Scientific Grand Challenges: Toward Exascale Supercomputing and Beyond

Vladimir Getov, University of Westminster

Scientific frontiers—astrophysics, whole genome sequencing, tectonic modeling, and so on—demand faster and bigger computers to analyze an avalanche of data and advance our knowledge. A quest for answers to these grand scientific challenges is the main motivation behind developing and building exascale supercomputers and beyond.

Substantially improving the speed with which we solve scientific problems has been the driving force of computational design and development since the dawn of computing. Indeed, John Vincent Atanasoff established electronic digital computing's core design principles and built the ABC machine in 1939 to more quickly and automatically solve systems of linear equations in physics. Several years later, John von Neumann proposed the stored program computer architecture as he tried to find a faster way of solving computational problems in nuclear physics. Since then, a new computer simulation–based approach has emerged and is increasingly adopted as one of the most successful modern methods for experimental scientific discovery.

The term "grand challenges"—originally introduced in 1989 by Nobel prize–winning physicist Kenneth G. Wilson—was coined and quickly adopted to describe major questions for which computers could apply revolutionary new methods to advance experimental scientific discovery. These methods included computer simulations and virtual environments requiring much more computational power than had been currently available.[1] Our ability to understand and address the nature of these grand challenges has been rapidly developing due to advances in computing speed and performance at a much higher scale and quality level. These advances are possible largely due to coordinated international efforts in recent years,[2] including progress in several global initiatives bringing together the US, China, Europe, Japan, and others.

Success in computational sciences is largely defined and driven by the rapid development of novel technologies, leading to the creation of complex and powerful high-performance computing facilities that also provide access to huge datasets and high-throughput communications. Sophisticated scientific instruments and systems, such as giant electronic microscopes, nuclear physics accelerators, or complex medical imaging equipment can be integral parts of extreme-scale computing infrastructures.

Combining these powerful technological forces has opened up new frontiers for exploration across many areas of research. Whole genome sequencing, astrophysics, national security, particle physics, and protein folding are just a small sample of the areas in which we have made rapid strides forward thanks to improved computational power and speed.

Indeed, the field of scientific grand challenges continues to open up. As we achieve additional astonishing research results, we will no doubt find even more areas for exploration. The unprecedented scale and quality of information that future exascale supercomputing could achieve will yield exabyte big data systems and spur on continued performance improvements and attendant research results.

IN THIS ISSUE

This special issue on grand challenges in scientific computing includes contributions that address a few of the main application domains in an attempt to provide an overview of the latest research results in this area. The four articles included in the issue give a flavor of the research activities in different scientific disciplines and, to a lesser extent, the supporting supercomputing infrastructures.

In their article "Numerical Weather Prediction with Big Ensemble Data Assimilation," Takemasa Miyoshi, Keiichi Kondo, and Koji Terasaki address the difficulties in maintaining and using increasingly large-scale data assimilation (DA) to improve numerical weather prediction (NWP). The problem is one of the most important and popular grand challenges, both in computational requirements and as a big data problem, in addition to the fact that it will achieve higher predictive resolution and accuracy for extreme weather events and thus help to minimize weather-related damage and loss of life. The article provides excellent and very accessible description of the state of the art in NWP, and shows the benefits of using a large number of lower-resolution samples for ensemble modeling with DA. The authors describe how this approach can improve predictive ability and help to manage cost in terms of both data size and computation using large-scale modern machines, such as the Japanese K computer used in the study. The authors' most salient point is the argument that more samples produces better predictions; this is achieved through effective use of DA approaches and localization of modeling in the simulations. The use of modern large-scale parallel machines enables scientists to continue to greatly improve NWP capabilities.

In the second article, Dylan Keon, Cherri M. Pancake, and Harry Yeh describe the difficulties in predicting natural disasters like tsunamis and storm waves. In "Protecting Our Shorelines: Modeling the Effects of Tsunamis and Storm Waves," they present their extensive experience in building and using such systems. The article contains a well-written description of a difficult problem—modeling tsunamis and storm surges—which requires a multidisciplinary approach to achieve a workable solution. The distributed data and the computational resources required to process the data pose significant challenges. The authors also carry over the model into the human aspects by examining what influences people's timeliness of response when they receive warnings. The authors cover all of these aspects and describe related work in this area along with examples of existing approaches. The article examines the intersection of high-performance computing (HPC), modern Web services delivery, user interface design, visualization, and the everyday lives of millions of people impacted by large-scale storm and tsunami events. Presenting an excellent example of a computing-enabled engineering practice that is relevant to readers on a human level (in terms of their safety and well-being), the article is also an impressive presentation of how computing theory, integrated across several disciplines, has been translated into best practice.

In "Quantum Molecular Dynamics in the Post-Petaflops Era," Nichols A. Romero, Aiichiro Nakano, Katherine M. Riley, Fuyuki Shimojo, Rajiv K. Kalia, Priya Vashishta, and Paul C. Messina provide a thorough overview of the research field and recent work in molecular dynamics (MD) that uses density functional theory (DFT) to generate the potential energy surface—including recent work in linear scaling and parallel replica dynamics approximations to DFT and MD. The article demonstrates the scaling of these methods to very large numbers of processors. Quantum-mechanical simulation has become a massive research field, although there are considerable challenges in extending such calculations to the exascale level and beyond. Thus the demonstrated efficient use of such large numbers of nodes is particularly impressive for this field of research.

In "The TOP500 List and Progress in High-Performance Computing," Erich Strohmaier, Hans W. Meuer, Jack Dongarra, and Horst D. Simon introduce and summarize supercomputer benchmarking using the TOP500 List and the High-Performance Linpack (HPL) benchmark. The article covers the advantages and limitations of HPL compared to other popular benchmarks like NASPB. It shows the performance trends and limitations toward the petascale level and supports these by looking at the Gordon Bell–awarded applications. The authors were original authors of the TOP500 project, and here they analyze performance trends of past, present, and future supercomputers and related technologies in a solid and measurable way. The article presents a long-term perspective of the Linpack benchmark and TOP500 list. At the same time, the authors are reasonably frank about the current system's weaknesses and suggest alternative benchmarks that could supplement Linpack and the TOP500 list. The most interesting information is the long-term trends of processor core counts and performance.

With numerous activities worldwide, the scientific grand challenges field continues to open new horizons for further work and interdisciplinary research at a global scale. Some of the most important areas that we could not cover in this special issue include simulations of the human brain, flying aircraft, automobile crash impact, cancer metastasis, seismic activity, and many others. Thus, the scientific grand challenges continue to bring us closer to revolutionary advances in state-of-the-art, extreme-scale computer technologies.

We hope you will enjoy reading the articles in this special issue, which could become a catalyst for further advances in the world of grand challenges in scientific computing.

REFERENCES

1. K.G. Wilson, "Grand Challenges to Computational Science," *Future Generation Computer Systems*, vol. 5, no. 2, 1989, pp. 171–189.
2. J. Dongarra et al., "The International Exascale Software Project Roadmap," *Int'l J. High Performance Computing Applications*, vol. 25, no. 1, 2011, pp. 2–60.

ABOUT THE AUTHOR

VLADIMIR GETOV is a professor of distributed and high-performance computing at the University of Westminster, London. His research interests include parallel architectures and performance, autonomous distributed computing, and high-performance programming environments. Getov received a PhD and DSc in computer science from the Bulgarian Academy of Sciences. He is a Senior Member of IEEE, a member of ACM, a Fellow of the British Computer Society, and *Computer*'s area editor for high-performance computing. Contact him at v.s.getov@westminster.ac.uk.