Published in final edited form as:

Comput Sci Eng. 2009 ; 11(5): 28–37. doi:10.1109/MCSE.2009.145.

Science and Engineering in the Petascale Era

Thom H. Dunning Jr.,

University of Illinois, Urbana-Champaign

Klaus Schulten,

University of Illinois, Urbana-Champaign

Jeroen Tromp,

Princeton University

Jeremiah P. Ostriker,

Princeton University

Kelvin Droegemeier,

University of Oklahoma

Ming Xue, and

University of Oklahoma

Paul Fussell

Boeing

Abstract

What breakthrough advances will petascale computing bring to various science and engineering fields? Experts in everything from astronomy to seismology envision the opportunities ahead and the impact they'll have on advancing our understanding of the world.

Petaflops computing is here, and sustained petaflops (petascale) computing is just around the corner. As discussed in this issue's introduction on page 7, petaflops computers are in place at the Los Alamos and Oak Ridge National Laboratories, and the Blue Waters sustained petaflops systems is under development at the University of Illinois, Urbana-Champaign. Such systems entail a dramatic increase in computing capability that will, in nearly all science and engineering fields, create breakthrough advances, including

- *Predictive models of biological systems.* Petascale computing will enable more accurate predictions of protein structures; higher fidelity simulation of systems that represent major units of life itself (such as viruses and ribosomes); and more comprehensive organization, comparison, and linkage of the vast amounts of data that researchers are collecting on fundamental life processes.
- *Design of new materials.* Detailed exploration of the structure and energetics of element and molecule combinations will illuminate principles underlying materials-by-design—a long sought goal in materials and molecular science.
- *Design of complex engineered systems.* Modeling the complex interaction of turbulent flow, along with a flame's chemical and physical processes, will create combustion devices that are more fuel-efficient and less polluting. It will also let engineers more fully capture the aerodynamics of flight, resulting in more fuel-efficient, quieter airplanes.
- *Prediction of global and regional climates.* By overcoming the limitations of today's climate models—such as inadequate resolution and simplified models of

physical, chemical, and biological processes—we'll achieve more reliable predictions on regional, continental, and global scales.

- *Emergence of precision cosmology.* Petascale computing will let us construct a virtual universe for detailed comparison with and understanding of the vast amount of data from a new generation of technologically advanced ground- and space-based optical and microwave telescopes.
- *Analysis of complex human behavior.* More detailed models of complex social, economic, and political behavior will help us better understand how pandemics spread and let policymakers and medical-service providers better estimate the effectiveness of different strategies to limit the spread of disease.

Petascale computers are impressive in their own right, providing extraordinary computing power, unprecedented amounts of memory, and massive on- and offline storage, but their true impact lies in the exciting advancements they make possible in many fields. In the following brief essays, experts in several such fields (see the bios for attribution) look ahead to possible petascale-era innovations and what they imply for our understanding of the world.

Seismology

In seismology, researchers will use the next generation of petascale computers for two main purposes: to simulate seismic wave propagation at a global scale and at unprecedented resolution, and to develop and implement adjoint/data assimilation methods. The latter aims to improve earthquake source models in conjunction with models of the Earth's interior; the first tantalizing images of the Southern California crust based on such adjoint tomography are just beginning to emerge.

The quality of tomographic images of the Earth's interior is closely tied to our ability to efficiently and accurately simulate 3D seismic wave propagation on global, regional, and local scales. In the past decade, seismologists have attempted to use asymptotic and numerical methods to address seismology's *forward problem*—that is, given a 3D Earth model, how can we accurately simulate the associated ground motions? At Princeton, we've collaborated with the University of Pau's Dimitri Komatitsch to simulate 3D acoustic, (an) elastic, and poroelastic wave propagation at unprecedented resolution and accuracy by taking advantage of modern numerical methods and harnessing parallel computers (see Figure 1). As a result of our efforts, we can now simulate 3D global seismic wave propagation on a modest PC cluster at periods of 20 seconds and longer, accounting for heterogeneity in the crust and mantle, topography, anisotropy, attenuation, fluid-solid interactions, self-gravitation, rotation, and the oceans. On the next-generation petaflops machines, we'll be able to reach a period of 1 second in global simulations.

The challenge now lies in harnessing these new-found forward-modeling capabilities to enhance image quality for the Earth's interior and the earthquake rupture process—that is, to address the seismological *inverse problem*. The objective here is to go beyond classical “travel time” (time-of-flight) tomography and to use information contained in entire seismic waveforms. On the face of it, this seems like a Herculean task, because such inversions involve hundreds or even thousands of model parameters. In principle, we can numerically calculate a seismogram's sensitivity with respect to the model parameters, but this would require a number of forward calculations equal to the number of model parameters (typically thousands). By drawing connections between seismic tomography, adjoint methods popular in climate and ocean dynamics, and time-reversal imaging, we've demonstrated that we can perform one iteration in tomographic inversions using just two numerical simulations for each earthquake: one calculation for the current model and a second adjoint calculation that

uses time-reversed signals at the receivers as simultaneous, fictitious sources. This has finally opened the door to solving the 3D inverse problem—that is, the problem of using the remaining differences between the data and the simulations to improve images of the Earth's interior.

Molecular Dynamics

Just as light microscopes gave scientists the first glimpse of cells, today the computational microscope provides a fine-grained look at the basics of life. Over the past two decades, as computers have grown more and more powerful, we've been able to conduct larger simulations and make more detailed discoveries. Twenty years ago, it was a triumph to model part of the cell wall with the correct physical properties—a simulation that involved tens of thousands of atoms. Today's terascale supercomputers perform trillions of calculations every second, enabling the simulation of systems of several million atoms. The Theoretical and Computational Biophysics Group that I lead at Illinois, for example, recently simulated an entire life form—the satellite tobacco mosaic virus (see Figure 2). Even for such a tiny form of life, the simulation involved 1 million atoms.

When Blue Waters comes online in 2011, its 200,000+ processors will increase the performance of the nanoscale molecular dynamics (NAMD) code by a factor of 50. With that burst of power, we'll be able to simulate larger cellular structures and to look at biomolecular processes for longer time scales. Today, we can simulate 10 microseconds of a process, but Blue Waters could deliver molecular dynamics' holy grail: the ability to simulate a full millisecond. This timescale would bring into focus a wide range of biological processes.

Atmospheric Science

Sixty years ago, researchers used the Electronic Numerical Integrator and Computer (ENIAC)—the first programmable, stored-memory digital computer—to create the first experimental, large-scale weather forecast. The grid consisted of 270 points, with a mesh spacing of 736 km, covering North America and parts of surrounding oceans. Integrating a single transport equation at one level of the atmosphere, the 24-hour forecast required 24 wall-clock hours and produced results that qualitatively mimicked nature. In stark contrast, today's computing power gives us operational weather forecast models that use 10-kilometer grid spacings on domains approximately $10,000 \times 10,000$ kilometers; they solve about a dozen highly nonlinear partial differential equations representing complex physical processes from below the ground surface to altitudes of about 50 kilometers. Processes of interest include short- and long-wave radiation, cloud and precipitation processes, land-surface exchanges, surface and subsurface water transport and heat conditions, subgrid-scale turbulence, and terrain effects. Forecast accuracy has increased considerably: today's 72-hour results are as accurate as 36-hour results were 20 years ago.

Although today's operational weather forecasts are impressive, their forecast models can't capture the severe local storms that routinely disrupt commerce and cause damage and loss of life. Research clearly demonstrates that grid spacings of 1 kilometer or less are required to adequately represent intense local weather and that we need many model runs (an ensemble) to quantify forecast uncertainty.

We've made considerable progress over the last several years in simulating intense local weather using terascale resources. As Figure 3 shows, we can now capture, in a manner consistent with detailed mobile radar observations, the entire 4D evolution of deep convective storms (including tornado development) with model grid spacings down to approximately 10 meters. However, even with idealized atmospheric environments and a

relatively limited treatment of physical processes, such experiments require several days of wall-clock time using a few thousand processors and produce output that can exceed a petabyte.

An even greater challenge is assimilating real observations into fine-scale models—such as using ensemble Kalman filtering with Doppler weather radar data—to yield realistic initial conditions for an actual prediction. Non-real-time experiments on terascale systems have even demonstrated the potential for predicting tornadoes up to two hours in advance. For example, Figure 4's forecast assimilated all available observations, including those from nearby next-generation Doppler radars (NEXRADs) and experimental radars from the US National Science Foundation Engineering Research Center for Collaborative Adaptive Sensing of the Atmosphere.

Consequently, simulating intense weather (by recreating past events, often via repeated experimentation with parameter tuning) and predicting it (by initializing models with current observations and predetermined model configurations) represent primary challenges for high-performance computing. Such simulations and predictions are a key driver for developing petascale systems such as Blue Waters.

Although scaling to petascale presents great opportunities, it also poses considerable challenges. We must significantly rethink our numerical solution techniques and possibly retool existing codes to take better advantage of emerging many-core architectures. Atmospheric prediction codes tend to be severely memory-bound, and it's considerably difficult to keep processing units fully fed; we thus need new approaches to utilize limited cache. With the domain decomposition strategy that scientists typically use for 2D (x, y) parallelization, load imbalance is often a major issue due to the disparate nature of physical processes occurring across the domain. High-resolution ensemble-based data assimilation involves large data movement, which is usually achieved through concurrent disk I/O. This will surely become a major bottleneck; a possible solution is to move most data through internodal networks, bypassing disk I/O as much as possible. For very large problem sizes, built-in fault tolerance becomes critical. Finally, observations needed for model initialization and forecast verification are generally incomplete on the physical scales that petascale systems make possible. This makes advanced data assimilation/data fusion techniques ever more important. At the same time, we must revise the treatment of key physical processes—including cloud microphysics, cloud radiation interaction, and fine-scale surface forcing—for the fine-scale predictions that petascale systems enable.

For storm-resolving models, the goal is to use their pinpoint forecasts to issue public warnings. This stands in sharp contrast to today's warnings, which are based on radar indications or spotter reports of storms and incipient or existing tornadoes, thus providing relatively short lead times to protect life and property. The model-based "warn on forecast" concept is a dramatic paradigm shift. It will require petascale systems for practicable demonstration and operational implementation—both for data assimilation and for generating multiple forecasts in an ensemble framework to quantify forecast uncertainty. If viable, warn-on-forecast could increase the lead-time for tornado warnings from the current US average of about 14 minutes to as much as two hours. Interestingly, such long lead times might vastly change the public response to weather warnings and lead social and behavioral scientists to engage in studying the warn-on-forecast concept.

Cosmology

Cosmology, one of the oldest subjects of human inquiry, has entered a new phase driven by three phenomena. First, vast amounts of quantitative data are available from ground- and space-based observatories that inventory the universe, from the vicinity of our own galaxy,

back to the era when the cosmic background radiation was released. Second, detailed quantitative theories describe the initial spectrum of Gaussian perturbations and let us calculate how these would grow approaching the current epoch. And finally, large-scale computational resources, complex physical modeling, and rapidly advancing computational algorithms let scientists

- quantitatively compute the nonlinear development of structure using these new quantitative theories, and
- compare the results to our rapidly developing knowledge from these observatories.

This process has been incredibly successful—so successful that we now know that the basic cosmological model is essentially correct. That is, we match innumerable, large-scale observational tests to a level of accuracy so that we can routinely use the comparisons between computed models and observed data to refine the accuracy of the cosmological model's input parameters. Facilities, physics, algorithms, and codes in the terascale era have let us say definitively that matter arrays itself—acting largely, but not entirely, through the force of gravity—into a “cosmic web” and, at the intersection of filaments, galaxies are formed at roughly the epochs, locations, and masses at which we find them. These large-scale simulations can say almost nothing about the internal distributions of gas, stars, dark matter, and so on within galaxies. Simultaneously, focusing on individual regions, we're beginning to be able to make initial simulations of specific galaxies and to include the range of physics needed so that we can tell whether our results match the real world.

But in the extreme, nonlinear limit (the density can be $>10^6$ times the global mean density), where the observables are the interior structures of the familiar galactic systems that have made up the visible universe since Hubble's time, lack of computational power has limited us to very rough and approximate treatments. Although the results of this “semianalytic” modeling have been reassuring, they are so constrained to fit the data that the reasoning behind them is somewhat suspect. Real *ab initio* hydrodynamic simulations using well-founded atomic physics and radiative transfer are just beginning to reach the level at which we can compare the outputs to reality's complex nature.

These improvements have technical consequences: we must compute many physical variables at each spatial position. In addition to the usual hydrodynamic variables of density, temperature, and velocity, we must store and update several variables for the gravitational field, chemical composition, and characteristics of the background radiation field, making the problem intrinsically “memory intensive.” And we certainly can't yet perform accurate calculations for a range of galaxy masses embedded in the full range of cosmic environments. In short, we can't reproduce the real universe or even determine if our current physical model for cosmology is consistent with the real universe of galaxies.

Also, because the input of *feedback*—the mechanical, thermal, and electromagnetic energy consequent to star galaxy formation—is driven by what happens in the minute fraction of the global volume within which gas condenses into stars and black holes (necessitating very high spatial and temporal resolution), we need efficient techniques for 4D domain decomposition, so we use very high small-scale resolution only where it's required. In addition, because both gravity and radiation transfer act over large distances, the communications requirements can be daunting. These computational requirements are of course multiplied dramatically when we want to use the new machines coming online to address the outstanding cosmological issues. As a result of these factors, we're putting great effort into algorithmic developments that will satisfy memory, domain decomposition, and communication requirements.

The current state of the art for the fixed mesh codes is $N = 1,768^3$ with full hydrodynamics or $N = 4,096^3$ for dark-matter-only simulations. Particle + hydrodynamics calculations (SPH) at 200^3 are routine, and we're moving rapidly to where we'll be able to perform hybrid computations, which we hope will combine the virtues of tree-particle and mesh approaches. At Princeton, we developed the widely used Tree Particle Mesh (TPM) algorithm for dark-matter-only computations. We're now working on a major upgrade, hydrodynamic TPM, which incorporates full hydrodynamics using mesh approaches on a background grid combined with numerous quasi-independent, embedded smaller regions followed with tree codes. The approach promises to be extremely computationally efficient on the next generation of machines.

So, what can we hope to reach with petascale machines having 100 to 1,000 times the capacity of current hardware?

- Dark-matter-only calculations that encompass the bulk of the visible universe having more than 2,000 megaparsecs box size at a resolution sufficient to reproduce mass scales down to small galaxies.
- Large-scale cosmological hydro simulations that can accurately test whether our models really do produce galaxies of all types at the right epochs, mass ranges, and environments (see Figure 5).
- Detailed hydro simulations of individual galaxies that can be compared with our detailed knowledge of real systems' internal structure (see Figure 6).

The excitement over what we might achieve in the petascale computing era is palpable. No doubt we'll face unforeseen impediments to reaching our goals. However, it does seem reasonable that within a decade we'll know whether or not our current cosmological paradigm can successfully reproduce the familiar world of galaxies discovered by Hubble and his successors.

Aeronautic Engineering

High-performance computing is central to Boeing's businesses. For example, we use tens of thousands of computational fluid dynamics simulations to evaluate designs and systematically explore design improvement possibilities for our aircraft. In designing the Dreamliner alone, CFD simulations reduced the number of wing designs tested in wind tunnels by a factor of seven.

We always use all the computational capability on our machine-room floors to solve the most significant challenges we face. And the bar is always rising; we know in advance how we'll use the next increase in computing capability—there's always more physics to include in the computational model and more design spaces to explore.

With petascale computing power, Boeing will be able to undertake much more complex simulations at much finer fidelity. Petascale simulations could include more extensive aerodynamics and would enable optimization that considers several physical processes, rather than just one. For example, simulations could include wing structure as well as aerodynamics, thus better capturing the complex interplay as the plane's wing bends and influences its aerodynamics, which affects the load and the wing's elastic response. With this computing capacity, Boeing could design better materials, consider structure and airframe, and study their response to dynamic loads.

Biographies

Thom H. Dunning, Jr., (Guest Editor) is director of the US National Center for Supercomputing Applications at the University of Illinois, Urbana-Champaign, where he also directs the Institute for Advanced Computing Applications and Technology and is the Distinguished interests include computational studies of molecular structure, energetics, and dynamics—most recently of a class of molecules exhibiting extended valence (hypervalence). Dunning has a PhD in chemistry from the California Institute of Technology. Contact him at tdunning@ncsa.illinois.edu.

Jeroen Tromp (“Seismology”) is Blair Professor of Geology and Professor of Applied & Computational Mathematics at Princeton University, where he is director of the Princeton Institute for Computational Science & Engineering (PICSciE) and conducts research in theoretical and computational seismology. Contact him at jtromp@princeton.edu.

Klaus Schulten (“Molecular Dynamics”) is the Swanlund Professor of Physics at the University of Illinois, Urbana-Champaign, where he directs the NIH Resource for Macromolecular Modeling and Bioinformatics and co-directs the Center for the Physics of Living Cells. Contact him at schulten@illinois.edu.

Kelvin Droegemeier (“Atmospheric Science”) is Regents’ Professor of Meteorology and Weathernews Chair at the University of Oklahoma, where his research focuses on the numerical simulation and prediction of severe storms. Contact him at kkd@ou.edu.

Ming Xue (“Atmospheric Science”) is a professor of meteorology at the University of Oklahoma, where he is director of the Center of Analysis and Prediction of Storms. He is also the Analysis and Prediction Thrust leader at the US National Science Foundation Engineering Research Center for Collaborative Adaptive Sensing of Atmosphere. He was the principal developer of the Advanced Regional Prediction System (ARPS) model. Contact him at mxue@ou.edu.

Jeremiah P. Ostriker (“Cosmology”) is director of the Princeton Institute for Computational Science and Engineering, where he works on the formation of galaxies and the growth of massive black holes. Contact him at ostriker@princeton.edu.

Paul Fussell (“Aeronautic Engineering”) is senior manager of mathematical modeling at Boeing Applied Math. Contact him at paul.fussell@boeing.com.

Figure 1. Snapshot from a simulation of the $M_w=9.2$ Sumatra-Andaman earthquake (26 Dec. 2004). Station PALK (Pallekele, Sri Lanka) recorded the vertical component seismogram.

Figure 2. Visualization of the satellite tobacco mosaic virus, simulated in full atomic detail by the University of Illinois, Urbana-Champaign's Theoretical and Computational Biophysics Group. The simulations help scientists determine what factors are important to the virus's structural integrity and how those factors might influence virus assembly inside host cells.

Figure 3.

A tornado funnel cloud, as visualized by the condensed cloud water, simulated at a 25-m horizontal resolution using the Advanced Regional Prediction System developed at the University of Oklahoma Center for Analysis and Prediction of Storms. The $3\text{ km} \times 3\text{ km}$ inset shows a low-altitude horizontal cross section of precipitation intensity (warm colors are associated with high intensity); the spiral precipitation pattern indicates strong rotation, which brings the precipitation toward the vortex center. The image also shows the horizontal wind vectors. Greg Foss of the Pittsburgh Supercomputing Center created the 3D cloud visualization using the simulation data.

Figure 4.

An 80-minute prediction of a tornado observed near Minco, Oklahoma, on 8 May 2007. The prediction assimilates operational as well as experimental radar data from the US National Science Foundation's Engineering Research Center for Collaborative Adaptive Sensing of Atmosphere. The observed Minco tornado is indicated by the blue triangle, while the concentrated black contours of vertical vorticity indicate strong rotation predicted by the model. Color shaping shows the predicted radar reflectivity field. Each tick mark on the axes represents the 400 m horizontal grid spacing used by the simulation. University of Oklahoma graduate student Alex Schenkman performed the numerical simulation.

Figure 5.

The cosmic web. Looking through a 100-light-year cube of a representative piece of the universe from a high-resolution hydrodynamic simulation, we see the distribution of baryons (normal chemical elements) in the “cosmic web” at redshift $z = 2$. Dark spots represent galaxies. The box contains $1,024^3$ cells, allowance for dark matter, and several chemical elements. (Image courtesy of R.Y. Cen and J.P. Ostriker.)

Figure 6. Synthesized elliptical galaxy. Galaxy formed in a very high-resolution simulation (200^3 particles) starting with cosmological initial conditions of a giant elliptical galaxy. Gas, dark matter, and stars are included. (Image courtesy of T. Naab & J. P. Ostriker.)