.
187. European Commission, *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions: High Performance Computing: Europe's Place in the Global Race* (European Commission, February 2012), http://www.eesc.europa.eu/?i=portal.en.ten-opinions.23266.
188. Thierry Van der Pyl, "The European HPC Strategy and Actions in Horizon 2020" (presentation, European Commission, February 5, 2014), http://www.apre.it/media/181410/van_der_pyl.pdf.
189. Earl Joseph, Steve Conway, and Bob Sorensen, "High Performance Computing in the EU: Progress on the Implementation of the European HPC Strategy" (presentation, IDC, 2015), 5, https://hpcuserforum.com/presentations/colorado-sept2015/IDC-EU-HPC-Denver-UF.pdf.
190. Ibid.
191. Ibid.

192. Ibid.
193. Robert Sorensen, "EU Consortium Gathers Core Capabilities to Build Exascale HPC Prototype," press release, February 25, 2016; "European Exascale System Interconnect and Storage: Partners," ExaNeSt, accessed April 20, 2016, http://www.exanest.eu/partners.html.
194. Sorensen, "EU Consortium Exascale HPC Prototype."
195. Top 500.org, "Current List: November 2015."
196. Joseph, Conway, and Sorensen, "High Performance Computing EU."
197. Innovate UK, (High Performance Computing (HPC) Market Map; accessed April 11, 2016), https://connect.innovateuk.org/documents/3108763/0/KTN%20Market%20Map%20v6?version=1.0.
198. Joseph, Conway, and Sorensen, "High Performance Computing EU," 9.
199. Ibid.
200. Ibid.
201. IDC, *High Performance Computing EU*, 22.
202. Shah, "Japan Chases World's Fastest Computer."
203. Dr. Bob Sorensen, in-person interview by Stephen Ezell, January 12, 2016.
204. Ibid.
205. Satoshi Matsuoka, "Japan's High Performance Computing Infrastructure (HPCI)," http://www.icri2014.eu/sites/default/files/presentations/Satoshi%20MATSUOKA.pdf.
206. Aloke Tikku and Vanita Srivastava, "India Set to Build Supercomputer Grid," *The Hindustan Times*, October 6, 2014, http://www.hindustantimes.com/india/india-set-to-build-supercomputer-grid/story-gPtagl8gOnKBcSvEwONlAP.html.
207. Tiffany Trader, "India to Launch $730M National Supercomputing Mission," *HPCWire*, October 6, 2014, http://www.hpcwire.com/2014/10/06/india-launch-730m-national-supercomputing-mission/.
208. Ibid.
209. Tikku and Srivastava, "India to Build Supercomputer Grid."
210. Trader, "India $730M National Supercomputing Mission."
211. Top 500.org, "Current List: November 2015."
212. Steve Conway, "Jysoo Lee Rates Progress on Korea's National Supercomputing Plan," *HPCWire*, December 8, 2014, http://www.hpcwire.com/2014/12/08/jysoo-lee-rates-progress-koreas-national-supercomputing-plan/.
213. Ibid.
214. Top 500.org, "Current List: November 2015."
215. "Russia Profile Weekly Experts Panel: 2009—Russia's Year in Review," *Russia Profile*, December 31, 2009.
216. Tiffany Trader, "Russia Aims to Close Supercomputing Gap," *HPCWire*, December 3, 2013, http://www.hpcwire.com/2013/12/03/russia-aims-close-supercomputing-gap/.
217. Ibid.
218. Guennadi Moukine, "Supercomputing Gap Seen as Threat to Economy," *The Moscow Times*, November 8, 2013, http://www.themoscowtimes.com/business/article/supercomputing-gap-seen-as-threat-to-economy/490405.html.
219. Seager, "Innovation of HPC."
220. Singer, "Federally Supported Innovations."
221. Ibid., 6.
222. National Research Council, *Continuing Innovation in Information Technology* (The National Academies Press, 2012), 14, http://www.nap.edu/catalog.php?record_id=13427.
223. Sorensen, Joseph, and Conway, "New U.S. Strategic Computing Initiative," 11.
224. Singer, "Federally Supported Innovations," 12.
225. Council on Competitiveness, *Solve*, 5.
226 Ibid., 12. Only 4 percent disagreed with the statement.
227. Ibid., 13.
228. The White House, "Executive Order—Creating a National Strategic Computing Initiative," news release, July 29, 2015, https://www.whitehouse.gov/the-press-office/2015/07/29/executive-order-creating-national-strategic-computing-initiative.
229. Trader, "New Elements of US Exascale Program."
230. Hansen, "Why the U.S. NSCI Matters."
231. Sorensen, Conway, Joseph, "New U.S. Strategic Computing Initiative."

232. The White House, “Executive Order—Creating a National Strategic Computing Initiative.”
233. Trader, “New Elements of US Exascale Program.”
234. Harmon, “HPC matters: Funding, collaboration, innovation.”
235. U.S. Department of Energy, “FY 2017 Congressional Budget Request—Budget in Brief” (U.S. Department of Energy, February 2016), 83, http://www.energy.gov/sites/prod/files/2016/02/f29/FY2017BudgetinBrief_0.pdf.
236. Hansen, “Why the U.S. NSCI Matters.”
237. Sorensen, Joseph, and Conway, “New U.S. Strategic Computing Initiative.”
238. Bob Sorensen, “Next Steps for the NSCI: Seeking to Ensure a Long and Lively Lifespan,” *HPC Wire*, January 27, 2016, http://www.hpcwire.com/2016/01/27/next-steps-nsci-looking-ensure-long-lively-lifespan/.
239. Robert D. Atkinson and Stephen J. Ezell, *Innovation Economics: The Race for Global Advantage* (New Haven, CT: Yale University Press, 2011).
240. Matthew Stepp et al., *Turning the Page: Reimagining the National Labs in the 21st Century Innovation Economy* (Information Technology and Innovation Foundation, Center for American Progress, and The Heritage Foundation, June 2013), 6, http://www2.itif.org/2013-turning-the-page.pdf?_ga=1.78960212.1554566957.1461032345.
241. Executive Office of the President, Office of Management and Budget (OMB), *Budget of the U.S. Government: Fiscal Year 2017*, (Washington, DC: OMB, February 2016), 27.
242. “China: Two 100 Petaflop Machines,” *Scientific Computing World*.
243. Robert Sorensen, in-person interview by Stephen Ezell, ITIF, January 12, 2016.
244. Stephen Wheat phone interview, March 1, 2016.
245 . Stephen J. Ezell and Robert D. Atkinson, “International Benchmarking of Countries’ Policies and Programs Supporting SME Manufacturers” (Information Technology and Innovation Foundation, September 2011), http://www.itif.org/files/2011-sme-manufacturing-tech-programss-new.pdf.
246. “About MGHPCC,” The Massachusetts Green High-Performance Computing Center, accessed April 12, 2016, http://www.mghpcc.org/about/.
247. Trader, "Why HPC Matters.”

ACKNOWLEDGMENTS

The authors wish to thank the following individuals for providing input to this report: Alex Key, John Wu, and Robert Sorensen. Any errors or omissions are the authors' alone.

ABOUT THE AUTHORS

Stephen J. Ezell is vice president, global innovation policy at the Information Technology and Innovation Foundation, where he focuses on science, technology, and innovation policy as well as international competitiveness and trade policy issues. He is the co-author of *Innovation Economics: The Race for Global Advantage* (Yale, 2012) and *Innovating in a Service-Driven Economy: Insights, Application, and Practice* (Palgrave McMillan, 2015).

Robert D. Atkinson is the founder and president of ITIF. Atkinson's books include *Innovation Economics: The Race for Global Advantage* (Yale, 2012), *Supply-Side Follies: Why Conservative Economics Fails, Liberal Economics Falters, and Innovation Economics is the Answer* (Rowman & Littlefield, 2006), and *The Past And Future Of America's Economy: Long Waves Of Innovation That Power Cycles Of Growth* (Edward Elgar, 2005). Atkinson holds a Ph.D. in city and regional planning from the University of North Carolina, Chapel Hill, and a master's degree in urban and regional planning from the University of Oregon.

ABOUT ITIF

The Information Technology and Innovation Foundation (ITIF) is a nonprofit, nonpartisan research and educational institute focusing on the intersection of technological innovation and public policy. Recognized as one of the world's leading science and technology think tanks, ITIF's mission is to formulate and promote policy solutions that accelerate innovation and boost productivity to spur growth, opportunity, and progress.

FOR MORE INFORMATION, VISIT US AT WWW.ITIF.ORG.