ASCI

Big engineering in the field of scientific computing

James S. Peery

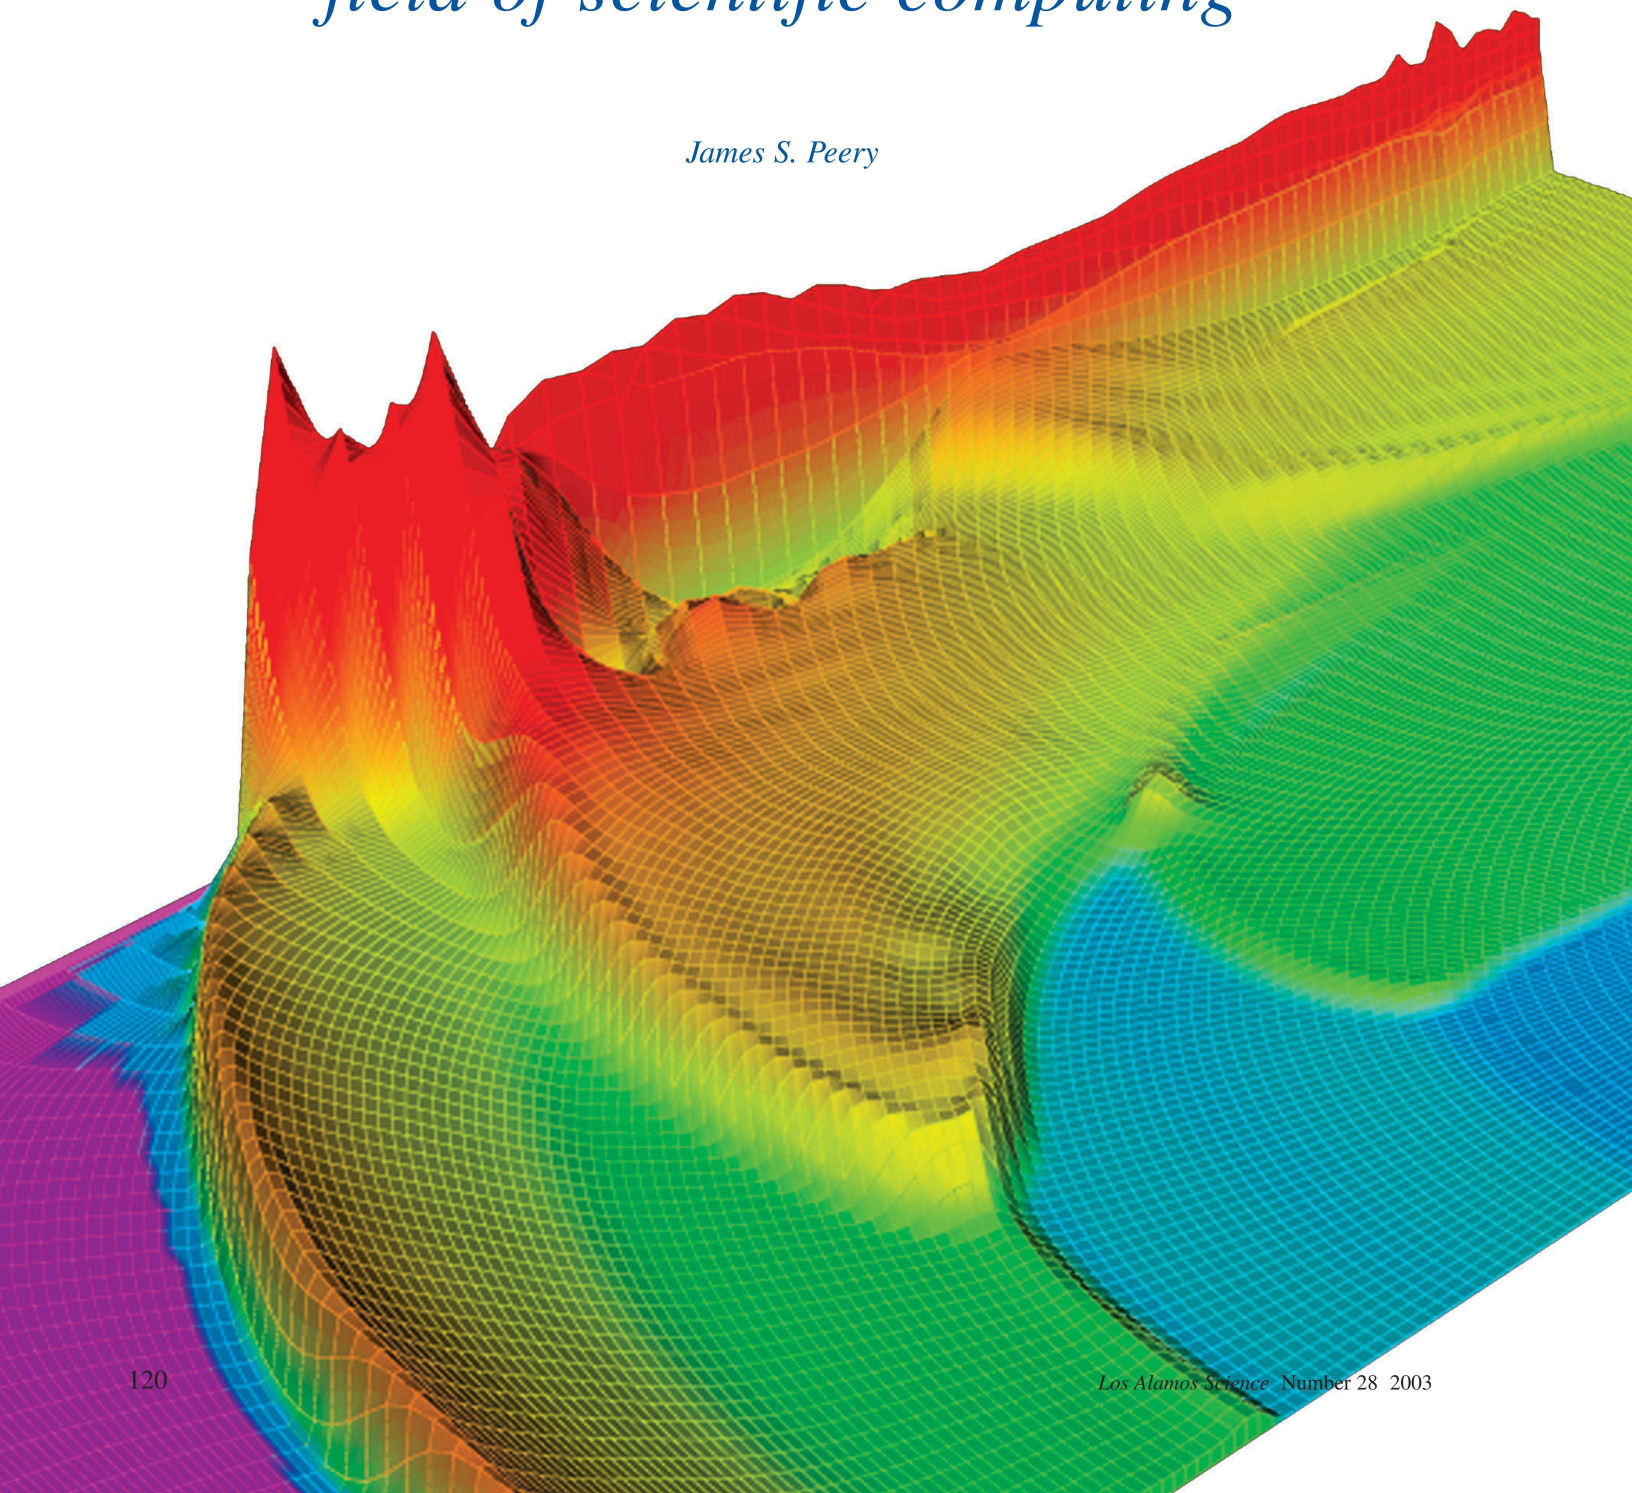

Imagine that, during routine inspection, we open up a weapons system and find a significant change that we did not expect. Or imagine the numerous possible orientations and conditions that a weapon might assume in a fire. These scenarios represent real possibilities. Will the weapon work as intended? Will the weapon be safe under almost preposterous conditions? These are the questions that the Los Alamos Advanced Simulation and Computing program will help to answer by providing the weapons designers with high-fidelity simulation capability on the world's most powerful computers.

The Advanced Simulation and Computing program evolved from the merging of the Accelerated Strategic Computing Initiative (ASCI), begun in 1996, and the ongoing stockpile computing program known as the Advanced Simulation and Computing Campaign. Continuing to use the acronym ASCI, this effort is perhaps the largest and most encompassing computational development program in the world. Its core mission is to provide simulation tools, including both the hardware and the software application codes, that enable the weapons designers to assess and certify the safety, performance, and reliability of the enduring nuclear weapons stockpile. As such, ASCI is a pillar of the Science-Based Stockpile Stewardship (SBSS) program. The success of ASCI, however, will have an even larger significance, by demonstrating that large-scale computational science can create potent tools to address many scientific challenges.

In a *Popular Science* report summarizing 15 years of big engineering, nine major construction projects were cited among which were the Toronto SkyDome, the Eurotunnel, and the Petronas Towers. These are multibillion dollar, multiyear projects involving multidisciplinary teams. ASCI is the first scientific software project to have a similar level of investment and a similar multi-institutional, multidisciplinary approach. Although the three labs involved, Los Alamos, Sandia, and Lawrence Livermore National Laboratories, develop their application codes independently, they work jointly on issues of computer science and hardware, on the testing of the application codes, and on developing visualization tools at scales never before attempted.

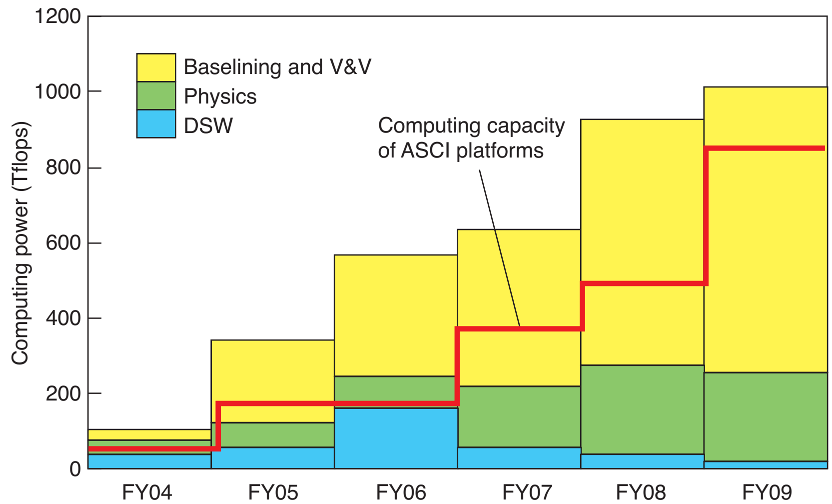

Figure 1. Computing Needs and the Integrated Delivered ASCI Platform Capacity

The anticipated computing needs for the Los Alamos and Livermore weapons program are divided into three categories: (1) direct stockpile work, or DSW (stockpile simulations, which assess stockpile issues, are part of refurbishment and annual certification); (2) physics (studies to increase understanding of weapons simulations); and (3) baselining and validation and verification, or V&V (creating 2-D and 3-D validated weapons models). The orange line represents the maximum capacity currently planned for the ASCI platforms. This integrated, delivered ASCI platform capacity is seen to fall short of the anticipated needs.

Although the ASCI program is well known for buying the world's most powerful computers, less than one-sixth of the total budget is spent on hardware. The major fraction of our effort goes into software development—simulation codes that faithfully model the end-to-end performance of a nuclear weapon. These multiphysics codes, validated through comparison with experiment and archival nuclear-weapons test data, represent the enduring product of the program.

Nuclear weapons are complex systems. During performance, materials change from solids to hot, dense plasmas, and physical processes operate on many different length and time scales. In order to produce predictive simulations of weapons performance, the codes must be built from accurate models of these physical processes and material behaviors validated through comparison with experiment. Further, the algorithms that represent these models must be both robust and computationally efficient, and they must be verified on simple problems by comparisons between numerical results and known solutions. Validation and verification are necessary to demonstrate the accuracy of these codes but are not sufficient to ensure their utility. The codes prove their usefulness when designers are able to set up problems rapidly, produce results in a reasonable time, and see the results in a form that can be easily and quickly interpreted. To satisfy these additional requirements, ASCI is making significant investments in developing visualization and other enabling tools and in production support for hardware and software.

At present, ASCI has responsibility for providing the computing resources (that is, cycles) for both the near-term needs of stockpile stewardship and the long-term development and application of high-fidelity simulation capability. More detailed physical models, coupled with higher resolution and three-dimensional (3-D) rather than two-dimensional (2-D) simulations, are projected to greatly increase the need for computing capacity. Figure 1 compares the anticipated computing needs for SBSS with the integrated delivered capacity based on the current ASCI computing-platform procurement schedule. Needless to say, the ASCI platforms alone will not provide the computer cycles required to meet the various demands of SBSS.

Beginnings of ASCI

At its inception in 1996, the ASCI program was conceived as an effort to accelerate the development of new, more-predictive weapons simulation tools. When supported by necessary computing resources, those tools would be able to support long-term stewardship of the stockpile in the absence of nuclear testing. To understand the magnitude of this undertaking, one needs to look at stockpile computing before ASCI. In the 1980s, coarsely resolved 2-D calculations might run for thousands of hours on the world's most powerful computers. Crays were the mainstays of production computing. After a decade of use, Crays had stable and well-understood vector architectures. Hundreds of those computers were in use around the world, although Los Alamos and Lawrence Livermore National Laboratories prided themselves on acquiring the first serial number of each latest model.

Before ASCI, the weapons codes (which are now referred to as "legacy" codes) were built by small development teams to support the day-to-day needs of the design community. The legacy codes matured by being applied to one-dimensional and 2-D problems whose timely solution was needed for planning and designing underground tests. Such tests had many goals, among which were certifying new designs, performing physics experiments, and confirming stockpile confidence. The heavy test schedule limited the time that could be spent on fundamental improvements; instead, legacy codes were calibrated to the underground test data with nonphysical parameters, sometimes termed knobs. This process produced useful engineering tools for interpolation, but their predictiveness for extrapolation was indeterminate. In other words, the success of that code development strategy depended on continued testing. The interaction of modeling and experiment is part of the scientific method. However, the political decision to cease nuclear testing required an immediate and urgent change of strategy—one result was ASCI.

Because the legacy codes can reproduce the results of underground tests, albeit, not from first principles, they are a direct link to the past and remain important to weapons designers. However, the design community also needs codes built from better physics models to assess the effects of aging components within the weapons, newly identified safety concerns, and other stockpile issues. The architects of ASCI understood that the new, more predictive codes would require huge increases in computing capability. Indeed, the program would have to revitalize the high-performance computer industry if high-fidelity simulations of the complex physics inside a nuclear weapon were ever to be practical. Initially, the program decided to focus on achieving long-term predictive capability at the expense of supporting short-term designer needs.

The vision was sold, and the planning began. The end goal of ASCI became the construction of new, high-fidelity, verified, and validated 3-D codes. High fidelity implies that the codes contain first-principles physics models and accurate, efficient numerical algorithms that produce converged solutions. Without fully understanding the magnitude of this vision, ASCI set out to develop 3-D codes capable of unprecedented resolution of physical processes in space and time. It was not long before the requirements were collected and the enormous complexity of the undertaking became clear. However, faced with the cessation of the underground testing and confronted by a rapidly aging weapons design community, management saw an urgent need to develop these more predictive tools and to train a new generation of designers as quickly as possible. Thus, the program grew at a rapid rate.

Developing Codes for Massively Parallel Computers

To simulate 3-D weapons system performance with high resolution and with reasonable turnaround times, one needs computers with 10^5 to 10^6 times more power than the Cray YMPs used at the end of the underground test period. The only type of architecture capable of delivering such power is a massively parallel computer in which at least 10,000 processors can be applied simultaneously to solving a problem. ASCI generated a multiyear, multiplatform plan to achieve that goal, which should be realized in 2005. The latest ASCI platform, the Q machine at Los Alamos, is approximately 10^4 times faster than a Cray YMP.

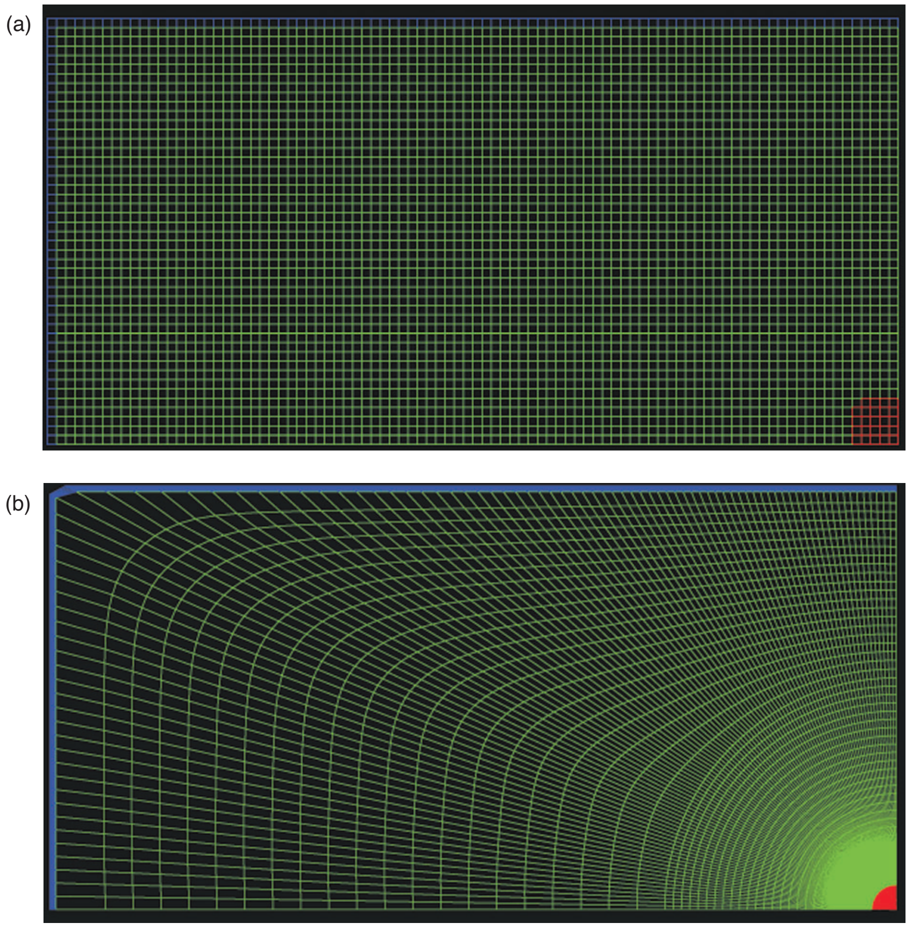

Figure 2. Eulerian and ALE Meshes for Simulating a Pulsed-Power Experiment

The meshes in (a) and (b) represent the computational domain in 2-D cylindrical geometry for a liner experiment at the Atlas pulsed-power facility. (See Figure 6 on page 74 in the article "The New World of the Nevada Test Site" for a description of such experiments.) The stainless steel containment vessel (blue) is separated by air (green) from the liner assembly (red). The regular Cartesian mesh in (a) is suitable for Eulerian hydrodynamics calculations. The boundary-fitted curvilinear mesh in (b) is suitable for ALE hydrodynamics calculations.

Effective use of massively parallel platforms demands new algorithmic strategies. ASCI began development of both a new generation of codes employing parallel algorithms and the associated setup and visualization tools. The more predictive physics models that provide the building blocks require the solution of nonlinear partial differential equations involving multiple scales of length and time. The equations of the individual models can rarely be solved in closed form. They must therefore be solved approximately on the computer. These approximations are based on discretization methods such as finite difference or finite elements. Discretization means that the computational domain is divided into discrete volumes, or cells, that are organized by a mesh (see Figure 2). Solution variables, such as density and temperature, are averaged over the cells. This operation effectively reduces the number of unknowns in the problem to a level that the computer can handle. Discretization represents a tradeoff between the accuracy and required completion time of a simulation. Accuracy increases as the square of the number of cells (for example, for second-order algorithms). The required work, which is proportional to the problem time, increases as the fourth power of the number of cells in three dimensions. Therefore, the accuracy divided by the work, or the efficacy, is a strongly decreasing function of the number of cells. Computer power can be traded off for longer run times, but the run times can quickly become unacceptably long if the problem is very large. Ultimately, the accuracy of a simulation is limited by computer speed and the time one is willing to wait for an answer. And, of course, the discretized problem must fit into the available memory of the computer.

From a physics viewpoint, the models depend on experimentally measured properties, and these imply scales of length and time that must be resolved if the simulation results are to be valid. Furthermore, the individual models are coupled to each other, and their collective behavior is more complicated than the sum of their individual behaviors. This complication is ignored in the legacy codes, a simplification termed operator splitting (see the article "Massively Parallel Multiphysics Code Development" on page 128), but recent research has indicated that this simplification is a poor approximation. In other words, resolving the individual physics models is part of verifying the algorithms, but that step does not guarantee getting the right answer. Verification must be followed by validating the full multiphysics code system against experimental data.

Setting up the mesh to represent the initial geometry of a weapons system in three dimensions can be a formidable challenge. There are two basic frameworks for solving the hydrodynamic equations that describe the motion of materials. Known as Eulerian and Lagrangian, they utilize different types of mesh (see Figure 2). Eulerian algorithms solve the equations on a mesh that remains fixed in space while the material flows through it. Lagrangian algorithms solve the equations on a mesh that moves with the material. Each method has advantages, and therefore ASCI set forth to develop both. Redundancy, in the sense of multiple independent approaches to code development, has long been a staple of the nuclear weapons program.

In general, we use computer-aided design (CAD) software to generate the 3-D geometries of weapons systems. But CAD software was designed for manufacturing applications; consequently, the CAD setups suffer from incompleteness, overlapping parts, and unnecessary detail and are therefore ill suited for ASCI set-up applications. To overcome these deficiencies, ASCI set forth to develop 3-D meshing algorithms for Lagrangian-based codes and volume filling techniques for Eulerian-based codes. Indeed, the regularity of the Eulerian meshes has already allowed us to create effective setup tools. Meshes for Lagrangian codes need to reflect both the initial geometry and the subsequent material motion. The tendency of imperfect meshes to tangle and thus bring the simulation to a premature end is a more difficult problem to overcome, and the lack of adequate setup continues to limit the use of Lagrangian codes in three dimensions. An example of a 2-D calculation using arbitrary Langrangian-Eulerian (ALE) techniques, a method that combines the advantages of Eulerian and Lagrangian approaches, is shown in Figure 3. In particular, a rezone and remap procedure is added to a Lagrangian algorithm.

In addition to being accurate and robust, our solution algorithms must scale on parallel architectures. In other words, at a minimum, if we increase the domain of the problem by two and the number of processors by two, we want the problem to run for the same time. Historically, scaling has been achieved in mesh-based algorithms (Eulerian and Lagrangian) by dividing the domain of the problem into small chunks or subdomains. Each chunk runs on the memory attached to one processor. When the equations call for information from more than one subdomain, the processors communicate through a communication network. The efficiency of a massively parallel computer depends on our ability to minimize this communication time. This requirement has mandated significant investments in high-speed network technology, effective domain decomposition software, and parallel algorithm development.

Only by continual investment in fundamental science can we create the realistic models and predictive capability that will enable code developers and weapons designers to address the problems presented by the aging nuclear stockpile.

The process of discretization described above transforms the partial differential equations into a large system of algebraic equations. Although standard techniques for solving these matrices have existed for many years, the systems resulting from 3-D ASCI algorithms are too large for these matrix solution techniques and thus too expensive in time for ASCI-class computers and architectures. Therefore, major efforts were initiated at all three national laboratories to develop new techniques that reduce the solution time. The results of these early efforts are beginning to pay off. In some areas of physics, new matrix solution techniques are not only effectively solving the large equation systems but also scaling well with problem size.

As ASCI identified new research areas such as those described above, it responded by allocating resources, forming research teams, and in many cases initiating trilaboratory collaborations. The scope of ASCI grew rapidly. Large teams of code physicists and computer scientists were assembled to write the physics codes. It is not unusual for a physics code team leader to represent a team of 20 or more developers and to interface with dozens of other teams who are producing software libraries or hardware relevant to the project. Teams of computational physicists, tasked with developing new algorithms and solution techniques, produced libraries for the physics codes. Teams of computer scientists developed tools, message-passing protocols, and encryption and system software for use on the massively parallel computers. Teams of engineers and computational scientists tackled the issues in problem setup and domain decomposition. Teams of hardware and software engineers and scientists developed tools to move vast quantities of data from disks to leading-edge, 3-D stereo display platforms for both office and custom-designed collaboratories and theaters (see the article "A Vision of Hidden Worlds" on page 135). ASCI responded quickly to the technical challenges, but the ensuing growing pains are still being felt.

ASCI Report Card

The mission of ASCI has evolved significantly since its inception. The first five years, from 1996 to 2001, were mostly directed at proof of principle. The goals were very ambitious,

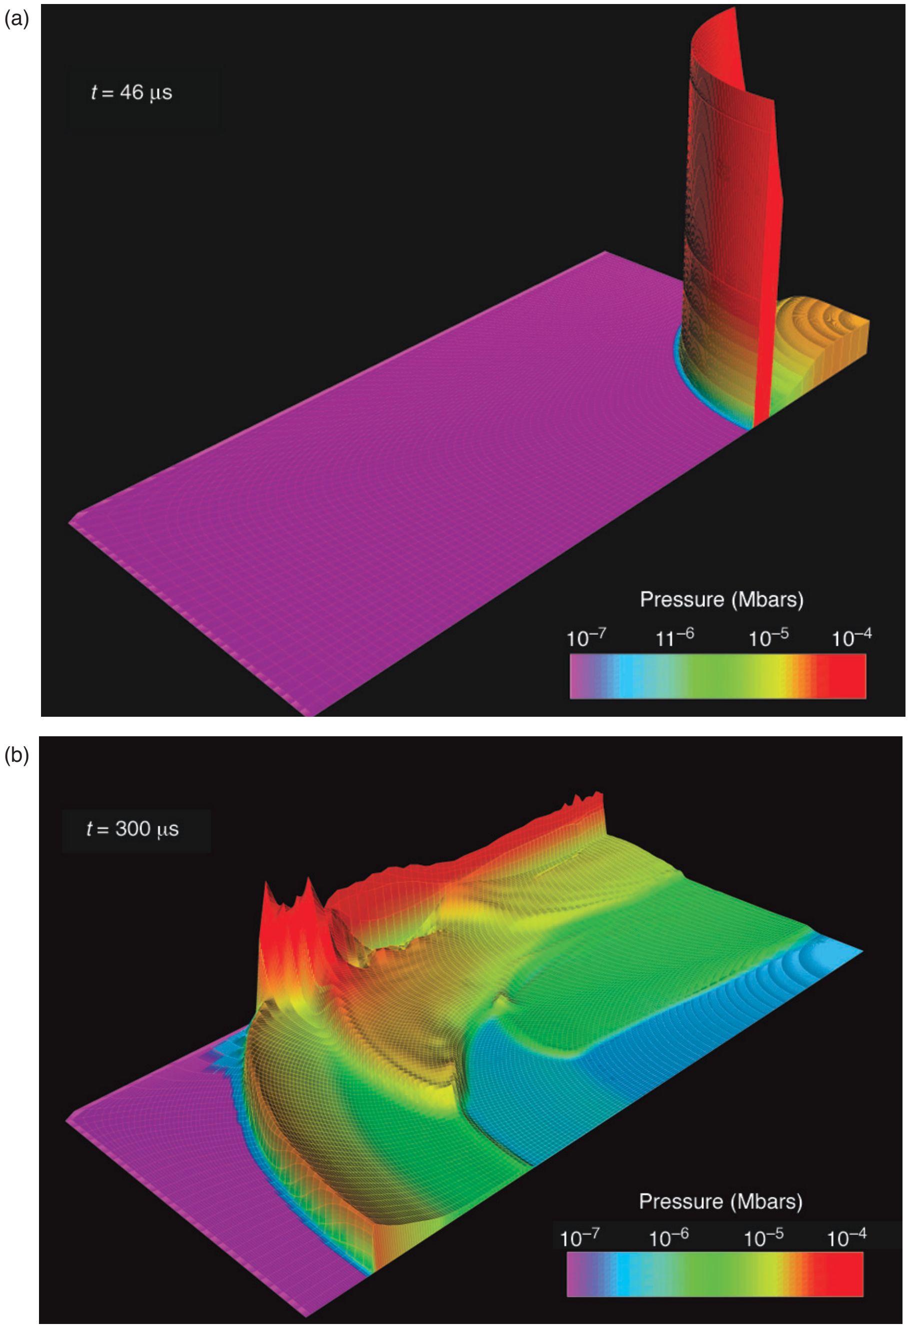

Figure 3. ALE Hydrodynamics Calculation of the Shock Structure inside Atlas Device

These pressure plots follow a pulsed-power liner experiment in the computational domain shown in Figure 2. At *t* = 0, 20 megajoules of energy, the maximum deliverable from the Atlas capacitor banks, is deposited into the liner assembly. (a) At *t* = 46 μ*s*, a sharp, spherical pressure shock is propagating outward in the air. (b) At *t* = 300 μ*s*, before the initial shock has reached the top of the containment vessel, multiple reflections of the shock have taken place between the sidewall of the containment vessel and the expanding liner assembly. These plots demonstrate the complex wave interactions and material flows that can be simulated using an ALE code as the mesh follows the material flows and is then readjusted to avoid tangling. The sizes and shapes of the mesh cells vary as the calculation proceeds.

but the program was short on requirements. Today, the focus has changed dramatically to deployment and to support of the new user-oriented tools. The end-to-end needs of the design community are driving program priorities and new activities. While ASCI continues its development of the new capabilities, it is also applying the simulation tools to immediate stockpile concerns.

Although the program is experiencing social engineering and project management tensions because of its rapid growth, it has engendered numerous technical success stories. At Los Alamos, the Crestone project has demonstrated unprecedented capabilities and geometric resolution through the first ever 3-D full system, end-to-end simulation of nuclear weapons performance. Moreover, codes of the Crestone project are used by more than half of the secondary design community. The Shavano project has provided a significant leap forward in its ability to model complex 3-D geometries and is gaining acceptance in the primary design community. The Blanca project has just recently completed a series of safety simulations and is being merged with the Shavano project. New physics models added to both the ASCI and legacy codes will continue to increase our predictive capability and add to our understanding of nuclear weapons. The Q machine and the Blue Mountain computer are delivering cycles to the design community and to the ASCI code development teams. Last, the computing infrastructure designed and deployed by ASCI, including both hardware and software (networks, computers, visualization displays), is facilitating the use of both the legacy codes and the new ASCI codes.

Los Alamos ASCI codes are now being used to close significant finding investigations, that is, to assess aging-weapons problems that are arising in the enduring stockpile. The ASCI codes also support the life extension programs for individual weapons systems by providing a means to evaluate the proposed steps for extending the shelf life of our present weapons systems. Significant efforts are under way to make the transition from legacy code calculations of baseline nuclear weapons performance to ASCI code calculations of those baselines. All these activities are enabled by the continuing operation and development of the supercomputing infrastructures at the national laboratories. Research continues on new techniques for storage, visualization, networking, and all aspects of the structure required by the modern generation of computing capabilities.

ASCI's goal of maintaining a healthy high-performance computing industry has been achieved. Although some vendors have exited the high-performance computing market, many have survived, new ones have emerged, and some have reengaged (for example, Cray). In addition, other federal agencies and universities have joined the push to maintain the U.S. high-performance computing industry. When ASCI started, many doubted that a teraflop computer (capable of performing 10^{12} floating point operations per second) could be built; now, through the efforts of many, ASCI has proved and enabled teraflop computing for all scientific communities.

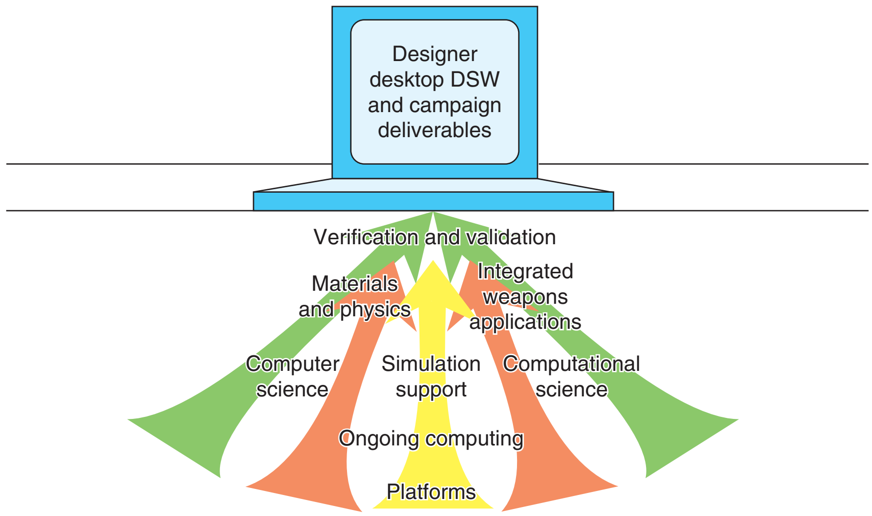

Figure 4. ASCI Program Elements
Each of the eight elements of the ASCI program is shown with a line of sight to the designer's desktop. The illustration suggests that stockpile certification and design requirements will continually guide the planning and execution of ASCI program elements.

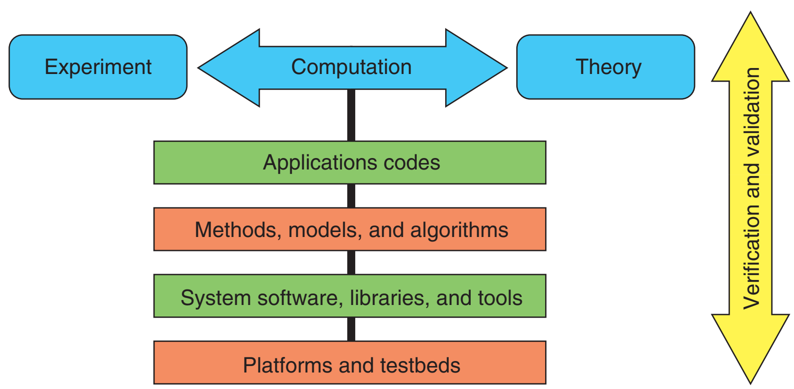

Figure 5. Elements of Increased Predictive Capability
The interaction between experiment and computation validates the codes and directs the course of new theoretical research.

Future of ASCI

ASCI has in place the foundations for 2- and 3-D codes based both on Eulerian and Lagrangian formulations and a computing infrastructure of more than 10 teraflops. The focus is now on integrating improved physics and engineering models into these codes and validating the codes against experimental data from both small-scale, nonnuclear integral tests and past underground nuclear tests. Once these codes can predict the baseline performance of nuclear weapons, they will become new repositories of expert designer judgment, as well as the best scientific tools for simulating the performance of the complex weapons currently in the stockpile as those weapons age or are modified. It is widely recognized that such simulation capabilities are essential if the National Nuclear Security Administration is to meet its statutory responsibility to assess and certify the stockpile annually. The ASCI codes will represent the ultimate integration of the theoretical and experimental efforts taking place within the stockpile stewardship program.

Inherent in the ASCI strategy is a tension between addressing the long-term simulation requirements of the weapons program and satisfying the weapons designers' needs for short-term improvements. For the last several years, the question has been, "When will ASCI deliver on its promises to the design community?" This question started in the hallways and is now the theme in Washington. The workload associated with these competing requirements has stretched the ASCI code teams almost beyond endurance. However, a new strategy based on a more realistic assessment of current computer resources is emerging. Heroic efforts have produced prototypes of 3-D weapons-system performance simulations, but those efforts have also shown that significant increases in computer power are required before the design community can routinely run high-fidelity 3-D simulations. Thus, until suitable platforms are available, the program will focus its efforts on developing and validating a production capability for 2-D spatially resolved simulations (see Figure 3). This goal is a better match to the current parallel platforms and computing infrastructure. Efforts to validate the codes and interactions with the design community will drive the development of more predictive material and physics models. As more-capable platforms become available, we will leverage those activities toward 3-D predictive capabilities.

The new strategy is directed at satisfying the current and anticipated designer requirements. ASCI is taking conscious steps to integrate its efforts more tightly with the ongoing work of the weapons designers. For example, we are currently aligning the ASCI milestones with the work that the code users must perform in support of stockpile assessment and certification. The ASCI milestones, which are reviewed periodically by an external review committee of experts in scientific computation, will continue to ensure steady improvement in the simulation capabilities for assessing and certifying a safe, secure, and reliable nuclear weapons stockpile.

Figure 4 shows the structure for implementing the new strategy. Each of the eight elements of the ASCI program is shown with a line of sight to the designer's desktop, meaning that stockpile certification and design requirements will continually guide the planning and execution of ASCI program elements.

Despite the modification in strategy, we still plan to deliver high-fidelity full-system physics characterizations of a nuclear weapon in 2009. At that time, we will also deliver a suite of validated codes, running on supercomputer platforms acquired through open procurement. Accompanying the codes will be user-friendly environments, advanced visualization tools for analysis, and the entire support structure to tie the components together. ASCI will also deliver high-performance storage, sophisticated solvers for linear systems, and high-bandwidth networks. In support of a true trilaboratory effort, ASCI continues to push the envelope in computing across platforms located at great distances from each other and in advanced encryption techniques and other approaches to ensure secure networking.

The process of quantifying margins and uncertainties in nuclear weapons systems will continue to influence ASCI priorities. In turn, ASCI's efforts to produce high-fidelity simulations will increase the predictive science capability and thus reduce uncertainties (see the article "QMU and Nuclear Weapons Certification" on page 47). The elements required to increase our predictive capability are shown in Figure 5. Building on Laboratory basic-research activities and external collaborations, ASCI will ensure that the tools needed to support the simulation of the most complex physics devices ever modeled will be ready when needed. Only by continual investment in fundamental science can we create the realistic models and predictive capability that will enable code developers and weapons designers to address the problems presented by the aging nuclear stockpile. ■

James Peery graduated from Texas A&M University with a Ph.D. degree in nuclear engineering in 1990. Before joining Los Alamos National Laboratory, James worked for almost 13 years at Sandia National Laboratories, where in his last assignment, he managed the Computational Solid Mechanics and Structural Dynamics Department. James was responsible for the development of state-of-the-art, massively parallel computational tools in the fields of transient dynamics, quasistatics, nonlinear implicit dynamics and structural dynamics. In addition, James was the ASCI program manager for STS Normal Environments and led coordination efforts for 12 code projects within the SIERRA architecture. Before the preceding assignment, James managed the Computation Physics Department and was responsible for the ALEGRA high-energy-density and CTH shock physics code projects. As a staff member, James was responsible for ALEGRA's project leadership. James' major research areas are in ALE algorithms and parallel algorithms. Currently, James is the deputy associate director of the Weapons Physics Directorate and is program director for the Los Alamos Advanced Computing and Simulation (ASCI) Program.