

The National Strategic Computing Initiative

Rudi Eigenmann
University of Delaware

Barry I. Schneider
National Institute of
Standards and Technology

The National Strategic Computing Initiative (NSCI) was introduced in 2015 by then-US President Barack Obama, with the goal of accelerating scientific discovery and economic competitiveness by maximizing the benefits of high-performance computing (HPC) research, development, and deployment. The initiative is led by three agencies—the National Science Foundation (NSF), the Department of Defense (DoD), and the Department of Energy (DoE)—with each agency focusing on the aspects in line with its

respective missions and priorities. There are also foundational and deployment agencies. The foundational agencies—NIST (<https://www.nist.gov>) and IARPA (<https://www.iarpa.gov>)—are responsible for the research and technology required to achieve the NSCI objectives, while the much larger set of deployment agencies are critical to expanding the scientific boundaries that result as a consequence of developing this new HPC capability. This special issue presents four articles that illustrate the state of the field of HPC and computational science, three years after the initiative's launch. The NSCI aims to address five strategic objectives:

1. accelerate delivery of a capable exascale computing system;
2. increase coherence between technology for modeling/simulation and data analytics;
3. establish a viable path forward in the “post-Moore’s law” era;
4. increase the capacity and capability of an enduring national HPC ecosystem; and
5. develop US government, industry, and academic collaborations to share the benefits.

The four articles chosen for this special issue relate to objectives 1 through 4, as argued by their authors.

Please note that an important element of objective 1—the international exascale effort—will be the topic of a forthcoming *CiSE* special issue (January/February 2019) and, therefore, is not covered explicitly here, although we will comment on the exascale thrust at the end of this introduction. Also, a description of the efforts to find a path forward, given the impending end of Moore’s law, appeared in the March/April 2017 special issue of *CiSE* (<https://www.computer.org/csdl/mags/cs/2017/02/index.html>), so is not repeated here.

The first two articles in the current issue describe computational applications in astrophysics and in computational biology, respectively. In “Simulating Stellar Hydrodynamics at Extreme Scale,” Paul R. Woodward, University of Minnesota; Falk Herwig, University of Victoria; and Ted Wetherbee, Fond du Lac Tribal and Community College, present a detailed description of what is required to take a large-scale physics code and extract performance on a current state-of-the-art computational instrument. In “Exascale Computing: A New Dawn for Computational Biology,” authors Christopher T. Lee and Rommie E. Amaro from the University of California, San Diego, argue that exascale computing will add some revolutionary capabilities to computational biology. Both articles demonstrate how advanced computational power can aggressively

push research frontiers by augmenting current computational models of physical or biological systems, leading to new insights not otherwise possible. Developing these application technologies is a key aspect of future exascale computing systems and thus pursues the NSCI objective 1. Both articles also describe how compute power and the processing of large data volumes need to be integrated, which is the goal of the NSCI objective 2.

The other two articles describe the current HPC ecosystem and discuss future needs, in line with NSCI objective 4. Nancy Wilkins-Diehr of the San Diego Supercomputer Center at the University of California, San Diego, and T. Daniel Crawford of Virginia Tech contributed “NSF’s Inaugural Software Institutes: The Science Gateways Community Institute and the Molecular Sciences Software Institute.” John Towns, from the University of Illinois’s National Center for Supercomputing Applications, and lead principal investigator of the NSF’s eXtreme Science and Engineering Discovery Environment (XSEDE), discusses the ecosystem’s current state and looks to the future with his article, “Toward an Open, Sustainable National Advanced Computing Ecosystem.” All of these authors are part of projects that represent some of the largest NSF investments devoted to building elements of a computational ecosystem; however, they were chosen for their scientific expertise and long-term involvement with HPC.

Of course, these four articles are only a small sample of the ongoing work related to NSCI, and they focus almost entirely on academic projects. The use of HPC in the commercial world is another very important aspect, one that is addressed by NSCI objective 5. We are aware of a number of technology transfer programs at NSF. NCSA’s Private Sector Program¹ is another example of how academia and the commercial world are working together. Many of the oil companies, BP being an outstanding example, use HPC for a variety of applications. The same has been true of the auto industry for decades. Without HPC for tasks such as crash tests, it would be difficult to competitively design and deliver today’s automobiles. Companies such as Proctor & Gamble use HPC to help design paper products, including napkins, tissues, and even baby diapers. The list is extensive. The major takeaway message is that HPC plays a key role in the US’s ability to compete economically in a world increasingly dominated by technology.

As mentioned earlier, an important aspect to be deferred to a forthcoming *CiSE* special issue is the actual construction and delivery of exascale machines. A worldwide effort is dedicated to delivering exascale computers in the next five years. This effort faces many new challenges. Moore’s law has been so good to us for so many years that it is difficult to accept the clear evidence of its impending end. We have seen the clock speed of chips decrease instead of increase, and new machines rely more and more on multicore designs, the use of accelerators, and other more recent computing technologies to achieve performance. Power consumption is having a major impact on the design of future supercomputers, and paradigms such as quantum and neuro-morphic computing are becoming the subject of more than just academic interest. In the US, the DoE is the largest player, simply because so many of its projects critically depend on advances in HPC. Needless to say, national security is at the top of the list. While the original DoE plan called for the delivery of an exascale computer in 2018, the deadline has shifted to 2021. Many challenges remain, as recently discussed by Horst Simon² at the NIST NSCI Seminar Series. Capable exascale computing is not simply a matter of building hardware platforms. Users must be able to effectively take advantage of the exascale capability for real-world problems. Doing so requires new algorithms, computational methods, and software architectures that can exploit extreme parallelism. Moreover, many of today’s most important problems require manipulating and/or analyzing very large data sets. This is a very different question than how many floating-point operations a machine can perform on a matrix–matrix multiply. All of these factors are converging, leading to a worldwide renewal of interest in HPC. The Chinese enthusiastically picked up the baton and, for a while, surpassed the US and other nations in building the biggest and fastest supercomputers. The NSCI can be seen as the US’s response, refocusing the nation’s energy on what has always been its key strength—technological leadership.

REFERENCES

1. *Worldwide Best Practices in Partnerships between HPC Centers and Industrial Users*, Nat'l Center for Supercomputing Applications, 22 August 2017; http://www.ncsa.illinois.edu/assets/pdf/industry/Industry_Report_2017.pdf.
2. H. Simon, "Ten Years of Exascale—Are We There Yet?," *NIST NSCI Seminar Series*, 25 January 2018; <https://extras.computer.org/extra/mcs2018050005s1.pdf>.

ABOUT THE AUTHORS

Rudi Eigenmann is a professor in the Department of Electrical and Computer Engineering at the University of Delaware. Contact him at eigenman@udel.edu.

Barry I. Schneider is a staff physicist in the Applied and Computational Mathematics Division at the National Institute of Standards and Technology. Contact him at bis@nist.gov.