How Big Can You Think?

Challenges at the Frontier

By Dimitri Kusnezov

WHEN I WAS A POSTDOC, A FRIEND TOLD ME THE FAMILIAR ANECDOTE OF HOW A THEORETICAL PHYSICIST APPROACHES RESEARCH: FACED WITH THE PROBLEM OF UNDERSTANDING A TABLE'S STABILITY, THEORISTS FIRST ANALYZE ITS STABILITY WITH ONE LEG AND AN INFINITE NUMBER OF LEGS, AND THEN SPEND THE REST OF THEIR CAREERS studying the table's stability with three legs. Physicists—and scientists in general—work hard to master specialties, achieving recognition for the depth of their focus but seldom for its breadth. Academic programs, funding agencies, and the career path from graduate student to endowed chair reinforce this approach. Having one foot in academia and the other in federal strategic planning and investment provides me with the ability to not only recognize the problems but also facilitate change. (In the past several years, several studies have emerged that support change, including some I've been involved with, but more than a decade of real execution at this frontier provides far more insight into the challenges we face.[1–3]) In my experience, I've found that our approach to large-scale computational science needs serious reexamination.

We currently compute at a scale larger than we could have imagined only a few years ago. We can do science well at fractions of a petaflop, on more than 10^5 processors, with data streaming into the petabytes. In a few years, we'll routinely use millions of processors on single calculations as multipetaflop systems find their way into the broader scientific community. Petaflops today are simply a matter of investment, and our hardware challenge is increasingly shifting to architectures in the exaflops. I no longer view the hardware as the dominant challenge or the driver of innovation. I believe the challenges lie in our approach to computational science's potential. We're at the point where computing is starting to outpace our ability to realize the rewards of computational science, largely due to our traditional approach.

The Role of Simulation

We're entering an era in simulation in which we'll be limited only by our imagination—by the size and breadth of our vision. It isn't too soon to think of the emergence of simulation as an essential economic driver, a means to preserve national security, enhance our global competitiveness, and maintain a leadership role in large scientific enterprises. Through investments we've made in supercomputers during the past decade, we've reached a position where we can apply enormous capability to broad-scale scientific needs, support true innovation in science, and address issues in national security and global economic competitiveness. This period is also marked by a need to complement the conventional academic approach—which focuses on well-defined problem spaces that a few specialized researchers can address—with a new approach that applies large, multidisciplinary teams to simulation of high-consequence problems that are less well-defined.

Several emerging multidisciplinary technical challenges will require innovative thinking to provide solutions with real impact. Examples include national security, entertainment, energy independence, and the challenges of global climate change and health security. Problems in these arenas don't generally have well-defined dynamical equations or precise quantitative questions to answer. For instance, questions with societal origins don't emerge as they do in scientific disciplines, where crispness and measurability are inherent—instead, the starting points are inexact, and the questions are imprecise. They tend to focus on the big picture, such as the effects of aging on weapons stockpiles or of global warming on the economy. Because the answers must inform high-consequence decisions, the simulation outputs must be credible and the uncertainties understood.

Given an unlimited amount of computing, it's natural to think that scaling up the problem of interest is the most interesting thing to compute. But scaling such problems to the computer's limit, although informative, isn't imaginative. The big contributions to innovation, science, and competitiveness will arise not from these types of problems but from challenges that require us to draw on diverse viewpoints and disciplines. Today, we're insufficiently prepared to move in this direction.

Two Problem Classes

Looking broadly at the frontiers of simulation, I see two classes of computing emerging. The first and more familiar is the one grounded in our academic background—using simulation to solve sets of equations to further our fundamental understanding of specific phenomena.[4,5] In this case, the endpoint is typically theory validation or the description of data from solving analytically intractable equations, which is the case for quantum chromodynamics and for looking at the spectrum or behavior of elementary particles.[6] In such calculations, we're concerned about the accumulation of error, the implementation of well-controlled or exact algorithms, and the extent to which we precisely capture the equations and their initial conditions. Because we have adequate control of the theory, our attention is principally focused on implementation methods and their efficiencies.

For this class, we find researchers who devote their careers to carefully moving our understanding forward. It's also marked by individual investigators or perhaps small teams of researchers who support the effort, typically all from the same discipline. Many scientific problems here aren't amenable to direct solution, maybe because, for example, the Hilbert space is growing too fast with dimension or the sampling of the path integrals is plagued with sign problems. These problems aren't unimportant, but I would suggest that solving them is less likely to have societal significance. Supercomputing's evolution certainly helps us solve problems at ever larger scales, but while potentially revealing, it's ultimately limited.

I won't discuss this class of problems further; other works highlight the importance of solving precisely posed problems and analyzing inherent ambiguities in the predictions.[7,8] Instead, let's focus on a second class because of the opportunities it brings, the successes we're seeing, and the consequences of credible simulations for key decision makers.

This second class drives simulations into an arena with less well-defined starting points, no clear Lagrangian or phase spaces, and often no precise measurables or end-points. The outcome's relevance isn't generally focused on whether we increased our fundamental understanding of the microscopic science—rather, the outcomes inform high-consequence, time-critical decisions. I believe that in this class of problems lies computational science's true promise, but this class is also fundamentally at odds with our conventional approach to the field, thus requiring us to move beyond the comfort zones defined by our sphere of expertise and toward broader views of simulation. Problems in this class don't typically have a home department, a set of technical journals, advancement infrastructure, or a funding agency.

High-Leverage Decisions

Whether maintaining a nation's nuclear deterrence and stockpile into the future with a commitment never to gather data again (that is, conduct a nuclear test) or trying to separate the causative effects of global climate change from the various possible drivers, a strong commonality of approaches comes to bear in applying scientific computing to the decision-making process. Here, the questions are high level, with no Hamiltonian to diagonalize, no clear basis set, and, most significantly, no clear technical approach to assessing uncertainties in the simulations. These considerations require critical examination of the outputs, particularly because the consequences of resting decisions on such complex simulations and blindly accepting the results can lead decision makers badly astray. Moreover, simulations shouldn't be pressured by political exigencies. To make good use of simulation results, those engaged in the enterprise must develop the skills to communicate to executive and legislative leaders, as well as managers in industry, the power of paying heed to this class of data and folding it in to the decision-making process. Graduate and professional schools, particularly in the business and management disciplines, would do well to include such training in their coursework.

One of the best examples of how computational science can inform decisions with strong international, economic, and political consequences involves the US nuclear weapons stockpile. In this case, simulation is transformational. Historically, nuclear weapons were designed according to what we could calculate. Most simulations were one-dimensional; only rarely did we explore the effects of two dimensions because such calculations were too machine intensive. Nuclear testing supported us as the final arbiter of our scientific and calculational ability.

Today's landscape is much altered. We don't have the control that we once had over the design space. Devices sitting in the stockpile age: they might corrode, cracks might develop, self-irradiation of radioactive materials can occur, and all the resulting features are 3D and generally of very small length scales. In short, our computational capability has to keep pace with features altered outside of our control. Maintaining a nuclear deterrent is a challenge that spans more than 15 orders of magnitude, is inherently nonequilibrium, and exhibits nonlinear behaviors. The extreme conditions aren't experimentally accessible, and the relevant timescales can range from nuclear/atomic reaction times to weeks, months, or decades, depending on the questions that arise. National laboratories have been ideal settings for such long-term programs, and today, 15 years after the US's last test, we're well on our way to maintaining the nuclear testing moratorium by using cost-effective simulations as surrogates.[9] This has enabled risk-informed decision-making through detailed scientific analyses and now allows broader consideration of national security problems, from nonproliferation to homeland security.

Other problems in this class are data rich, enabled by sensor arrays of various forms or assortments of inverse events. Such cases require robust predictive capabilities—the migration of pathogens or the spread of rodent-carried diseases, for example, might be relevant to understanding the impact of a global climate model. In these cases, discerning true causative events and reconstruction of mechanisms becomes critical. The breadth of expertise that would drive this understanding might range from physical and computer scientists to mathematicians, economists, social scientists, and anthropologists. Yet, beyond mechanisms, we're ultimately interested in the impacts on society or the environment. Integrated activities at this scale can take a decade to mature and must include strong team efforts, from researchers to funding agencies.

Business Models

Computational science's integration into business models is starting to gain momentum as scattered examples emerge on the return on investments. But the timescale is long: with the demise of long-term research and development in the corporate world, the advantage of sustained investment in supercomputers or even code-development teams isn't entirely obvious.

Back in 1993, Goodyear partnered with computational science and engineering groups at Sandia National Laboratories in a decade-long project to develop a virtual toolset to design and test new ideas for tires.[10] In a state of fiscal crisis in 2004, the company leadership laid down the challenge that a distinctive new product needed to be developed that would capture the market and be done in less than 12 months. Using advanced mechanics tools, parallel solvers, material models, and validation, researchers developed an integrated tire design and virtual testing environment to simulate tire response to high-speed rolling, cornering, hitting bumps, and braking during wet, dry, and icy conditions. Goodyear's costs as a percentage of sales decreased by 2.6 percent the year its first fully simulated tires hit the market, and its R&D costs for tire testing and design decreased by 25 percent; a new tire's time to market decreased from two years to as little as nine months, and sales grew from US$15 billion in 2003 to $20 billion in 2005. Success required a heavy lab–industry effort, followed by a strong push from above to align the vision with the old way of doing business, but the benefits are striking.

Although not as integrated in its scope, Chevron's fall 2006 announcement of its discovery of a massive oil field in the Gulf of Mexico was the result of supercomputing.[11] The company's willingness to take a US$125 million risk with a billion-dollar robotic drill-string was informed via high-end simulation. Chevron had looked in this area before, but simulation applied to seismic imaging of 10^{15} points—each requiring a significant amount of processing—finally brought the field into focus through the miles-thick canopy of salt. Real-time adjustment of algorithms discovered an oil field 300 miles long and 100 miles wide.

The Future

The problems that present the greatest challenges require the largest collection of diverse viewpoints because the science involved is always multiscale and hence multidisciplinary. They also typically involve some degree of urgency—the solution's timeliness is coupled to policy, economic, or scientific decisions. Given that they often feature large data sets for which suitable analysis methods are still wanting, ascertaining the veracity of predictions and developing testable hypotheses are central to success. Here, we can build on existing experience in verification, validation, and uncertainty quantification, but placing uncertainties on predictions requires tremendous effort. Because we don't have sufficient existing methodologies, investment here is a priority for the next decade. In most of the relevant problems, we can't conduct full-scale experiments to validate conclusions, so we need an approach to determine what types of experimental tests are required and at which scales (to minimize uncertainty) as well as how to deal with a mix of phenomenology, aleatory, and epistemic uncertainties.

The past decade of experience has taught us many of the elements needed to realize the potential of high-end simulation. The core attribute in the stockpile stewardship program has been to focus multidisciplinary teams, or teams of teams, and direct their energy and intellect to large, specific, technical challenges. These teams have the support of a remarkable computing infrastructure—one that we believe will increase in power and capability by many orders of magnitude in the next 10 to 15 years. We've already seen a difference in our ability to develop techniques for multiscale simulations and the analysis of extremely complex coupled systems. As we expand on what we've learned, our techniques and timeliness will only improve. The tools produced in our mission-driven simulations have provided a decision-making framework that's starting to replace instinct and belief with solid science.

I see a similar benefit for other areas, including climate change, the spread of diseases, viable energy systems, and production and distribution. A coherent effort to address a complex grand-challenge problem is the single most important element for success in harnessing scientific diversity, and the lessons we've learned about how to do this in national security can help enhance the quality of life across the world.

Although academia offers the necessary expertise, it will be hard to develop this model in a university setting. Academic institutions provide a diversity of viewpoints, with departments offering a much broader range of disciplines than national laboratories—researchers in medical, biological, and mathematical sciences can intermingle with physical scientists, economists, and sociologists. However, several factors severely limit such settings: immersion in a single broad-scale problem for a protracted time is fundamentally at odds with traditional faculty advancement, cross-departmental projects that aren't properly pigeonholed into conventional areas of specialization rarely find homes, and funding agencies don't naturally align with this model, so there's no financial leverage to help compel universities to rethink their structures. Moreover, with rewards based on a person's unique contributions and broad-scale simulation as an integrated venture beyond the confines of a department, visionaries will be challenged to find homes. Co-location of disciplines on campus does not, by itself, remedy this.

For the past 10 years, the US National Nuclear Security Administration has been pushing a focused, mission-oriented approach in select universities by funding multidisciplinary and multidepartmental teams to deliver on single computational science problems, sometimes in concert with industry and national laboratories. Such efforts have had a remarkable cultural and scientific impact (see Figure 1) on the field at these sites,[12] but I've found that institutional deans and presidents still don't always see the inherent value of this new thinking as an integral part of their responsibilities.

The past decade has been punctuated by the nuclear security enterprise's hardware successes—in both revitalizing the US supercomputing industry and fielding the most aggressive technologies at remarkable scales. Today, these technologies are more widely available—almost commodity products—and we're beginning to see the large-scale adoption of such systems in the US and abroad. Sufficient demand should carry us all forward into exaflops.

Intellectually, the rewards will come from how we plan our scientific and engineering research in the context of broad-scale simulation. We'll likely begin to see the decline of empirically based phenomenology as the need to postulate more complex models will be replaced by the ability to drill down to the appropriate level of fundamental interactions. The ephemeral glory of fielding supercomputers shouldn't overpower the only real legacy of these systems: the new innovation, insight, and science they help us discover. We're limited only by how big we're willing to think.

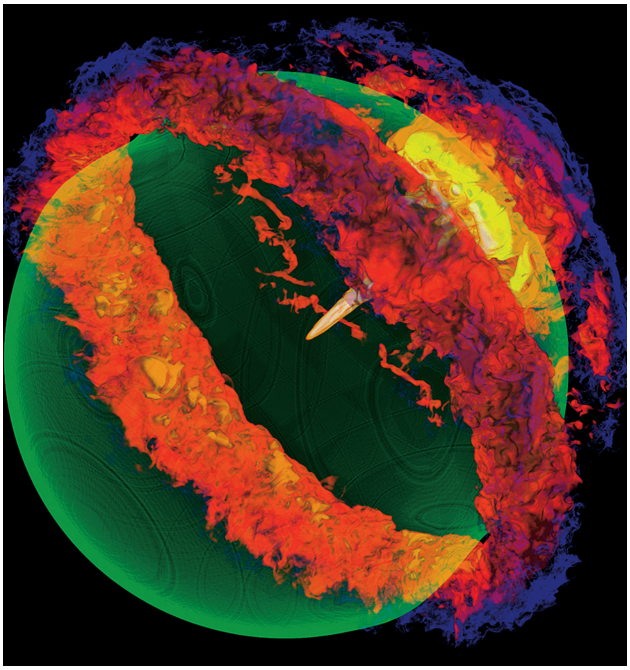

Figure 1. New mechanisms have been discovered for Type 1a supernova explosions at the US National Nuclear Security Administration's Advanced Simulation and Computing FLASH Center at the University of Chicago. This 3D simulation identifies the moment in time when hot ash from a burning bubble breaks through the surface of a star and begins to sweep over it.

References

1. M.R. Benioff and E.D. Lazowska, *Computational Science: Ensuring America's Competitiveness*, President's Information Technology Advisory Committee (PITAC) Report, June 2005; www.nitrd.gov.
2. J.T. Oden, *Simulation-Based Engineering Science*, US Nat'l Science Foundation Blue Ribbon Panel Report, Jan. 2003; www.nsf.gov/pubs/reports/sbes_final_report.pdf.
3. K. Holbrook and D.E. Shaw, *HPC Users Conference Report*, Council on Competitiveness, July 2005; www.compete.org.
4. S. Yip, "Synergistic Science," *Nature Materials*, vol. 2, Jan. 2003, pp. 3–5.
5. O. Yaşar and R.H. Landau, "Elements of Computational Science and Engineering Education," *SIAM Rev.*, vol. 45, no. 4, 2003, pp. 787–805.
6. U.M. Heller, "Recent Progress in Finite Temperature Lattice QCD," *Proc. Science* (LAT2006), vol. 11, 2006, pp. 1–15; http://arxiv.org/hep-lat/0610114.
7. L.P. Kadanoff, "Excellence in Computer Simulation," *Computing in Science & Eng.*, vol. 6, no. 2, 2004, pp. 57–67.
8. R.B. Laughlin, "The Physical Basis of Computability," *Computing in Science & Eng.*, vol. 4, no. 3, 2002, pp. 27–30.
9. R.L. Wagner Jr., "Science, Uncertainty and Risk: The Problem of Complex Phenomena," *APS News*, vol. 12, no. 1, Jan. 2003, p. 8.
10. *Sandia and Goodyear Scientists Develop State-of-the-Art Models*, tech. report SAND2006-4051P, Sandia Nat'l Labs., 2006.
11. S. Conway, "HPC Keys Gulf of Mexico Discovery for Chevron," *HPC Wire*, vol. 15, no. 37, 2006.
12. F. Reddy, "Why Stars Explode," *Astronomy Magazine*, Mar. 2007, pp. 38–43.

Dimitri Kusnezov is the director of the Office of Advanced Simulation and Computing at the US National Nuclear Security Administration and a senior research scientist in the physics department at Yale University. He has published numerous papers in research areas that include theoretical nuclear and many-body physics, nonequilibrium statistical mechanics, group theory, and mathematical physics. Kusnezov has a PhD in theoretical nuclear physics from Princeton University. Contact him at dimitri.kusnezov@nnsa.doe.gov.