LAWRENCE LIVERMORE NATIONAL LABORATORY

LLNL-JRNL-420382

Extreme Scale Computing to Secure the Nation

D. L. Brown, J. R. McGraw, J. R. Johnson, D. Frincke

November 18, 2009

SciDAC Review

Disclaimer

Extreme Scale Computing to Secure the Nation

Contributors: David L. Brown and Jim McGraw, Lawrence Livermore National Laboratory; John R. Johnson and Deb Frincke, Pacific Northwest National Laboratory

Since the dawn of modern electronic computing in the mid 1940's, U.S. national security programs have been dominant users of every new generation of high-performance computer. Indeed, the first general-purpose electronic computer, ENIAC (the Electronic Numerical Integrator and Computer), was used to calculate the expected explosive yield of early thermonuclear weapons designs. Even the U. S. numerical weather prediction program, another early application for high-performance computing, was initially funded jointly by sponsors that included the U.S. Air Force and Navy, agencies interested in accurate weather predictions to support U.S. military operations. For the decades of the cold war, national security requirements continued to drive the development of high performance computing (HPC), including advancement of the computing hardware and development of sophisticated simulation codes to support weapons and military aircraft design, numerical weather prediction as well as data-intensive applications such as cryptography and cybersecurity

U.S. national security concerns continue to drive the development of high-performance computers and software in the U.S. and in fact, events following the end of the cold war have driven an *increase* in the growth rate of computer performance at the high-end of the market. This mainly derives from our nation's observance of a moratorium on underground nuclear testing beginning in 1992, followed by our voluntary adherence to the Comprehensive Test Ban Treaty (CTBT) beginning in 1995. The CTBT prohibits further underground nuclear tests, which in the past had been a key component of the nation's science-based program for assuring the reliability, performance and safety of U.S. nuclear weapons. In response to this change, the U.S. Department of Energy (DOE) initiated the Science-Based Stockpile Stewardship (SBSS) program in response to the Fiscal Year 1994 National Defense Authorization Act, which requires, "in the absence of nuclear testing, a progam to:

1. Support a focused, multifaceted program to increase the understanding of the enduring stockpile;
2. Predict, detect, and evaluate potential problems of the aging of the stockpile;
3. Refurbish and re-manufacture weapons and components, as required; and
4. Maintain the science and engineering institutions needed to support the nation's nuclear deterrent, now and in the future".

This program continues to fulfill its national security mission by adding significant new capabilities for producing scientific results through large-scale computational simulation coupled with careful experimentation, including sub-critical nuclear experiments permitted under the CTBT.

To develop the computational science and the computational horsepower needed to support its mission, SBSS initiated the *Accelerated Strategic Computing Initiative,* later renamed the *Advanced Simulation & Computing (ASC)* program (sidebar: "History of ASC Computing Program Computing Capability"). The modern 3D computational simulation capability of the ASC program supports the assessment and certification of the current nuclear stockpile through calibration with past underground test (UGT) data. While an impressive accomplishment, continued evolution of national security mission requirements will demand computing resources at a significantly greater scale than we have today. In particular, continued observance and potential Senate confirmation of the Comprehensive Test Ban Treaty (CTBT) together with the U.S administration's promise for a significant reduction in the size of the stockpile and the inexorable aging and consequent refurbishment of the stockpile all demand increasing refinement of our computational simulation capabilities. Assessment of the present and future stockpile with increased confidence of the safety and reliability *without* reliance upon calibration with past or future test data is a long-term goal of the ASC program. This will be accomplished through significant increases in the scientific bases that underlie the computational tools. Computer codes must be developed that replace phenomenology with increased levels of scientific understanding together with an accompanying quantification of uncertainty. These advanced codes will place significantly higher demands on the computing infrastructure than do the current 3D ASC codes.

This article discusses not only the need for a future computing capability at the exascale for the SBSS program, but also considers high performance computing requirements for broader national security questions. For example, the increasing concern over potential nuclear terrorist threats demands a capability to assess threats and potential disablement technologies as well as a rapid forensic capability for determining a nuclear weapons design from post-detonation evidence (**nuclear counter-terrorism)**. Preventing the spread of nuclear materials and devices (**nuclear non-proliferation**) requires the development of new proliferation-resistant materials and the ability to distinguish between natural and man-made seismic events. Reduction of biological, chemical and cyber threats will require capabilities for massive data analysis and simulation of complex systems (sidebar "Getting Ahead of the Threat: Cyber Security Science at the Exascale Fontier"). Safety analysis for weapons systems and for warfighter protection systems can also be supported by high performance computing (sidebar: "Safety analysis through Multi-Physics HPC Simulation"). Extreme scale computing – in particular for computational simulation and analysis – is the integrating element that will enable us to meet all of these national security challenges. While they will never replace experiment nor the physical intuition of expert scientists, computer codes with increasingly strong science bases will help form and test that intuition in a more fundamental way. Within the next decade, we expect to see the development and deployment of computers and the attendant programming models and software environments at the exascale (10^{18} operations per second and beyond). This computing horsepower will be an essential contributing element to the significant simulation capabilities that are required for the broader national security mission.

Exascale Needs for Stockpile Stewardship

The first ten years of the SBSS program focused largely on understanding the existing weapons stockpile by developing computational tools for predicting, detecting and evaluating problems in the aging weapons systems. The tools developed were sufficient to support "lifetime extension programs" (LEPs) for the weapons, which included the ability to refurbish and re-manufacture weapons and components as required, while maintaining the expertise and institutions necessary to support the nuclear deterrent into the future. Computational simulation has provided SBSS the ability to make scientific analyses without directly relying on underground testing (sidebar: "How as computational simulation is created").

The nuclear stockpile was initially designed using one-(1D) and two-dimensional (2D) codes. These codes were strongly calibrated to the underground test base and while sufficient to "interpolate" the test results were not adequate to extrapolate significantly beyond the conditions of the original tests. To make predictions and develop understanding for conditions beyond those of available test data, the ASC program recognized that significantly better mathematical and physical fidelity would be required. The initial focus of ASC simulation code development was on improving the geometric or dimensional fidelity of the simulations to obtain capabilities not encumbered by the symmetry-assumptions inherent in 1D and 2D simulations. Particularly as one considers hydrodynamic instabilities or material variations that occur due to aging, the phenomena of interest are inherently three-dimensional. The ASC 3D codes now provide the nation a set of tools that can assess and certify the current stockpile.

The ASC program's simulation codes fall into three main categories. The *Integrated Design codes* (IDC) are the main workhorses used by weapons designers and analysts. These are large, complex multi-physics codes that have hundreds of person-years invested in their design and implementation. The IDCs include the classified codes used by designers and analysts to simulate the nuclear safety, performance, and reliability of stockpile systems. These are complex, integrated hydrodynamics, radiation-hydrodynamics, and transport codes for application to Stockpile Stewardship, design and analysis of experiments, general purpose hydrodynamics and radiation-hydrodynamics problems, and analyzing radiation and particle transport problems for a variety of applications. These codes are principally used to understand performance, safety, surety and weapons system response issues. The *Material Property codes* (MPC) are used to develop material property databases that are essential inputs to the integrated design codes. These include codes to compute constitutive models (relations between physical quantities that approximate the response of a material to external forces), codes that implement models for material fracture used for multi-scale modeling, and codes for simulating high-explosive detonations. As computing resources increase, results from these codes continue to enhance the physical fidelity of the IDCs. Finally, the *Research codes* identify future paths for the IDCs and MPCs, developing new physical models and numerical algorithms that have the potential to improve future code capabilities. These are typically single-physics codes, focused on improvement of one type of simulation, e.g. atomistic, turbulence, dislocation dynamics or molecular dynamics, but may also include codes designed to study, understand and improve multi-physics modeling capabilities for the future.

As the demands for mathematical and physical fidelity continue to increase, even the current 3D ASC codes will be insufficient to predict with confidence the safety and reliability of nuclear weapons as they age. As time passes, systems change due to aging or improvements made through the LEPs, and the state of the systems drifts ever further from the conditions they were in when originally tested. Consequently, simulations based on calibration to data become increasingly unreliable. In addition, as the number of weapons is decreased to 1500 as is currently planned, the degree of confidence in each weapon must grow. Ultimately, a *predictive simulation* capability is required that can give science-based predictions with quantified uncertainty without the necessity of calibration to underground tests (sidebar: Uncertainty Quantification). Reaching this goal will require significant additional developments in simulation code capability that increase the use of first principles calculations to determine physical parameters, address multi-scale phenomena and rigorously quantify uncertainty. Full simulation of inherently 3D phenomena such as turbulence and material failure will be required. Improved physical models (for example, not based on interpolation of experimental data) will be required for key areas such as energy balance and radiation damage. Concurrently, it is expected that three to five orders of magnitude improvement in computer performance will be required to consolidate the understanding of weapons physics sufficiently for certification without reliance on underground tests. This will require not only the development of new computer hardware platforms at the exascale, but also the software environments that support them. Simulation codes at this scale will be able to perform 3D predictive simulations with uncertainty quantification, including the study of wide-ranging weapons safety and security scenarios.

Computational Support for Broader National Security Issues

For more than a decade, the ASC program has been at the forefront of HPC in development of computer platforms and computational and scientific tools to support the stockpile stewardship mission. These capabilities are not, however, all limited in application to stockpile-related problems. Applications of ASC computational tools and capabilities benefit broader national security missions and other problems of national interest, and as the ASC program matures, there will be increasing dependence on its capabilities from the larger national security community. Areas outside stockpile stewardship for which ASC capabilities and tools are applied include nuclear forensics, nuclear counterterrorism, seismic modeling for nonproliferation, radiation hardening and survivability for microelectronics, vulnerabilities of critical infrastructure, weapon effects and foreign assessments. In the past, ASC resources have also been used in the space shuttle Columbia investigation and in missile defense simulations.

Nuclear Counterterrorism. The nuclear threat has posed two main challenges, namely the need to eliminate the threat of a nuclear attack, and the minimization of the impact of such an attack should one occur. Furthermore, the nuclear threat has evolved in recent years. While the development of the U.S. stockpile was originally in response to the existence of large hostile nation-states, we are now concerned with proliferation of nuclear technology to other nation-states as well of to non-state actors (e.g. terrorist organizations). Not only do we require technology to develop and maintain nuclear weapons, but also to counter potential terrorist threats. A nuclear counterterrorism program must be able to detect the existence of nuclear weapons elsewhere and to determine the source of nuclear materials ("attribution"). Exascale computing capabilities will be required for the assessment of threats, the assessment of weapons disablement technologies, and for a rapid response capability to determine a nuclear weapons design from post-detonation debris should an attack occur.

There are a number of factors that drive the computational technology needs for nuclear counterterrorism. Analysis of a potential proliferant or improvised designs can be significantly more computationally intensive than modern stockpile systems since the systems are expected to be much larger physically and hence will require larger computational meshes to simulate (see sidebar "How a computational simulation is created"). In addition, they often contain artifacts that necessitate 3D modeling and the diversity of potential designs is extensive. The need to be more proactive in analyzing the risk posed by all likely potential designs could require thousands of simulations. The assessment of disablement technology effectiveness is another area that will require large scale computing. For example, a single simulation of a shape charge jet using a realistic target will require an estimated 5 petaflop-days. As hundreds of such calculations would be required to study all potential weapons designs and disablement technologies, this is easily an exascale challenge.

In the event of a nuclear detonation, there will be tremendous pressure to provide as much information as quickly as possible on the type of device detonated to help enable identification of the perpetrators. This involves a backward engineering exercise where post-detonation diagnostics are matched to possible designs, again requiring numerous simulations. Using an example time frame of three days for turn-around of such an analysis, one can show that we quickly exceed the capacity of the current ASC computing capability as we move from running all simulations in 1D to an all 2D suite, with an all 3D suite extending about 6 orders of magnitude beyond current capacity.

Nuclear non-proliferation is another area where HPC simulation technology will be essential. In evaluating intelligence information, we must be able to determine existing and emerging weapons designs of our adversaries, locate and quantify existing nuclear materials, and evaluate the proliferation impact of the 'civil' nuclear renaissance (i.e. the development of nuclear power production technology in hostile states). In order to monitor possible nuclear tests in other parts of the world, we must be able to differentiate nuclear tests and other man-made explosions from naturally occurring events such as earthquakes. Computational seismology is an essential technology that supports this determination. Modern simulation tools and high performance computing allow for the modeling of explosion-generated seismic waves from the detonation underground, through the 3D earth to the observing seismic station. Current terascale applications are routine and petascale calculations have been demonstrated. However, to model seismic wave generation and propagation with sufficient fidelity for modern national security requirements, the development of code coupling and giant advances in computational power will be required. (sidebar: "Detonation-to- Detector, Simulation for Nuclear Explosion Monitoring").

Summary

Computational simulation has been and will continue to be an essential element of U.S. national security programs. The ASC program has been driven since its inception by the need to ensure the safety, reliability and performance of the nation's nuclear weapons stockpile without nuclear testing. This has been done by emphasizing the development of high-fidelity, three-dimensional simulation codes, the creation and deployment of the required computational capabilities and supporting infrastructure. As computational simulation has become an essential element of SBSS, the capabilities enabled by ASC have put computational science on an equal footing with theoretical and experimental science as a tool for studying basic issues of weapons science and for scientific discovery. While the scientific base of ASC simulation capabilities has progressed significantly, the next decade of the program will increasingly emphasize the development of a deeper understanding of the underlying science, a continual improvement in the theoretical models that provide the scientific basis for the weapons codes, and an increased and more quantitative understanding of their limitations.

In the past, nuclear weapons designers had access to full-system experiments (i.e. underground nuclear tests). Designs relied on codes containing physically incomplete models that were calibrated to measured data using adjustable parameters and combined with many simplifying assumptions. Now that the mission is no longer design of new weapons, but understanding the aging of existing weapons systems, detailed three-dimensional simulation and analysis has become a requirement. A science-based predictive capability is an essential replacement to extrapolation based on calibration and expert judgment. The new models must provide the capability to credibly extrapolate from past underground tests into new physical regimes. An essential component of this technology is the ability to calculate, measure and understand the uncertainty in the predictions. The combination of a reduction in phenomenology, replaced by higher fidelity science and the accompanying quantification of uncertainty drives the need for the continued advancement of the computing hardware towards and beyond the exascale.

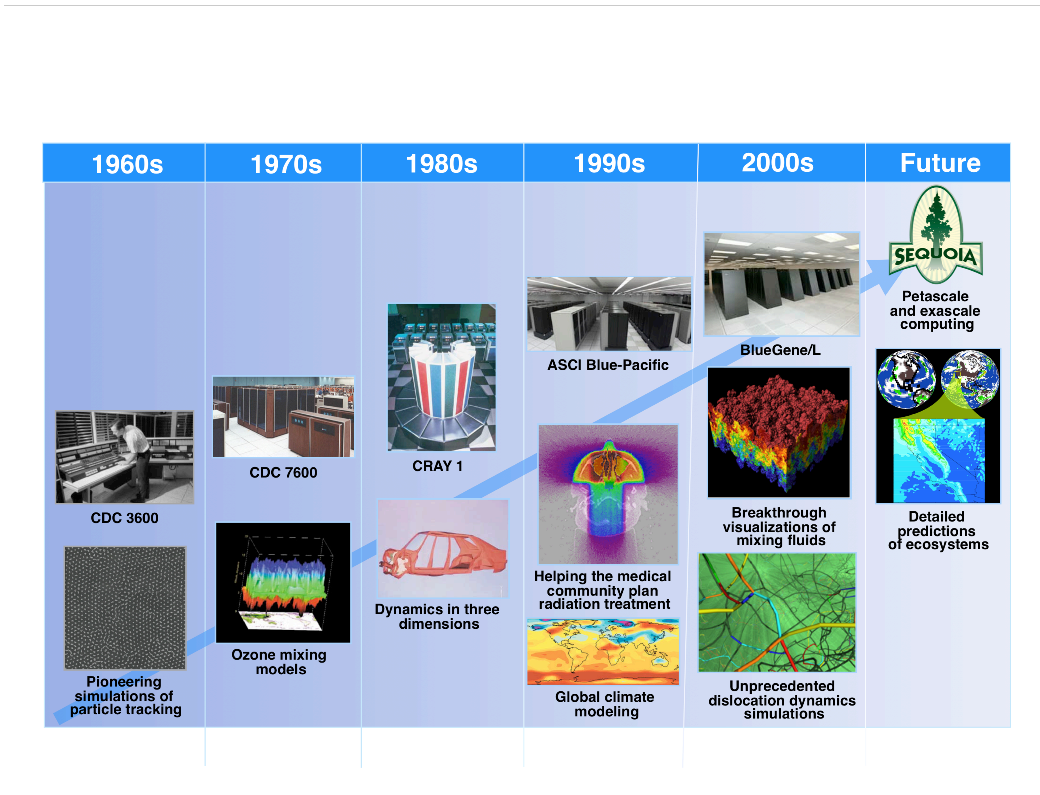

Caption: Each new generation of computing hardware expands the horizons computational science.   From early simulations involving 870 particles that contributed to the understanding of matter, to recent turbulence simulations on BlueGene/L involving over 10^9 particles, each of the simulations pictured has added new understanding and insight of complex physical phenomena.

Sidebar:
History of ASC Program Computing Capability

The *Accelerated Strategic Computing Initiative* was initiated in 1995 as part of the DOE's Science-Based Stockpile Stewardship (SBSS) program to develop the computational science and the computational horsepower needed to support the nuclear weapons mission. It was designed as a "balanced and integrated program of computational simulation, fundamental scientific research and improved nuclear and non-nuclear experiments dedicated" to support understanding of the enduring stockpile, prediction, detection and evaluation of potential problems in the *aging stockpile,* essential support of Lifetime Extension Programs (LEPs) and to maintain the science and engineering institutions needed to support the deterrent into the future. The Advanced Simulation and Computing Program (ASC, the successor to ASCI) had the role of providing the high-end computational simulation capabilities, both hardware and software, to enable the efforts within the SSP. In the absence of Underground Testing (UGT), ASC became the *integrating element* of the SSP.

The previous approach used underground nuclear tests to validate theory, providing the science-based analysis of nuclear weapons reliability, performance and safety. With the cessation of underground tests, validation under SBSS required expertise from weapons design scientists and engineers slated to retire within ten years, and hence, while previous computational requirements had been fulfilled by one and two-dimensional simulation codes, ASC would require three-dimensional codes with greater fidelity and physical realism to reproduce, replace and improve upon the human expertise that would be lost. The challenge was immense: an increase of nearly five orders of magnitude in computer performance combined with the development of new three-dimensional simulation codes to provide high-fidelity scientific predictions would be required to meet the initial ten-year objectives of the program.

Lawrence Livermore, Los Alamos and Sandia National Laboratories, the ASC "tri-labs", met this challenge by partnering with the computer industry to produce and acquire a new class of computers—large-scale machines that relied not only on faster processors to increase the speed of computation, but also on massive parallelism in which a very large number of operations could be performed simultaneously on a computer made from many processors connected together in a large closely-coupled network. The first decade of ASC high-performance computers culminated with the delivery in 2005 of the 100 TeraFLOPS IBM "Purple" computer, together with the concurrent development and deployment of significant three-dimensional simulation capabilities, augmenting and replacing the "legacy" one-dimensional and two-dimensional simulation codes that had been used in the past.

At the same time that Purple was delivered to the ASC program, the 360 TeraFLOPS IBM BlueGene/L was also acquired. This computer was the first of a new generation that anticipated the future need to deliver high-performance computational capability with far less power per FLOPS than its predecessors. BlueGene/L accomplished this feat by using slower (and hence less energy-demanding) processors, combined with a significantly larger total number of processors. While the peak performance of this machine topped the world's supercomputers for several years, BlueGene/L was significantly more difficult to program, and the development of new, more effective programming models and supporting software continues to be a challenge for this class of extreme scale computers.

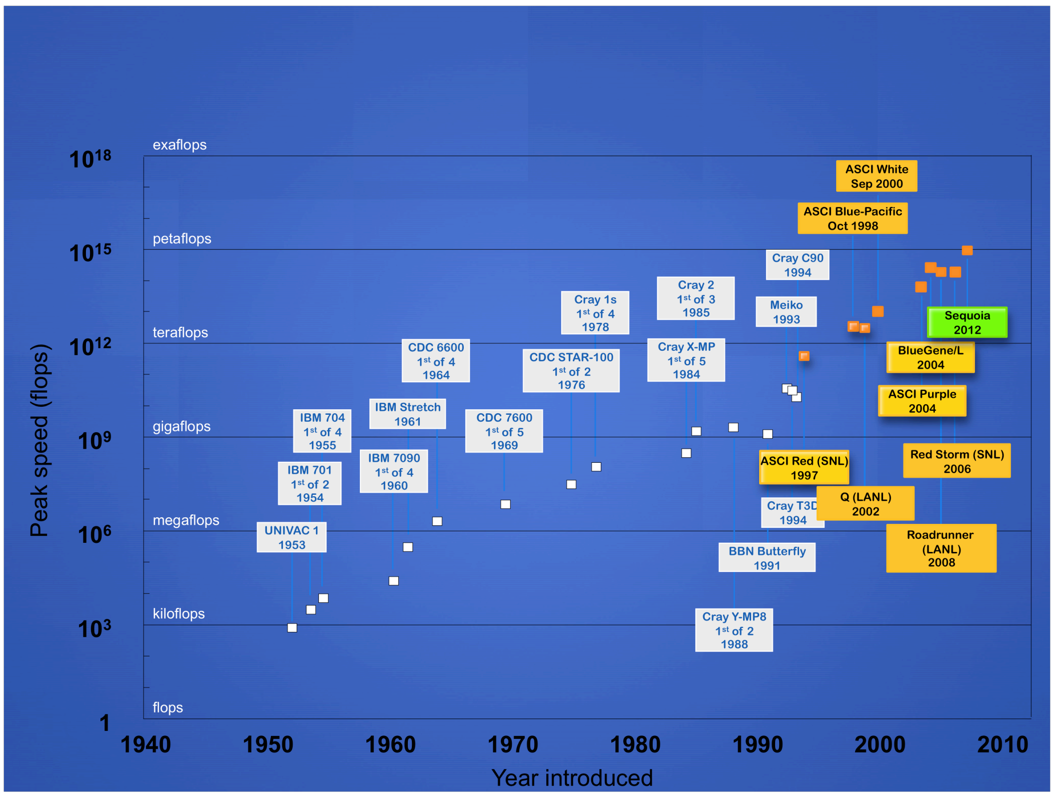

Figure: The ASC program was responsible for an increase in supercomputer performance of nearly five orders of magnitude in just 10 years. The peak performance of each of the ASC platforms (in orange) is shown plotted against delivery date. For reference, the exponential growth of supercomputer performance deployed for national security programs at Lawrence Livermore National Laboratory from 1953 until the beginning of the ASC program are also shown (white).

Sidebar

Cyber Security Analysis – a rapidly emerging driver for exascale computing

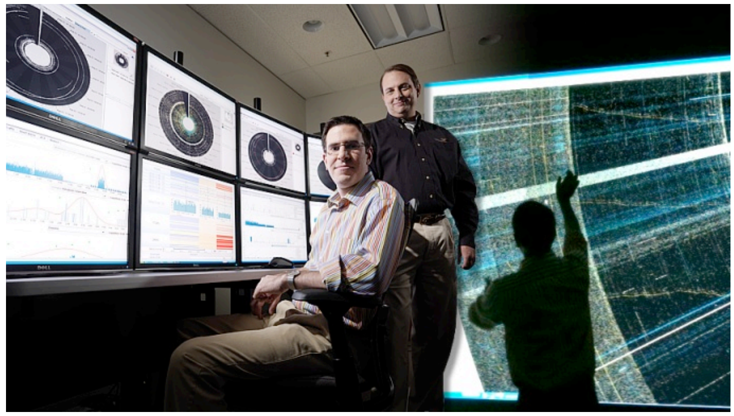

Figure 1 Bill Pike and Jeff Mauth analyzing network traffic with PNNL's Traffic Circle

Threats to the nation's cyber infrastructure, i.e. the large networked information systems on which business, energy and defense infrastructures increasingly depend, have become a serious national security concern. Understanding and addressing these threats will require computational capabilities that rival or exceed those for simulation and modeling of the physical systems of importance to national security.

While modern cyber security operators at the high end tend to work with terabytes or petabytes of data, increased needs for complex computations and continually expanding datasets are already pushing the cyber security field towards exascale computing. For example, a single month of cooperative enterprise perimeter data collected by one government agency is known to contain tens of millions of distinct IP addresses – flows between source and destination – representing communication between one party and another. A consolidated 2D matrix representation of the implied connectivity graphs for this data would have a few quadrillion elements for each value of interest (e.g. first packet arrival time, port used, total bytes transferred). This simple connectivity matrix would be sparse (a few billion entries). A richer cybersecurity flow representation would involve representation of additional information: intermediate hops within a flow (devices besides the source/destination pair with access to the flow), dynamic and static address translation (allowing one IP address to map to different computers at different times, or allowing one IP address to represent multiple computers at a given time), and the need for detailed supplemental data about activities within any or all of the computers or devices along the flow, including intermediate hops. Combining these factors yields a complex multidimensional database of exascale size.

While real-time and near-term forensic analysis is important, so is longer-term retrospective analysis. Retrospective analysis extends the complexity through not just the address space, but across time. As an example, data collected at one-second time intervals for a sophisticated eighteen month cyber attack against the enterprise mentioned above would introduce over 25 billion new edges and nearly a million time steps. For either deep forensic analysis or “real-time” (aggregated to 1 minute intervals) cyber analysis and response, even the most basic analytical computations on a dynamic graph of over a quadrillion elements and a million time steps is intractable on today’s leadership class machines. And again, this is only considering network flows themselves, and not the critical interplay of activities within the hosts involved in the communication.

Emerging trends are to consider predictive as well as retrospective approaches to cyber defense. An example of a predictive approach is to use modeling and simulation to determine the effect of a new cybersecurity remedy on an Internet that includes IPv6 (Internet Protocol v6) traffic. IPv6, described as the successor to IPv4, has several advantages, among them increasing the directly addressable nodes on the Internet to about 3.4×10^{38}. Even with a simplified discrete network model that instantiates 340 “virtual machines” per physical node, simulation of this system would require a supercomputer with 10^{36} nodes –well beyond the exascale range.

Advancing cybersecurity analysis for very large scale and collaborative networks, particularly if the goal is to move from a “catch and patch” approach to one that is more proactive and predictive, will require fundamental advances in our understanding of the structure, mechanics, and dynamics of complex cyber networks. Progress is hindered by the enormous scale of the problem and the lack of scalable computing resources to perform high fidelity modeling, simulation, and analysis. Both operators seeking to protect systems, and researchers seeking to understand solutions, face computational complexity barriers. The ability to model, simulate, and analyze these problems is an important scientific challenge for global security – and, given the scale of the problem, one that will require extreme computing resources.

Sidebar:

Safety analysis through multi-physics HPC simulation

Safety analysis is an important element of national security. It is essential to know that the systems we design containing high explosives and propellants are inherently safe should an accident occur. We must also assure that the protection systems we provide our warfighters are maximally effective against enemy insult. High performance multi-physics simulations can provide critical insights for these questions.

One example of a simulation tool used for safety analysis is the ALE3D code at Lawrence Livermore National Laboratory, which enables simulation for problems that must couple physics across a wide range of length and time scales. Example applications include high-explosive modeling from detonation to thermal cook-off, modeling of composite armor, penetrator mechanics, conventional weapon safety and performance, metal forging, concrete fracture rebar response, understanding how to protect of our forces through analysis of urban canyon blast, and safety systems such as armored vehicles and personal armor designed to protect against traumatic brain injury.

The figure demonstrates one small example of how we're continuously improving the fidelity of the physics, and the size and complexity of the systems that can be modeled, thanks to the rapid increase in computing power provided through programs like ASC. The ultimate goal – much like that of the stockpile stewardship program, is to provide predictive simulation capabilities that can replace much of the empirical work currently done through expensive experimentation and testing.

The ALE3D code has grown in capabilities over the last 20+ years due to a combination of new algorithms and the growth in raw computational power of the largest available supercomputers. Analysts can build models on a laptop or workstation, and then easily scale their problems up to use the largest supercomputers available in the DOE and DoD. Large computing resources allow analysts to not only run problems that were intractable on previous generations of computers, but also to run hundreds or thousands of calculations for parameter studies and uncertainty quantification.

Yesteryear

Early work in safety analysis of explosives and propellants focused on using simplified 2D models of high explosives detonation. These calculations were used to design higher fidelity models validated through comparison with experimental results.

Today

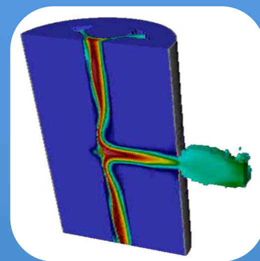

Today we can model safety scenarios such as bullet or fragment impact upon a rocket motor or conventional munition. Improved physics allows for a more predictive response to mechanical insult, such as whether a propellant or high explosive will deflagrate or detonate.

Exascale

Exascale computing will be required for full system designs using the advanced higher fidelity models being developed today and in the future. This can save millions of dollars by giving designers a virtual proving ground for analysis of safety features.

Increases in compute power allow for higher fidelity physics, larger calculations encompassing full system designs, and the large number of calculations required to improved uncertainty quantification and the ultimate goal of predictive simulation.

Sidebar:

Creating a computational simulation

Advanced simulation techniques lie at the heart of many of the nation's most pressing scientific challenges, including understanding our changing climate, designing safe and efficient energy sources, and managing the nation's nuclear stockpile. For example, when designing next-generation nuclear reactors, many new or modified designs can be evaluated using computer simulations before the best designs are chosen for further study. These simulations are built using advanced mathematical models that describe the underlying physical phenomena, and sophisticated software tools that allow scientists to examine solutions for many different scenarios.

To build the simulations, research scientists first devise a mathematical model of the physical process they would like to study. This results in one or more equations that approximate physical processes, along with a description of what is occurring on the boundary (boundary conditions) and at the beginning of the simulation (initial conditions). In addition, scientists must develop a computer representation of the computational domain. The geometry of the domain can be as simple as a rectangular box or sphere or as complex as one can imagine when studying advanced scientific devices. In most cases, the mathematical equations describing the physical phenomenon cannot be solved analytically on the domain of interest. Instead, the domain is decomposed into a collection of simpler geometries—a mesh—typically comprising triangles or quadrilaterals in two dimensions and tetrahedrons or hexahedrons in three dimensions. Once the mesh has been generated, the mathematical equations are approximated on that mesh resulting in a system of algebraic equations that is easier to solve on a computer than the original equation is. Once the solution of this system of equations is obtained, it is extensively analyzed and, when possible, validated against experiments to ensure the solution is correct. This process is repeated with adjustments made to the mathematical model, the computational domain, the mesh, or the numerical solution process until the scientific goal in question is achieved.

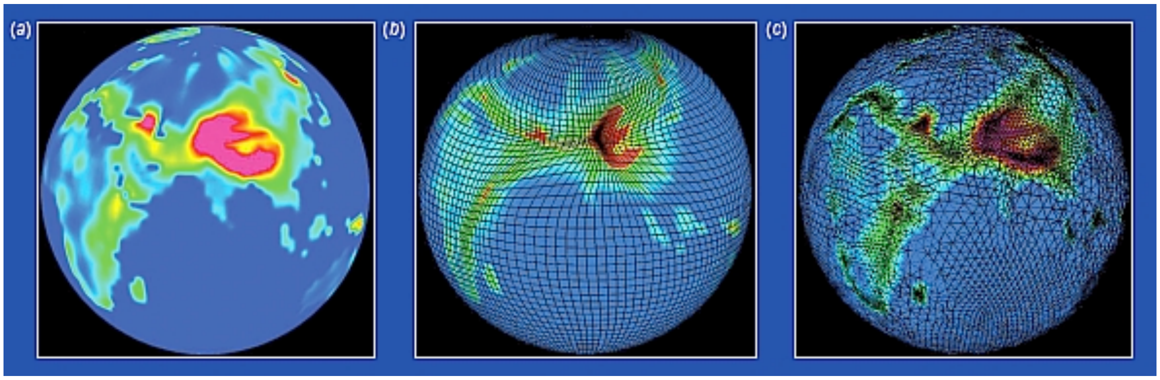

Figure: Different types of meshes can be used to represent geometry in a simulation. This figure represents (a) orography (height of the terrain) with color, and uses structured (b) and unstructured (c) meshes to concentrate mesh points in regions of interest.

Adapted from "ITAPS: Advanced Simulation Technologies for Application Scientitsts", SciDAC Review, Issue 13, Summer 2009 with permi ssion of the author.

Sidebar

Uncertainty Quantification

Today, researchers in many different disciplines turn to numerical simulations to predict or explain complex physical phenomena. For simulations of physical phenomena to be meaningful and allow for accurate prediction, it becomes crucial to provide estimates of the uncertainty associated with the numerical simulation. Examples of different disciplines with implications for national security that rely heavily on numerical simulation with quantified uncertainty are the assessment of safety, surety and performance of the aging/evolving stockpile without nuclear testing and predicting climate response to energy technology strategies. This includes evaluating the range, area and limits of predictability; regional impacts on agriculture; predicting responses to mitigation strategies; monitoring and managing greenhouse emissions and concentrations, and the impact on energy policy.

Uncertainty Quantification (UQ) is the end-to-end study of the reliability of scientific inferences. Ideally, UQ results in

- a quantitative assessment of that reliability;
- an understanding of possible sources of error and uncertainty in the
- inferences and predictions;
- identification of the sources of error and uncertainty;
- a clear understanding of the assumptions on which the assessment is based.

An obvious question arises: how accurate are the predictive simulations? Ideally, the optimal approach for validating simulations would be direct comparisons to the real experiments. Unfortunately, in many circumstances we cannot run the experiment, because of time, cost, or inability to gather enough experimental data to validate the simulation. In the examples given above, we cannot perform a nuclear test and we cannot create a climate for comparisons with numerical simulation. Thus we are left to quantify and bound the uncertainties in the simulations. This process of UQ provides a measure of the variability of the simulation results in terms of error bars or probability density functions (PDFs) and also helps identify the most significant sources of uncertainty. For example, outputs of numerical simulations may be sensitive to small variations of the many independent input parameters.

A brute force approach for quantifying the full range of uncertainty in the simulation is not possible to execute for the national security applications mentioned above. In such an approach, the simulation would be performed for each of the different combinations of the values that the parameters can take on and then analyze the results to compute the variability of answers. If an individual simulation requires

20 parameters and takes 24 hours to execute (this is a relatively small problem for the examples above), then if we also assume that due to uncertainty in the precise values of the parameters, each of the 20 parameters can be represented approximately by just three different values (e.g. low, high and intermediate) over the full range of the parameter, this implies the need for 3^{20} or about 9 billion simulations, which would take several million years to calculate on today's fastest available supercomputer. Consideration of the full range of uncertainty involving 100 parameters would require more time than the age of the universe on a massively parallel supercomputer to explore all possible permutations.

Researchers at Lawrence Livermore National Laboratory are taking a two-pronged approach to improving the state of the art in uncertainty quantification. The first approach will explore "intelligent" strategies to adaptively navigate the complex terrain of the uncertainty space of problems for which the number of uncertain parameters is high (a high dimensional uncertainty space) and develop new approaches that achieve a reduction in the high dimensionality of this space making the uncertainty quantification analysis tractable. A simple schematic example for a 2 dimensional parameter space is shown in the figure. The second approach is a longer-term strategy that examines the possibility of including the propagation of uncertainty directly into the key numerical algorithms used in specific scientific simulations. Propagating uncertainty estimates of the underlying algorithms can produce reliable and quantifiable estimates of the error bars induced by a simulation. This approach will require significant new research, and will involve the use of, for example, adjoint methods for computing solution variability as a function of interesting parameters.

Figure : White circles in part A identify simulation tests run varying two parameters. Based on those answers, we can predict results for nearby values of the two parameters and compute a predicted accuracy for those values. In this case, purple is highest predicted accuracy. Based on this result, adaptive sampling recommends more samples in the upper right area, thus improving the overall accuracy shown in part B.

The contours of a response model's prediction accuracy

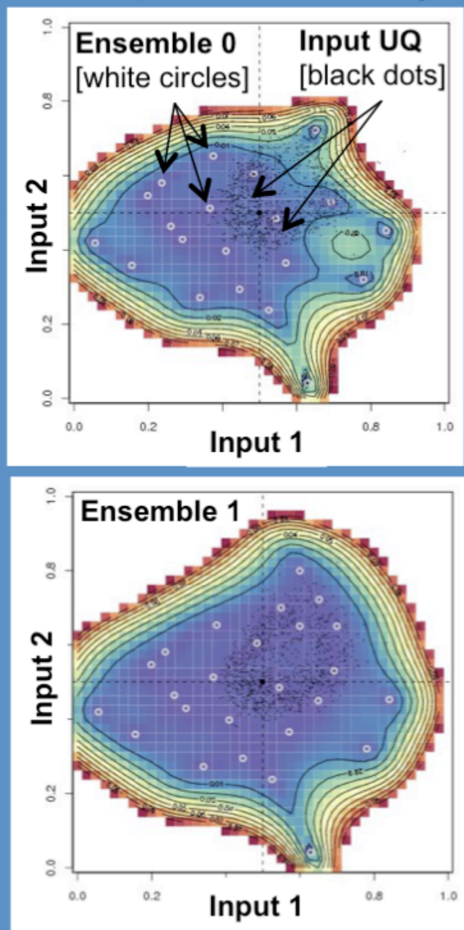

Sidebar

Detonation-to-Detector Simulation for Nuclear Explosion Monitoring

Seismology provides the best and most timely data to detect, locate and characterize a weapons system test (e.g. identify as an explosion, discriminate from background earthquakes, estimate the yield). When a nuclear device is detonated underground it vaporizes the rock immediately surrounding the device, creates a cavity and sends a shock wave into the earth within a few tenths of a second. This shock wave excites seismic waves that can be observed several minutes afterward at large distances, depending on the explosive yield, propagation pathway through the earth and observation conditions (e.g. background noise at a station). For example a 1 kiloton fully coupled (buried) explosion in strong rock can result in a Richter magnitude equivalent 4 event and be observed at distances of 1000's km. However, the level and character of ground shaking are strongly dependent on the geologic material in the immediate vicinity of the detonation and the path the waves travel through the earth from source to sensor. The earth is heterogeneous on all scales, that is its composition and physical properties vary in three-dimensions (3D) from the mineral grains that compose a single rock to the tectonic plates that form the outer layers of our planet. Not surprisingly seismic waves are strongly impacted by this 3D variability, especially in the near-surface crust and uppermost mantle (0-100 km depths). This makes every source-sensor path through the earth unique and introduces tremendous uncertainty in seismic modeling.

During the Cold War, nuclear explosion monitoring was focused on a few well-known test sites and empirical analysis of seismic recordings from past events. Now national security requirements are concerned with proliferation and emerging nuclear states and this requires broad area monitoring without reliance on past explosions. Equally challenging is the need to monitor some areas devoid of earthquakes for comparison. These requirements stimulate the development of simulation and modeling tools to gain a predictive capability for seismic waves emerging from explosions and/or earthquakes from any source location to any sensor, especially in areas of proliferation concern. Modern simulation tools and high performance computing allow for the modeling of explosion-generated seismic waves from the detonation underground, through the 3D earth to the observing seismic station (sensor). Current terascale applications are routine and petascale calculations have been demonstrated. However, in order to model seismic wave generation and propagation to sufficient fidelity across the broad range of spatial (10^{-2} – 10^{6} m) and temporal scales (10^{-6} – 10^{3} s) requires development of code coupling and giant advances in computational power.

Full-scale 3D end-to-end (device-to-detector) modeling is the goal for computational nuclear explosion seismology and will require exascale computing. This will enable seismologists to better detect weak signals from small distant events, provide improved estimates of event location and explosive yield as well as reduce uncertainty in source type discrimination. Toward this goal we are modeling elements of the full-scale problem and working to piece them together. Presently we can model shock wave generation in different rock types. Similar to the SBSS program, calibration of material models with past nuclear test data is essential to build confidence in simulation results for the shock waves generated by underground nuclear tests. To improve propagation models on continental scales, earthquake data can be used to image 3D earth structure. New methods based on fully 3D simulation of earthquake motions and time-reversal (adjoint waveform tomography) are being applied to resolve detailed sub-surface structure. While these methods have been demonstrated, they will require exascale computing to perform the large number of high-resolution calculations for 100's of earthquakes and iterations over model optimizations. The future is bright for exascale computing to enable major advances in seismic nuclear monitoring.

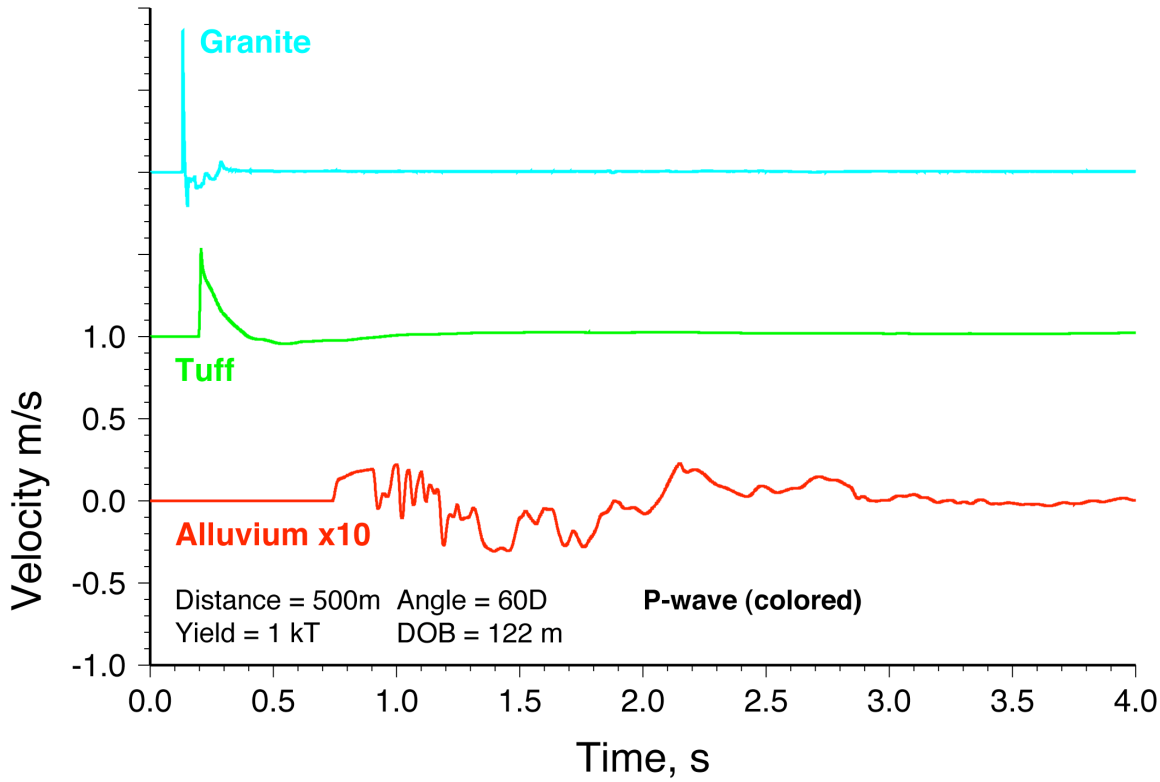

Figure NP.1. Modeling of explosion generated shock waves in different material models (granite, volcanic tuff and sedimentary alluvium) results in vastly different wave motions. Material models were generated to reproduce observations from a large set of legacy ground motion data from past nuclear tests.

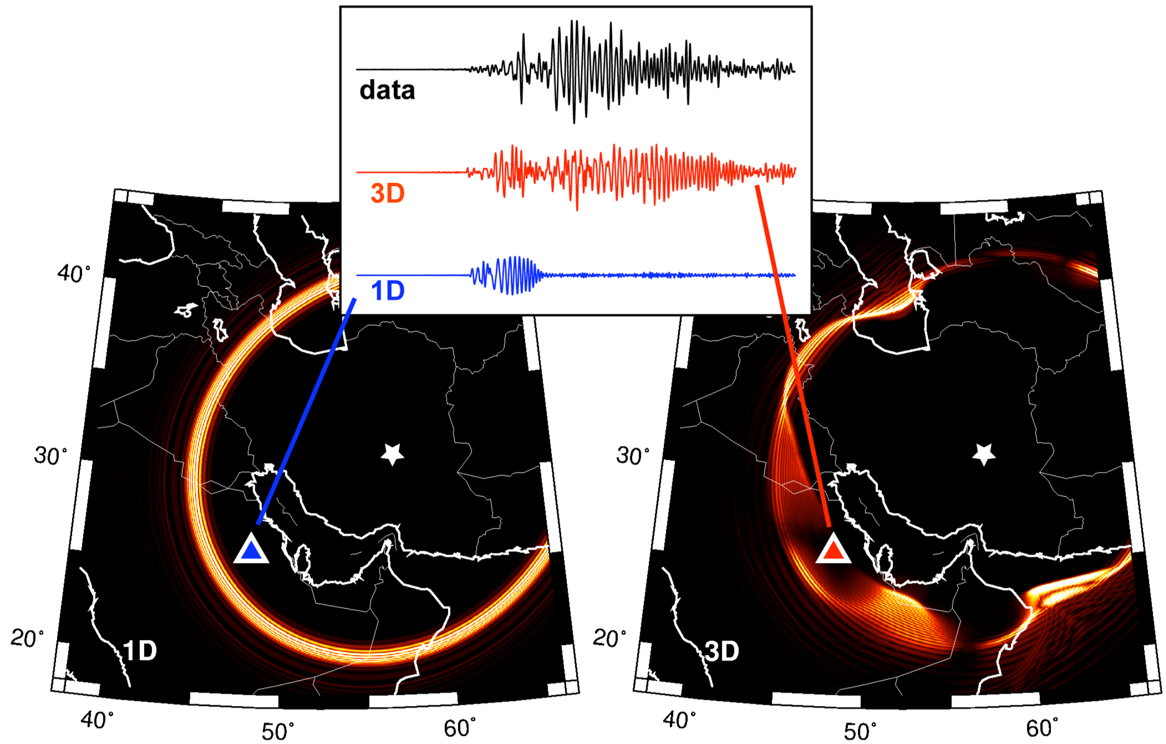

Figure NP.2. Simulations of a seismic event in the Middle East using an unrealistic one-dimensional (1D, left) and three-dimensional (3D, right) model. Comparison with an observed seismogram (top, black) shows that the 3D model (red) predicts the long-duration character of the data and the 1D model (blue) is not at all like the data.

Sidebar: Predicting the Bizarre Properties of Plutonium[1]

Plutonium is arguably the most complex element known, and it is one of the least well understood. Before it liquefies, plutonium exhibits six solid material phases that vary considerably in density. Plus, a seventh phase may appear when the radioactive metal is under pressure.

To understand material phases, think of carbon and its most common solid phases: soft graphite and hard diamond. Both are made of carbon atoms, but the bonds that form between the atoms create two very different materials. Many elements have two or more solid phases, but most have no more than four. With six phases, solid plutonium is highly unusual. The material's peculiarities do not stop there. Experiments over the years have demonstrated other anomalous properties, including an almost complete absence of magnetism and highly unusual resistivity.

In an effort to explain some of plutonium's strange behavior and better understand results from past experiments, a team of Lawrence Livermore scientists and international collaborators used Livermore's Atlas supercomputer to perform some of the most precise predictions yet of delta-phase plutonium. For these simulations, the team combined density functional theory (DFT) and dynamical mean field theory (DMFT) to calculate plutonium's delta-phase electronic structure, specifically its lack of magnetic "susceptibility." While DFT is useful for explaining the energy and interactions of many electronic systems, it may break down for certain properties of strongly correlated systems. Previous research combining the two theories to simulate delta-phase plutonium could obtain only approximate solutions to the DMFT equations. Recent advances in the continuous-time quantum Monte Carlo methods – and the power of the Atlas supercomputer – allowed exact solutions for the first time.

The team's simulations predicted that at room temperature with delta-phase plutonium at its equilibrium volume, the *f* electrons are delocalized. That is, they easily move about the lattice and are not associated with one particular atom. Under these conditions, the material's magnetic susceptibility is very low above 600 Kelvins and only slightly higher at lower temperatures.

A change in temperature does not affect the magnetism of conventional metals such as platinum or molybdenum. In contrast, most strongly correlated materials other than plutonium exhibit a distinct relationship between temperature and magnetic susceptibility. Their magnetic susceptibility is high at low temperatures and lower at high temperatures. Experiments demonstrate, however, that plutonium's

---

[1] The material in this section is condensed from an article appearing in Science & Technology Review, Lawrence Livermore National Laboratory, November 2008, and is used here with the permission of LLNL and the author.

magnetic susceptibility is unlike that of other strongly correlated materials. In fact, it behaves more like a conventional metal, although it exhibits slight temperature dependence.

If the plutonium volume is expanded—if the lattice is stretched so that the *f* electrons are farther apart—the magnetic susceptibility of plutonium changes. The Livermore simulations showed that as the lattice expands, the *f* electrons become heavier, or localized. That is, they are more associated with one particular atom and thus cannot easily hop through the lattice. The plutonium is then a more strongly correlated material. The transition from delocalized to localized behavior occurs at increasingly lower temperatures as the lattice volume continues to expand. With greater distance between the electrons, delta-phase plutonium begins to behave more like other strongly correlated materials. As temperature drops, the metal's magnetic susceptibility increases.

The team intends to tackle plutonium's alpha phase next. Predicting the behavior of alpha-phase plutonium will be more challenging than the delta-phase simulations. The smallest individual crystal in delta-phase plutonium contains one atom. In the alpha phase, 16 atoms make up the smallest crystal. Another challenge will be to explain the role that other materials play in stabilizing delta-phase plutonium. Only with the aid of powerful supercomputers can researchers answer plutonium's many riddles.

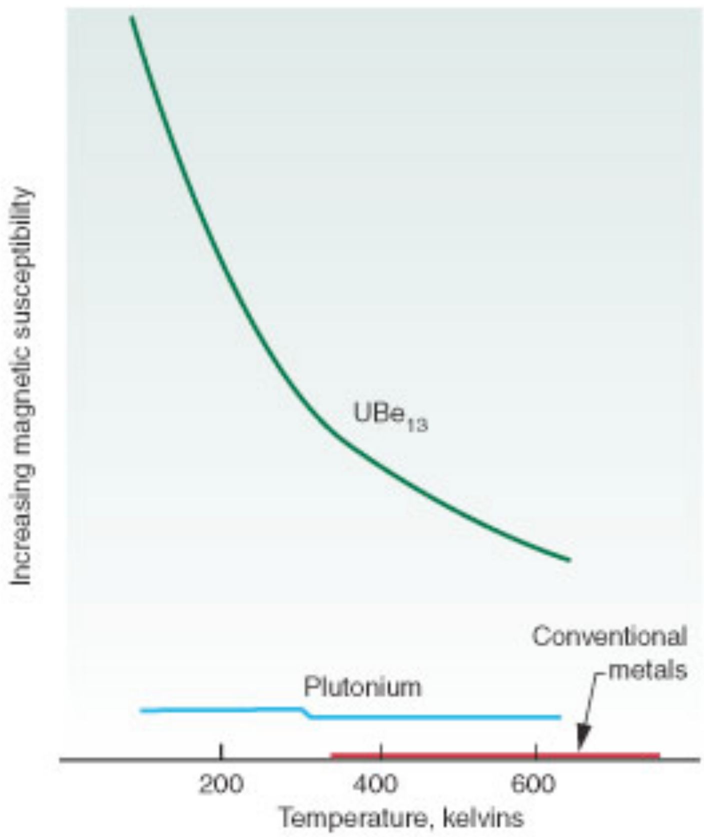

[Caption] The magnetic susceptibility of plutonium is unusual. For conventional metals such as molybdenum, titanium, and platinum, magnetism does not change with temperature (red curve). Strongly correlated materials, however, are more magnetic at low temperatures than they are at higher temperatures. The green curve shown for uranium–beryllium-13 (UBe_{13}) is typical. Plutonium's magnetic susceptibility (blue curve) lies between these cases.

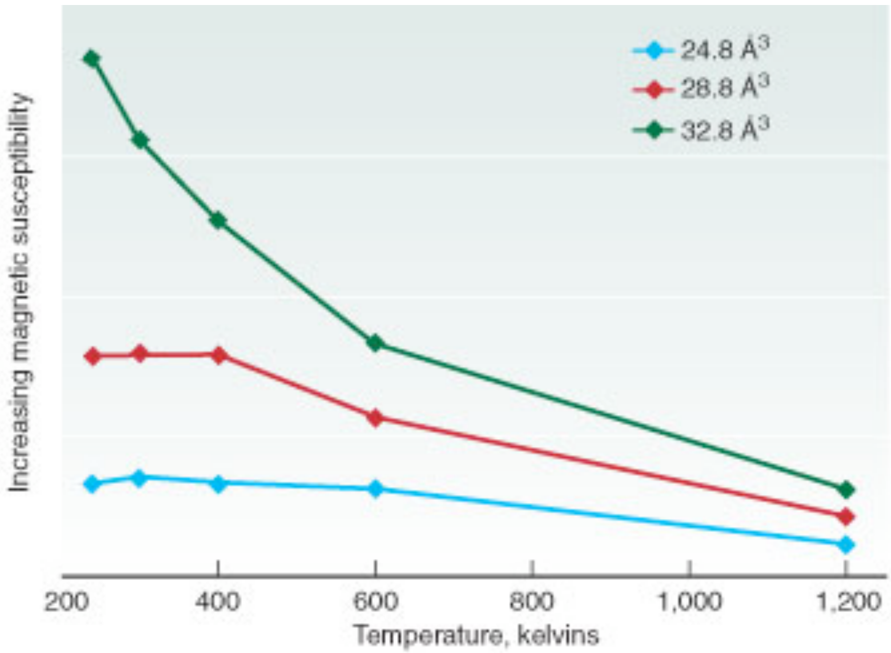

[Caption] In the Livermore simulations, the volume of plutonium was increased, pulling the plutonium atoms farther apart. Only under these circumstances and at the largest volume does plutonium's magnetic susceptibility begin to mimic the temperature dependence of other strongly correlated materials such as UBe_{13}. ($Å^3$ = cubic angstroms, where 1 $Å^3$ equals 1×10^{-30} cubic meters.)

Sidebar:

Computations for High Energy Density Physics[2]

Energy density, the amount of energy stored in a given volume, can take many forms. Aeronautical engineers want a fuel with maximum energy density for rocket liftoff. High-energy-density foods are vital to endurance athletes, such as cyclists in the Tour de France. As the object with the highest energy density ever created by humankind is an exploding thermonuclear weapon, it is not surprising that the study of high energy density physics is important science question for national security.

Recently, in a set of high-energy-density (HED) experiments, researchers at Lawrence Livermore National Laboratory used ultrahigh-intensity lasers with ultrashort pulses to zap one side of a reduced-mass target. The goal was to get a dense target as hot as possible in the presence of a large magnetic field. The 2007 experiments performed on the Callisto laser in Livermore's Jupiter Laser Facility showed that the target did get hot, but unexpectedly, a large number of protons were ejected from the entire surface of the target. This surprised Livermore physicists, who expected behavior similar to when a laser zaps a larger (millimeter size) target, resulting in a beamlike pattern of protons that blows off the back of the target.

In order to understand these unexpected results, the researchers used Livermore's 9216-processor 44 Teraflop Atlas supercomputer to simulate the results of these laser-driven HED experiments. The simulations turned out to be very important, as they helped explain the physics behind what the surprising experimental observations and measurements.

The Atlas simulations used a particle-in-cell code specifically designed for studying electrons in a high-energy plasma. This allowed the researchers to study laser–plasma interactions in reduced-mass targets at full scale from first principles. In a two-dimensional (2D) simulation of a large target, electrons accelerated by the laser generated an electric field on both the top and bottom of the target. In a simulation of a smaller, "finite" reduced-mass target, large electric fields developed on the sides of the target as well, which explained the signal detected all around the target in the actual experiments. This simulation showed that shrinking a target to a smaller size does not increase target temperature as one might expect, but instead increases the total number of ions accelerated from all of its surfaces. This was precisely the

---

[2] The material in this section is condensed from an article appearing in Science & Technology Review, Lawrence Livermore National Laboratory, January-February 2009, and is used here with the permission of LLNL and the author.

pattern seen in the experiments. A 3D simulation also predicted maximum proton energies out the back of the target to be about 5 mega-electronvolts, agreeing with experimental results. Results from the Atlas simulations indicate that smaller targets may be more efficient ion accelerators than larger targets, which could make fast ignition using proton beams competitive with hot electron–based fast ignition.

Simulation on powerful supercomputers is often the only way to both understand the results of experiments in detail and to develop physical insight into complex experimental processes. Simulations can parse the physical constituents that affect the whole and examine microscopic details not easily detected during an experiment. In addition, computer simulations can explore regimes of temperature, density, and pressure that experiments cannot yet achieve, serving as a guide for future experiments. Ultimately, scientists must depend on both experiments and simulations working in tandem to advance HED physics research.

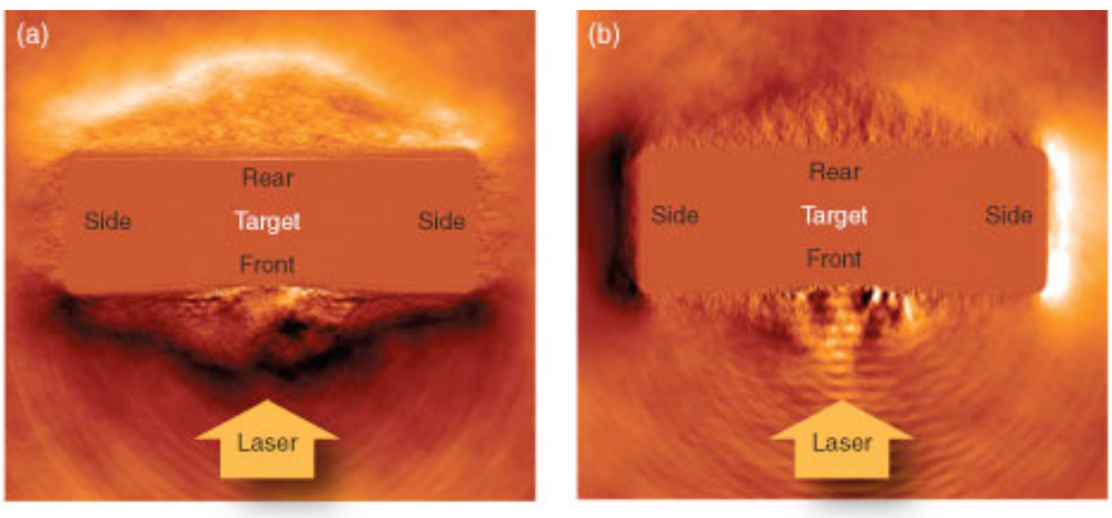

[Caption] Two-dimensional simulations on Atlas incorporate full physics and help explain the Callisto laser experimental results. Electrons accelerated by the laser generate an electric field (a) on both the front and rear of the reduced-mass target and (b) on the target sides as well.

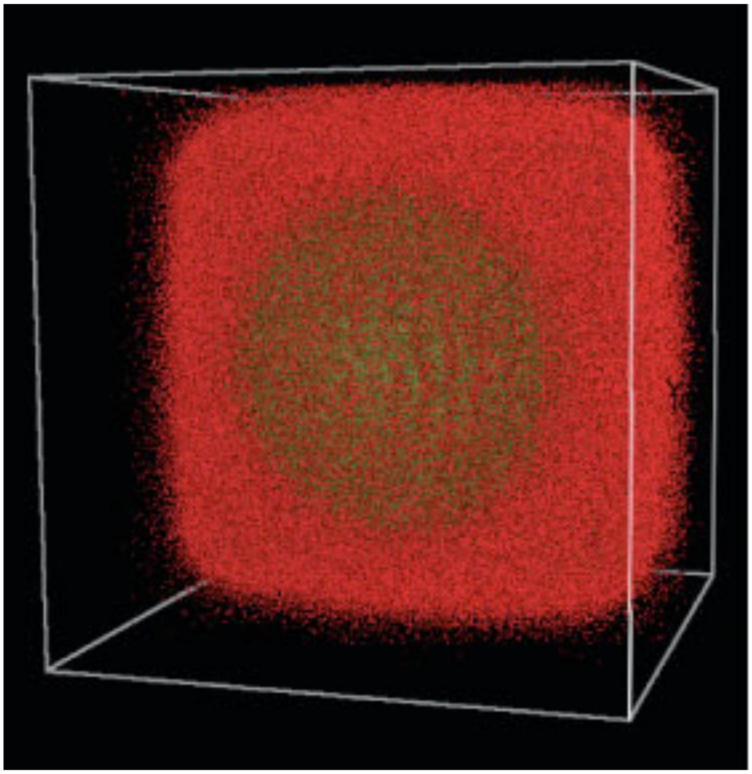

A three-dimensional simulation of a reduced-mass target shows vertical striations in electron density due to a laser polarization effect—a phenomenon that can only be observed in three dimensions.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.