NATIONAL SECURITY

SCIENCE

But will it work?

Also in this issue

Roadrunner—On the Road to Trinity
Big Data, Fast Data
Killing Killer Asteroids
Strategic Deterrence in the 21st Century

Los Alamos NATIONAL LABORATORY

FROM TRINITY TO TRINITY

A New Frontier in Supercomputing

CHARLES McMILLAN, Laboratory Director

WELCOME to this issue of National Security Science.

This year, the Laboratory is celebrating its 70th anniversary, and this issue's principal topic—advanced supercomputing—is one that is dear to my heart: not only can today's supercomputers trace their origins back to computers used in the Manhattan Project, but I was privileged to work with colleagues in the early 1990s on the formation of the Accelerated Strategic Computing Initiative (ASCI) program. ASCI launched the modern generation of supercomputers; it was a key component of America's plan to move from reliance on full-scale underground nuclear tests to much more sophisticated computer-based simulations of full-scale weapon detonations.

When the United States stopped underground nuclear testing in 1992, a mission of the nuclear security labs fundamentally changed: we went from designing, building, and testing nuclear weapons to using our science and engineering capabilities to ensure that the stockpile of current weapons remained safe, secure, and effective into the future.

So far, we have succeeded in this mission. Last September, along with the directors of the other national security laboratories, I sent my Annual Assessment Letter to the secretaries of Energy and Defense and to the Nuclear Weapons Council. My letter states that after a thorough technical evaluation by Laboratory staff, I can give my assurance that the weapons in the stockpile that Los Alamos stewards have no reliability or safety concerns that must be resolved with underground testing. My predecessors have sent similar letters of assurance every year for the past 17 years.

Without supercomputing, we could not do the technical work that underwrites the Annual Assessment Letter. The weapons in the stockpile are built of thousands of components that are now beyond their expected lifespan. These aging components must be continuously studied and evaluated. Any problems that are found must be identified and addressed.

In the absence of continued real-world underground nuclear testing, the national security laboratories have increased their reliance on simulations, which have demonstrated increasing fidelity from year to year. This increased fidelity has placed growing demands on our supercomputers. These supercomputers are unique, designed and built at the laboratories in partnership with commercial venders. These new designs have found their way into commercial and academic applications. It is not an exaggeration to say that the needs of Stockpile Stewardship have driven the evolution of the world's supercomputing capability.

Because the weapons detonate in 3D, we need high-resolution 3D simulations to observe their performance. This is a much greater challenge than the one faced by the Manhattan Project; it is a challenge to build a nuclear weapon that works, but it is much harder to understand how and why it works—we need this level of knowledge to simulate it in high-resolution 3D.

When we at ASCI first estimated what we would need by now in high-performance computing, we underestimated. In my view, we must continue to advance the power and resolution of our computers to do our mission; the ongoing weapon life-extension programs and our annual assessment of the deterrent depend on it.

This means a new frontier in supercomputing, one we are calling Trinity. With Trinity, we come full circle: the Trinity Test of 1945 was the first full-scale, real-world test of a nuclear weapon; with the new Trinity supercomputer, our goal will be to provide the computing power to explore one of the most challenging puzzles remaining from nuclear testing; a puzzle that has eluded solution for almost 70 years.

I hope this issue of National Security Science *leaves you with a better understanding of why and how supercomputing is key to the Annual Assessment Letter and why our supercomputing capabilities must continue to grow.*

About the Cover

The U.S. nuclear stockpile is aging. Is it still reliable? Los Alamos uses supercomputer simulations to help assess the reliability of the nuclear deterrent.

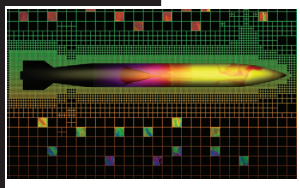

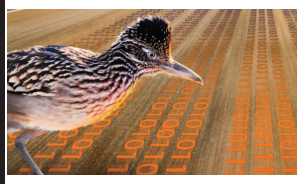

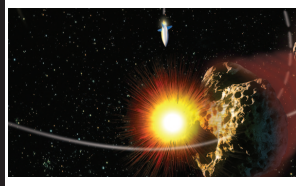

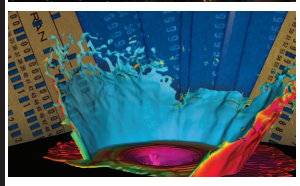

2 — *But Will It Work?*

7 — *Roadrunner—On the Road to Trinity*

14 — *The Tip of the Iceberg*
*Say, what **is** "under the floor" of a supercomputer?*

24 — *Big Data, Fast Data*

30 — *Killing Killer Asteroids*

34 — *Punched Cards to Petaflops*

42 — *Strategic Deterrence in the 21st Century*

52 — *Interview with Bob Webster*
Associate Director for Weapons Physics

57 — *Reflections*

But will it work?

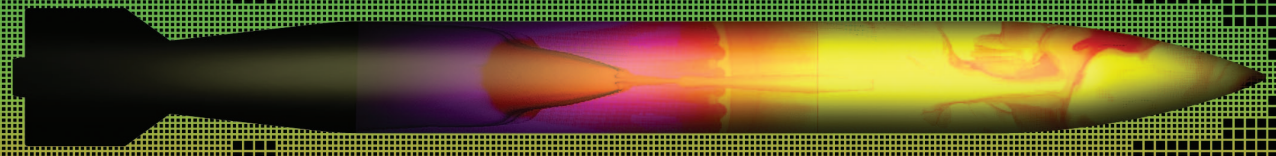

The threat is real. The security challenges confronting the United States in the 21st century are complex and multifaceted, and they demand the best science, technology, and engineering.

North Korea, despite international sanctions, continues to pursue both nuclear-weapon and ballistic-missile technology. On January 24 of this year, following action by the United Nations condemning the December launch of a missile, the North Korean Defense Commission issued the following statement: "We are not disguising the fact that the various satellites and long range rockets that we will fire and the high level nuclear test we will carry out are targeted at the United States."

Shortly thereafter, North Korea conducted its third nuclear weapons test.

Equally intransigent, Iran recently informed the International Atomic Energy Agency that it would introduce new centrifuges to its main uranium enrichment plant near Natanz. The new IR-2m units will allow Iran to enrich uranium at higher rates.

What could go wrong with a nuclear bomb or warhead?

Will It Work?

The end of the Cold War marked a turning point in how the United States maintained the safety, security, and reliability of its nuclear deterrent. From 1945 until 1992, the United States routinely replaced nuclear weapons with new systems that had been designed, tested, and fielded using a series of simulations, experiments, and tests at the laboratories, including underground nuclear weapons tests in Nevada. That stopped with the last underground test (Divider) on September 23, 1992.

So what could go wrong with a nuclear bomb or warhead? The nation has maintained nuclear weapons for more than half a century, so we must know all there is to know. After all, they are pretty much just an explosive with a detonator system inside a metal casing, right?

Not Just Another Bomb

Compared with that of conventional weapons, the effectiveness of modern nuclear weapons is a remarkable feat of physics and engineering; nuclear weapons have the almost unbelievable capability, using plutonium and uranium, to convert a few pounds of these elements into the explosive equivalent of thousands or millions of tons of TNT. This is about 100 times more destructive power than was released by the bombs that devastated Hiroshima and Nagasaki.

To achieve this kind of yield while meeting demanding safety and security standards is not easy. The complexity of a nuclear weapon is profound. Whereas conventional bombs and warheads use a simple design of a high explosive and a relatively simple detonator, nuclear weapons use complicated, high-precision mechanisms.

By itself (not including its delivery system, such as a missile), a nuclear weapon consists of many thousands of highly engineered, precision-crafted components, including complex electrical systems. Components can be made of steel, aluminum, silicon, and even plastic. These components must be made small enough and precise enough so that once assembled, the entire system will fit inside an 11.5-foot-long by 1-foot-diameter bomb or so that several can ride inside the nose cone of a missile. To manage this feat, components are sometimes manufactured to tolerances many times smaller than a human hair.

More to the point, to be successful, the interactions of the weapon components have to mesh precisely to initiate the most complex of *natural physical processes* and to stimulate them to work together, synergistically. These complex processes include chemistries**,** solid-state physics, plasma physics, and nuclear and thermonuclear reactions—the energy source of the Sun and stars.

Individually, many of these components and processes are far from being completely understood, even today. How and why these all work together to create a successful nuclear weapon explosion remains elusive.

The nuclear weapons in the U.S. stockpile were designed and built to be replaced every 10 to 15 years, or sooner. These weapons have lived beyond their expected lifespan, and their components continue to age.

Weapons with Crow's Feet

The nuclear weapons in the U.S. stockpile were designed and built to be replaced with new designs and builds every 10 to 15 years, or sooner if the U.S. defense strategies required it. Now these weapons have lived beyond their expected lifespans, and their components continue to age.

Over time, plastics become brittle and break down, releasing gases. Electronic systems based on antiquated technologies like vacuum tubes and plastic-coated copper wire corrode. Adhesives bonding key components weaken. Metal coatings naturally deteriorate. Metals and metal joints corrode and weaken. Vibrations from transportation and deployment further impact components. Many of these issues are faced by every car owner. With years of environmental changes in temperature, pressure, and humidity, and in the presence of radioactive elements like plutonium and uranium, components degrade and may eventually fail to work.

In short, aging materials change and in so doing, change their physical, chemical, and mechanical properties; aged materials no longer behave as they once did. New behaviors can be unpredictable.

Nuclear weapons must work, and perfectly, *only* if the president of the United States has authorized their use—and they must *never* work if the president has not authorized their use. Can an aging stockpile meet these demanding requirements?

For example, using hypotheses, theories, and experimental trial and error, the swordsmiths of feudal Japan learned that when specific combinations of high- and low-carbon steel, along with other materials, were precisely processed by controlling temperature and carefully folding, welding, and quenching them, a superior sword was born: the katana, a.k.a. the samurai sword, the most feared and revered weapon of their time. They did not fully understand *why* these materials and processes worked in this way, but they could describe what materials to use and explain how to process them, and they could predict a consistent outcome. For the job at hand, this level of knowledge was sufficient.

Scientists could say with some confidence that they understood some of the weapon physics. But by no means did they claim to understand all of what they did and saw.

But explaining and predicting phenomena using testing does not necessarily mean *understanding* phenomena.

Turning Knobs

Real-world trial-and-error experimentation is sometimes called the "engineering approach" because it uses hands-on building and testing of theoretical concepts. This was largely the approach by which nuclear weapons were invented and by which they evolved for 47 years. Like Japanese swordsmiths, weapons scientists hypothesized, theorized, and experimented, using this and trying that with different materials, processes, and designs in very successful efforts to meet the requirements established by the U.S. military.

For the job at hand, this level of knowledge was sufficient and was codified in weapon simulation computer programs. A deeper understanding was certainly desired and sought out, but it was not necessary in order to accomplish the Cold War mission. Regardless, better tools were needed to better understand thermonuclear weapons.

Still, scientists' ability to explain and predict weapons phenomena got stronger, and an amazing body of knowledge grew, so they could say with some confidence that they understood some of the weapon physics. But by no means did they claim to understand *all* of what they did and saw. In fact, it was not uncommon for test results to contradict scientists' best predictions and call into question what they thought they understood.

The nation needed a new way to assess the stockpile. The answer would be the science-based Stockpile Stewardship Program.

When reality did not match their predictions, the scientists were often forced to adjust their calculations until these matched test results, but without really understanding why these adjustments worked. It was like the early days of radio, when the scientists and engineers understood the principles of the device but lacked the predictive power to design the radios to respond exactly. Radio response was "tuned" (often by turning a knob) to achieve the final high-precision match required so that radio transmitters and receivers could work together. Indeed, the practice by nuclear scientists of massaging calculations until they fit their real-world test results was called "turning knobs." The knobs were embedded in the weapon simulation computer codes. Thus, testing was done not only to see if a weapon worked, but also to try and eliminate particularly troubling knobs by gaining a better understanding of the weapon physics.

But then real-world underground testing in Nevada and deployment of new systems went away. What would take their place? Would the military retain confidence in systems aboard the submarines and planes and in missile silos? Could the president be assured that systems were safe, secure, and effective? Would allies and adversaries be convinced of the effectiveness of America's nuclear deterrent?

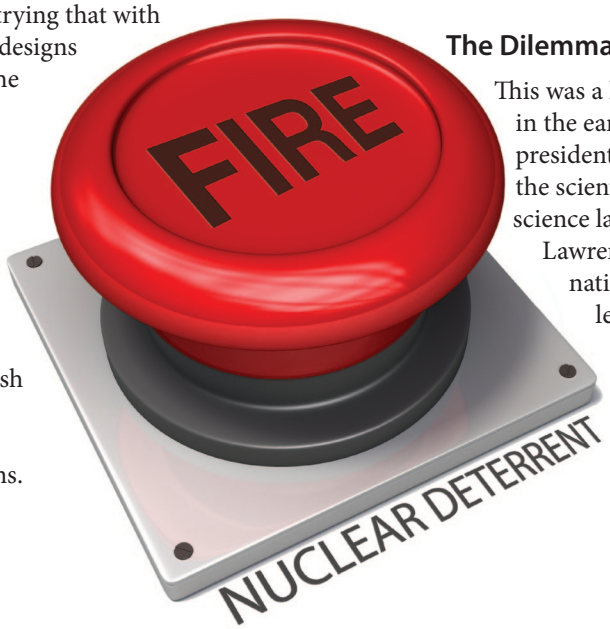

The Dilemma

This was a huge dilemma facing the nation in the early 1990s. Members of the president's cabinet, members of Congress, the scientists at the three national security science laboratories (the Los Alamos, Lawrence Livermore, and Sandia national laboratories), and military leaders debated whether it was wise to end weapons development and underground nuclear testing without having a satisfactory alternative in place. Without new production and with a ban

on underground testing, the nation needed something to ensure that stockpiled weapons would continue to work into the future, perhaps for decades. In other words, the nation needed a new way to assess the stockpile. The answer would be the science-based Stockpile Stewardship Program (SSP).

Today the SSP is applying the best experimental, computational, modeling and simulation, and engineering capabilities to provide the scientists and engineers at the laboratories with the tools to understand what is happening to the nation's deterrent. These tools are allowing the laboratory directors to successfully execute life-extension programs in support of Navy and Air Force systems, to resolve issues that arise in these aging nuclear systems. At the end of each fiscal year, the laboratories are required by law to report to the president of the United States, through the secretaries of Energy and Defense, on the state and health of the nation's deterrent.

Supercomputer Simulations

From the beginning of the Manhattan Project, Los Alamos has relied on experimental data and weapon simulations running on state-of-the-art computers when designing weapons. During the Cold War the national security laboratories continued to use the most powerful and advanced computers for weapons simulations. The Department of Energy's Accelerated Strategic Computing Initiative (ASCI) program—now called the Advanced Simulation and Computing (ASC) program—was established in 1995 as a pillar of the Stockpile Stewardship Program (SSP). The goal was to enable high-resolution 3D simulations of nuclear weapons by 2005. The idea of executing high-resolution 3D simulations of a nuclear weapon was, in 1995, revolutionary.

The idea of executing high-resolution 3D simulations of a nuclear weapon was, in 1995, revolutionary.

Computing science in the early 1990s was not up to the task. This was the era of floppy discs, Apple's Newton and Macintosh computers, and Windows 95. In fact, the notion of being able to build computers that had the power, speed, and memory needed to accurately model and simulate a nuclear weapon explosion from first principles, much less in 3D, was thought nearly impossible by many.

For example, a standard unit of computer speed and power is floating-point operations per second (flops). An initial calculation done at this time by Lawrence Livermore National Laboratory concluded that more than 100 teraflops (*100 trillion* flops) would be required to execute the high-resolution 3D simulations, with sufficient accuracy, for the SSP. But at the time, Livermore's most powerful computer provided only 13.7 gigaflops (*13.7 million* flops). This meant that its computing power would need a 7,000-fold increase in less than a decade, implying a technological growth rate many times that given by Moore's Law—the industry standard for predicting increases in computing power. To run the models and simulations then envisioned, the laboratories needed significantly larger, faster, and more powerful computers than Moore's Law allowed.

High-resolution 3D weapons simulations would require vast leaps in supercomputer design, development, programming, and computing power. It would also require unprecedented levels of electric power. This meant needing large new infrastructures to provide power and cooling.

Avatars Won't Work

It may be difficult for most people to grasp the difficulty of creating 3D simulations for the SSP. After all, computer-generated 3D graphical representations of nuclear events can be made for the movies. Hollywood produces simulations that *appear* to be real but do not need to reflect reality; in contrast, the SSP needs simulations that reflect how nature

really works. Weapons scientists must produce a high-resolution representation of real events, as nature would unfold them.

To accomplish this, they must rely on the quality and accuracy of their experimental data, models, and programs. There is no tolerance for any "garbage in, garbage out" dynamic in nuclear weapons science. Whenever possible, those elements have been rigorously tested—and verified with the highest levels of confidence and validated against experimental data—before being used in simulations that will represent the real world.

A Choice of One

Despite the challenges of developing the advanced computing platforms and the codes, there was no other choice. The scientific basis for assessing the stockpile is formed by the ASCI tools, in partnership with new experimental tools like the Dual-Axis Radiographic Hydrotest facility at Los Alamos and the National Ignition Facility at Livermore, and with data from the more than 1,000 nuclear weapons tests conducted by the United States up until 1992.

There is no tolerance for any "garbage in, garbage out" dynamic in nuclear weapons science.

The demands of the SSP and the evolution of the ASC would eventually make computer modeling and simulation *an integral part* of science in general, changing the centuries-old way of doing science. Even for fields of science that can still perform real-world testing, virtual-world computer modeling and simulation have become a normal part of the scientific process. Today, computer simulations are a regular, key element—alongside theory and real-world testing—of practicing the scientific method.

The SSP is a successful and evolving effort that continues to push the state of the art. Incredibly, in 2005, the 100-teraflops goal was reached. But the SSP's supercomputing needs were *vastly* underestimated. Consequently, in 2008, Los Alamos unveiled the world's first petaflop (1.0 *million billion* flops) supercomputer, Roadrunner. Since Roadrunner, another petascale machine, Cielo, has come online at Los Alamos, and Livermore has stood up Sequoia at almost 16 petaflops. Los Alamos is proceeding with Trinity, a 30- to 40-petaflop machine. The machines at Los Alamos and at Livermore are working 24 hours a day, 7 days a week, 365 days a year. The demand for time on these and other machines by scientists is never-ending.

The SSP has successfully resolved problems related to aging, even problems in the original designs and manufacturing of some weapons, and has enabled corrections. It has made the life-extension programs for weapons a success. It has successfully resolved the need for some knobs in the weapon codes. It has provided simulations of nuclear weapons and their subsystems in 3D. It is important to understand that these 3D simulations are often referred to as "hero calculations," given the amount time (weeks, months, and in some cases, years) required to set up the code and run the calculation, even on petascale machines.

Supercomputers and the weapon codes have played a key role in all these successes. They will become even more important as the stockpile continues to age and the nation continues its moratorium on conducting real-world, full-scale tests.

Supercomputer simulations have allowed scientists and engineers to discover phenomena previously hidden to real-world experimentation, making it necessary to ask more questions, change some theories, and explore new directions. As a result, the need for high-resolution 3D simulations is clear. Indeed, some weapons issues can be accurately addressed *only* in 3D. However, high-resolution 3D simulations require vastly more powerful supercomputers than 2D simulations do.

Bigger Than Manhattan

To help put the SSP effort into perspective, it took almost $26 billion in today's money and two years for the scientists of the Manhattan Project to build the first atomic bombs, relatively simple devices compared with today's deterrent. But to understand how and why the weapons work remains a work in progress; the nuclear weapons community continues to pursue a complete understanding of nuclear weapons, almost seven decades after the Manhattan Project ended.

When the nation and our allies are banking on the reliability, safety, and security of the aging nuclear deterrent to protect them from ever more dangerous and unpredictable aggressors—but without detonating a weapon to absolutely, positively know the stockpile works—is there ever enough science to be done? As the stockpile shrinks and ages and as weapons rely more on replaced and rebuilt components, more questions come to the surface, and more science, not less, is required for future annual assessments.

In a world without continued real-world testing, experimental data coupled with high-performance computing, modeling, and simulation are *the* game in town.

The SSP is reliant on supercomputing, and the SSP is the *only* way to answer the question, *But will it work?*

~Clay Dillingham

ROADRUNNER
ON THE ROAD TO
TRINITY

For almost 50 years, nuclear weapons were tested explosively, first at the Pacific Proving Grounds and until 1992, at the Nevada Test Site. Now nuclear weapons are simulated on supercomputers using codes in which computational physicists have captured the essence of weapon performance.

The simulations are part of Stockpile Stewardship, the National Nuclear Security Administration (NNSA) program established as a means of assessing the safety, security, and reliability of the United States' stockpiled nuclear weapons in the absence of nuclear testing. The assesments drawn from simulations enable the directors of the three U.S. national security laboratories—Los Alamos, Lawrence Livermore, and Sandia—to annually inform the president of the United States, through the secretaries of Energy and Defense and the Nuclear Weapons Council, that the stockpile of current weapons remains, in their view, safe, secure, and effective.

Roadrunner was a step on the road to Trinity, and it was a huge step—one for the record books.

Providing the input that underwrites such an important assesment is a daunting task, so the supercomputers that weapons scientists use are always on computation's leading edge. Los Alamos National Laboratory's Roadrunner supercomputer won world fame for breaking a computing speed record, but Los Alamos is pushing its computing capabilities to newer, higher levels to meet the growing needs of Stockpile Stewardship. So the Laboratory is looking forward to its next big computer, a machine already named but still to be fully conceived: Trinity. Trinity will be designed and procured by the New Mexico Alliance for Computing at Extreme Scales, a partnership between Los Alamos and Sandia. Roadrunner was a step on the road to Trinity, and it was a huge step—one for the record books.

The Need for Speed

Like its namesake, the state bird of New Mexico, Roadrunner was certainly speedy. In June 2008 it took over first place on the Top500, the international list of the world's fastest supercomputers. It was the world's first petascale supercomputer, that is, the first to reach a sustained quadrillion (thousand trillion)—1,000,000,000,000,000—calculations (called "floating-point operations") per second: petaflops.

Roadrunner was first in another way as well. It had a unique design (architecture) for a supercomputer. Cheryl Wampler, the Laboratory's deputy director for the Advanced Simulation and Computing (ASC) program—the NNSA program that oversees the development of supercomputing technology for Stockpile Stewardship—describes Roadrunner as "a historic machine, a pioneer." Then she adds, "It was also controversial."

The controversy was about that unique architecture. Roadrunner was a "hybrid" in that it combined two different kinds of processors (chips that read and carry out program instructions). It had 6,563 dual-core general-purpose processors (AMD Opterons)—the kind found in almost all computers, except that each was actually two processors in one because of its two compute cores. Each core was then linked to a special graphics processor (PowerXCell 8i) called the Cell. The Cell was an enhanced version of a specialized processor originally designed for the Sony PlayStation 3. For Roadrunner the Cell was adapted specifically to support scientific computing.

Although hybrid computers had existed before Roadrunner, no one had ever tried that approach on such a large scale, and many doubted that a hybrid supercomputer would work. So for Los Alamos and IBM, who collaborated to design the computer, Roadrunner was a leap of faith—but a leap with purpose.

Roadrunner's processors were housed in 17 rows of "racks," the tall black cabinets shown here. The racks covered almost 6,000 square feet of floor space. Roadrunner's "hybrid" architecture made it the first petaflop supercomputer and gave it exceptional energy efficiency. (Photo: Los Alamos)

High-speed calculation was certainly part of the purpose, as Jim Lujan, in the Laboratory's High Performance Computing Division, explains: "We wanted to take the next major step in computing. Computers get faster every year, but we needed an exponential increase to solve the physics problems we have in the NNSA."

The physics problems associated with nuclear weapon simulations are both big and complex, especially since the simulations need to be of an entire weapon, not just its individual components. Only a very fast supercomputer can complete all the calculations involved in a reasonable amount of time.

Computers get faster every year, but we needed an exponential increase to solve the physics problems we have in the NNSA.

"We're looking at all the physics involved with a weapon, from initiation to completion of a detonation," says Lujan. "With slower, less powerful computers, we were only able to do pieces—primaries, secondaries, or parts of the ignition. Now we're looking at the whole system, end-to-end. We could have done it 10 years ago with the computing systems we had then, but it would have taken years upon years upon years, and that doesn't make it a tractable problem."

It's All in the Details

Although simulation only *predicts* weapon performance, when nuclear testing *demonstrated* it, simulation pays a dividend in detail—detail about all the internal processes that feed into performance.

The scientists who develop weapons codes lay the groundwork for that dividend by including equations and data representing all aspects of a weapon: the size and shape of components, the chemical makeup and behavioral tendencies of constituent materials (from plastics to plutonium), and the phenomena acting on everything inside a weapon during detonation. Data collected during the days of nuclear testing inform the code developers, as do new data drawn from experiments on the many materials used in nuclear weapons and on data from *non*nuclear tests of weapon components.

When a computer is fast enough to work all of a code's equations in a reasonable time span, the resulting simulations deepen scientists' understanding of weapon behavior. The faster the computer, the more equations it can handle and, therefore, the more detail its simulations can provide about smaller pieces of the whole.

That small pieces are critical to the fate of an entire system is just as true for nuclear weapons as it was for the space shuttle Challenger on January 18, 1986, when the failure of an O-ring just 0.280 inch in diameter led to the space craft's destruction and the death of everyone aboard.

So scientists use simulation to learn how all parts of a weapon behave during detonation and use what they learn to predict the continued viability of the weapons in the stockpile. The information harvested from each simulation, along with the data from continued experiments and nonnuclear tests, is then used to improve the codes so that future simulations, provided they are run on a fast-enough computer, will be even more detailed and more accurately predictive of performance.

Matching Code to Computer

Roadrunner's speed was derived from its architecture. The two processors shared functions, with the Cell taking on the most computationally intense parts of a calculation to hurry the work along. Essentially, the Cell acted as a computational accelerator. But the division of tasks was not automatic; it was achieved through programming, so preexisting codes had to be revised by their developers. In fact, the code work began years before Roadrunner was developed.

Rewriting a code for Roadrunner required developers to puzzle out the best way to divide calculations between the processors and to carefully consider how to distribute data so that it would be in the right place when needed. That kind of rethinking was a significant challenge, but there were rewards for successful code revision. In work for Stockpile Stewardship, Roadrunner took on a troubling, long-standing gap in understanding about how energy flows in a detonation and how yield is dependent upon that flow. The big computer made a significant contribution to that understanding.

As code developers grappled with the new machine and its unfamiliar architecture, they were hearing a warning shot. They were being alerted to upcoming changes in high-performance computing.

General-science researchers also had time on Roadrunner, during the computer's first six months of work—its "shakedown" period before it was transitioned to classified status. Scientists using the "ML" code used Roadrunner to study the genetic sequences of HIV and to build an HIV family tree that they hope will help in the search for a vaccine for the virus. Other codes were used to produce breakthroughs in fields such as materials science, astronomy, and laser-plasma interaction.

Brian Albright of the Computational Physics Division was one of several researchers using the VPIC (vector particle-in-cell) code for projects that included the laser-plasma interaction studies. He noted that adjusting the code for Roadrunner had far-reaching benefits. "We set up several code teams to adapt VPIC for Roadrunner, and the changes made VPIC a more general code, not tied to a single computer architecture."

A Shot across the Bow

As demonstrated by successes with VPIC, code work done to accommodate Roadrunner will be generically applicable to emerging new types of supercomputers. That fact highlights another significant purpose behind Roadrunner: demonstrating that the weapons codes need to be broken free of their dependence on computer architectures that are rapidly becoming antiquated. As code developers grappled with the new machine and its unfamiliar architecture, they were hearing a warning shot. They were being alerted to upcoming changes in high-performance computing—changes that everyone working with Stockpile Stewardship is now embracing.

As Wampler explains it, "Roadrunner was an 'Advanced Architecture' supercomputer, a category defined by the ASC program. Advanced Architecture machines are meant to evolve computing into the future."

As expected with any ahead-of-the-curve technology, Roadrunner could seem to be, as Wampler describes it, "a little flaky," but that was because of its evolutionary nature; it was not necessarily intended to be used, or used comfortably, by all codes. However, Roadrunner did exactly what it *was* intended to do regarding the weapons codes: it got the codes moving toward new architectures. The codes will *have* to run on new architectures . . . and soon. Roadrunner was challenging because the supercomputing future will be challenging.

Cielo—End of an Era

Things are already changing in the Laboratory's Nicholas C. Metropolis Center for Modeling and Simulation (also known as the Strategic Computing Center, or SCC), where Roadrunner was installed. Roadrunner's time has run out, and it is going away. The big machine will soon be dismantled and will assume its rightful place in history as the focal point of dramatic change at the start of a new era in computing history.

What remains behind at the SCC is an equally powerful supercomputer named Cielo, which will be active for the next few years and is already hard at work; it has been overlapping with Roadrunner since January 2011. Cielo is more like previous Stockpile Stewardship workhorses (not a hybrid) while operating at Roadrunner scale: 1.43 petaflops. It was built to accommodate the present-day weapons codes—no

An explosion containment vessel is shown being prepared for a test at the Laboratory's Dual-Axis Radiographic Hydrotest (DARHT) facility. DARHT is one of several Los Alamos facilities where tests and experiments generate data that Los Alamos scientists incorporate into codes that are run on supercomputers like Roadrunner for simulations of full-scale nuclear weapon tests. (Photo: Los Alamos)

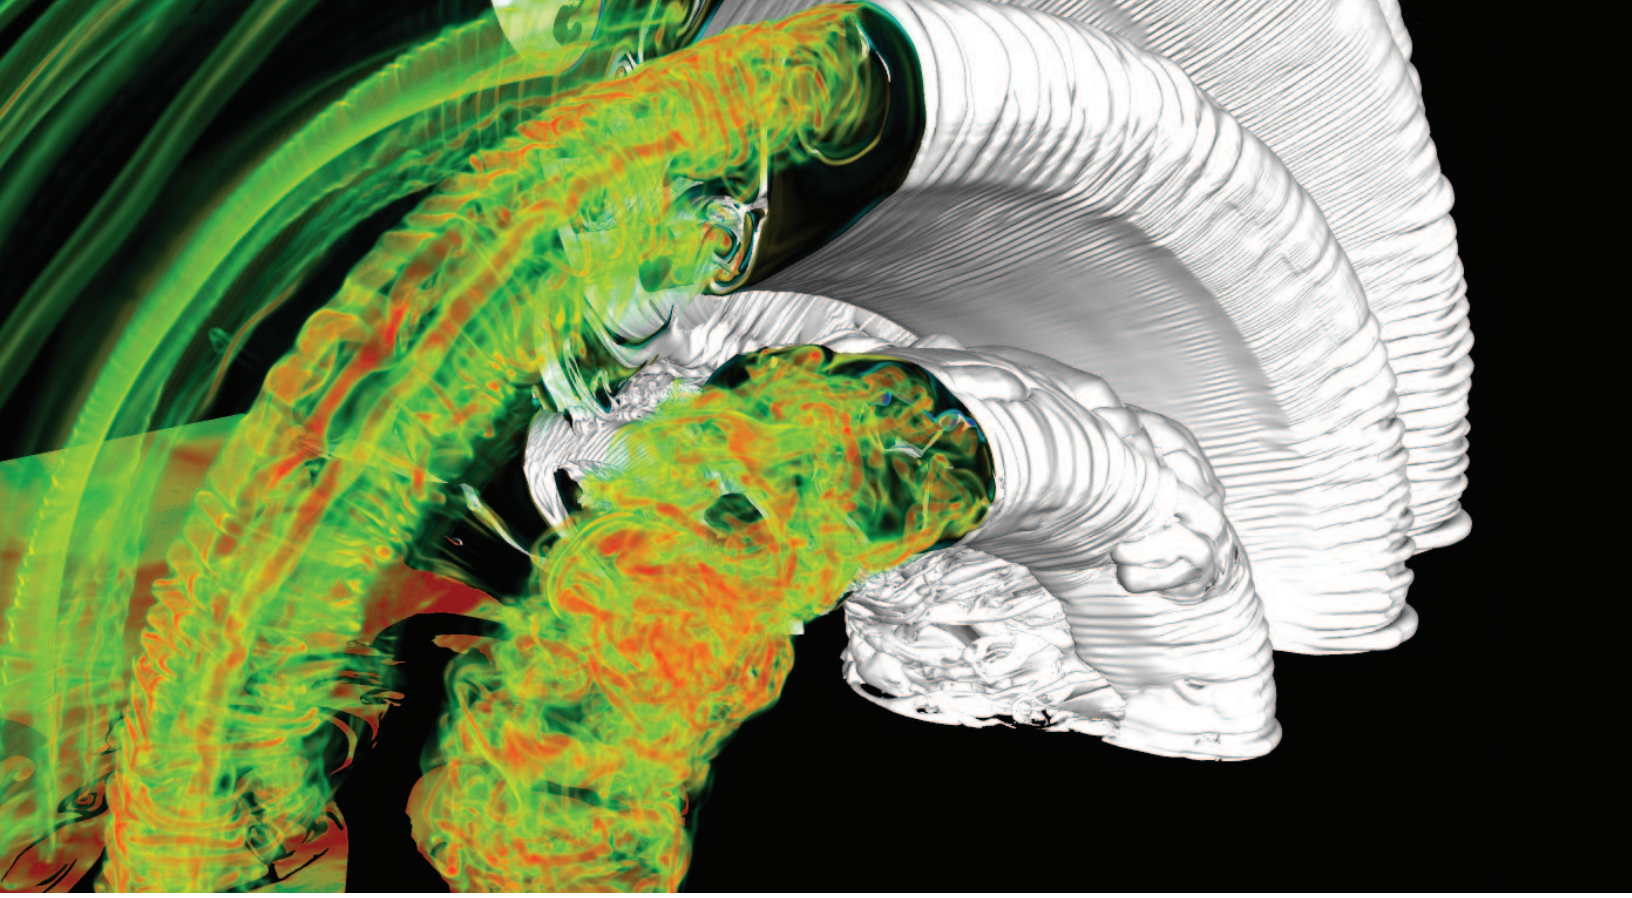

At Los Alamos, simulations such as those run on Roadrunner and currently on Cielo can be seen as full-color 3D visualizations such as this one showing the growth of turbulence in the gas inside an inertial confinement fusion capsule. This image is a time shot from a full simulation done on Cielo. (Photo: Los Alamos)

extensive adaptations needed—and its continuity with previous systems will give *all* the codes time to be prepared for future machines.

After Cielo, we won't be able to say, "Give me another one just like the last one." High-performance computing is moving in a new direction, and we have to move with it.

The future is coming quickly because Cielo will be on its way out by 2015. And when it is gone, it will *really* be gone. Cielo is the last of its kind.

"After Cielo," says Wampler, "we won't be able to say, 'Give me another one just like the last one.' High-performance computing is moving in a new direction, and we have to move with it."

Keeping up is essential because the national laboratories do not typically have their supercomputer's hardware built for them from scratch. Instead, their computers are built with commercially available technology as the starting point. That approach is more economical. Besides, the most rapid changes in computer technology are occurring in the private sector; the national laboratories must track commercial technology trends if their supercomputers are to be on the cutting edge, easily maintained, and able to minimize downtime.

Market trends are driving changes, with companies trying new architectures and developing new technologies such as different kinds of processors, memory units, and interconnects, and all those innovations add up to an entirely new software environment.

But the needs of Stockpile Stewardship are also driving change. The program's demands are growing and require greater and greater computational power because nuclear testing, and the certainty it provided about weapon reliability, is receding further and further into the past. Indeed the stockpiled weapons are not the same as they were when nuclear tests produced as much information as we were able to capture. The weapons' components, materials, and systems have aged, possibly changing in ways that make them behave differently. Some of the components have been remanufactured. As complex as modern nuclear weapons already are, they become even more complex as they age, so time is raising new questions that are waiting for answers.

Trinity is expected to be the first platform large enough and fast enough to begin to accommodate finely resolved 3D calculations for full-scale, end-to-end weapons calculations.

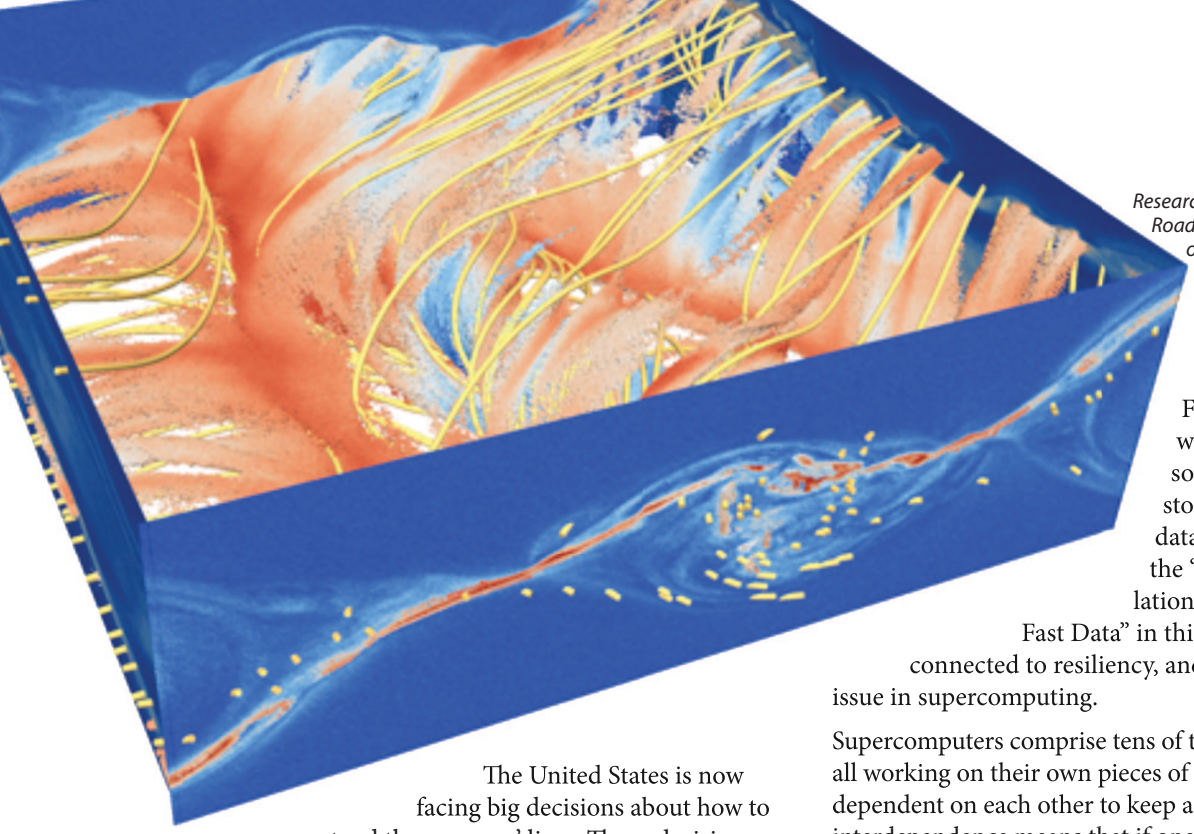

Researchers used the VPIC code on Roadrunner to produce this simulation of magnetic reconnection. VPIC was also used to simulate laser-plasma interactions.

The United States is now facing big decisions about how to extend the weapons' lives. Those decisions will depend on how much scientists know about what makes a weapon work . . . or not work. As a result, the national security labs need even bigger and more detailed simulations with superior resolution, in 3D, for capturing fine-scale features. And the simulations will need to cover all the physics associated with a full-scale weapon.

In other words, the laboratories will need simulations that are vastly more predictive than what today's supercomputers can provide. So the national security science labs have to take dead aim on the supercomputing future.

Forward to Trinity

Eventually, maybe as soon as 2020, supercomputers will reach exascale—one quintillion (1,000,000,000,000,000,000) calculations per second—making them 1,000 times faster than Roadrunner. Such speed bodes well for the needs of the national security labs. But along with being fast, future supercomputers will need to be energy efficient if they are to be welcome in an increasingly energy-conscious world.

Bigger systems have many, many more processors, each of which can fail, causing the whole machine to stop.

Roadrunner made great strides in energy efficiency. Although bigger and older than Cielo, Roadrunner used significantly less energy to achieve essentially the same speed. Future supercomputers will need to emulate that success . . . and improve on it if speed and power are to continue increasing.

Future supercomputers will also need new memory solutions for handling and storing the vast amounts of data used in simulations and the "big data" those same simulations generate (see "Big Data, Fast Data" in this issue). Memory is closely connected to resiliency, and resiliency is a very big issue in supercomputing.

Supercomputers comprise tens of thousands of processors all working on their own pieces of a calculation and all dependent on each other to keep a calculation going. That interdependence means that if one processor fails, they all have to stop. And with more and more processors coming into play as systems grow, the number of failures is growing as well.

Supercomputing codes mitigate the losses from such failures by "checkpointing," halting at regular intervals to save the current state of a calculation to storage (long-term memory). Then when a failure occurs, the latest checkpoint can be recalled and the calculation restarted at that point. The faster a checkpoint can be saved and the faster it can be recalled after a failure, the less downtime there will be—and that is resiliency. The designers of tomorrow's supercomputers will need to look hard at ways to make them as resilient as possible.

Getting to better efficiency and resiliency, and reaching exascale, will not happen all at once. It is happening in stages, and the next stage of the journey is Los Alamos' next addition: Trinity.

Roadrunner got everyone thinking in new ways about how to build a supercomputer.

Trinity will be an Advanced Technology System—a new ASC category replacing Advanced Architecture. This new machine, tentatively projected for installation in 2015–2016, could be 40 to 50 times faster than Roadrunner and Cielo. Although it is expected to be, like Roadrunner, a significant break with the past, it will have to serve, like Cielo, as the working platform for *all* the Stockpile Stewardship codes that need to run at large scales. As such, Trinity is expected to be the first platform large enough and fast enough to begin to accommodate finely resolved 3D calculations for full-scale, end-to-end weapons calculations.

Exactly what Trinity will be like is still under discussion. Trinity's designers are thinking in particular about what Gary Grider, of the Laboratory's High Performance Computing Division, calls "being smarter about the use of silicon"—silicon being the key material that makes up the transistors on the computer's processors (chips).

When Roadrunner's specialized processor, the Cell, assumed some functions that the Opteron processors would otherwise have done, it was taking over operations that were not only computationally complex but also computationally *expensive*—requiring a great deal of time and power and many floating-point operations per second.

Using the Cell for complex and expensive functions was Roadrunner's example of using silicon "smartly." If all a computer's functions are performed on a general-purpose processor (like the Opteron), there will be times when only parts of the processor are in use, even though all parts of the processor, even the inactive ones, are being energized. That adds up to wasted energy, a "dumb" use of silicon.

So Roadrunner's division of labor enabled not only greater speed but also greater energy efficiency. It was a strategy that landed Roadrunner on the Green500 list, a measure of world supercomputer efficiency, in addition to being on the speedy Top 500 list.

"Roadrunner wasn't the first time anyone had thought of using specialized processors that way," says Grider, "but it was the first time anyone had done a really big demonstration of what that means. So we produced a petaflop machine using 3 megawatts of power. At the same time, the Jaguar supercomputer at Oak Ridge used 8 megawatts of power to get a petaflop because they were doing things the old-fashioned way—using only general-purpose processors."

Another old-fashioned way to use silicon involves the movement of data to and from processors. Says Grider, "Computer programs today are written on the model that you move data to the processors, you do calculations, and then you write the results out to memory [storage]. Then later you move data back out of memory to be returned to the processors for new calculations. We can't afford to do that anymore because moving data around requires a lot of power—it's not very efficient. In addition, the time it takes to move data is wasted. It's time spent *not* computing."

There is more than one possible solution. Why not have some processing done right there in memory? Or, why not store data closer to processors? The Laboratory's computer scientists began to think of both possibilities as Roadrunner's huge computations raised questions about efficient versus inefficient data movement.

The Smart Supercomputing Future

Roadrunner got everyone thinking in new ways about how to build a supercomputer. Specialized processors are already being included; Lawrence Livermore's Sequoia, the fastest supercomputer in the world in June 2012, and Oak Ridge's Titan, the current fastest, are examples of that. "So our demonstration with Roadrunner," Grider concludes, "caused everyone to pay attention."

Trinity is on the way to exascale—and who knows what else lies beyond.

Certainly Trinity's designers are paying attention. They have put a lot of new ideas on the table, not only a greater use of specialized processors but also the creation of layered memory, with a "flash" memory to improve efficiency in data movement between processors. There are a lot of decisions to be made before the machine becomes a reality, and the final picture is still out of focus.

But if the exact nature of Trinity is still uncertain, what *is* certain is that it will not do what Cielo does—provide a comfortable environment for current, unadapted weapons codes. So the code developers at Los Alamos and at all the national security labs are already working to radically change their codes to fit the radical architectures to come.

Trinity is on the way. Exascale—and who knows what else—lies beyond. Roadrunner's warning came just in time for everyone to get ready.

~Eileen Patterson

THE TIP OF THE ICEBERG

Say, what is
"under the floor"
of a supercomputer?

Every year for the past 17 years, the director of Los Alamos National Laboratory has had a legally required task: write a letter—a personal assessment of Los Alamos–designed warheads and bombs in the U.S. nuclear stockpile. This letter is sent to the secretaries of Energy and Defense and to the Nuclear Weapons Council. Through them the letter goes to the president of the United States.

The technical basis for the director's assessment comes from the Laboratory's ongoing execution of the nation's Stockpile Stewardship Program; Los Alamos' mission is to study its portion of the aging stockpile, find any problems, and address them. And for the past 17 years, the director's letter has said, in effect, that any problems that have arisen in Los Alamos weapons are being addressed and resolved without the need for full-scale underground nuclear testing.

When it comes to the Laboratory's work on the annual assessment, the director's letter is just the tip of the iceberg. The director composes the letter with the expert advice of the Laboratory's nuclear weapons experts, who, in turn, depend on the results from another year's worth of intense scientific investigation and analysis done across the 36 square miles of Laboratory property.

One key component of all that work, the one that the director and the Laboratory's experts depend on to an ever-increasing degree, is the Laboratory's supercomputers. In the absence of real-world testing, supercomputers provide the only viable alternative for assessing the safety, reliability, and performance of the stockpile: virtual-world simulations.

I, Iceberg

Hollywood movies such as the *Matrix* series or *I, Robot* typically portray supercomputers as massive, room-filling machines that churn out answers to the most complex questions—all by themselves. In fact, like the director's Annual Assessment Letter, supercomputers are themselves the tip of an iceberg.

Without people, a supercomputer would be no more than a humble jumble of wires, bits, and boxes.

Although these rows of huge machines are the most visible component of supercomputing, they are but one leg of today's supercomputing environment, which has three main components. The first leg is the supercomputers, which are the processors that run the simulations. The triad also includes a huge, separate system for storing simulation data (and other data). This leg is composed of racks of shelves containing thousands of data-storage disks sealed inside temperature- and humidity-controlled automated libraries. Remotely controlled robotic "librarians" are sent to retrieve the desired disks or return them to the shelves after they are

The most important assets in the Laboratory's supercomputing environment are the people—designing, building, programming, and maintaining the computers that have become such a critical part of national security science. (Photo: Los Alamos)

played on the libraries' disk readers. The third leg consists of the many *non*-supercomputers at the national security laboratories. The users of these computers request their data, over specially designed networks, from the robotic librarians so they can visualize and analyze the simulations from afar.

The Los Alamos supercomputers are supported by a grand infrastructure of equipment used to cool the supercomputers and to feed them the enormous amounts of electrical power they need. They also need vast amounts of experimental data as input for the simulation codes they run, along with the simulation codes themselves (also called programs, or applications), tailored to run efficiently on the supercomputers. In addition, system software is necessary to execute the codes, manage the flow of work, and store and analyze data.

People are the most vital component of any supercomputer's supporting infrastructure. It takes hundreds of computer scientists, engineers, and support staff to design, build, maintain, and operate a supercomputer and all the system software and codes it takes to do valuable science. Without such people, a supercomputer would be no more than a humble jumble of wires, bits, and boxes.

The computer room's vast floor space is 43,500 square feet, essentially an acre—90 percent of a football field.

Supercomputers That Fill a Stadium

At Los Alamos, supercomputers, and the immense amount of machinery that backs them up, are in the Nicholas C. Metropolis Center for Modeling and Simulation, known pragmatically as the Strategic Computing Center (SCC).

Roadrunner, the world's first petaflop computer, joined other supercomputers in the SCC's computer room in 2008. It is a *big* machine, containing 57 miles of fiber-optic cables and weighing a half-million pounds. It covers over 6,000 square feet of floor space, 1,200 square feet more than a football field's end zone. But that represents only a portion of the computer room's vast floor space, which is 43,500 square feet, essentially an acre—90 percent of a football field (minus the end zones). (Roadrunner has finished its work for the Laboratory and is currently being shut down.)

What is really amazing, however, lies *beneath* the supercomputer room floor. A trip straight down reveals more vast spaces crowded with machinery that users never see.

The computer room is the SCC's second floor, but that one floor is actually *two*, separated by almost four feet. That 4-foot space hosts the miles of bundled network cables, electrical power lines inside large-diameter conduit, and other subfloor equipment the supercomputers rely on. The double floor provides enough room for engineers and

The Metropolis Center, also called the Strategic Computing Center, is the home of Los Alamos' supercomputers and the vast infrastructure that supports them. (Photo: Los Alamos)

maintenance staff, decked out like spelunkers in hardhats and headlamps, to build and manage these subfloor systems.

Below this double floor, on the building's first floor, is another acre-size room, a half-acre of which holds row upon row of cabin-size air-conditioning units. These cool the air and then blow it upwards into the computing room, where it draws the heat off the hard-working computers. The now-warmed air then rises to the third floor (basically an acre of empty space), whereupon it is drawn back down, at the rate of 2.5 million cubic feet per minute, to the first floor by the air coolers so the cooling cycle can begin again.

An additional half-acre of floor space stretches beyond the cooling room and holds the SCC's electric power infrastructure, the machines that collectively keep the supercomputers running. There are rows of towering power distribution units (PDUs), containing transformers and circuit breakers, and for backup power, rotary uninterruptible power supply (RUPS) generators. Each RUPS uses motor generator technology. Electricity fed into the RUPS is used to build kinetic energy in a 9-foot-diameter flywheel that, in turn, generates electricity.

Supercomputers are cooled by chilled air circulating at the rate of 2.5 million cubic feet per minute.

That bit of extra electricity evens out the flow of power to the supercomputers in the case of a power surge from, for example, a lightning strike, a common occurrence in summertime Los Alamos. In the case of a power outage, there is enough kinetic energy built up in the flywheel to provide 8–12 seconds of electricity to the supercomputers. Those few seconds are long enough for data about the current state of a running calculation to be written to memory, reducing the loss of valuable data.

The PDUs transform the incoming high-voltage electric power feed into lower voltage and distribute it to the supercomputers according to each machine's particular voltage needs, for example, 220 volts for Roadrunner and 480 volts for the supercomputer named Cielo.

The Guardians

Because the Laboratory's supercomputers work on national security problems 24 hours a day, 365 days a year, they require dedicated overseers who stay onsite and collectively keep the same exhausting schedule. The members of the SCC's operations staff, a team of 22, are the experts who keep things running and make sure anything that goes wrong gets fixed, right away.

Divided into three shifts, the members of the operations staff tend monitoring equipment and keep watch from inside the Operations Center, a high, windowed nerve center that overlooks the computer room. The staff's tasks are many and varied, as they are charged not only with overseeing the computer hardware and software but also, for example, with keeping tabs on the cooling system. The computer room's environment must stay cool enough to prevent damage to the valuable computers; too much heat is a major threat.

These dedicated guardians are expected to be able to fix both hardware and software problems in about an hour. For software problems requiring additional support, a team of 30 software administrators, also stationed onsite, backs them up. If a software problem occurs outside regular business hours, the administrators can be called in and must report to the SCC within two hours.

Evolution to Revolution

Looking for all the world like row upon row of large gym lockers, a supercomputer is visibly very different from a personal computer (PC). But the real difference is in the work supercomputers do and the way they do it.

The guardians are expected to be able to fix both hardware and software problems in about an hour.

Today's supercomputers are collections of tens of thousands of processors housed in "racks," cabinets holding the processors and supporting equipment. The large number of processors is needed because supercomputers run immense calculations that no PC could do. The calculations are divided into smaller portions that the processors work on concurrently. This is parallel computing or actually, for a supercomputer, *massively* parallel computing.

A new supercomputer for Los Alamos can take years to create. The process begins with an intense collaboration between commercial computer companies, like IBM,

The Laboratory's Luna supercomputer can be accessed by all three national security labs (Los Alamos, Livermore, and Sandia), making it a "trilab" machine. Roadrunner and Cielo are also trilab machines. (Photo: Los Alamos)

Cray, Hewlett-Packard, etc., and Los Alamos' computer experts, who have extensive experience both operating and designing supercomputers. Los Alamos computer personnel involve themselves in the creation of each new Laboratory supercomputer from the generation of the first ideas to the machine's delivery . . . and after its delivery. Once it is built and delivered, before it is put to work, a supercomputer is disassembled, inspected, and reassembled to ensure that it can handle classified data securely and can be fixed and maintained by Laboratory staff.

Once it is built and delivered, a supercomputer is disassembled, inspected, and reassembled to ensure that it can handle classified data securely.

As a practical and economic necessity, each new Los Alamos supercomputer takes advantage of commercial technological advances. And in the 21st century, beginning with Roadrunner, technology from the private sector is being evolved in innovative ways that are, in effect, a reinvention of how a supercomputer is built. Roadrunner, for example, used video game technology originally conceived for the Sony PlayStation 3, and with that technology, it became the world's first hybrid supercomputer, with an architecture that linked two different types of processors to share computational functions. This particular evolutionary step in supercomputer architecture let Roadrunner surge onto the global stage as the world's first petaflop computer.

Architectures are still evolving, so the next generation of machines will be radically new, even revolutionary, as will Trinity, Los Alamos' next supercomputer, projected to arrive in 2015–2016. On Trinity, Laboratory designers and their industry partners will be trying out numerous innovations that will directly affect supercomputing's future. So Trinity will be unlike any other computer Los Alamos researchers have used. And by the way, it will be 40 to 50 times faster than Roadrunner.

The exact form Trinity will take is still being decided, as design discussions are still underway, but whatever the final design is, it will be a means to an end. The form each new supercomputer takes is dictated by what the Laboratory needs the machine to do. In general that always means it must answer more questions, answer new kinds of questions about new and bigger problems, compute more data, and compute more data faster.

Los Alamos' specific need, however, is focused on the stockpiled nuclear weapons and the continuous analysis of them. Laboratory supercomputers are already simulating the detonation of nuclear weapons, but Trinity and the computers that will succeed it at the Laboratory will need to simulate more and more of the entire weapon (button-to-boom) and in the finest-possible detail. Design efforts for Trinity will be aimed at that goal, and a great deal of effort will go into creating the many new and complex subsystems that the computer will need.

Saving Checkpoints Is the Name of the Game

At the system level, some design requirements remain the same from supercomputer to supercomputer, even when the next one is as fundamentally different as Trinity will be. For example, while a PC serves one user at a time, Laboratory supercomputers must serve many users simultaneously—users from the Laboratory's various divisions and from the other national security labs far beyond Los Alamos. The computer they use must be designed not only to accommodate that multitude of users but also to provide ultra-secure access for the protection of classified data.

Every Los Alamos supercomputer must also be designed to enable an operator to quickly and easily identify and locate which component within the computer's 6,000 square feet (or more) of equipment needs repair. And repairs will always be needed because of the ever-increasing size and speed of supercomputers. As these machines get larger and faster, they naturally become more and more subject to breakdown.

Think about this: If a PC crashes once year and a supercomputer is equal to at least 10,000 PCs, one might expect to see 11 failures *per hour* on a supercomputer. Consider what such a failure rate could mean for an extensive computation. At Los Alamos, a nuclear weapon simulation can take weeks or even months to be completed, and those weeks and months are already costly in terms of computer time filled and electrical power used. In addition, successful simulations require a large collaborative effort between, for example, the weapons scientists, computer designers, computer code developers, and members of the supercomputer operations team. A breakdown equals time and money lost.

With downtime being a supercomputing inevitability, it is commonplace to mitigate the loss by "checkpointing," which is like hitting "Save." At predetermined times—say, every four hours—the calculation is paused and the results of the computation up to that point (the "checkpoint") are downloaded to memory. Returning the simulation to the closest checkpoint allows a simulation (or other type of calculation) to be restarted after a crash with the least amount of data loss.

Unfortunately, the compute time lost even to checkpointing is becoming dearer as supercomputers grow larger and therefore more prone to periodic crashes, so Trinity's designers are working on new checkpointing methods and systems that will maintain a higher level of computational productivity. Los Alamos is working closely with industry to develop this kind of defensive capability.

An Itch That Needs Scratching

PCs are all fundamentally the same, similarly designed to do the same tasks. Users can just go out and buy the software they need for their brand of PC. But supercomputers are different. Designed and built to fill a specific need, each one scratches a hard-to-reach itch. At Los Alamos, the special need is scientific computing and simulation, and a super-computer's users need specially written codes for each project.

Who develops the advanced codes used on Los Alamos supercomputers—the codes for weapon simulation or for general science research? Those highly specialized programs are created in-house, and for many years, the Laboratory's successive supercomputers have had enough in common that existing codes adapted well to them. Trinity's architecture and performance characteristics, however, will presage a complete upheaval. The codes will need to be overhauled, not just adapted: more of a "build it from scratch" compared with just an updating.

The developers are already busy making codes "Trinity friendly" and doing so without having anywhere near the variety and amount of resources the giant commercial computer companies have available. For this work, developers depend on partnering with a range of Laboratory scientists, who provide the unique algorithms for solving the basic physics equations governing how the dynamics of a complex system play out over time. This is true whether the system being studied is the climate or a nuclear explosion. The nature of the scientists' algorithms and the new data generated as a system changes with time determine how the code developers design and build a code to make efficient use of the supercomputer and its data storage and networking connections. In this age of "big data," building programs that efficiently generate unbelievably massive datasets on a supercomputer and make them useful has become a grand challenge. (See the article "Big Data, Fast Data—Prepping for Exascale" in this issue.)

A Titanic Achievement

Designing, building, operating, and maintaining a supercomputer are completely different experiences than working with Word or Excel on a PC at home or at the office. That is true today and will be true, in spades, tomorrow. Computer architectures continue to evolve, leading to the upcoming Trinity and eventually to machines still unimagined.

The Laboratory's supercomputers cannot exist without massively complex and expensive infrastructures, which are often unacknowledged and unappreciated, and without the effort and creative thinking of hundreds of scientists, engineers, and technicians. Working together, they meet the challenge of providing the most-advanced supercomputing environments in the world and then use them to perform the national security science that makes the director's Annual Assessment Letter possible.

It is hard work, and it is certainly worth it.

~ Clay Dillingham

Using supercomputers, scientists can interact with simulations of everything from nuclear detonations to protein synthesis or the birth of galaxies. These simulations can boggle the mind—and at the same time provide clarity. Scientists wear special glasses to view simulations in 3D at extremely high resolution. They can even manipulate the simulations, as the viewer shown here is doing. (Photo: Los Alamos)

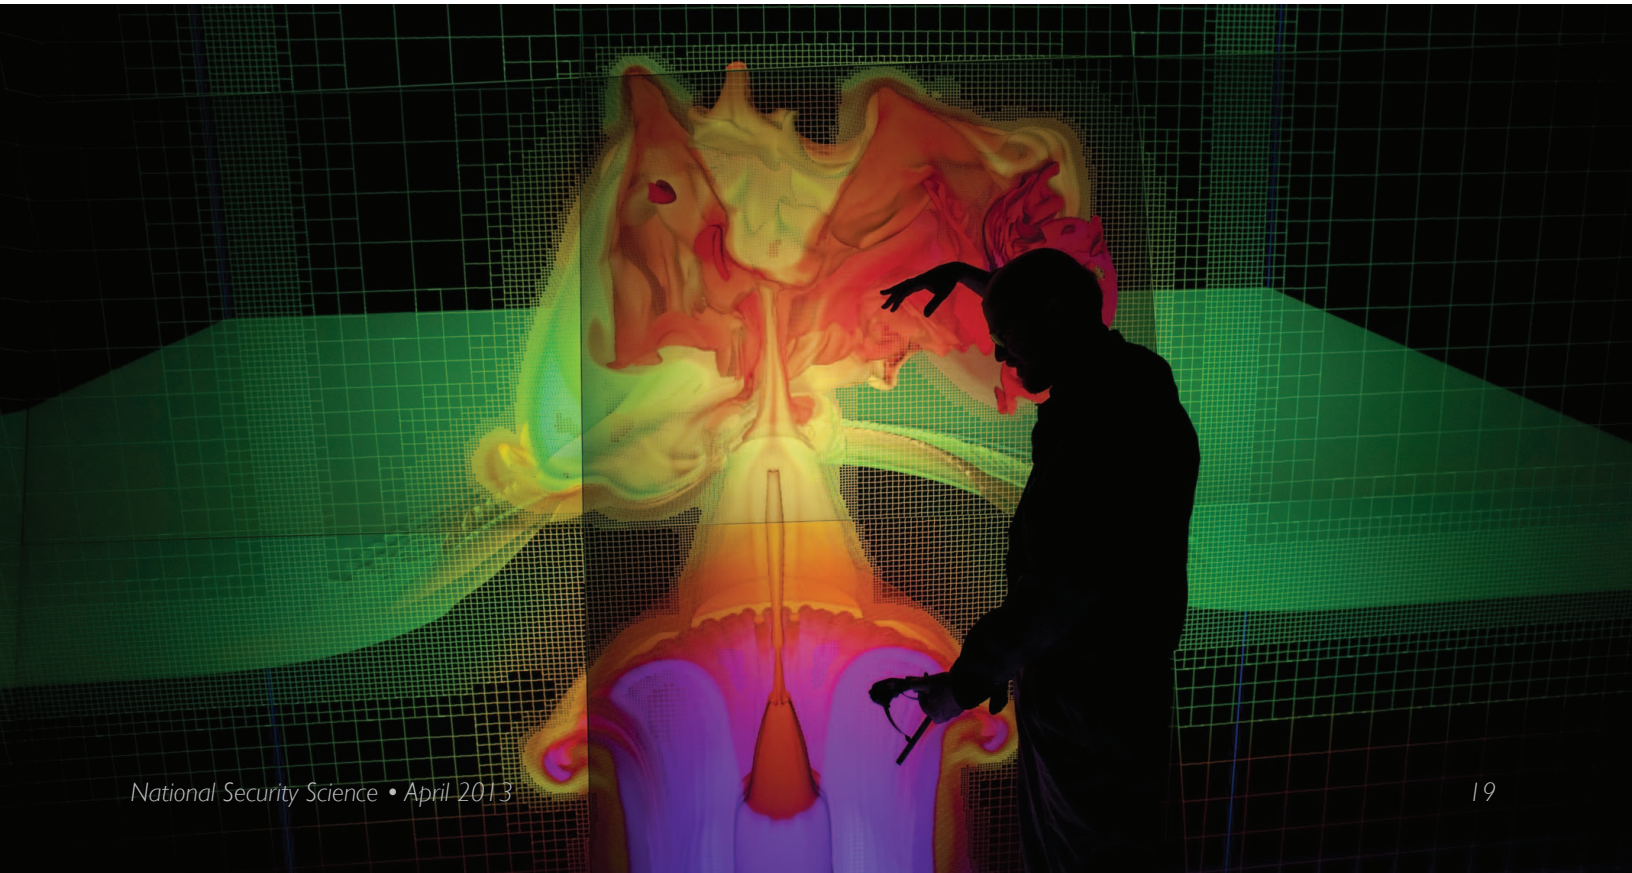

What it takes to run a roadrunner

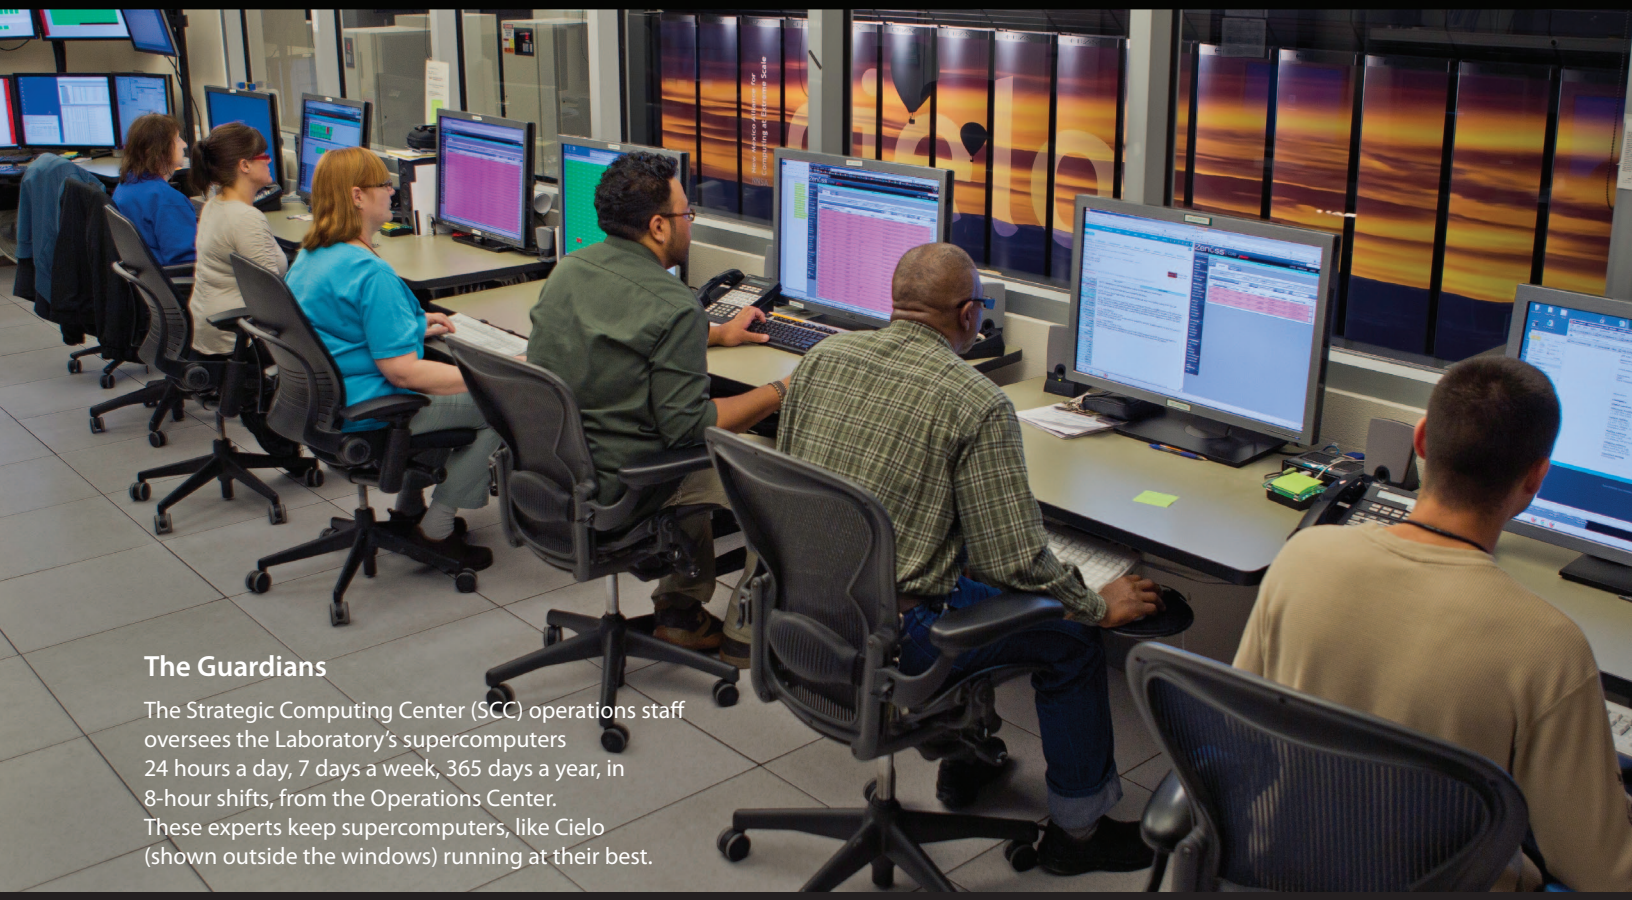

The Guardians

The Strategic Computing Center (SCC) operations staff oversees the Laboratory's supercomputers 24 hours a day, 7 days a week, 365 days a year, in 8-hour shifts, from the Operations Center. These experts keep supercomputers, like Cielo (shown outside the windows) running at their best.

Floor Space

The SCC is a 300,000-square-foot building. The vast floor of the supercomputing room is 43,500 square feet, almost an acre in size.

Electric Power

The amount and cost of electric power required to run a supercomputer are staggering.

Today, a megawatt (MW) of power costs $1 million per year. Roadrunner uses 2 MW per year. Cielo, the Laboratory's newest supercomputer, is a 3-MW machine. Trinity will be a 12-MW machine.

The combined supercomputing facilities at Los Alamos use $17 million per year of electricity.

Using all that electric power means that supercomputers generate lots of heat. If not kept cool, a supercomputer will get too hot and overheat, causing processors to fail and the machine to need costly, timely repairs.

An electrician, wearing personal protective gear, works on a 480-volt breaker inside a power distribution unit.

Managers at the SCC inspect the double floor beneath the acre-size supercomputing room. Several of the giant air-cooling units are visible in the foreground and behind the managers.

The high winds blowing beneath the supercomputer room are generated by the massive air-cooling units.

Air-Cooling

Beneath the acre-size supercomputing room in the SCC is a 1.5-acre floor that houses 84 giant 40-ton air-cooling units. Together, these units can move 2.5 million cubic feet of chilled air per minute through the supercomputing room above.

The air-cooling units use water, cooled by evaporation, to chill the air before it is blown upward to circulate around the supercomputers.

The air, now heated by cooling the supercomputers, is drawn back down to the lower floor and back into the air-cooling units. This process transfers the heat from the air to the water, which is then recooled by evaporation.

To capture the dust and dirt that might otherwise blow into the super-computers, the 84 giant air-coolers use 1,848 air filters. It takes two staff members an entire month to change the filters.

Water for Cooling

The amount of water required to cool the air that, in turn, cools a supercomputer is also staggering. The SCC uses 45,200,000 gallons of water per year to cool its supercomputers. This amount of water costs approximately $120,000 per year.

By the end of the decade, as supercomputers become more powerful and require more cooling, the SCC is predicted to double its water use to 100,000,000 gallons.

The SCC has five evaporative cooling towers. These towers evaporate water to dissipate the heat absorbed by the water in the air-cooling units.

There is room to add an additional cooling tower as the supercomputing needs of the Stockpile Stewardship Program increase.

How the Laboratory is turning Big Data into Fast Data and making it useful...

Big Data

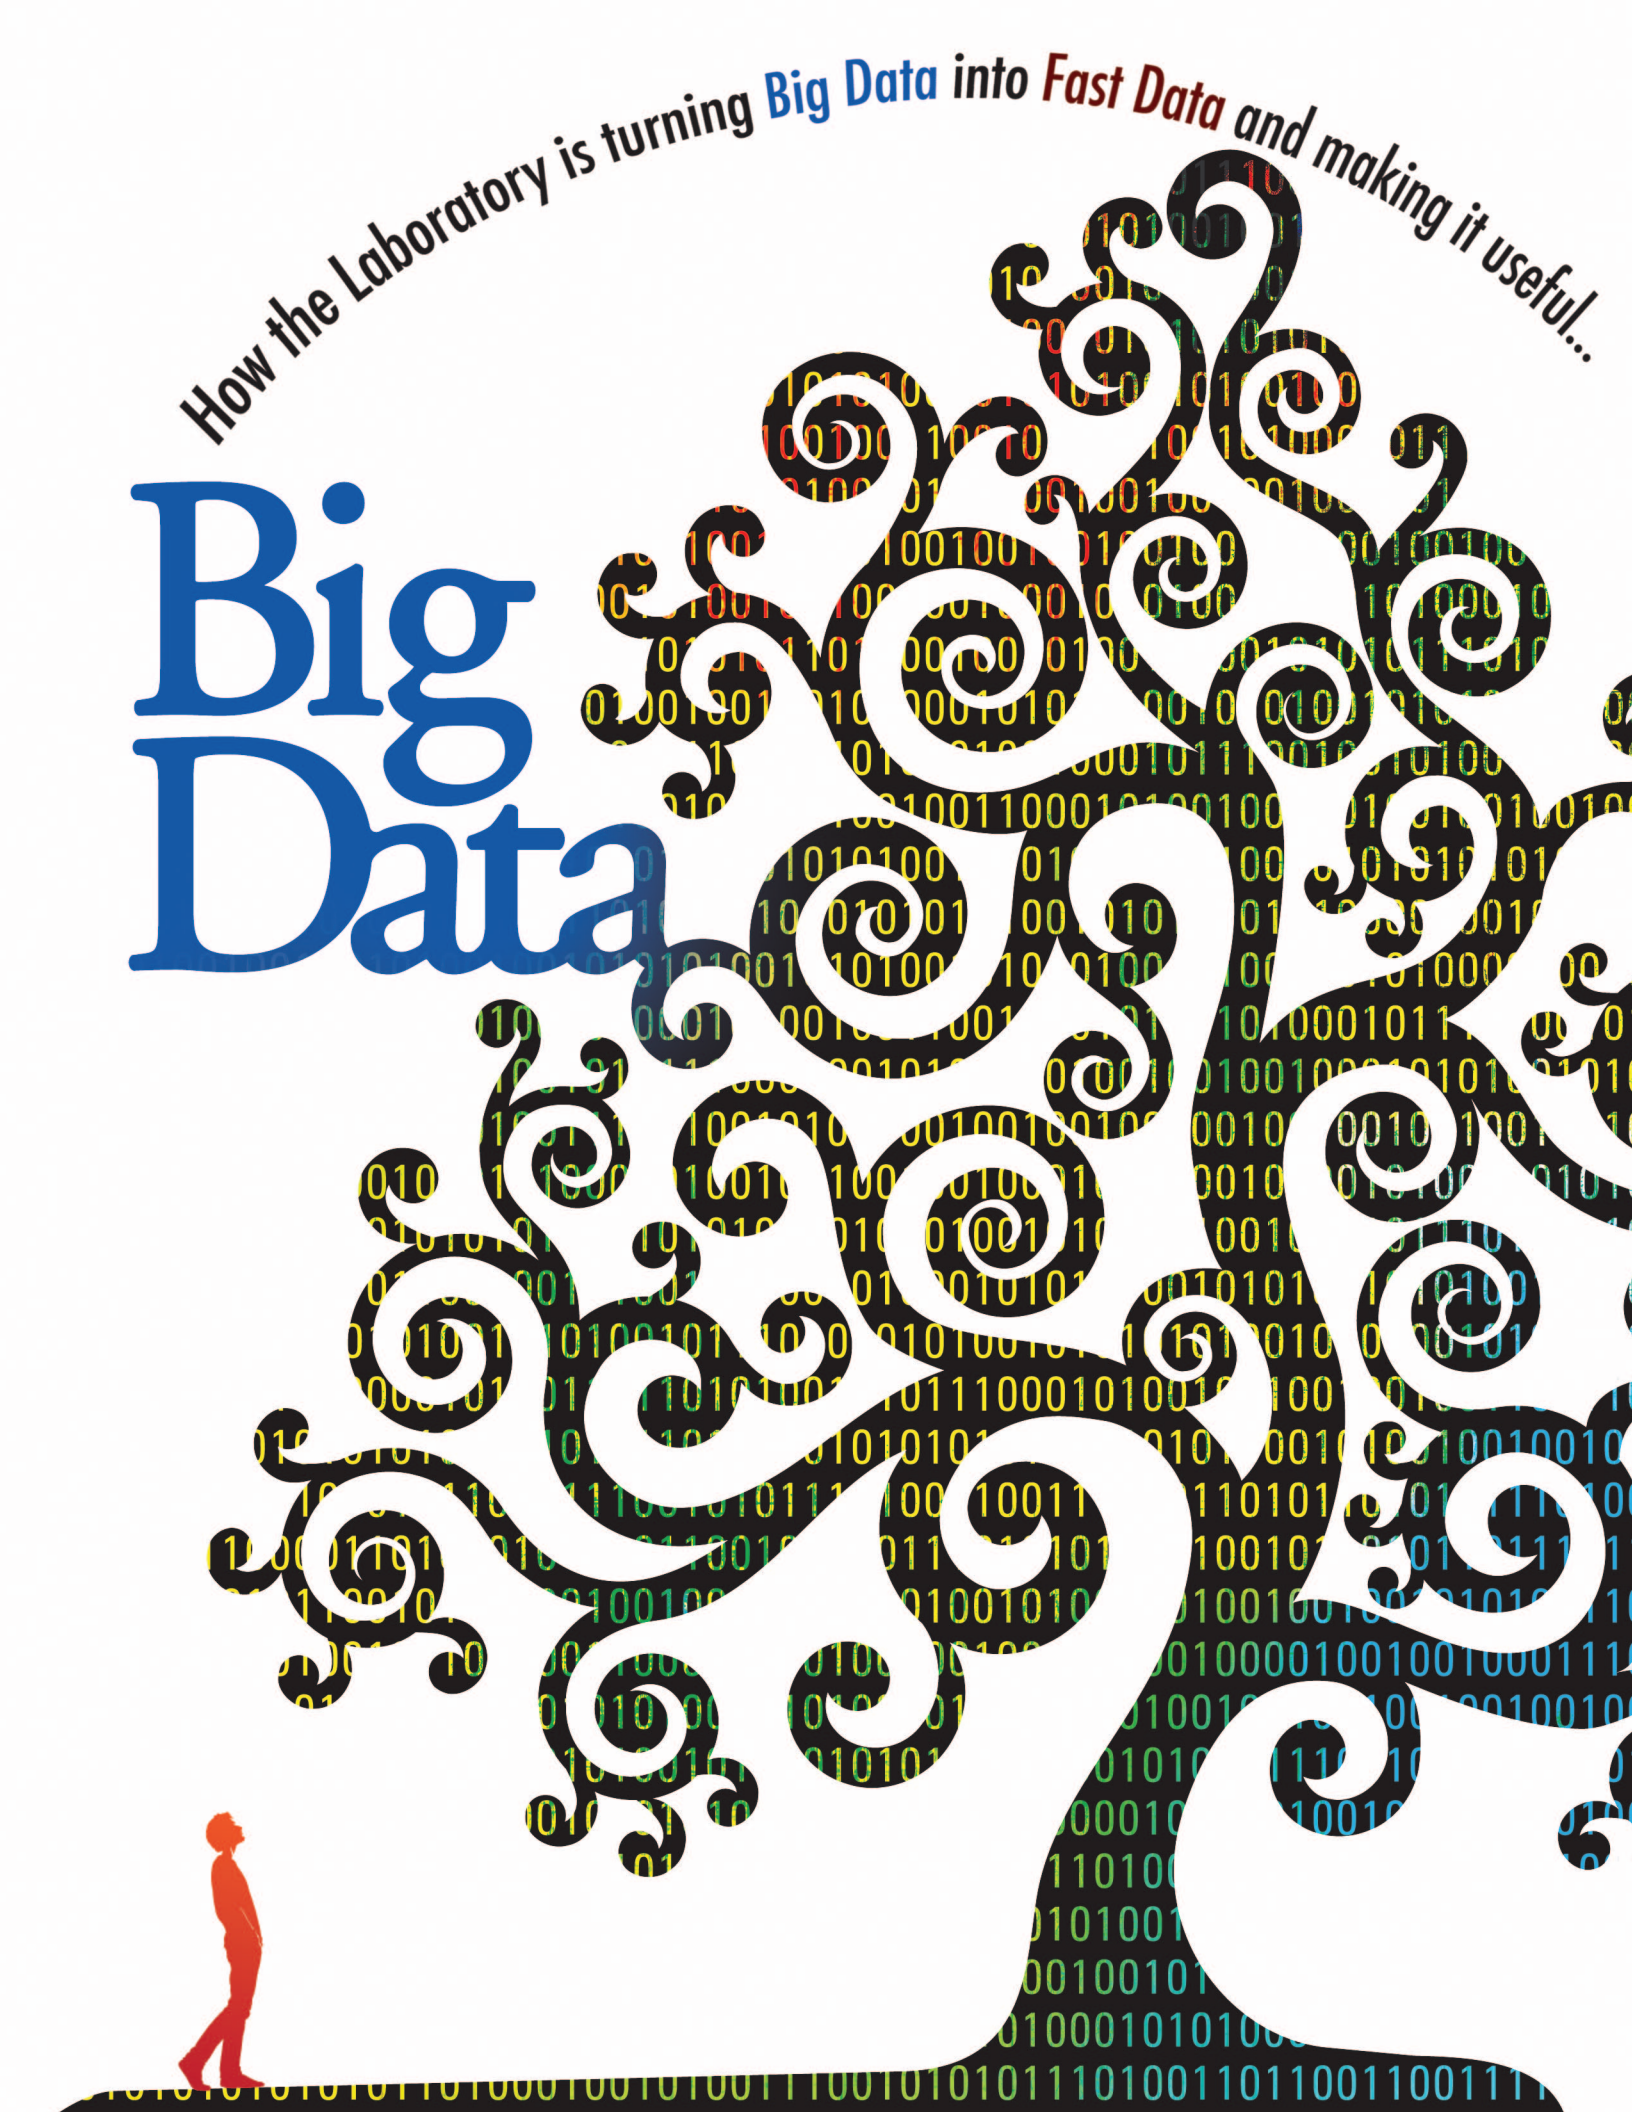

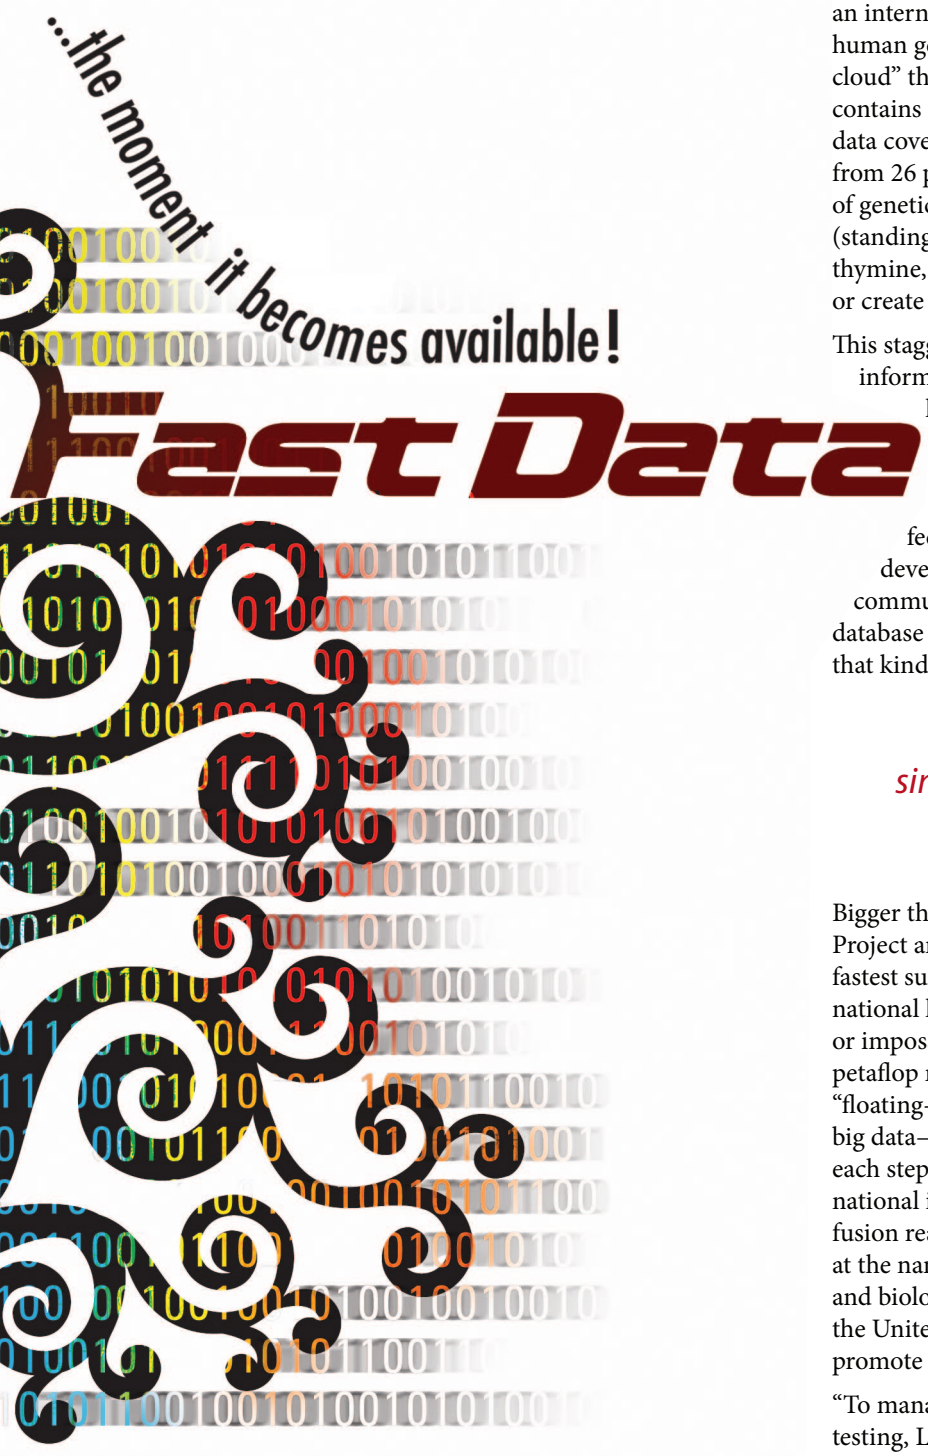

Big data is everywhere. Massive sets of digital data are being collected or generated for science, medicine, astronomy and cosmology, national security, cyber-security, situational awareness for our warfighters, social networking, financial markets, and more. And those datasets are big on a scale that boggles the mind.

A good example of big data collected from nature is the recently released database from the 1,000 Genomes Project, an international effort to establish a detailed catalog of human genetic variation. Made publicly available on "the cloud" through Amazon Simple Store Services, the database contains 200 terabytes (200 trillion bytes) of DNA sequence data covering the complete genomes of close to 2,000 humans from 26 populations. If printed as text, these endless strings of genetic code, written in only four letters A, T, C, and G (standing for the four nucleotide bases of DNA: adenine, thymine, cytosine, guanine), would fill 16 million file cabinets or create a paper stack the height of a skyscraper.

This staggering pile of data is a potential gold mine of information for studying such things as differences in human disease resistance and drug metabolism. But can the medical community mine the gold? Does it have the necessary infrastructure and analysis tools for the job? Only recently, because of a $200 million federal big data initiative, were the necessary tools developed and made available to the medical research community for accessing and analyzing the 1,000 Genomes database for insights into human health and disease. It takes that kind of effort to convert big data into *valuable* data.

The national laboratories simulate systems that are otherwise difficult or impossible to test.

Bigger than the dataset collected by the 1,000 Genomes Project are the datasets *generated* by today's largest and fastest supercomputers, which are being used by the national laboratories to simulate systems that are difficult or impossible to test. The laboratories' supercomputers are petaflop machines that achieve more than a quadrillion "floating-point" operations a second (petaflops) and generate big data—hundreds of terabytes of new data—to simulate each step in the dynamic performance of complex systems of national interest. Those systems include the changing climate, fusion reactors and advanced fission reactors, new materials at the nanoscale (one billionth of a meter), complex chemical and biological systems, and nuclear weapons systems, which the United States has not tested since 1992 in order to promote the goals of the Comprehensive Test Ban Treaty.

"To manage the U.S. nuclear weapons stockpile without testing, Los Alamos and Livermore simulate weapons rather than blowing them up, and to achieve the highest-fidelity

simulations possible, we use the largest computers available and generate big data at an ever-increasing scale," explains Gary Grider of the High Performance Computing Division at Los Alamos. "The problem of big data is always about value—about trying to learn something from the data. At that level, we're the same as Google: we want to turn big data into useful information in an affordable and reliable way. And that way must also be scalable—remaining affordable and reliable as datasets continue to grow exponentially."

But are the national labs getting the most out of this big weapons simulation data from the latest supercomputers? And are they ready with the data management and analysis tools to handle the much larger datasets that will be produced by the next generation of machines?

To achieve the highest fidelity simulations possible, we use the largest computers available and generate big data at an ever increasing scale.

The current answer is a big NO! Unlike the 1,000 Genomes Project big data initiative, the initiative for big weapons data is nowhere near complete, but it has been going on quietly behind the scenes at Los Alamos for almost a decade.

The Big Data Bottleneck

For the past 20 years, supercomputers have generated ever-more simulation data at ever-faster speeds, but those data are not useful until they are selected and moved to permanent storage, organized into files, and then accessed by auxiliary computers that analyze the data and create visualizations of the simulated systems. All those data-handling steps are being challenged by big simulation data, but the biggest challenge is the growing mismatch between the rate at which supercomputers generate data and the rate at which those data can be transferred from the supercomputer to magnetic disk storage, the best permanent storage around. Like cars trying to exit a five-lane highway by way of a narrow ramp, big simulation data of the future will hit a big bottleneck in the transfer path between the supercomputer and storage (see figure below).

Without a solution, computing in 2020 will see crippling data traffic jams in which exaflop supercomputers are idle half the time.

To be specific, the supercomputer world is racing to increase calculation speed a 1,000-fold by 2020—from petaflops to exaflops (a quintillion operations a second)—whereas data-transfer rates to disk storage are expected to increase only 30-fold by that year. Without a solution to this growing mismatch, computing in 2020 will see crippling data traffic jams in which exaflop supercomputers are idle half the time, bloated with data stuck at the bottlenecks separating data generation from data storage and analysis.

Computing at the exascale has often been viewed as a holy grail. For the national security labs, that is because exascale is the scale at which high-fidelity, 3D weapon simulations become practical (see "Will It Work?" in this issue). But the closer supercomputing speeds get to the exascale, the larger the specter of big data becomes. To prepare for the next-generation computers and ensure that they live up to their promise, Grider and colleagues are working closely with

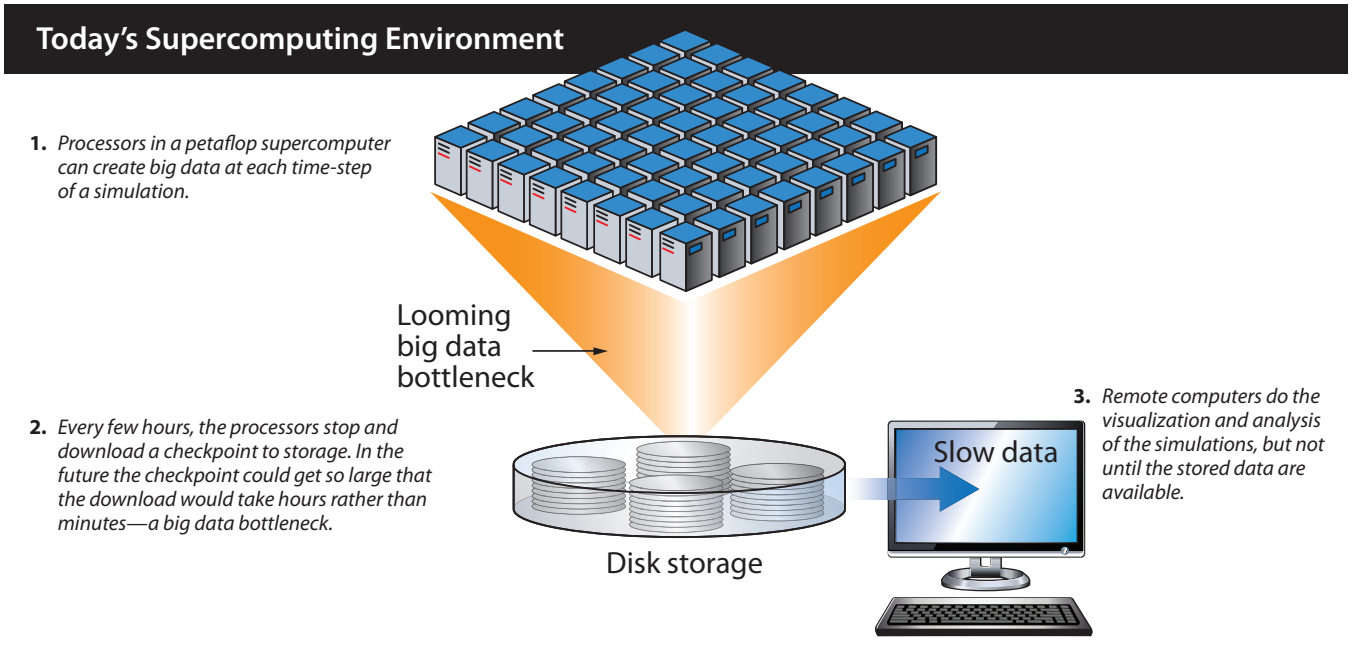

industry and coming up with affordable, scalable solutions. These will not only relieve the big data bottlenecks to disk storage but presage a more effective approach for managing big data simulations at the exascale and beyond.

How Big Simulations Get Done

To better understand these big data solutions, you have to know how today's high-performance computers work. These machines are massively parallel: they can contain more than a million processors, and all million-plus of them work in tandem on tiny bits of the same simulation.

Suppose the simulation is needed because a killer asteroid, one the size of the Rose Bowl, is on a collision course with Earth, and the government wants to know if a nuclear detonation can destroy it. This scenario cannot be tested in a laboratory. But it could be simulated on a supercomputer to help predict whether a nuclear detonation would succeed (see "Killing Killer Asteroids" in this issue).

The simulation might run from weeks to months. And the work is never smooth going.

To do the simulation, a model of an asteroid is placed in a computational box (a way to specify the 3D coordinates of every point in the asteroid model). In this case, the supercomputer is to simulate the entire event, that is, compute all the heating, vaporizing, fracturing, and accelerating, along with the final trajectories of the asteroid fragments, that result from the blast wave from a nuclear detonation hitting the asteroid.

To simulate that event on a modern supercomputer, the computational box is divided into a 100 million smaller cubes of equal size, just as a Rubik's cube is divided into smaller cubes. Groups of the small cubes are assigned to different processors, and each processor solves the physics equations describing what the blast wave does to the material in its set of cubes. The event's duration is divided into discrete time steps (say, several microseconds long), and together, the processors simulate the event one time step at a time. When a processor computes that fragments of rock and vaporized rock in one of its assigned cubes are crossing into a neighboring cube, the processor must pass its latest data about their position, density, temperature, velocity, and so on to the processor for the neighboring cube.

Even though all the processors are sharing the computational load, each processor must solve complicated sets of physics equations for each of the hundreds of thousands of time steps, so the simulation might run for weeks to months to reach completion. And the work is never smooth going. A petaflop computer has millions of parts connected by miles of cable, and a processor fails on average every 10 to 30 hours,

How a Simulation Is Done

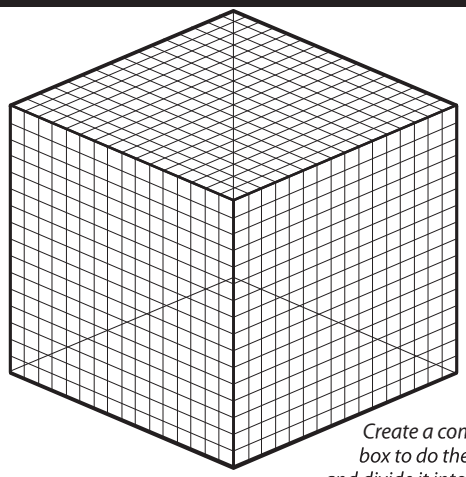

Create a computational box to do the simulation and divide it into 100 million cubes.

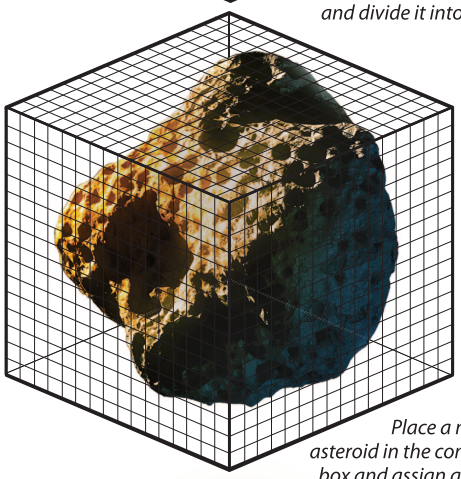

Place a model of the asteroid in the computational box and assign groups of the cubes to different processors.

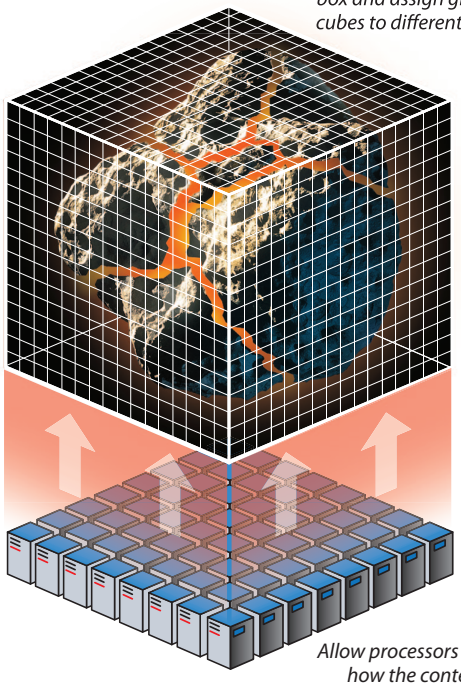

Allow processors to compute how the contents of each cube evolve.

corrupting some of the data needed for the next time step. And because what happens in one cube depends on what comes in from and goes out to neighboring cubes, all the processors must work cooperatively. The failure of one processor has a domino effect: when one stops, all the rest must stop. Does that mean the simulation must return to "start" each time a failure occurs? That would be like writing a document and never using the "Save" command—a very dangerous strategy.

Instead, a supercomputer has to play defense. Every 4 hours, it stops and creates a checkpoint, the analog of pressing "Save" or taking a snapshot of the simulation. All the processors stop at the same simulation time step; update the data describing the temperature, pressure, position, velocity, and so on of materials in their cubes; and send the data to the storage system, which is outside the main computer. Thus, whenever one or two processors fail and the computer crashes, the computer automatically stops, retrieves the data from the nearest checkpoint, and resumes computing at that point. These reference checkpoints not only provide a backup but also record the calculation's progress.

A supercomputer has to play defense. Every 4 hours, it stops and creates a checkpoint, the analog of pressing "Save."

Storing checkpoints sounds simple, but a petaflop supercomputer must save as many as 50 to 100 terabytes of data for each checkpoint, so this kind of "Save" can be very costly in time. Grider explains, "The disk drive in your computer at home might have 1 terabyte of storage capacity, and it would take you about 11 hours of writing to fill that up. We need to transfer all 50 to 100 terabytes in about 5 minutes because while we're writing to memory, we're not getting any science done. So we need 10,000 disk drives hooked together to transfer the checkpoint data to all the storage disks in parallel and get the job done in minutes." On today's petaflop machines, the job does get done, but barely.

Years ago, Los Alamos anticipated that its next big development after Roadrunner, the first petaflop machine, would be Trinity, which, at a speed of 40 to 100 petaflops, would need to store 2 or 3 petabytes of data at each checkpoint. That would require buying 30,000 disk drives at a cost of $30 million, or 20 percent of the machine's cost, and they would be difficult to maintain. An exascale machine would need about 100,000 disk drives, costing 40 to 50 percent of the machine's cost; that would be unaffordable. Without those disk drives, it would take an hour or two to dump the data at each checkpoint, so a major fraction of the computing time would be lost to defensive storage. Neither option was acceptable and both would get worse over time. "Our only course," says Grider, "was to initiate research and development with government, academia, and industry and find an affordable, scalable way around the big data bottleneck."

Burst Buffers—From Big Data to Fast Data

The bottleneck problem that Los Alamos is solving with industry is two-fold: *decreasing* how long processors remain idle when transferring checkpoint data to storage and *increasing* how quickly checkpoint data is fed back to the processors when they fail.

The solution that is in the works capitalizes on flash memories—solid-state storage devices that can write (store) data about 10,000 times faster than disk drives can. If flash memories are placed between the processors and the disk storage, they can "buffer" the mismatch between the burst of checkpoint data needing to be downloaded very quickly and the disk drives, which write data slowly. Grider coined the name "burst buffer" to describe the device that will hold this rapid-writing flash memory and have the right connections to both the supercomputer and the disk storage.

Grider explains, "The concept of the burst buffer is to have the burst of data written onto flash very quickly and then have it written from flash to disk slowly. That way you don't need so many disk drives, and you use the storage disks for what they're good at, namely capacity storage—storing large quantities of data securely."

Imagine racks of processors that are doing the simulation and beneath them the permanent storage system. Each of the million processors is connected to one of many thousands of burst buffers that together act as a staging area to hold checkpoint data before they are sent to permanent storage (see figure, opposite page). An entire checkpoint in the form of a huge petabyte data stream—a burst of data—is downloaded from all the processors in parallel and is absorbed in seconds by the flash memories in the buffers, the processors then resume the simulation. Later, the checkpoint is drained from the burst buffers to disk storage, but at the much slower rate that the disk drives can handle. That means that the processors are stopped so briefly for the downloading to flash that they run almost continuously, with data being written from flash to disk in the background while the processors keep doing science.

We'll be able to watch the simulation as it's happening and intervene if we see something that needs changing. This is truly big data becoming fast data.

Further, if one adds two flash memory units to each burst buffer, one of those units could hold onto the most recent checkpoint data for hours, and download it to disk storage *only after* the second flash unit had received data for a new checkpoint from the processors. Because data downloading would toggle between units, a complete checkpoint would always be available in the burst buffers, ready to be fed back

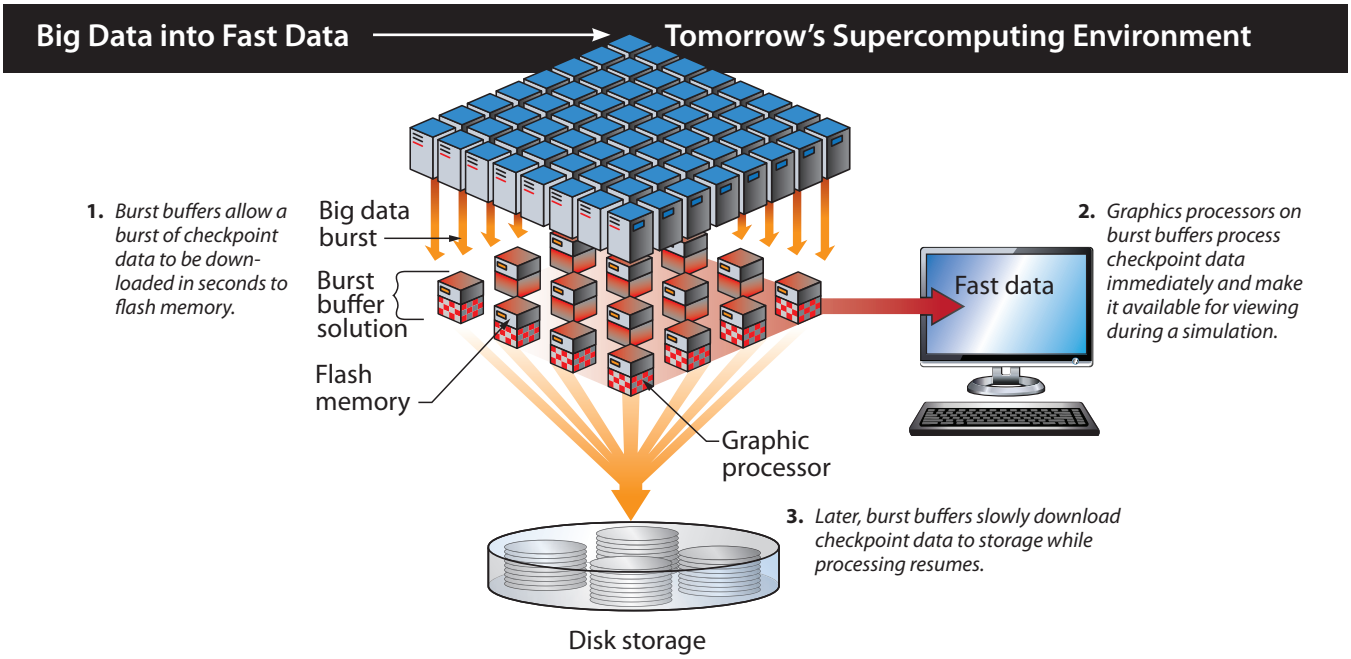

to the processors if a failure required the simulation to be restarted. Flash would virtually eliminate delays caused by both a processor failure and a slow-moving "Save."

Once burst buffers have enough flash memory units to temporarily store checkpoint data, it becomes possible to add graphics processors to each burst buffer. Then, instead of waiting until the end of a run for a visualization of the completed simulation, the current situation, the checkpoint data could be processed into a visualization while the simulation was in progress!

Los Alamos knew back in 2006 it needed to innovate, and it came up with a winner.

"That means we'll be able to watch the simulation as it's happening and intervene in the middle of a run if we see something that needs changing," says Grider. This is truly big data becoming fast data—useful at the moment it becomes available.

"This is the beginning of a big story," continues Grider. "Adding graphics processors to the burst buffer is an example of what's called 'process-in-memory'—processing data where it's most appropriate. Today we move the data to the processors, do the math (addition, multiplication, whatever it is), and then write the results back out to memory [storage]. But the time it takes to move the data is wasted because it's time in which no computing is going on. It may take less time to ship the process to where the data is, and that's what we'd be doing by shipping analysis and visualization to a processor in the burst buffer. So the big story is that processing in the future could go on wherever there's data—in memory, in flash, near disk, near tape. That way some of the processing for a big simulation can take place off the main computer." That is how big data will become fast data.

For weapons simulations, the burst buffer idea is great because it not only allows the downloading or uploading of big data in a few minutes, but it also enables big data to be processed during the simulation, making it useful data.

Race to the Exascale

The Laboratory was driven to develop the burst buffer so it can to do high-resolution 3D simulations of nuclear weapon detonations at the exascale by 2020, but it also needs it because of the constraints of performing exascale simulations affordably and within practical time limits. Los Alamos knew back in 2006 that big data at the exascale would lead to big data bottlenecks and make the old way of doing supercomputing unaffordable. It knew it needed to innovate, and it came up with a winner.

"The burst buffer with its flash memory is the only way we'll be able to build a cost-effective exaflop machine in the 2020 time frame," explains Grider, "and we'll be trying it out on Trinity in the 2015–2016 time frame. Then, when we really need it, we'll have it working. And even as early as Trinity, we'll be testing burst buffers with processors that can analyze and distill the data while the simulation is running." And that's not all. According to Grider, Trinity will be a testbed not only for the burst buffer, but for debugging some of the software Los Alamos will need to keep an exaflop machine running smoothly.

Los Alamos is doing serious prepping for the exascale.

~Necia Grant Cooper

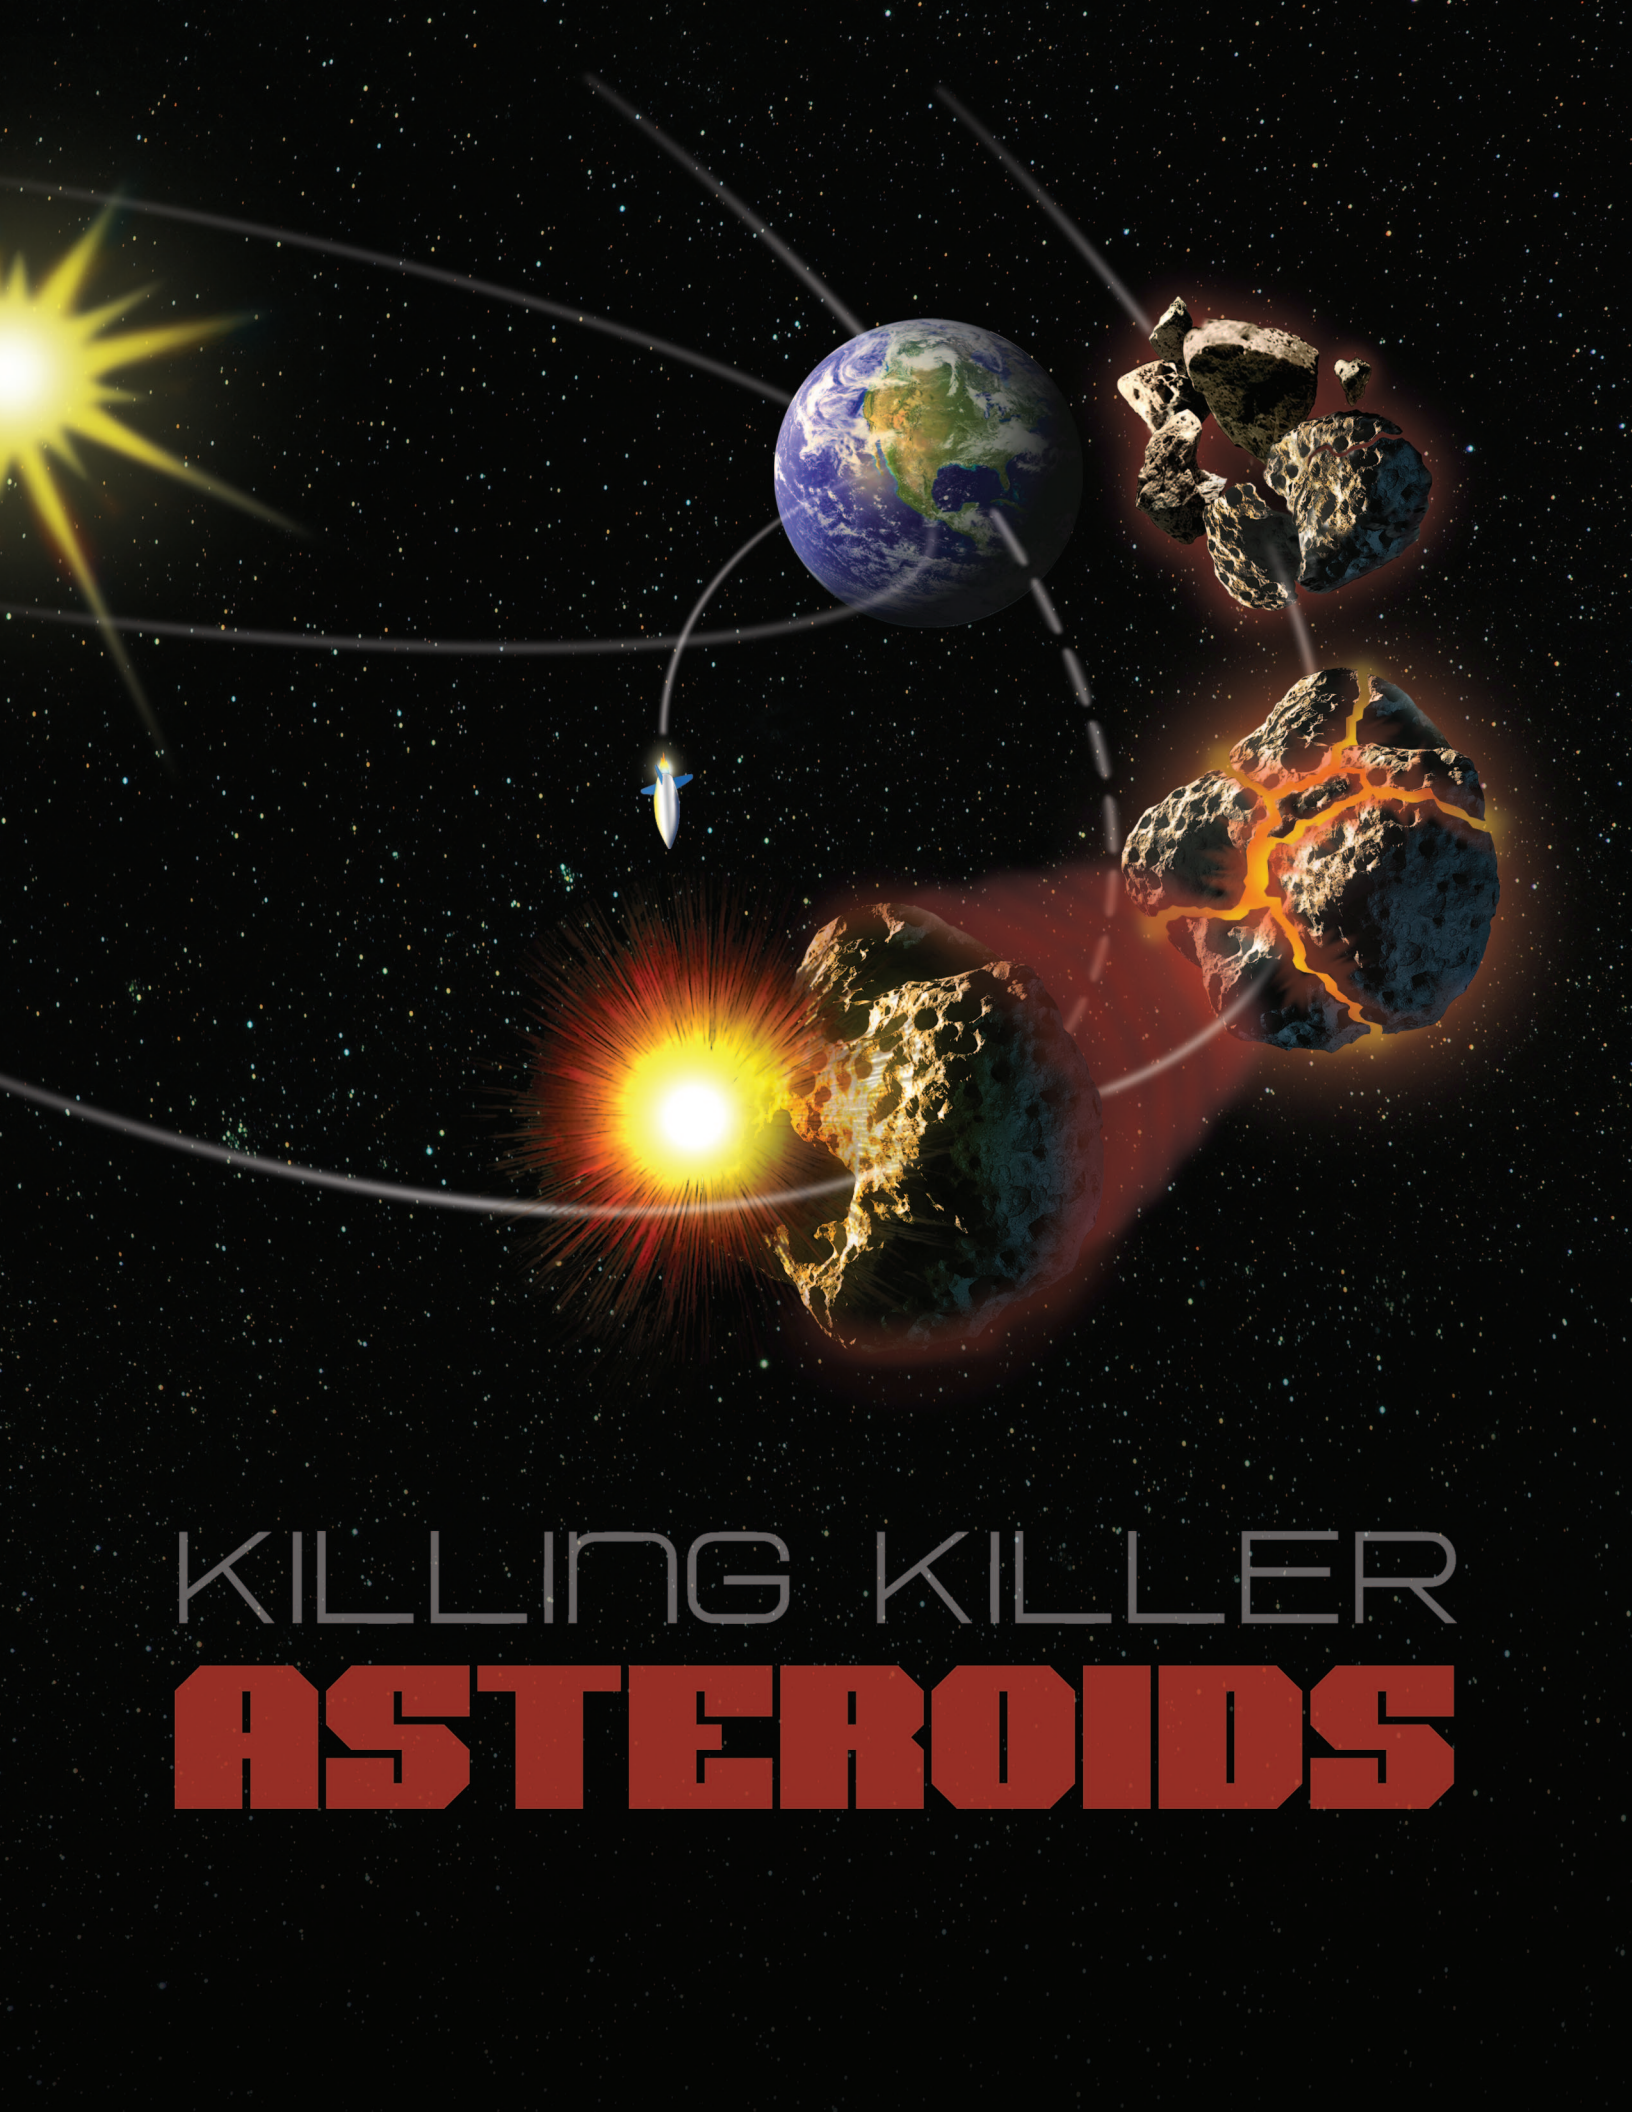

KILLING KILLER ASTEROIDS

In 2004 an alarm went around the globe that a very large near-Earth asteroid, about three football fields in diameter, had a frighteningly high chance (1 in 37) of striking the planet in 2029. Named Apophis for the Egyptian god of darkness and destruction, this space rock would pack a gigantic wallop if it actually struck Earth, releasing the energy of 500 megatons (million tons) of TNT, or 10 of the largest hydrogen bombs ever tested.

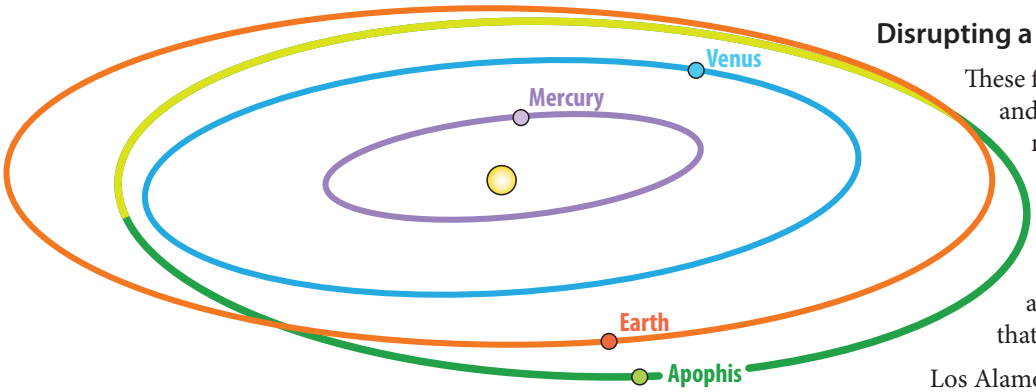

The orbit of the asteroid Apophis is so close to Earth's that we can expect many close encounters in the future.

As more observations accumulated, the Apophis threat was dramatically downgraded. Apophis was expected to pass relatively close to Earth in 2029, at a distance closer than the geosynchronous communication satellites that keep us all in touch with one another, but it would not be on an impact trajectory. However, should it pass through a small region of space called a "gravitational keyhole," the killer asteroid would return seven years later on a collision course and strike Earth on February 13, 2036!

Then during January of this year, scientists used NASA's giant Goldstone radar dish to track Apophis as it passed within 9 million miles of Earth and, from the results, recalculated its future orbits. Mercifully, its chance of passing through the keyhole in 2029 is now zero, and its return in 2036 will be at a very comfortable 14 million miles away.

Any near-Earth object greater than a half-mile in diameter can become a deadly threat.

Whew! We can all temporarily breathe a sigh of relief. However, the likelihood that one day a killer asteroid will be on a collision course with Earth is very high. Under a 2005 congressional mandate, government-sponsored surveys using ground and space-based telescopes have discovered 9,500 near-Earth objects; 1,300 of these, are deemed potentially hazardous. New asteroids and comets can be expected to enter Earth's neighborhood as the gravitational pull of passing stars and collisions between asteroids do their work to alter the orbits of these (mostly) Solar-system residents.

Also, we know with certainty from many fields of study that 63 million years ago, a 6-mile-diameter asteroid collided with Earth, striking Mexico's Yucatan peninsula, releasing 10 million megatons of energy, creating a huge crater, and causing the extinction of the dinosaurs, a major change in climate, and the beginning of a new geological age. Any near-Earth object greater than a half-mile in diameter can become a deadly threat, potentially causing a mass extinction of us.

Disrupting a Killer Asteroid

These facts keep many professional and lay astronomers busy monitoring the sky. Recognizing the risk, astrophysicists are working on ways to intercept a killer asteroid and disrupt it in some way that will avert disaster.

Los Alamos astrophysicist Robert Weaver is working on how to protect humanity from a killer asteroid by using a nuclear explosive. Weaver is not worried about the intercept problem. He would count on the rocket power and operational control already developed by NASA to intercept a threatening object and deliver the nuclear device. NASA's Dawn Mission has been able to place a spacecraft in orbit around Vesta, a huge almost-planet-size asteroid in the asteroid belt between Mars and Jupiter, and the NASA Deep Impact mission sent a probe into the nucleus of comet 9P/Tempel. In other words, we have the technology to rendezvous with a killer object and try to blow it up with a nuclear explosive. But will it work?

Simulations on Los Alamos' powerful Cielo supercomputer suggest that a 1-megaton nuclear blast could deter a killer asteroid.

Weaver's initial set of simulations on Los Alamos' powerful Cielo supercomputer demonstrates the basic physics of how a nuclear burst would do the job. The simulations suggest that a 1-megaton nuclear blast could deter a killer asteroid the size of Apophis or somewhat larger.

By far the most detailed of Weaver's calculations is a 3D computer simulation of a megaton blast on the surface of the potato-shaped Itokawa asteroid. Visited by Japan's Hayabusa asteroid lander back in 2005, Itokawa is a conglomerate of granite rocks, a quarter of a mile long and about half as wide, held together by self-gravity (the gravitational attraction among its constituents). Weaver used the most modern, sophisticated Los Alamos codes to predict the progress of

a megaton nuclear blast wave from the point of detonation through the asteroid.

"A big plume coming out of the asteroid in the simulation [see image on opposite page, bottom left] is the effect of all that heated rock in the vicinity of the explosion being expelled from the asteroid at high velocities," Weaver says. "The shock wave from the explosion transfers kinetic energy to the individual rocks, and then as the rocks move, they hit other rocks, causing more rock-to-rock kinetic energy transfers. These rock-to-rock interactions propagate the energy from the surface all the way through to the opposite end of the asteroid, totally disrupting these rubble piles." A YouTube video of the 3D simulation can be seen at http://www.youtube.com/watch?v=hOcNbAV6SiI.

A calculation of this complexity requires the next generation of supercomputers, the so-called exascale computers.

Computing Limitations

Los Alamos' Cielo supercomputer is a 1.43-petaflops machine—meaning it performs just over a quadrillion (million billion) arithmetic computations per second. It is one of the most powerful computers on the planet. Cielo is composed of 32,000 independent computers that work "in parallel"; that is, they work on separate parts of the calculation simultaneously.

Even with Cielo's massive computing power running for a full month, the 3D calculation simulated the detonation's progress for only 30 milliseconds, at which point the blast wave had traveled through only about 25 percent of the asteroid's volume. To reach completion the simulation needs to run 10–60 seconds past detonation, following the blast wave through the entire asteroid, computing the breakup into rocks that then collide with each other, and finally following the trajectories of the individual pieces resulting from the breakup. That would take Cielo about three years of running time. To be practical, a calculation of this complexity should take only a few days, and that requires the next generation of supercomputers, the so-called exascale computers that would calculate a billion billion computations per second, or 1,000 times more calculations than a 1-petaflop supercomputer.

The 2D Results

To complete the simulation on Cielo, Weaver made some drastic simplifications to the asteroid model so that it could be run in 2D instead of 3D. The lumpy asteroid became a simple, smooth cylinder made up of smooth cylindrical rocks. The asteroid model's symmetry meant that the outcome of the blast could be calculated in just a couple of days on Cielo, compared with three years for the full 3D calculation.

Weaver was very encouraged by the results. "In my 2D calculations, I'm seeing velocities of meters per second imparted to expelled rock on the side of the asteroid opposite the detonation point," Weaver says. "The escape velocity [the velocity needed to escape the self-gravity of the asteroid] for an asteroid like Itokowa is only fractions of a centimeter per second, so the expelled rocks have over 100 times the escape velocity and can therefore overcome the forces of gravity tending to reassemble them into a loose pile of rock. That was a surprise to me and gave me some confidence that a nuclear blast really would be an effective mitigation technique. The asteroid would not re-collect, and it would not pose a hazard of a bunch of smaller rocks hitting the Earth."

Some astrophysicists had predicted that fragments from a nuclear blast would move very slowly, so slowly that they

This photograph, taken in 2005 by a Japanese space probe, shows the elongated, lumpy shape of the asteroid Itokawa. The asteroid was the basis of Weaver's simulations, performed on the Cielo supercomputer.

This 3D simulation shows the asteroid Itokawa during the initial impact of a blast wave (red and white) from a nuclear detonation on its surface. Itokawa is depicted as a random distribution of rocks 5–50 yards in diameter.

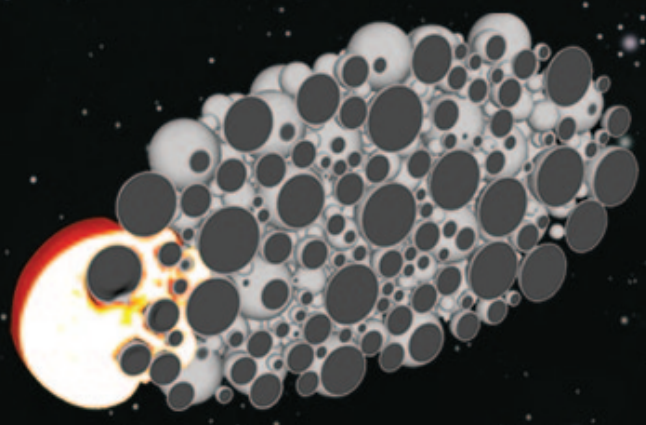

would recondense into a bunch of large rocks, large enough to hit Earth's surface with damaging impact. The fact that the simulated fragments had speeds well beyond the escape velocity refuted that prediction. Moreover, pointing the blast in a direction perpendicular to the asteroid's motion would make the rocket effect of the blast (which heats material at the asteroid's surface and creates a blowoff opposite the direction of impact) force the fragments to move in the same direction as the blast impact, out of the asteroid orbit and away from Earth. In other words, a nuclear blast could act like a propulsion system, directing the asteroid fragments in a desired direction.

Weaver will next turn to simulating larger and larger rocks of varying compositions up to the size of a "dinosaur killer."

"All this depends obviously on exactly where the intercept is done, how far away from the Earth it is, how much time we have left—and all of these are unknowns until we discover a threatening asteroid," Weaver says. "What I think I'm bringing to the table for the first time are truly validated simulations of these nonuniform, nonspherical compositions that will hopefully give policy makers a better understanding of what their options are." Weaver will next turn to simulating larger and larger rocks of varying compositions up to the size of a "dinosaur killer" (about 6.2 miles across). To that end, Weaver and Los Alamos will soon begin a collaboration with Lawrence Livermore National Laboratory that will pool computational and funding resources to take this kind of asteroid mitigation exploration to the next level, assessing a range of potential threats.

Nuclear versus Alternative Options

In suggesting a nuclear energy source for asteroid mitigation, Weaver says he is being practical. Nonnuclear options could prevent impact by deflecting the incoming asteroid through the use of gravitational tractors (spacecraft that travel alongside the asteroid for a decade or two and have enough mass to pull the object off its collision course with Earth) and impactors (rockets that make direct hits on the asteroid and throw it off course). Weaver believes these nonnuclear options would need a decade of planning and development before they could be deployed. In addition, they would need to be deployed many years in advance of the impending collision.

But objects can appear with little warning, he explains. The most likely ones are extraplanetary comets, objects not bound in orbit around the Sun that travel toward us in a plane different from the Earth-Sun plane. If we have only a 6-month lead time, the most practical option is a nuclear device. "From my perspective," he says, "the nuclear option is for the surprise asteroid or comet that we haven't seen before, one that basically comes out of nowhere and gives us just a few months to respond," says Weaver.

As if to illustrate Weaver's point, Earth recently got a violent demonstration from one of his "out of nowhere" objects. On February 15 a meteor blazed through Russian skies and exploded, generating a brilliant flash and a shower of meteorites. Fifteen hundred people were injured by the broken glass and debris resulting from the shock wave.

With such a graphic example in people's minds, the pros and cons of alternatives are being hashed out at the next biannual Planetary Defense Conference. There, scientists of all persuasions discuss the best mitigation strategies and the international agreements that must be put in place before any of the strategies can be implemented.

~ Necia Grant Cooper

As the blast wave travels through the asteroid, it creates a backward-flowing plume of hot material (red, orange, and white in this simulation), propelling the asteroid in the opposite direction.

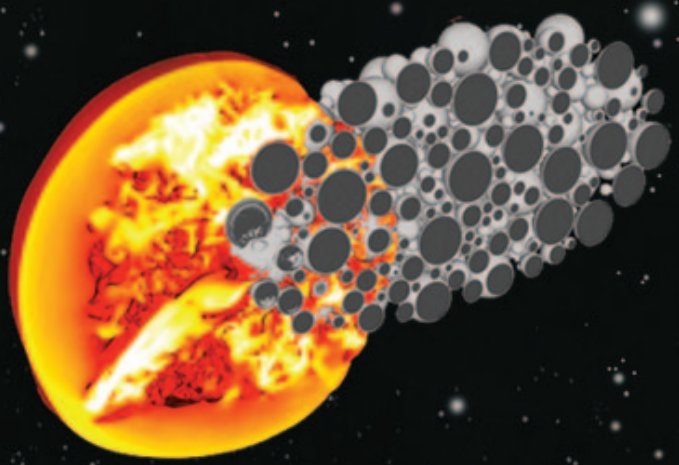

Alternative proposals for altering an asteroid's path include using the gravitational pull between the asteroid and an unmanned space "tractor" to shift the asteroid to a new orbit.

Punched Cards to Petaflops

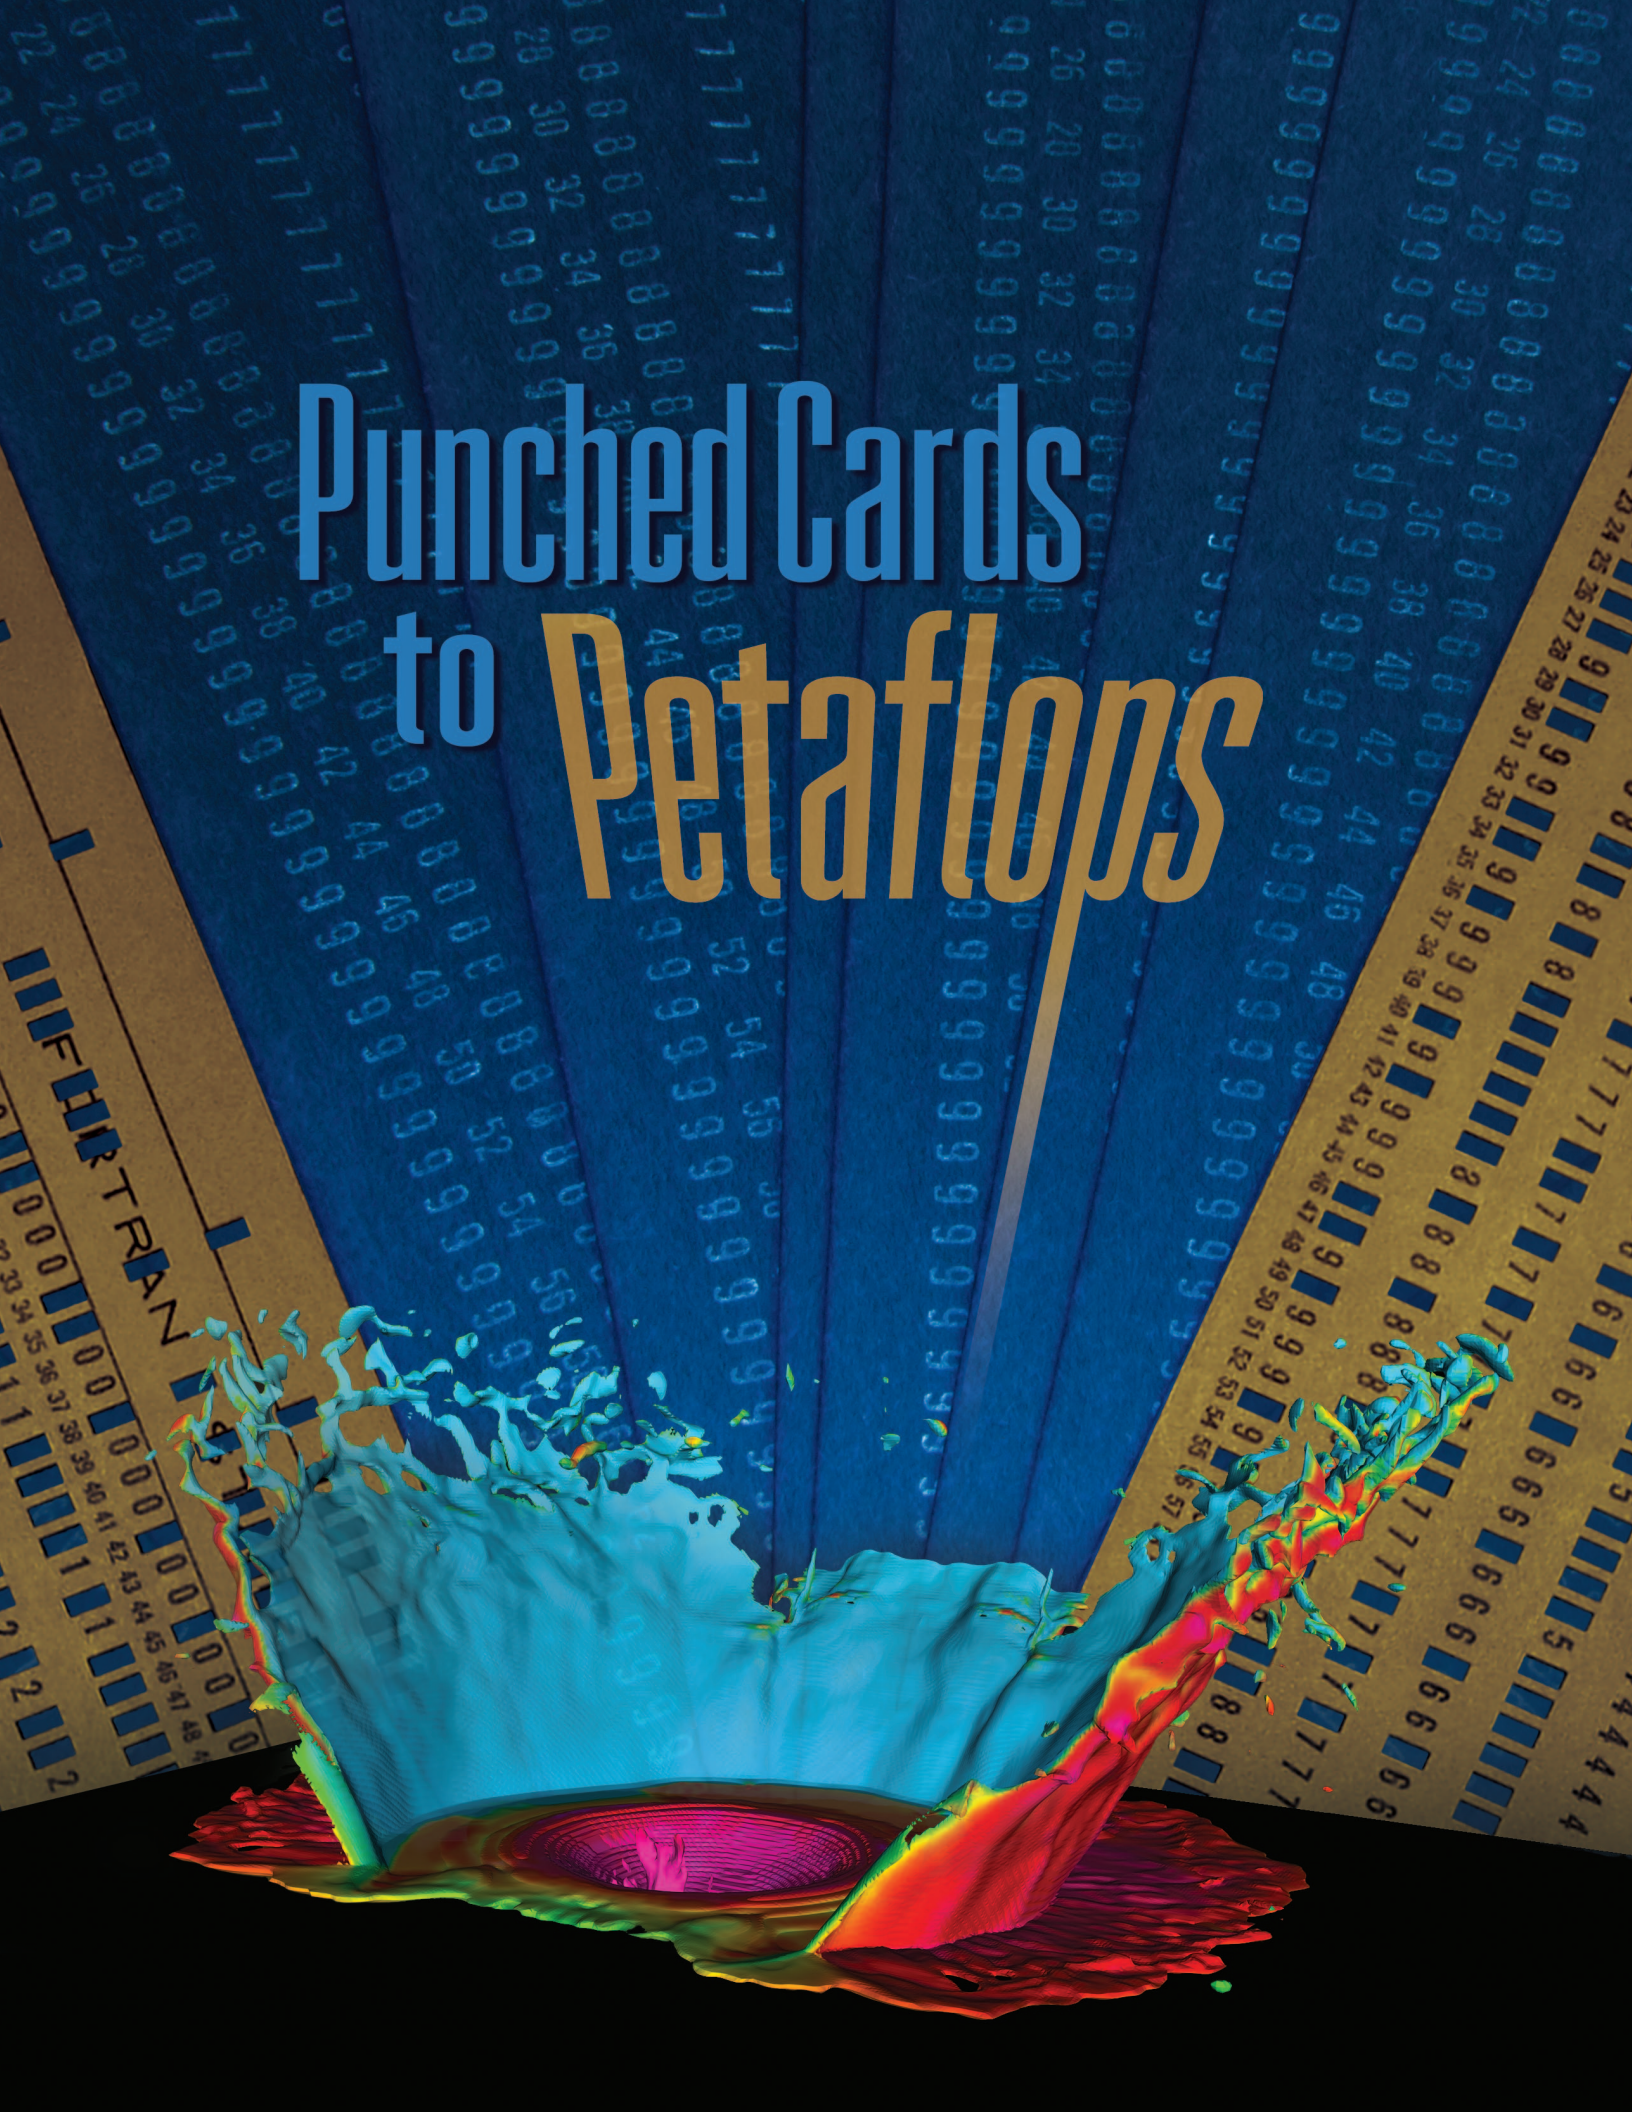

Using the most-advanced computers of that time, by 1952 Los Alamos built and tested Ivy-Mike, the world's first full-scale hydrogen bomb. It achieved a yield equivalent to 10,400,000 tons of TNT. (Photo: Los Alamos)

In May 2008, Los Alamos National Laboratory's Roadrunner became the most powerful supercomputer in the world. The coveted title of world's most powerful supercomputer changes hands often, but Roadrunner was not just another fast machine. Its pioneering architecture allowed Roadrunner to break the petaflop computing barrier by performing more than a thousand trillion floating-point operations (calculations) per second. In doing so, Roadrunner sparked a technological revolution.

Computing has provided the tools for the solutions of many problems in nuclear science that would otherwise have been either intractable or much delayed.

Los Alamos enjoys a rich history of innovation in many fields, including supercomputing. As Laboratory Fellow Jack Worlton wrote in 1978, "Nuclear science has provided technical motivation and much of the funding for large-scale scientific computing, and computing has provided the tools and techniques for the solutions of many problems in nuclear science that would otherwise have been either intractable or much delayed." Decades before the petaflop barrier was broken, the Laboratory relied on mechanical desktop calculators and punched-card machines to perform calculations necessary for building the first nuclear weapons. This relationship, between national security and computing, is no mere coincidence.

Early Computing and the Manhattan Project

For millennia, humans have needed calculating tools to perform an assortment of tasks, including basic arithmetic, records management, and timekeeping. In the 17th century, important devices such as the slide rule and the first mechanical calculator were invented, but it was not until the late 19th century that computers capable of interpreting data, such as advanced punched-card machines, were developed. Punched-card technology remains with us today, but it gradually fell out of favor as a platform for state-of-the-art computing in the early 20th century.

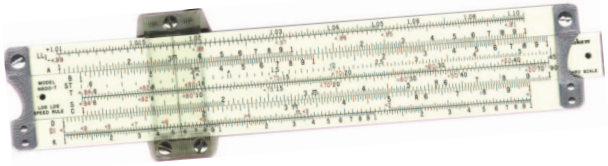

The slide rule is a simple, mechanical analog computing device that uses distance to represent real numbers.

In the decades leading up to World War II, complex analog computers rose to prominence. Analog computers use measurable physical entities, such as distance or voltage, to represent numerical data. Although analog devices, such as the astrolabes used by early navigators and astronomers, have been around for thousands of years, analog computers remained relatively simple machines until

Early punched-card technology was used for many applications, including controlling machines such as the player piano.

Mechanical calculators, such as the Marchant, were used heavily at the wartime Laboratory. (Photo: Los Alamos)

the early 20th century. The development of advanced analog computers culminated with MIT's differential analyzer, a machine named for its ability to solve complex differential equations. The differential analyzer was invented by Vannevar Bush and his student Harold Hazen.

As the 1930s drew to a close, both Bush and his machine were drafted for defense projects. On the eve of the country's entry into World War II, research on critical defense technologies, such as radar, was performed at the MIT Radiation Laboratory using tools that included the differential analyzer. Bush managed several important programs as head of the National Defense Research Committee (NDRC) and later accepted a new position as head of the Office of Scientific Research and Development.

The NDRC programs included the germinal atomic bomb project, which Bush took a personal interest in. After the Japanese attack on Pearl Harbor, the bomb project grew rapidly and, with Bush's concurrence, was transferred to the Army Corps of Engineers. The project came to be called the Manhattan Project, and the weapons design laboratory, sited northwest of Santa Fe, New Mexico, was known as Project Y.

From the very beginning, Los Alamos National Laboratory relied heavily on computers and computing machines.

From the very beginning, Project Y (which eventually became Los Alamos National Laboratory) relied heavily on computers and computing machines to design the world's first atomic bombs. When the Laboratory commenced operations in the spring of 1943, several mechanical calculating machines were purchased. These devices, the most useful of which was the Marchant desktop calculator, were primarily used to perform calculations in support of the gun-assembled uranium weapon program (Little Boy). Relatively powerful IBM punched-card machines soon followed, thus enabling more-complex computing in support of the implosion-assembled plutonium weapon program (the Trinity device and Fat Man).

Digital computers, such as the ENIAC, use numerical digits to represent data. (Photo: Los Alamos)

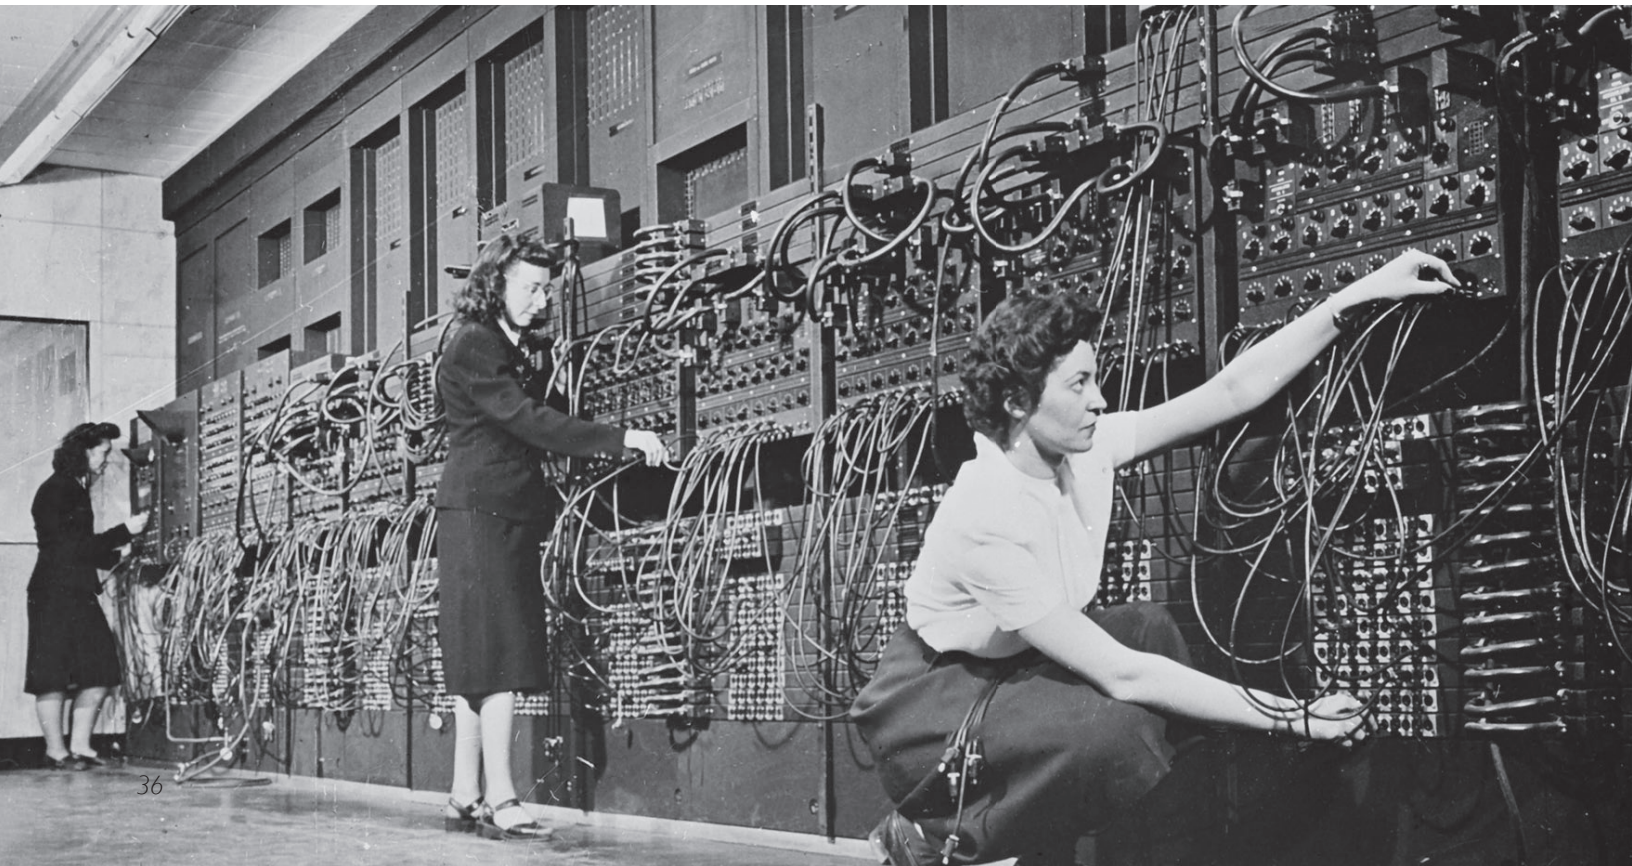

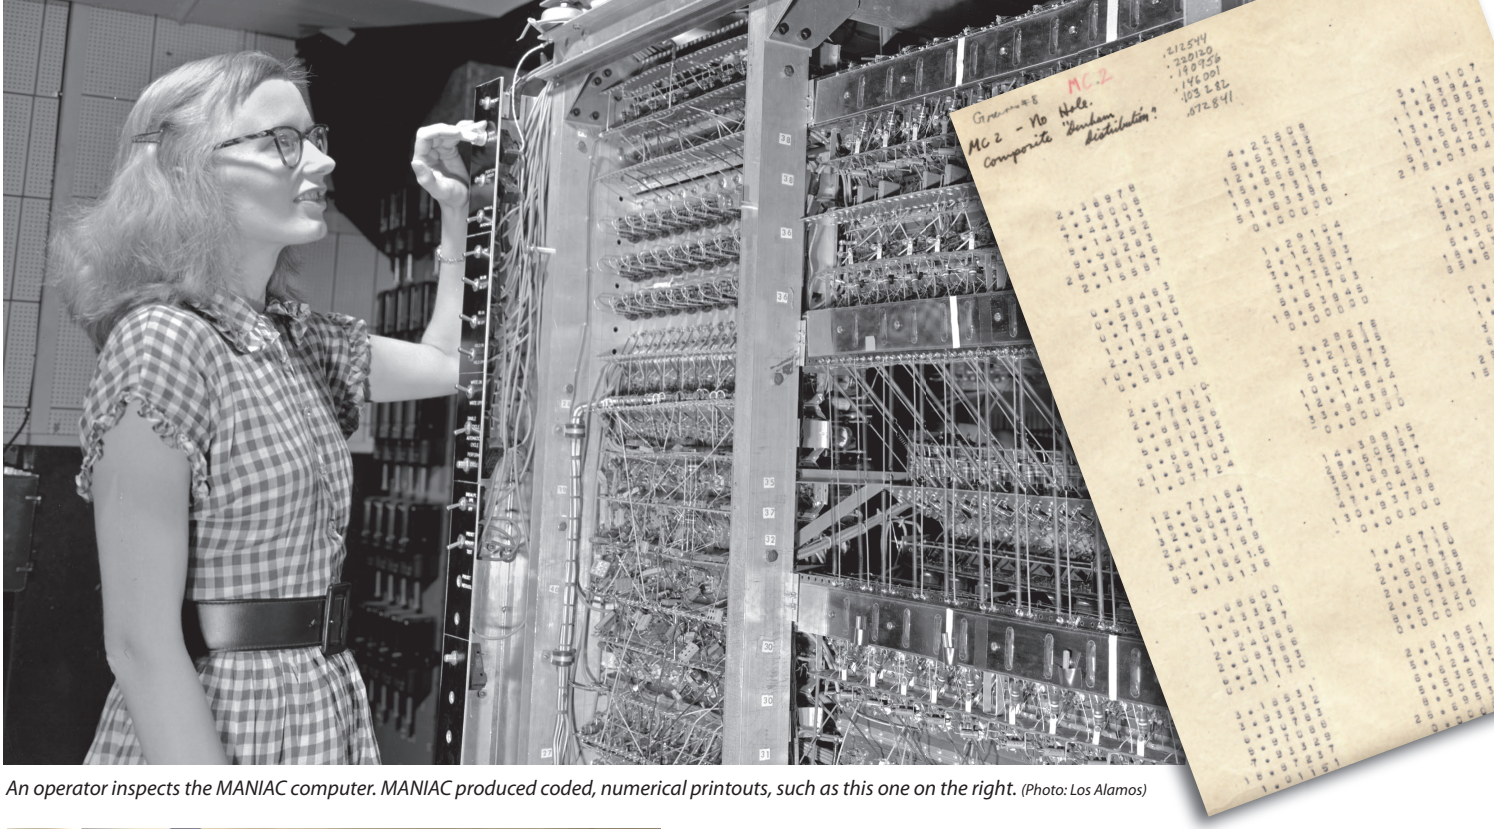

An operator inspects the MANIAC computer. MANIAC produced coded, numerical printouts, such as this one on the right. (Photo: Los Alamos)

In a gun-assembled device such as Little Boy (foreground), a uranium projectile is propelled at a uranium target to initiate the fission chain reaction. Plutonium requires a more complex method of assembly known as implosion. In implosion devices, such as Fat Man (rear), a sphere of plutonium is compressed by high explosives to initiate the fission chain reaction. (Photo: Los Alamos)

Early on, the wartime computing machines at Los Alamos lacked mechanical reliability and, largely as a result, required routine repairs and often yielded inaccurate results. But the early computing program at Los Alamos boasted several notable scientists, Richard Feynman (who would win the 1965 Nobel Prize in Physics) and Nicholas Metropolis among them. Feynman and Metropolis decided to personally start repairing the Marchant and punched-card machines. As Metropolis and his Los Alamos colleague Eldred C. Nelson said after the war, "When the [Laboratory] administration discovered this extracurricular activity, some critical eyebrows were raised and service was interrupted. Then, as the number of working computers dwindled, criticism turned to pleas to restore the status quo."

Despite the lack of reliability, the early computing technology became nearly indispensable, especially as Feynman and Metropolis grew more adept at maintaining them, enabling the scientific staff to model complex experiments. The data produced in these models helped scientists understand the physics of implosion. Likewise, computing enabled scientists to accurately predict other physical scientific phenomena, such as the weapon's explosive yield, pertaining to the Trinity test of July 16, 1945.

A few weeks after the Trinity test, atomic bombs were used to help bring World War II to an abrupt and victorious conclusion.

ENIAC, still unknown to the public because its existence was classified, had quietly ushered in the age of modern computing.

The eminent Hungarian mathematician, John von Neumann, also played an important part at Project Y as a consultant. Von Neumann introduced the Los Alamos staff to many cutting-edge computing technologies, including the world's first electronic digital computer, the ENIAC (Electronic Numerical Integrator And Computer), which was under construction at the University of Pennsylvania. ENIAC was designed to make calculations for the Army's artillery firing tables, whose data helped gunners accurately aim their weapons. ENIAC's versatile architecture also enabled it to perform calculations in support of early hydrogen

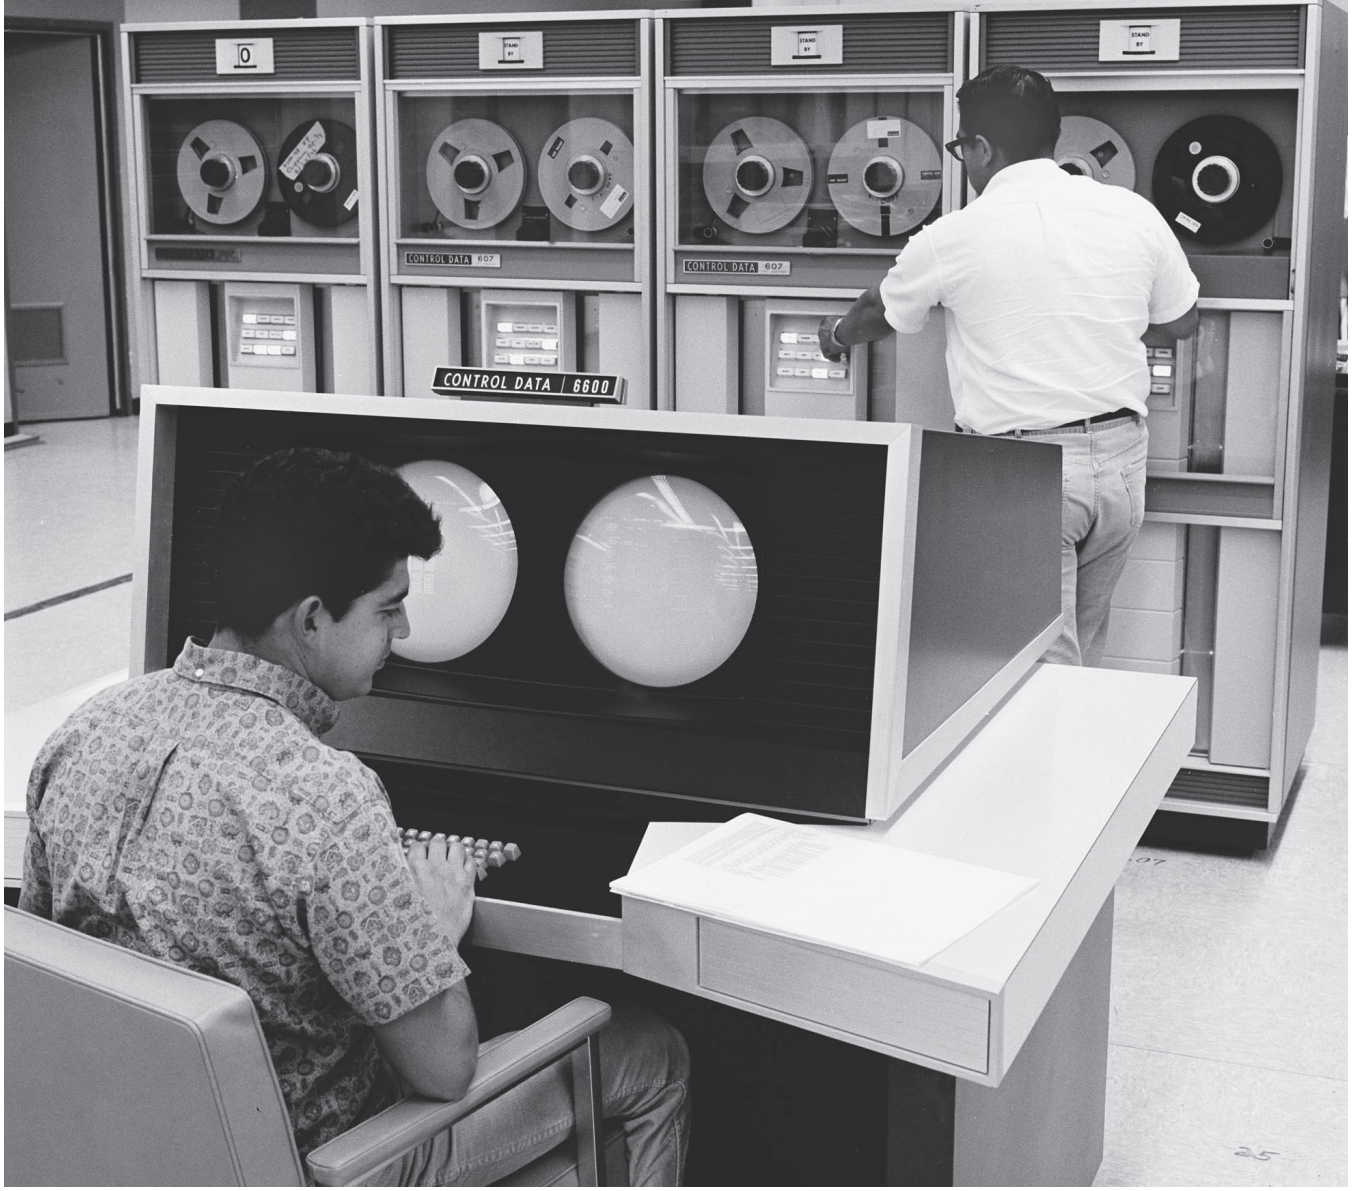

The CDC 6600, the first computer to break the megaflop barrier, is generally regarded as the world's first supercomputer. (Photo: Los Alamos)

bomb research. In fact, ENIAC's first job was a hydrogen bomb calculation for the Los Alamos staff. ENIAC, then still unknown to the public because its existence was classified, had quietly ushered in the age of modern computing.

Marchant desktop calculators and IBM punched-card machines continued to see service at the Laboratory for years after the war. But the quest for more-complex and more-powerful weapons called for more-complex, and more-powerful computers. As Metropolis and Los Alamos physicist Frank Harlow remembered, the "experience of the war years was enough to excite the involved scientists and engineers to the power of mechanized calculations and acted as a tremendous spur to the postwar development of the modern computer."

Cold War Computing

In the months and years following World War II, scientists at Los Alamos refined fission weapons and explored the feasibility of building the hydrogen bomb, a weapon many orders of magnitude more powerful than Little Boy and Fat Man. Von Neumann arranged for the ENIAC to run some of the early hydrogen bomb calculations in Pennsylvania, but it soon became clear that Los Alamos needed its own modern computer. Metropolis, who was working at the University of Chicago, accepted an invitation to return to Los Alamos to build such a machine.

The ENIAC had spawned the development of several similar computers. Metropolis, who consulted with von Neumann, studied several of these computers and designed the Los Alamos version to be more powerful and user friendly. He called it the Mathematical Analyzer, Numerical Integrator, And Computer, or MANIAC for short. As construction started in 1948, research on the hydrogen bomb progressed steadily. In the months after the Soviets conducted their first atomic bomb test, in August 1949, work on the hydrogen bomb accelerated.

In early 1952, the MANIAC was completed. Several months later, on October 31, the hydrogen bomb—the world's first full-scale thermonuclear device—was tested at Enewetak Atoll in the Pacific. The test, dubbed Ivy-Mike, unleashed a blast equivalent to nearly 500 Fat Man–type bombs and completely vaporized the small island it was conducted on.

The innovative Connection Machine, CM-5, was the first massively parallel supercomputer at Los Alamos. It was built by the Thinking Machines Corporation. (Photo: Los Alamos)

Computers played no small role in the timely success of Ivy-Mike. As the Laboratory's second director, Norris Bradbury, stated in 1954, "Computers are an essential part of any thermonuclear computation. They have a very great task to play because the computations in this field are not things you make with a slide rule or a small pad of paper." He concluded, "Only recently, with the development of machines such as the MANIAC, the computer at Princeton, [and] IBM computers, have we had the machines which even begin to attack the problem which was confronting us."

Throughout much of the '50s and '60s, the Laboratory managed to double computing capacity every two years.

The MANIAC line of computers enjoyed a long and successful run at the Laboratory and beyond. The MANIAC II, which was completed in 1956, was easier to operate than the original. After more than 20 years of service, MANIAC II was shut down in 1977. Metropolis also constructed a third machine, MANIAC III, at the University of Chicago in the early 1960s. But in the 1950s, as digital electronic computing technology became less expensive, more reliable, and more powerful, commercially produced computers started to gradually displace ENIAC's handmade descendants.

The Laboratory purchased its first commercial computer, an IBM 701, in 1953. The acquisition of the 701 opened a new era in Los Alamos computing, which would be dominated by commercial machines and custom computers developed jointly with corporate partners. Throughout much of the 1950s and 1960s, the Laboratory managed to double computing capacity every two years. This remarkable achievement was made possible through partnerships with private companies and breakthroughs in computing technology.

The Laboratory purchased its first commercial computer, an IBM 701, in 1953.

The most significant advancement during this era was the development of transistors. Up to that point, computers relied on vacuum tubes, which produce heat and require routine replacement, to control electric currents. The ENIAC, for instance, contained over 17,000 vacuum tubes. Transistors, on the other hand, were smaller, more reliable, cheaper, and less complex. To meet the growing computing needs of the weapons program, in 1961 the Laboratory received an IBM Stretch, the company's first transistor computer**.** Although the Stretch never achieved the lofty performance goals set by IBM, it retained the title of world's fastest computer into the mid-1960s.

Los Alamos scientists next looked to Control Data Corporation (CDC) for machines with even more power. CDC delivered by producing the world's first supercomputer, the model 6600. The 6600s, which were the first computers capable of performing a million floating-point operations per second (megaflops), were soon supplemented by even faster CDC 7600 models.

Seymour Cray, the CDC designer who led the development teams that produced the 6600 and 7600, left the company to start his own in 1972. His company, Cray Research, completed its first design, the revolutionary 160-megaflop Cray-1, in 1975 and delivered it to Los Alamos the following year. The Cray-1 used integrated circuits (individual chips containing numerous transistors) to improve performance and an innovative Freon cooling system to ensure the machines did not overheat. Seymour Cray also used revolutionary "vector" processing, which enabled the Cray-1 to process infor-mation far more efficiently than any other computer of its day. During the 1980s, the Laboratory purchased additional Cray computers, most notably the X-MP. From 1982 to 1985, the X-MP, which used multiple "vector" processors, reigned as the world's fastest computer.

Seymour Cray with the Cray-1. The Cray-1's elegant shape increased performance by decreasing the length of signal-carrying wires. (Photo: Cray Inc.)

The women who operated the wartime Laboratory's desktop calculators and punched-card machines were themselves were called "computers." They often were the wives of Los Alamos scientists. (Photo: Los Alamos)

Supercomputers would change almost as rapidly and as drastically as the global political landscape of the early 1990s.

As the 1980s drew to a close, Los Alamos remained a key driver in the evolution of computing by once again partnering with IBM and starting a collaboration with the Thinking Machines Corporation. Thinking Machines' massively parallel Connection Machine series, which used thousands of microprocessors to perform numerous calculations simultaneously, would take Los Alamos into the gigaflop era (a billion floating-point operations per second), which had already been opened by the Cray-2 elsewhere. But the fortunes of Thinking Machines, despite its innovative lineup of supercomputers, would change almost as rapidly and as drastically as the global political landscape of the early 1990s.

Computing Since 1992

As the Cold War came to an abrupt end, government funding for supercomputers shrank. These cutbacks played a role in bankrupting Thinking Machines in 1994 and Cray Computer Corporation, an offshoot of Cray Research, the following year. But just as these companies went out of business, Congress created the Science-Based Stockpile Stewardship Program "to ensure the preservation of the core intellectual and technical competencies of the United States in nuclear weapons, including weapons design, system integration, manufacturing, security, use control, reliability assessment, and certification." Specifically, the new law called for "an increased level of effort for advanced computational capabilities to enhance the simulation and modeling capabilities of the United States with respect to the detonation of nuclear weapons." As such, the Accelerated Strategic Computing Initiative (ASCI) was launched to rapidly develop the much more powerful computers necessary to sustain the Stockpile Stewardship Program.

The success of ASCI came largely as a result of unprecedented levels of cooperation between the national laboratories and private industry. At Los Alamos, the partnership with Cray resumed with a trio of machines in the mid-1990s.

The evolution of the computer between the 1940s and 1960s, as told by Patsy, Gail, Milly, and Norma. (Photo: U.S. Army)

The T3D, which arrived first, marked Cray's entry into massively parallel computing, and the T90 and J90 machines, which came soon after, pushed vector technology to its limit. But for ASCI to meet its national security goals, computing would need to make the technological leap from billions to trillions of floating-point operations per second (teraflops). A new architecture, parallel clusters, would make it possible.

The system could perform 3 trillion floating-point operations per second, making it the world's third-fastest computer in 1999.

As its name implies, a parallel-cluster computer is actually many computers that function together as a single unit. The Laboratory's first parallel cluster machine, Blue Mountain, was a collection of 48 Silicon Graphics computers. The system could perform 3 trillion floating-point operations per second, making it the world's third-fastest computer in 1999. Blue Mountain trailed only two ASCI counterparts at Sandia National Laboratories and Lawrence Livermore National Laboratory, respectively, as the millennium came to a close.

ASCI Q, the collaborative product of a partnership between Hewlett-Packard and Los Alamos, emerged as the world's second-fastest computer in 2002. Q eventually achieved a speed of 20 teraflops. But it would take the revolutionary hybrid cluster technology of Roadrunner to finally break the petaflop barrier in 2008, thus giving Los Alamos the world's fastest computer again for the first time since one of its Connection Machines, the CM-5, in 1993.

National Security and the Future of Scientific Innovation

Throughout the Cold War, many of the world's most powerful computers were developed for national defense purposes, in particular for applications pertaining to the development and maintenance of nuclear weapons. During the Laboratory's 40th anniversary year of 1983, Frank Harlow and Nicholas Metropolis stated, "It is a stunning tribute to Los Alamos bomb designers and their colleagues that many of the most powerful procedures for taming computers to the myriad tasks of modern science and technology were developed right here."

Throughout the history of the Laboratory, computers have been specifically developed for the nuclear weapons program. Today, as the Laboratory turns 70, more-powerful computers enable more-detailed weapons simulations. More-detailed weapons simulations, supported by the Laboratory's experimental data, produce greater certainty in assessing the nuclear stockpile for safety and reliability. Greater certainty in assessing the nuclear stockpile ensures the nation will be able to maintain a credible nuclear deterrent well into the future.

But throughout the Laboratory's history, computers have also been used to further many fields of scientific endeavor. Biological explorations into the human genome, which continue today, can be traced all the way back to the 1950s, when Los Alamos scientists attempted to decode DNA sequences using MANIAC. Through the years, Laboratory computers have also been used for mineral exploration, basic science, and energy research. As we move deeper into the 21st century, Laboratory computers will continue to produce ever more detailed and accurate models for understanding global climate change, the spread of pandemics, and the nature of our universe, as well as the state of the nuclear weapons stockpile.

When Roadrunner broke the petaflop barrier, it inspired a new generation of supercomputers worldwide. These versatile machines enable scientists to perform research in a wide range of fields, but national security applications continue to play a significant role in driving the development of computing technology itself. In fact, as of November 2012, Department of Energy laboratories possess three of the four most powerful supercomputers in the world, including the fastest, Titan, a Cray machine at Oak Ridge National Laboratory. Roadrunner changed history, but it was not the first Los Alamos computer to do so. It will not be the last.

– Alan B. Carr

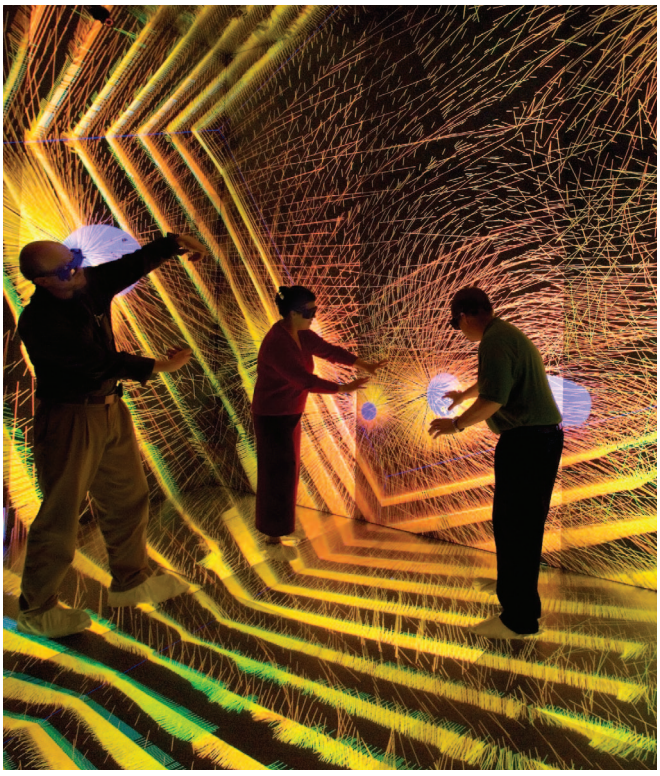

Today, Los Alamos' supercomputers enable scientists to run highly detailed 3D interactive simulations to solve problems in national security and general science. (Photo: Los Alamos)

Strategic Deterrence

in the 21st Century

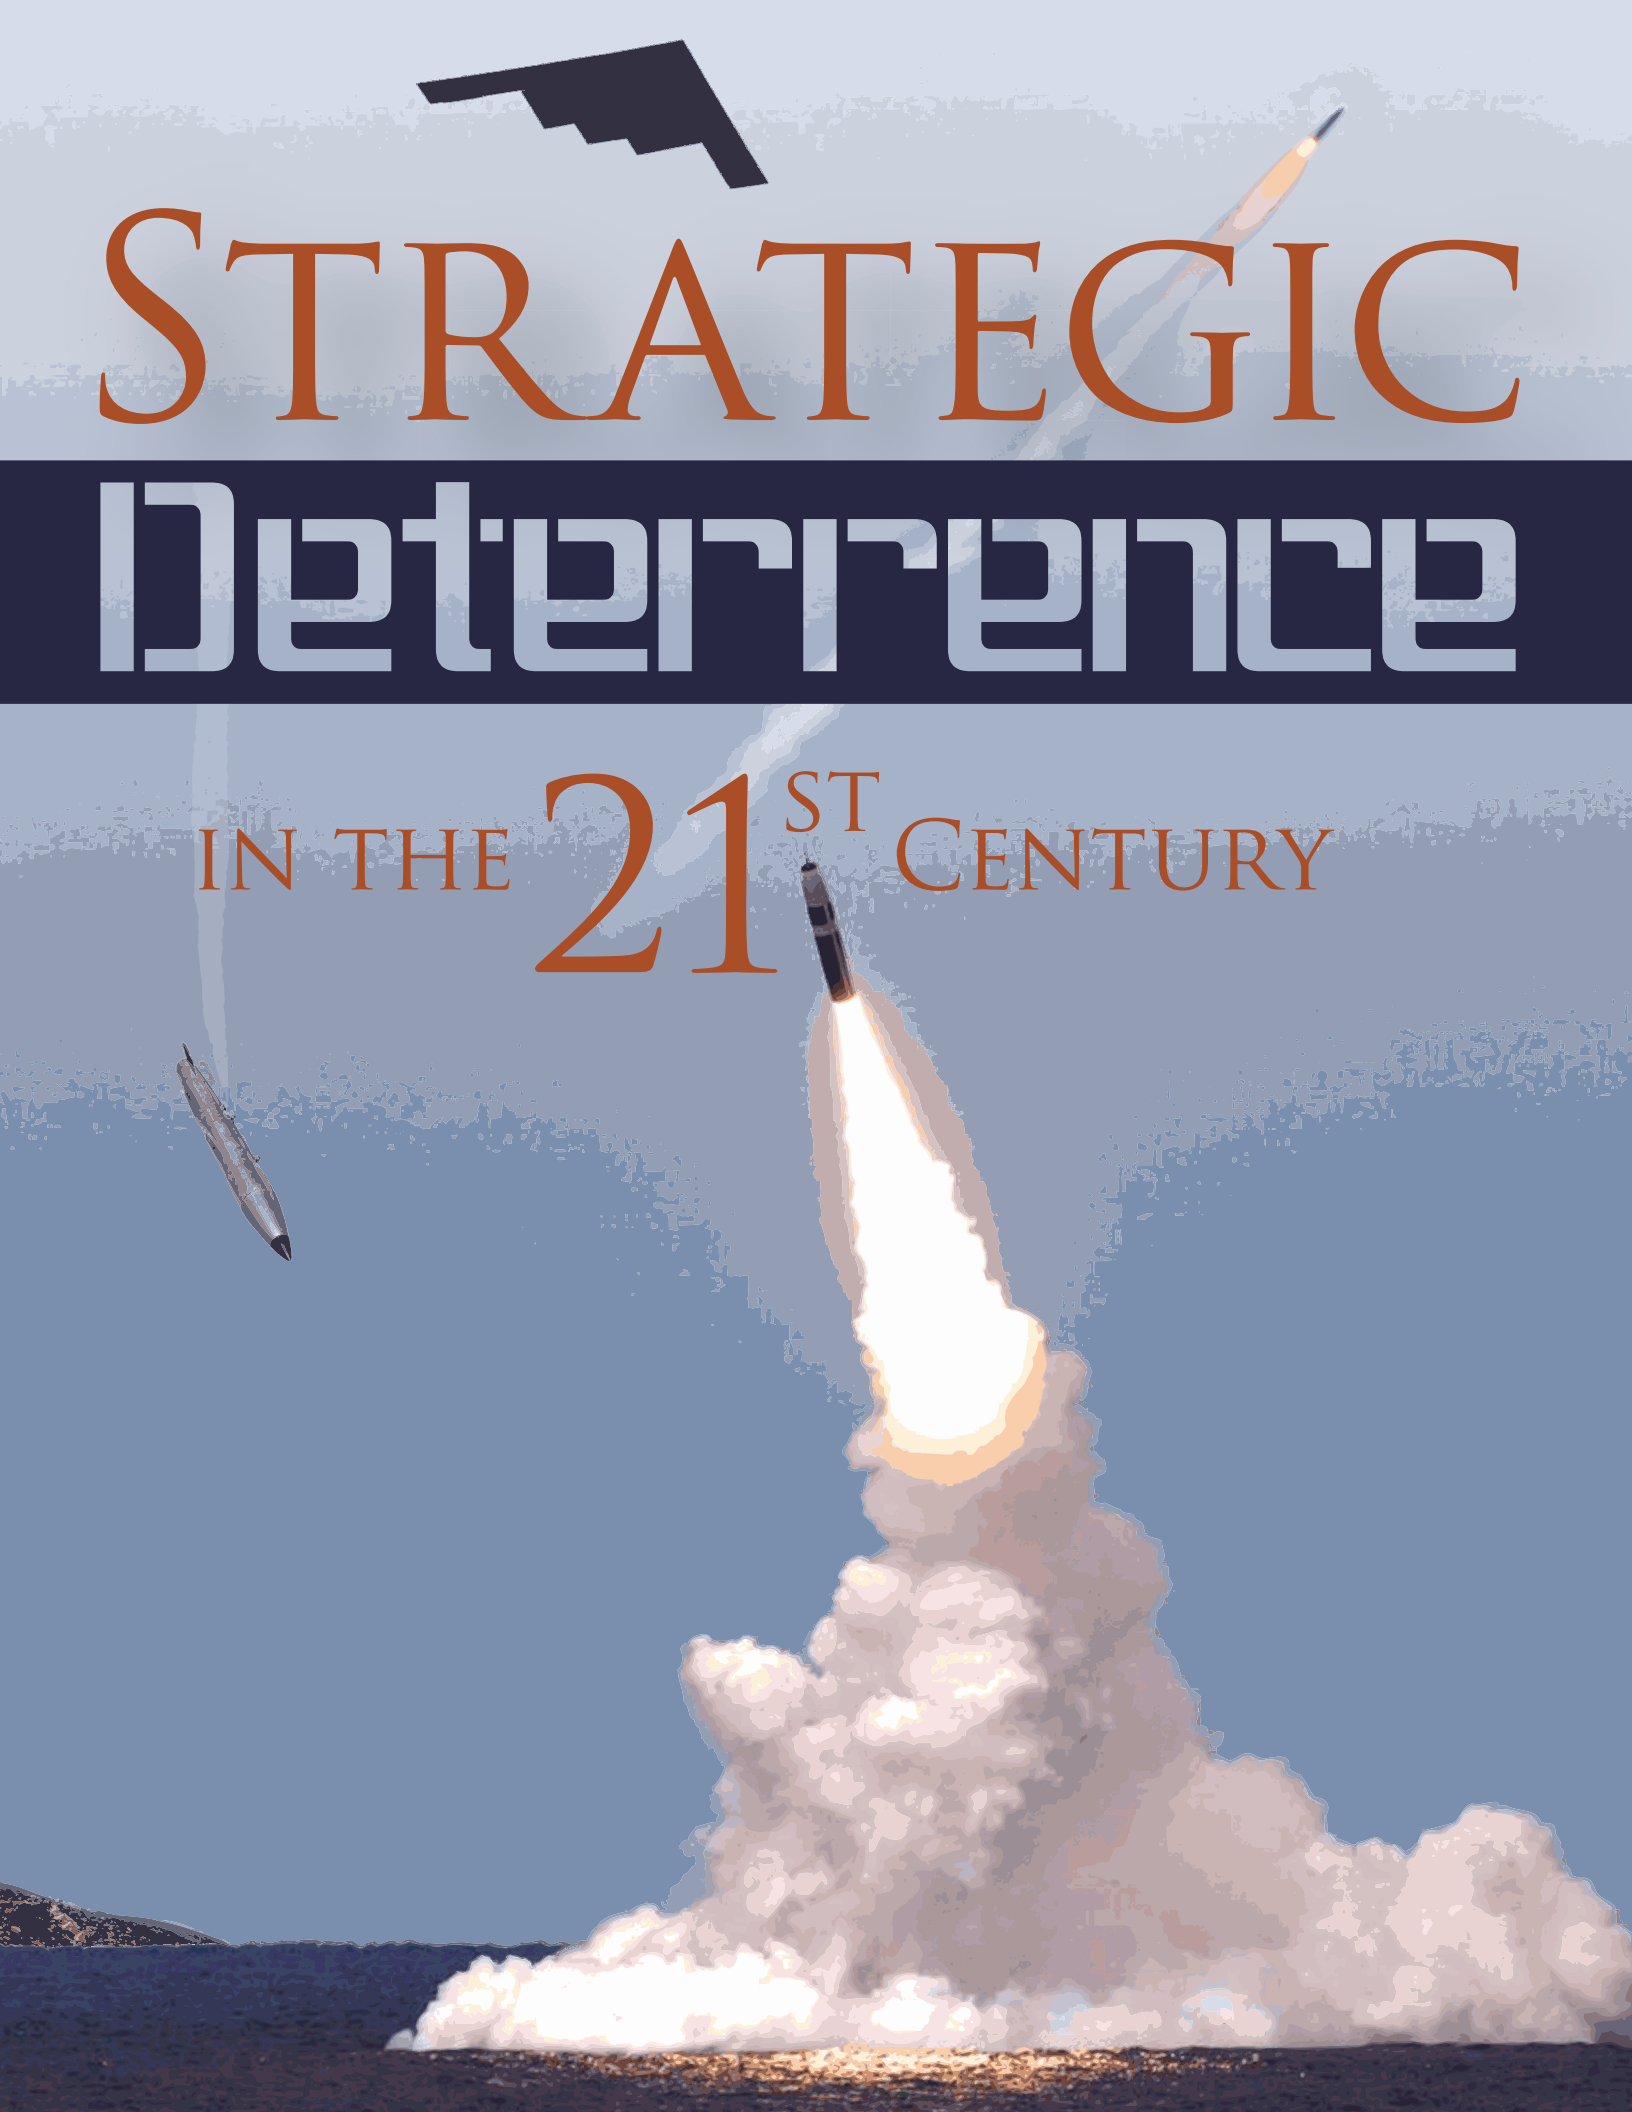

Editor's note: As a service to our readers, NSS *provides this article (originally published in* Undersea Warfare *magazine, Spring 2012) because it presents, in detail, many of the key issues in the debate over the role of nuclear weapons in today's security environment. The author, Admiral Richard W. Mies (retired), is a former commander, Submarine Forces, and former commander, U.S. Strategic Command. He is chairman of the Naval Submarine League Board of Directors.*

The viewpoints expressed in this article are the author's and do not reflect the viewpoints of Los Alamos National Laboratory.

No discussion of deterrence strategy in the 21st century can be meaningful without a clear understanding of how nuclear weapons have revolutionized and transformed warfare. In a small book written at the dawn of the nuclear age, a group of scholars drew some profound and prescient conclusions about the significance for human warfare of what they termed "the absolute weapon." The authors recognized that the atom bomb was revolutionary and fundamentally different from conventional weaponry. Pound for pound, nuclear weapons were several million times more potent; no adequate defense against them was known or foreseen to exist; and some proliferation of nuclear weapon technology to other nations was inevitable, barring international control.[1] One of the most insightful, fundamental conclusions they reached reflected the atom bomb's revolutionary nature:

"Thus far the chief purpose of our military establishment has been to win wars. From now on its principal purpose must be to avert them."[2]

Nuclear weapons have extended the potential of warfare to a level where classical warfare concepts cease to have meaning—to the reductio ad absurdum[3] of warfare. In parallel, they have also come to be seen as different not just by their potency, but "by convention—by an understanding, a tradition, a consensus, a shared willingness to see them as different."[4] And this revolution in warfare—the virtually unlimited capacity to harm each other—is likely to be with us forever, since the knowledge to build nuclear weapons cannot be erased.

The Transformation of Warfare

Because of their revolutionary nature, nuclear weapons are, first and foremost, instruments of national policy, as opposed to instruments of military operations. Nuclear weapons serve as a deterrent against major war, a hedge against an uncertain future, a guarantee of our security commitments to our allies and friends, and a disincentive to those who would contemplate developing or otherwise acquiring their own nuclear weapons. They are primarily weapons of war *prevention*, as opposed to war *fighting*, although war prevention and war fighting cannot be totally disassociated. Nuclear weapons deter by the possibility of their use and by no other means.

Deterrence strategies, which evolved during the Cold War, recognize that the greatest utility of nuclear weapons is in their non-use—in the diplomacy derived from the threat of their use. In that sense, *nuclear weapons are used every day.* The concepts of deterrence, assurance, and dissuasion associated with nuclear weapons differ fundamentally from classical military strategy in that they deal with the exploitation of *potential* force rather than the application of force.

A ground-based interceptor lifts off from Vandenberg Air Force Base in California. The launch is a test of the Ballistic Missile Defense System, which successfully intercepted a long-range target launched from Kodiak, Alaska. (Photo: DoD)

They are intended to shape behavior and, as such, they share some common elements of inducements—of threats and/or promises, explicit or implicit—to either prevent or promote an action. Their primary purpose is to influence potential adversaries' intentions far more than their capabilities through two interrelated means—the power to hurt and the power to deny.[5] These powers are most successful when held in reserve and their non-use, their potential, exploited through diplomacy. The most successful threats are the ones that never have to be carried out. As Sun Tzu noted, "To subdue the enemy without fighting is the acme of skill."[6]

Flexible Response

The great paradox of nuclear weapons is that they deter conflict by the possibility of their use, and the more a potential adversary perceives the credibility of our capabilities and will, the less likely they are to challenge their use. The converse of that proposition is also true. To be credible, capabilities and plans have been developed since the early 1960s to provide the president with as broad a range of options as considered prudent to enable the president to respond with the *minimum use of force* sufficient to deny an adversary's objective.

Nuclear deterrence ultimately depends on the threat of retaliation—not on our capability to strike first, but on the assurance we always have the capability to strike second.

This has been the nature of the concept of "flexible response" and the core of U.S. and NATO targeting doctrines. To argue that this has made nuclear weapons more useable is to ignore their central paradox and their fundamental difference from conventional weapons. To allow nuclear weapon use to become incredible would increase, not lessen, the risk of war.

And because nuclear weapons are primarily designed for war avoidance, nuclear deterrence ultimately depends on the threat of retaliation—not on our capability to strike first, but on the assurance we always have the capability to strike second. In my experience, our strategic forces have always been viewed by our leaders as weapons of last resort, to be employed only when deterrence has failed and all other means to counter aggression or coercion have failed.

From a war-fighting perspective, nuclear weapons have historically been regarded as the nation's "ultimate insurance policy"—de facto weapons of last resort—the least-preferred option, short of surrender, to protect vital national interests.

Strategic Force Evolution

During the past decade, our strategic forces have been on a journey of reductions that was charted in the 2001 and 2010 Nuclear Posture Reviews (NPR) and codified in the Moscow Treaty and, more recently, the New START Treaty. The journey began out of recognition that U.S. nuclear doctrine and forces needed to have lower salience and a less adversarial character, most directly as a result of our changed relationship with Russia, and also out of recognition that deterrence was likely to be more complex and perhaps less reliable, particularly against non-state actors, *although not necessarily less relevant*. I emphasize that this is about a journey rather than a destination because the journey is far more important than the destination.

Simultaneously, *since the end of the Cold War, we have experienced significant erosion in our strategic deterrent capabilities well documented in a number of reports.*[7] In spite of the rhetoric of the past two NPRs and the National Defense Strategy, there has been a paucity of thinking by senior-level decision-makers about the role of our strategic deterrent, and particularly the role of nuclear weapons in the 21st century. Many reasons are given for this, such as the Global War on Terror, operations in Afghanistan and Iraq, unchallenged U.S. conventional superiority. Nevertheless, the result is a glaring mismatch between the rhetoric of national strategy and the resources committed to our national strategy objectives.

Despite recent actions to arrest some of this erosion, our strategic forces appear to be adrift—paralyzed by inaction and a lack of consensus. The fundamental underlying cause has been a lack of attention to nuclear weapon issues by senior leadership—both civilian and military—across both present and past administrations. This lack of senior leadership attention has resulted in public confusion, congressional distrust, and a serious erosion of advocacy, expertise, and proficiency in our nuclear forces.

Our Aging Nuclear Enterprise

While we have made great progress in the drawdown of our strategic forces, progress to modernize our strategic deterrent enterprise has been inadequate to meet our national security needs. If one thinks about our strategic capabilities as an enterprise, it really resembles a pyramid, as Figure 1 depicts, whose foundation is the scientific and technological expertise resident in our nuclear complex employees and in our strategic operating forces. That foundation is growing increasingly thin and brittle—through both an aging workforce and difficulties recruiting and retaining the best and brightest.

And while many have spoken eloquently about the importance of science and technology programs as critical underpinnings of the Department of Energy's (DOE) portion of the nuclear enterprise, there are really few, if any, programs on the Department of Defense (DoD) side that are analogous to DOE's science-based stockpile stewardship program or the advanced computing initiatives. We have raised a whole generation of *war-fighters* within DoD who have received virtually no professional education in the theories of deterrence, assurance, and dissuasion, and who consequently often fail to think in *war-prevention*

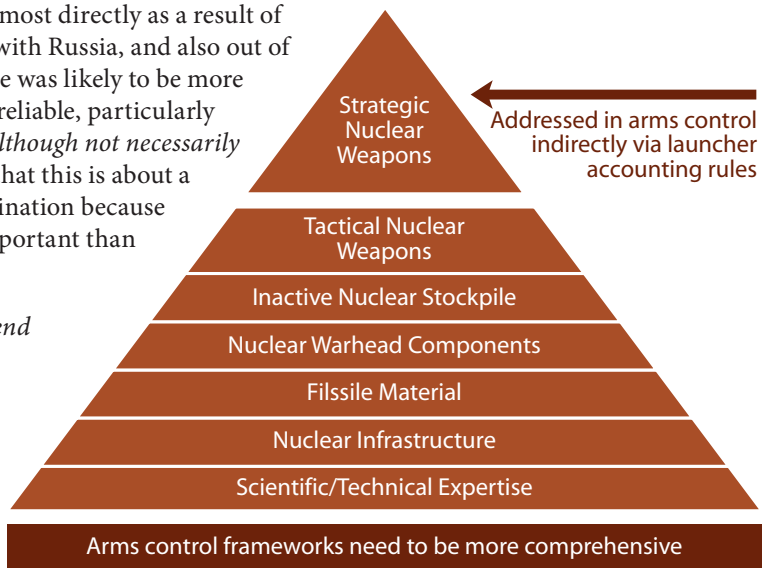

Figure 1. The U.S. nuclear enterprise is aging on all levels.

terms. Additionally, there has been until recently little, if any, programmatic advocacy within the Office of the Secretary of Defense, the Joint Staff, and the military services for the strategic nuclear enterprise.

Several points are worthy of mention with respect to this enterprise pyramid. Foremost, deterrence depends on the health of the entire pyramid, not just any one element. We can't deter with just a strong foundation—a "virtual deterrent" is simply not credible. Second, the distinction between tactical and strategic nuclear weapons is an outmoded, treaty-derived distinction that relates more to delivery platforms than actual warheads. There is little significant difference in the design and capabilities of our tactical and strategic warheads. The principal distinction is in the delivery platform; any tactical nuclear weapon can be used with strategic effect.

Despite these factors, our focus on the enterprise tends to be disproportionately narrow—driven to an over-emphasis on the very top of the pyramid—to strategic weapons—and even then indirectly—because of our captivation with strategic warhead numbers.[8] As a consequence, we often fail to view the enterprise in a more comprehensive way.

Deterrence depends on the health of the entire pyramid, not just any one element.

Figure 2 illustrates the aging of our legacy Cold War stockpile and our lack of robust design and production capability. We have lost people with unique skills as well as design and production knowledge. Many of our warheads are beyond their design lives and lack desirable safety and surety features we are now capable of incorporating into replacement designs. Our legacy warheads are sophisticated machines,

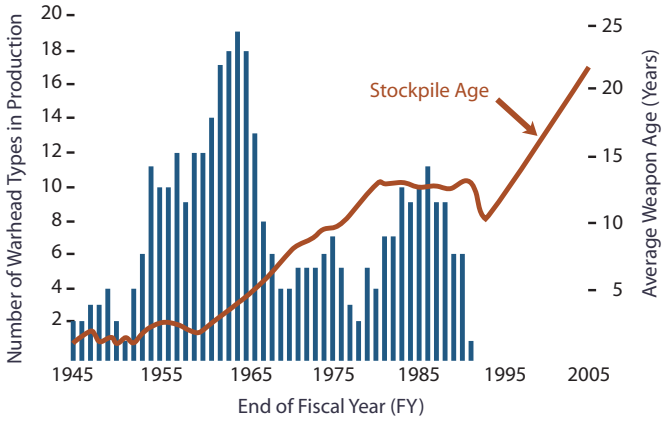

Figure 2. The aging of the nuclear stockpile inherited from the Cold War.

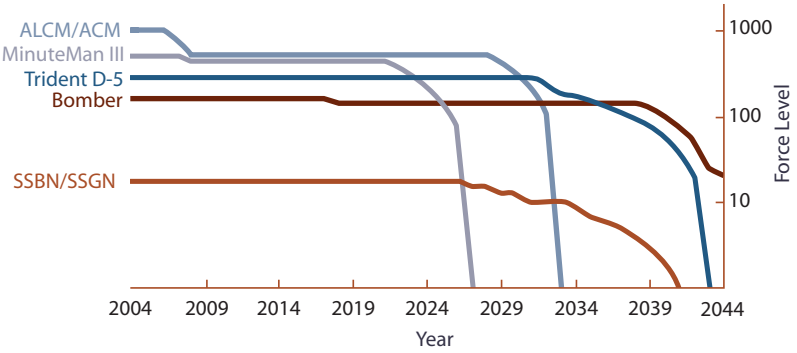

Figure 3. The aging of U.S. strategic nuclear delivery platforms. All are approaching end-of-life.

similar to a 20th century Rolls Royce, with as many as 6,000 intricate parts and complex chemical interactions. Because of their sophistication, some warhead performance margins are extremely narrow. And unlike wine, the reliability of sophisticated machines doesn't improve with age. The best we can do is to extend their lives. Needless to say, reestablishing design and production capabilities remains a very complex and lengthy process.

Figure 3 complements the previous one. Not only is our warhead stockpile aging, all of our strategic delivery systems are aging and approaching end-of-life in an austere and potentially adverse fiscal environment. Contrast this with other key nuclear-capable nations who are modernizing substantially their strategic forces.

Risks and Uncertainties of Strategic Force Reductions

As we contemplate further reductions in our nuclear forces beyond the New START Treaty to lower levels consistent with our national security needs, we will inevitably encounter several risks related to the national security concepts of deterrence, assurance, and dissuasion.

A smaller arsenal may appear to be a more tempting and easier target for preemption, breakout, or a race to parity.

First, some of our allies may seriously question the credibility of our extended nuclear deterrent, so instead of promoting non-proliferation, our reductions may have the perverse, opposite effect. Decades ago, British Prime Minister Denis Healey explained the difference between extended deterrence and assurance with the observation that, "it takes only 5 percent credibility of American retaliation to deter the Russians, but 95 percent credibility to reassure the Europeans." By this, he meant that assuring allies may be more challenging than deterring foes, that there are different measures of adequacy for these two different goals.[9]

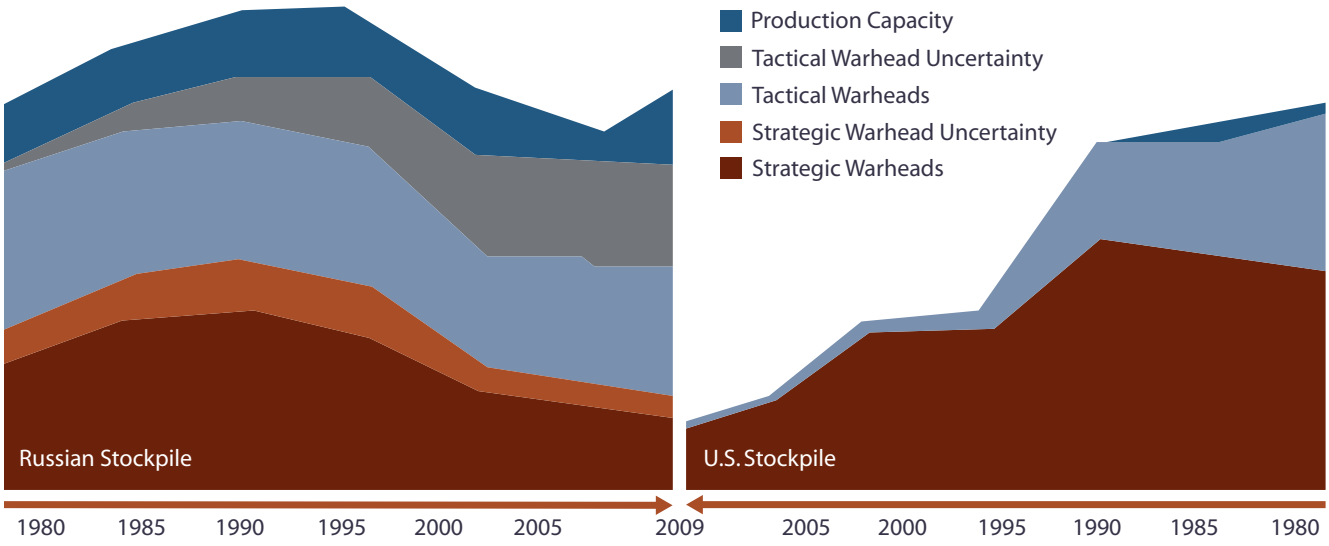

Figure 4. U.S./Russian stockpile comparison. Note the decline of the U.S. nuclear weapons stockpile compared with Russia's.

Second, below certain levels, potential adversaries may be encouraged to challenge us. A smaller arsenal may appear to be a more tempting and easier target for preemption, breakout, or a race to parity.

Third, at some level, it will become more difficult and economically impractical to sustain the present strategic triad. While there is nothing sacrosanct about the triad, numerous analyses and studies have repeatedly reaffirmed the wisdom of preserving the complementary capabilities of land-based intercontinental ballistic missiles (ICBMs), submarine-launched ballistic missiles (SLBMs), and strategic bombers.[10]

Each leg of the triad contributes unique attributes that enhance deterrence and reduce risk, such that the whole is greater than the sum of the parts. ICBMs provide a prompt response, the potential to launch under attack[11], and a hardened, geographically-dispersed target base. Additionally, single-warhead ICBMs are considered stabilizing, since they are less attractive as targets than multiple-warhead ICBMs because the ratio of weapons required to destroy them is greater than one. Missile submarines provide survivable, assured response and the mobility to adapt missile over-flight to targets. Strategic bombers provide great flexibility in force posturing, signaling intentions, route planning, and recall-ability.

Together they comprise a robust deterrent capability that complicates a potential adversary's offensive and defensive planning and a synergistic force that provides protection against the failure of a single leg.

Our unilateral disarmament initiatives have done little to promote similar initiatives in our potential adversaries; they have reduced our arms control negotiating leverage.

A fourth risk concerns the asymmetries in U.S. and Russian nuclear stockpiles. Figure 4 is a *relative* comparison of the U.S. and Russian nuclear stockpiles over the past three decades. (Note that both stockpile charts start from the outside and work toward the center.)

This comparison raises several noteworthy points. First, we have dramatically and unilaterally drawn down our tactical nuclear forces in contrast to Russia. To my knowledge, our unilateral disarmament initiatives have done little to promote similar initiatives in our potential adversaries, and at the same time, they have reduced our arms control negotiating leverage. In that sense, the lead part of the "lead and hedge" strategy—the idea that if we lead, others will follow—has proven illusory.

The U. S. has sought to maintain a nuclear weapons capability "second to none." Are we in danger of allowing our nuclear preeminence to become "second to one"?

Second, and similarly, the NPR's promises of a responsive infrastructure remain largely unfulfilled. In contrast to Russia, we have had virtually no warhead production capability for the past two decades and have little likelihood of developing a robust one within the coming decade.

Finally, because of the difficulties and our lack of leverage in expanding treaty negotiations to include tactical nuclear forces and production capability, if we jointly agree to reduce our strategic nuclear forces to even lower levels, the asymmetries in our respective stockpiles will become even more pronounced. As stated earlier, the artificial and inappropriate distinction between strategic and tactical nuclear weapons is cause for concern.

As Ambassador Robert Joseph has written, "Since the start of the atomic age, from Harry Truman to George W. Bush, the United States has sought to maintain, in the words of John F. Kennedy, a nuclear weapons capability 'second to none.'" Are we in danger of allowing our nuclear preeminence to become 'second to one'?[12]

Those who advocate nuclear abolition need to answer some fundamental questions about the logic of zero.

A fifth risk concerns strategic targeting doctrine. Figure 5 is a notional chart intended to illustrate several of the dilemmas of strategic targeting. The curve on the right represents our present and long-standing targeting doctrine of flexible response—a doctrine designed to hold at risk our potential adversaries' military forces, war-supporting industry, command and control capabilities, and military and national civilian leadership, while minimizing to the maximum extent possible collateral damage to population and civilian infrastructure. It is a doctrine designed to provide the president the widest range of options using the minimum level of force intended to achieve our objectives. The curve on the far left illustrates that if we adopted a counter-population targeting strategy, we could achieve significantly more damage with fewer weapons. But at what cost and credibility?

As we reduce the number of available weapons, that flexible response curve moves to the left, which will diminish the robustness and flexibility inherent in a moderately sized arsenal (a few thousand, as compared to a few hundred). Greater stress will be placed on the reliability and survivability of our remaining forces. As stated earlier, at some level, it will become more difficult and economically impractical to sustain the present strategic triad.

And of greatest concern, it will reduce the range of flexible response options designed to provide the president with minimum use of force. Ultimately, below a certain level, to remain credible our targeting doctrine and policies would have to shift away from our traditional flexible response targets to counter-population targets, as depicted by the two curves on the left, which represent the range of counter population options. This transition would be counter to our historical practice, politically less tolerable, and morally repugnant. Although I am not an international lawyer, I would also argue that such a transition is in violation of the Law of Armed Conflict and the Theory of Just War.

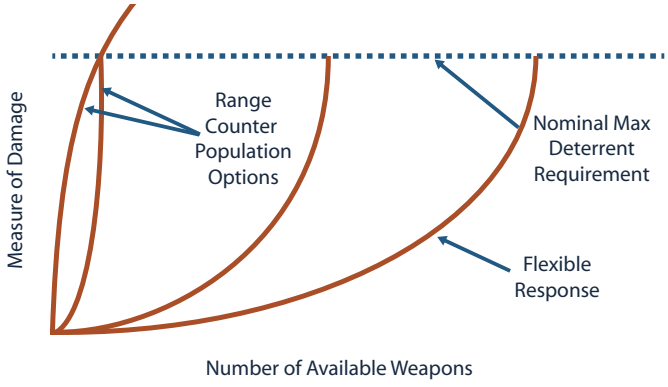

Figure 5. The strategic targeting doctrine dilemma: the relationship between warhead numbers and strategic doctrine.

The Illogic of Zero[13]

In light of the aforementioned transformation of warfare, the widely publicized initiative to eliminate nuclear weapons deserves critical review. Theories and concepts abound on the political, strategic, and military significance of nuclear weapons, but we should be mindful of their limitations. We lack sufficient hard evidence about the consequences of nuclear weapon abolition. In the words of an experienced practitioner:

> *"The resulting limitations in our knowledge ought to instill in all who make predictive statements about these issues a degree of humility not always evident... There is no substitute for looking at the merits of what is said [rather] than the eminence of who said it … the means for creating a world without actual nuclear weapons would have to be of a basic political kind, not a matter of technical arms control. Secure nuclear abolition would be consequence, not cause; and in the journey it has to be cart, not horse... Better unquestionably, pending political transformation, to have nuclear weapons but not war than to have war but not nuclear weapons."*[14]

If biological terrorism remains a major threat despite the abolition of biological weapons, why do proponents believe that the abolition of nuclear weapons will significantly reduce the nuclear threat?

If, as another experienced statesman has stated, "Nations don't distrust each other because they are armed; they are armed because they distrust each other,"[15] shouldn't our focus be on the more fundamental, underlying causes of distrust instead of disarmament? Hence a significant burden of proof rests upon *those who advocate nuclear abolition. They need to answer some fundamental questions about the logic of zero.* Without compelling answers to these questions and achievable actions, I believe their vision will prove counterproductive, promote unrealistic expectations, and serve as justification to keep the strategic enterprise adrift—paralyzed and frozen in time.

First: Is it feasible? If so, what detailed, specific actions must be taken by individual nations and the international community, and in what time frames? How do you achieve those reductions and avoid the risks and uncertainties

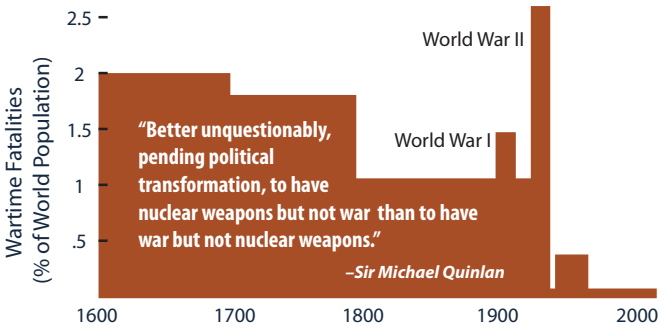

Figure 6. Strategic deterrence –a transformation in warfare. The historical trend of wartime fatalities as a percentage of world population from 1600 to 2000, showing the effect of nuclear deterrence after World War II.

outlined previously? Which other nuclear nations share the abolitionist vision and have actually demonstrated a commitment to work in concert toward that goal?

Second: Is it verifiable? If so, by whom, and with what means? How would compliance be enforced? Considering the examples of North Korea and Iran, is such an intrusive and comprehensive verification regime achievable in our existing geopolitical framework?

Third: If it is both feasible and verifiable, is it inherently stabilizing, and hence sustainable? Since the knowledge to build nuclear weapons cannot be erased, and many nations will have latent nuclear capabilities, what disincentives will preclude cheating or breakout? If biological terrorism remains a major threat despite the abolition of biological weapons, why do proponents believe that the abolition of nuclear weapons will significantly reduce the nuclear threat? What means will exist to prevent a terrorist from acquiring fissile material, which will still be in abundant supply? What means will exist to prevent a rogue nation from aspiring to become a nuclear superpower in a non-nuclear world? As a former professor of mine has written, under abolition, present nuclear powers would actually be latent nuclear powers—hardly "former nuclear powers." If the atom bomb could be invented from scratch during World War II, imagine how quickly the nuclear genie could be conjured back into action now.

> *"In summary, a world without nuclear weapons would be a world in which the United States, Russia, Israel, China and half a dozen or a dozen other countries would have hair-trigger mobilization plans to rebuild nuclear weapons and mobilize or commandeer delivery systems, and would have prepared targets to pre-empt other nations' nuclear facilities, all in a high-alert status, with practice drills and secure emergency communications. Every crisis would be a nuclear crisis; any war could become a nuclear war. The urge to preempt would dominate; whoever gets the first few weapons will coerce or preempt. It would be a nervous world."*[16]

Lastly, if nuclear weapon abolition can be achieved and sustained, is it really desirable? How can we be sure we are not making the world safe for conventional war? And while it may be imaginable to envision a world without nuclear weapons while we are the world's superpower, how safe and secure will we be as a nation when, at some future, inevitable time, we no longer enjoy that distinction? To me these are the most fundamental questions the abolitionists blithely ignore.

Figure 6 reinforces this last question. As this graph of wartime fatalities as a percentage of world population illustrates, conventional warfare took a devastating toll throughout history before the advent of nuclear weapons. However, since the advent of nuclear weapons, the transformation of warfare has been dramatic. The fact that there has not been a war between major powers in almost 70 years is without historical precedent. In contrast, the idea that conventional weapons can credibly deter as effectively as nuclear weapons lacks historical evidence.

As Margaret Thatcher has reportedly stated, "There is a memorial to the failure of conventional deterrence in every town and village in Europe…. A thousand years of European history prove that conventional weapons do not deter."[17] What evidence do those advocating disarmament and nuclear abolition proffer that illustrates how disarmament has made the world more peaceful?

Nuclear forces are in reality very cost-effective relative to conventional forces and historically have consumed less than 5 percent of the DoD budget.

Conventional deterrence can obviously complement strategic deterrence; but, there is no evidence it can supplant it. Regardless of force superiority, conventional weapons are contestable both temporally and geographically; in contrast, nuclear weapons are not contestable. Whereas in the past, nations sought to achieve strategic objectives through war, nuclear weapons have created a strong restraining force among nations to avert war. And that has contributed to a remarkable, revolutionary transformation in warfare.

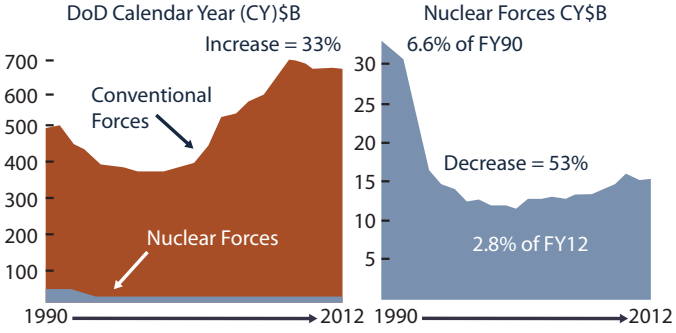

Figure 7. Nuclear force costs (1990–2012). Relative cost trends for nuclear and conventional weapons.

An F-22 Raptor and B-2 Spirit deployed to Andersen Air Force Base, Guam, fly in formation over the Pacific Ocean April 14, 2009. The F-22 Raptor and the B-2 Spirit deployment to Andersen marks the first time F-22 Raptors and B-2 Spirits, the key strategic stealth platforms in the Air Force inventory, were deployed together outside the continental United States. (Photo: U.S. Air Force)

Misperceptions About Nuclear Weapons

There is a common misperception that nuclear forces are disproportionately expensive—a rich "cash cow" that can be milked with further reductions to free up funding for other priorities. As the graph on the left of Figure 7 illustrates, nuclear forces (including dual-capable forces like bombers) are in reality very cost-effective relative to conventional forces and historically have consumed less than 5 percent of the DoD budget. Most of this cost is driven by over-head and infrastructure, such that warhead reductions will not result in meaningful savings. The graph on the right of Figure 7 is an expanded view of the nuclear force costs in the left graph. Considering their role in war prevention, one should think of our nuclear forces much like we think personally about health and life insurance. Their cost, as a small percentage of the DoD budget, is a very reasonable premium for the nation's "ultimate insurance policy."

There is also a naïve and mistaken belief that the "nuclear danger" is directly proportional to the number of nuclear weapons, and accordingly, lower is axiomatically better. However, disarmament is not inherently stabilizing. One can envision many scenarios where small numbers breed instability.

In addition, there is a common fallacy about deterrence that holds that nuclear weapons deter only nuclear weapons. To accept that, one has to accept that nuclear weapons have played no role in the remarkable peace among the nuclear powers during the past six decades despite periods of significant tension and East-West confrontation. While it is impossible to prove a negative, how else does one reasonably justify the precipitous change depicted in Figure 6?[18]

And it would be equally fallacious to assume, that without some fundamental change in the political configuration of the world, nuclear weapons have no relevance for the future. Deterrence is about preventing all major wars, not just nuclear ones, since major war is the most likely road to nuclear war. As such, a policy of "weapons of last resort" makes sense. A policy of "no first use" of nuclear weapons, if believable, weakens deterrence of major conventional war and rests upon a false strategic premise.

Finally, the oft-cited characterization that our strategic forces are on "hair trigger" alert is a scare tactic routinely used to justify proposals to lessen the potential responsiveness of our strategic forces. In fact, multiple stringent procedural and technical safeguards are in place to guard against accidental or unauthorized launch and to ensure the highest levels of nuclear weapon safety, security, reliability, and command and control. Robust reconstitution capabilities are in place to survive sufficient forces, command and control systems, and national leadership to enable us to "ride out" an attack and not rely upon "launch on warning." In peacetime, our strategic forces are not even targeted against potential adversaries. The U.S. trigger is built so we can always wait.

Guiding Principles for Strategic Force Reduction

There are a number of fundamental principles that should guide further strategic force reductions.

Because we have neither new delivery platforms nor new warheads in development, we must not be hasty in taking irreversible steps to reduce our capabilities and flexibility.

First, we should continue to focus on arms control measures that directly and demonstrably enhance stability and reduce the risks of war. Stability—the lack of an incentive on either side to initiate major aggression or conflict, the assurance against being caught by surprise, the safety in waiting—rather than numerical parity is the most important criterion in assessing force structure and posture options. As Albert Wohlstetter wrote many years ago, "Relaxation of tensions, which everyone thinks is good, is not easily distinguished from relaxing one's guard, which everyone thinks is bad."[19] Deterrence ultimately depends not on our capability to strike first, but on the assurance, we always have the capability to strike second.

Stability rather than numerical parity is the most important criterion.

Second, we must preserve sufficient deterrent capabilities to respond to future challenges, to provide a cushion against imperfect intelligence and technological surprises, and to provide a reconstitution capability as a hedge against unwelcome geopolitical developments. As we reduce our nuclear forces to lower levels, numbers alone become less important. Attributes such as survivability, reliability, transparency, accountability, reconstitution, force asymmetries, production infrastructures, and verifiability become more and more important. It is ultimately the character and posture of our forces, as well as those of our allies and adversaries, more than just numbers, that makes the strategic environment stable or unstable. Preservation of our capability to adapt our deterrent forces to a rapidly changing and unpredictable strategic future is critical. Because we have neither new delivery platforms nor new warheads in development, we must not be hasty in taking irreversible steps to reduce our capabilities and flexibility.

Third, strategy must be the starting point—it should drive numbers rather than the reverse. A number of people have declared with unwarranted certitude that we can successfully reduce our operationally deployed forces to some lower number (for example, 500 or 1,000) without ever formulating or articulating what changes in national strategy, objectives, capabilities, force structure, and force posture would be required. Instead of threat-based or capability-based deterrence underpinned by rigorous analyses, war-gaming and risk assessment, they seem to be advocating a form of faith-based deterrence.

Strategy must be the starting point for rigorous analysis with a logic path akin to the following:

- Whom do we want to deter, and under what circumstances might we need to simultaneously deter more than one potential adversary?
- What do those potential adversaries hold that they value most?
- What kinds of capabilities do we need to hold what they value at risk under the most stressful of scenarios?
- What kinds of capabilities do we need to meet our extended deterrence commitments to our allies and friends?
- How do we hedge those capabilities against technological surprise and imperfect intelligence?

An Ohio class ballistic missile submarine. (Photo: U.S. Navy)

- What form of strategic reserve, supporting infrastructure, and reconstitution capabilities are required to maintain those capabilities?
- How do we posture those capabilities to promote stability—for example, to discourage any potential adversary from preemption, to avoid a "use them or lose them" situation, and to ensure we always have the capability to strike second?
- And finally, what numbers of various capabilities, based upon rigorous analyses, are required to hold at risk a sufficient amount of what our potential adversaries value without accepting undue risk ourselves, while providing the president the widest range of options using the minimum level of force intended to achieve our objectives?

Fourth, we need to view reduction as a means to an end—national security—and not as an end itself. Given the clear risks and elusive benefits inherent in additional deep reductions, those who advocate them bear the burden of proof to demonstrate exactly how and why such cuts would serve to enhance national security.

Summary

An early strategist's metaphor that nuclear planners are like homebuilders remains true today. A wise architect does not design only for benign environments, but for the worst weather conditions one can reasonably anticipate. We have to consistently maintain a 'building code' for our strategic forces to ensure they can weather the most stressing scenarios we can reasonably postulate.[20]

None of the foregoing discussion is intended to discourage reductions in our nuclear arsenal *that promote greater stability*, but it is essential to recognize that the journey is far more important than the destination, and that *the overriding goal is not reductions for disarmament's sake, but increased international stability and, most important, the avoidance of war*. We need to carefully manage the risks and uncertainties we face in this new strategic era. Our strategic enterprise, and particularly our force structure and doctrine, needs to be robust, flexible and credible. We must always maintain the ability to both reassure our allies and convince potential agressors to choose peace rather than war, restraint rather than escalation, and conflict termination rather than continuation.

~ Admiral Richard W. Mies, United States Navy (retired)

Footnotes

1. Frederick S. Dunn, et al., The Absolute Weapon (New York: Harcourt, Brace and Co, 1946), pp 21-107.

2. Ibid, p 76.

3. Michael Quinlan, *Thinking About Nuclear Weapons,* Whitehall Paper 41 (London: The Royal United Services Institute for Defence Studies, 1997), p 8.

4. Thomas C. Schelling, Arms and Influence (New Haven: Yale University Press, 1967), p 134.

5. See Schelling, ibid, for a fuller discussion.

6. Sun Tzu, The Art of War, translation by Samuel B. Griffith (Oxford University Press, 1963), p 77.

7. Numerous classified reports, including two Defense Science Board Reports separately chaired by Adm. Henry Chiles and Gen. Larry Welch, the Schlesinger Task Force Report, the End-to-End Review chaired by Lt. Gen. Brent Scowcroft, and the Nuclear Comprehensive Report that I chaired, serve as examples.

8. Albeit misleading, because strategic war-head numbers are tied to counting rules associated with delivery platforms due to practical limitations in our monitoring and verification capabilities.

9. Keith Payne, "Future of Deterrence: The Art of Defining How Much Is Enough," *Comparative Strategy: an International Journal,* Vol. 29, Issue 3, 2010, p 219. Some refer to this as the Prime Minister Healy paradox.

10. In every STRATCOM force structure analysis I've been involved with over the years, there were two general truths: 1) For the same force levels, a triad performs better than a dyad, and a dyad performs better than a monad. Diversity affords a hedge against single-point failures and significantly complicates a potential adversary's offensive and defensive planning considerations. 2) There is a tyranny in low platform numbers that greatly restricts the flexibility, survivability, and resiliency of the force. Fewer weapons in more delivery platforms fare far better than too many weapons in too few platforms.

11. While it is not U.S. policy to depend upon launch under attack, the ambiguity associated with the potential to launch under attack complicates any adversary's preemption calculations.

12. Robert G. Joseph, "Second to One," *National Review,* October 17, 2011.

13. For a complementary rationale, see Bruno Tertrais, "The Illogic of Zero," *The Washington Quarterly,* April 2010.

14. Michael Quinlan, *Thinking About Nuclear Weapons,* Whitehall Paper 41 (London: Royal United Services Institute, 1997), pp 5, 41.

15. Salvador de Madariaga, Morning without Noon (Farnborough, Hampshire, U.K.: Saxon House, 1974), p 48.

16. Thomas C. Schelling, "A World Without Nuclear Weapons?" *Daedalus,* September 2009, p 124-129.

17. Although these alleged statements are widely quoted, no definitive source for them has ever been cited.

18. For a more complete and compelling rationale, see James Schlesinger, "The Impact of Nuclear Weapons on History," *The Washington Quarterly,* Autumn 1993.

19. Albert Wohlstetter, "The Delicate Balance of Terror," originally published in *The New Republic,* September 1, 1958 (revised November 6, 1958), RAND Corporation publication P-1472.

20. Thomas C. Schelling, in discussion with the author.

Admiral Mies, when serving as STRATCOM's Commander. (Photo: DoD)

Interview with
Bob Webster
Associate Director for Weapons Physics

The new associate director discusses his viewpoints regarding the nature of the science and scientists at the Laboratory... and why he misses using his slide rule.

Recently, Bob Webster was named the Laboratory's Associate Director for Weapons Physics. He oversees the Computational Physics and Theoretical Design divisions, as well as the Laboratory's Advanced Simulation and Computing (ASC) program. The Laboratory's ASC capabilities are inextricably woven into the work of weapons physics and design. Webster recently spoke with *National Security Science* (*NSS*) about his new role at the Laboratory.

NSS: Where in the evolution of computing did you start your career?

Webster: I think I was in the last class at Case Western Reserve University that used slide rules in the exams. Slide rules were abandoned between my freshman first semester and the spring semester that year; the university *finally* let us use calculators. So initially many of us were still carrying "slip sticks" to engineering classes. That was a different era in terms of how we thought about solving physics problems because we didn't have computers the way we have them now.

NSS: So do you think anything was lost by leaving the slide rule behind?

Webster: I think there was some value in the way that we had to think about the problems when we were still using slide rules, a way that we could reintroduce into the system right now. At the same time, though, there is a tremendous opportunity presented by leaving slide rules behind.

If you look at the last 20 to 25 years, there's a fundamental shift in how supercomputing underwrites our evaluation of scientific problems. In the '70s, supercomputing, or high-performance computing, which wasn't very "high performance" by today's standards, was sometimes viewed as a crutch.

Today it's an integral part of synthesizing theories—we can evaluate very complex scenarios that we can't actually test. For economic, political, and risk factors, we can't always employ the classic, direct scientific experimentation that we were taught to do. I think that's something we need to get out to folks—supercomputing is integral. We can't separate it from doing the experiments and doing the analytic theories anymore.

But there was loss there with leaving the slide rule behind. We started to leave experiments behind more than we should have. Experiments got very expensive, so there's a tendency

to try to compute your way around a problem. If all you have is a slide rule, you must use experiments to inform how you think and reason and internalize the uncertainties and the possibility of error in your calculations.

It's different with a computer—you can run a simulation and get an answer with less consideration of the interplay of the different pieces of physics every step of the way along the solution path. That's a seductive feature of the computer that could bite you. When using a slide rule or a calculator, such considerations couldn't be ignored.

NSS: So one should miss the slide rule because it compelled scientists to approach the experimental process in some very useful and enlightening ways?

Webster: There was a feeling, where I went to school, that higher math was something that required people to learn very complex functions that you could use to represent the solution. The truth of it was, you could only solve problems that were under certain spotlights for those kinds of theoretical approaches. That's something that has changed; you used to have to recast the problem that you were trying to solve so that you could evaluate it with known analytic or semi-analytic solution techniques.

For example, we would frequently have to treat something as spherical when we knew it wasn't, just so we would have a solution technique available. With a computer, one doesn't have to do that. So, we used to get a more exact answer to a more approximate view of the problem, and that has now flipped with the computer. Today, you can actually go get the answer, an approximate answer, to a more exact posing of the problem.

NSS: Do you see a trend away from experimentation because of the economics of it?

Webster: I think there are several factors. It's the economics, but also there's the perceived risk. People can get hurt when they do experiments. We're afraid of that.

If we think about the Stockpile Stewardship Program over the next 15 years, without experiments, how are we going to develop the trust we need to have in scientists?

NSS: Particularly experiments in the Weapons Program?

Webster: It's particularly true in the Weapons Program, but not just the Weapons Program. It's also true in the Energy Program. The Lab used to have a magnetic fusion research division, and we did experiments that were in some cases, by today's standards, perhaps dangerous experiments. Today, the country, and people in general, seem to be less willing to take those kinds of risks. In some cases, that response is justified.

Bob Webster (right) being interviewed by NSS *staff in the lobby of the Laboratory's Strategic Computing Center (SCC). The SCC is home to the Laboratory's suite of supercomputers used to solve problems in national security.* *(Photos: Los Alamos)*

For example, we are less willing to do experiments that present risks to the environment. As we've moved that way, a scientist's ability to develop judgment based upon experiments has been diminished. If we think about the Stockpile Stewardship Program over the next 15 years, without experiments, how are we going to develop the trust we need to have in scientists? If we don't expose them to situations where they have to make decisions, and then watch how they react—both when they get the answer right and when they get the answer wrong—how do we evaluate them? How do we know that they're stewarding things well? We need to give scientists the opportunity to pose and solve the types of problems relevant to stewardship.

NSS: It's a conflict. Science is based on taking risks—that's how it moves forward—but at the same time, there's a countervailing weight to be conservative. As an associate director, how do you help manage the two extremes?

Webster: I'm struggling with that a little bit right now. To really have a balance—people developing the self-confidence to make decisions, take the reasonable risks, and develop wisdom but not take imprudent risks—requires a lot more interaction and a lot more opportunity to design something, to create and execute an experiment. It's important to build something, see how it works, and be humbled if *and when* you get it wrong because we're all going to get it wrong sometimes.

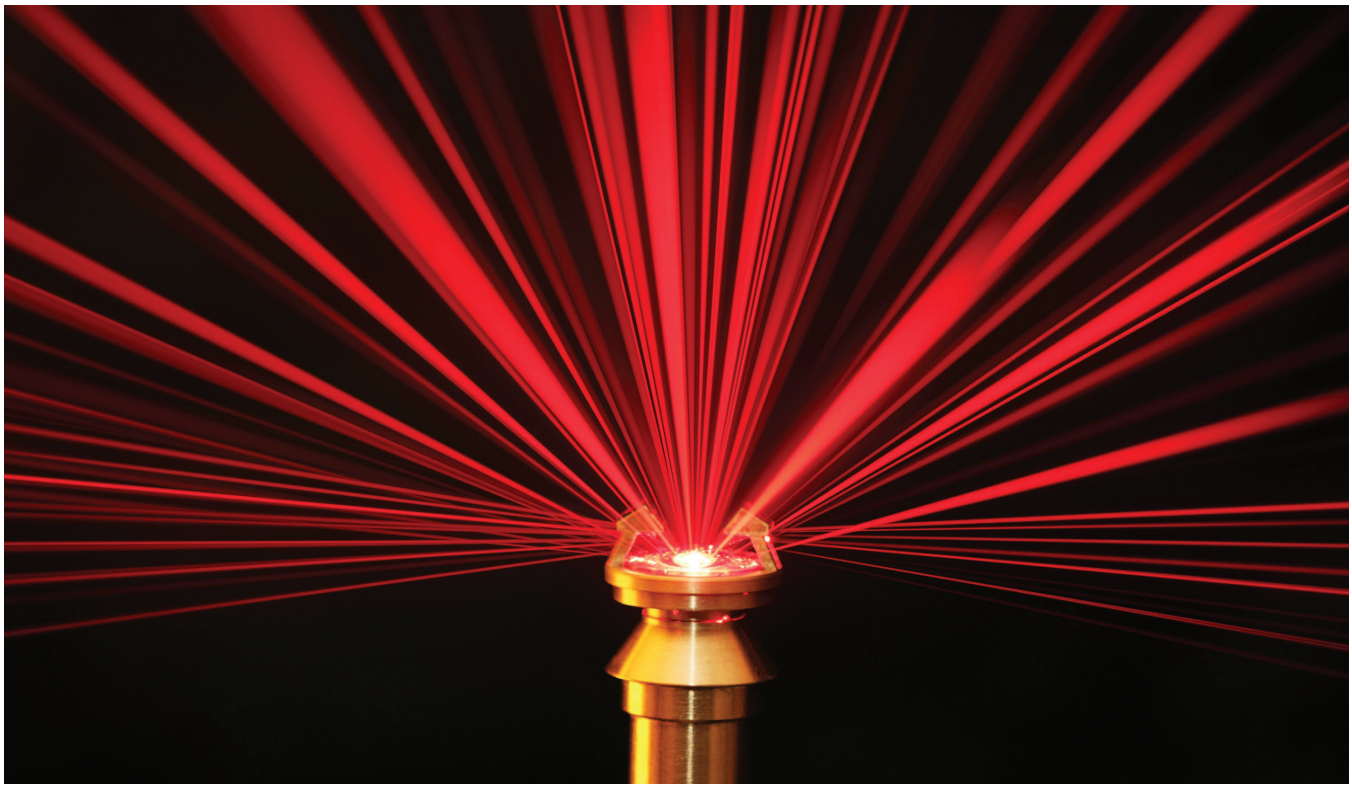

In December 2012, Los Alamos and its partners National Security Technologies and Sandia National Laboratories successfully executed a subcritical plutonium experiment in the Gemini series at the underground U1a facility of the Nevada National Security Site. In this experiment, a novel optical diagnostic (shown here) measured the motion of a plutonium surface along more than 100 rays, providing orders-of-magnitude more data than similar past experiments. These data will challenge models of plutonium behavior, ultimately increasing our confidence in the computer simulations that ensure the safety and effectiveness of the nation's nuclear stockpile. (Photo: National Security Technologies)

So I am searching right now within the directorate for different experimental, scientific, and program areas to assign to people to spur them toward these opportunities. I'm looking for an increase in how our folks perform what I think of as "cold physics experiments": moving metal with high explosives. How do I get an active group of people working in that area, whether it's for Stockpile Stewardship or for global security concerns?

NSS: Where do you see your directorate going in the next 5 to 10 years?

Webster: We've got a couple of real challenges that are going to come up within the Stockpile Stewardship Program. There are some changes that are likely to be required for the stockpile that will require active decisions from the Laboratory, including from our directorate. In a number of cases, they won't be easy decisions to make. We want to be certain that if we make a change to a system, the system is going to function.

When we're touching systems on intervals of 30 years or so, we need to remember that industry changes a lot in 30 years.

Clearly, we'll need to increase the amount of computing that we're using. The solution techniques that we have available to us today pretty much demand that we increase the computing that we've got. We'll need that computing power so we can take on the problems we anticipate and be really confident that we have the right answers.

I also see that we're going to have to revitalize a number of components in the experimental programs, partly to get the data to answer the questions we need to answer. Equally important, and maybe even more important, we need to give future stewards the opportunity to experiment, to test themselves against nature, to demonstrate that they're actually capable of predicting what nature will do. That's fundamentally what we are going to be doing.

NSS: Are these challenges being brought on by the age of the stockpile or by changes to components in the stockpile? What's driving these challenges?

Webster: Aging is certainly a concern. Part of aging is the evolution of the manufacturing environment in which we do our work. It's not clear that we can simply rebuild what we have because what we have now was built with a certain set of manufacturing processes. And those may or may not be the processes that we have available today if we need to rebuild a system.

When we're touching systems on intervals of 30 years or so, we need to remember that industry changes a lot in 30 years. The processes we have available to us change. Processes that stockpile stewards will be using to maintain the systems we deploy over the next 5 to 10 years are going to change, too. We have to respond to those realities. We need to have a notion of how we're going to qualify the materials and the parts that are manufactured by the available processes today, so we have a way to monitor them tomorrow through a surveillance program. We also need to know we'll have the skills and capabilities on hand to deal with any issues we find during surveillance.

We challenged our designers to go through the whole process: fielding an explosive experiment and making a prediction—which is an important step when we're developing scientists with good judgment.

NSS: Would you comment on the Laboratory's recent success with the Gemini experimental series in Nevada, which included a subcritical experimental shot?

Webster: That's a good example of evaluating judgment, which goes beyond the good science. That experiment—that shot—was seen as being a success because we got some data that challenges our thoughts. In fact, we got more data out of that than we ever got from shots of that class in the past. It's remarkable the huge steps that the Gemini experiment represents and the data that came out of it.

But equal to that, and perhaps more important, we benefitted because we challenged our designers to go through the whole process: fielding an explosive experiment and making a prediction—at the risk of looking silly, which is an important step when we're developing scientists with good judgment.

Gemini also tested the judgment of people who assembled the shot. We exercised the entire spectrum of taking plutonium, manufacturing it, shipping it, assembling it, and measuring data from it.

So, we got far more than just science data. We improved the judgment skills of everyone involved—the scientists, the engineers, and the technicians.

NSS: From your perspective, what sort of global security questions might be addressed with supercomputing?

Webster: The global security twist on things throws a wrinkle into the use of high-performance computing in general. In the weapons world, we have fairly exquisite knowledge about what our weapon systems look like and how to construct very high-fidelity computer models of them to understand how they perform.

At Los Alamos, simulations such as those run on Roadrunner and currently on Cielo can be seen as full-color 3D visualizations displayed on a screen or in a viewer-surrounding environment. Visualizations allow scientists to watch dynamic phenomena as they evolve. (Photo: Los Alamos)

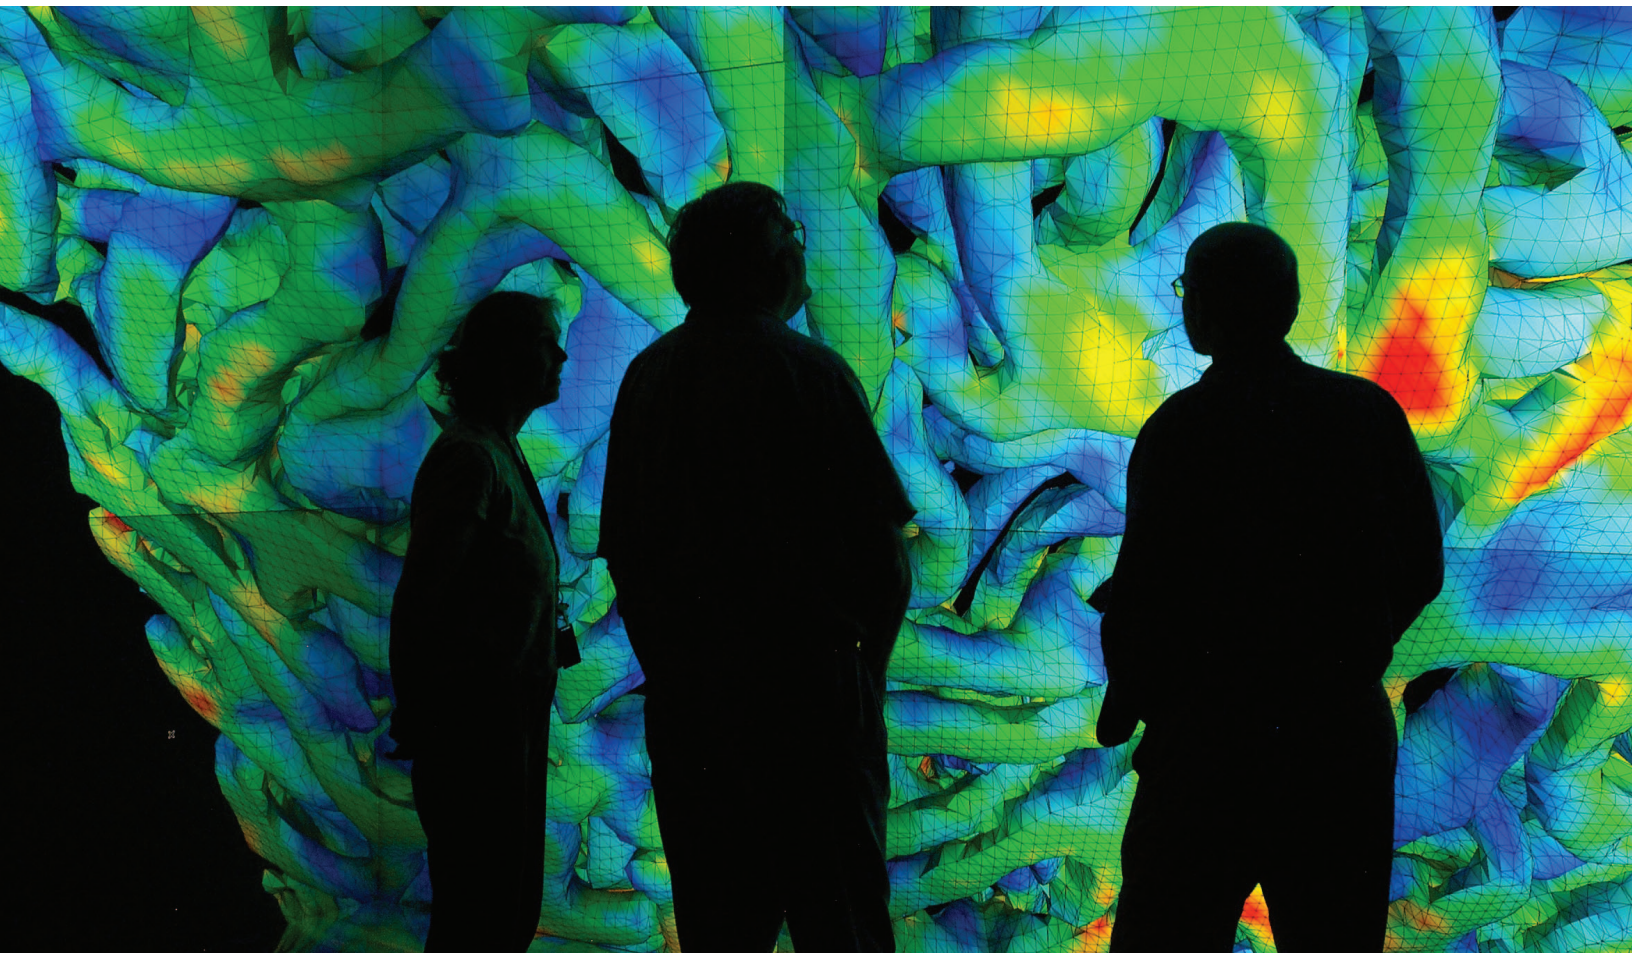

In global security, we don't always have that level of detail because we're imagining what *somebody*—a nation or terrorist group, for example—could be doing. In that sense, the high-fidelity driver for predicting their behavior is not so strong. But we still use high-performance computing, for example, if we're trying to predict how a group might use a particular material, given the material's properties, for terrorist activities. So there are aspects of high-performance computing for global security that play a fundamental role, and that could play a *more* fundamental role, in predicting global security threats.

NSS: What kinds of people do you think the Laboratory should be attracting?

Webster: The Laboratory is not sustainable if we don't attract really creative, talented people who know when to take a risk and when to be conservative. When we make stockpile decisions, we need to be conservative, but when we're searching for new solutions to other kinds of problems, we need to be willing to take risks. A person who can balance conservatism with risk is a complicated person to try to find.

If you really love science, the Lab is a place where you can come and retool yourself. You can be curious. You can go find an expert and learn and have impact on a world-changing scale. This is a scientist's "place to be."

NSS: How do you convince people to come work for the Laboratory?

Webster: There are very few places in the country or in the world where, if you have a passion for science, you can apply that science to informing, for example, national policy. Weapons research is one such area. Climate research is another.

If you really love science, the Lab is a place where you can come and retool yourself. You can be curious. You can go find an expert and learn and have impact on a world-changing scale as you apply your science. You might not make as much money as you would if you were using those same skills as a quant [a researcher on Wall Street]. They use some of the same scientific techniques to predict financial markets, and you might make a lot of money at it. But if you love basic science and want to have a career as a scientist, the Lab is a place where you can impact the world in a positive way. This is a scientist's "place to be."

NSS: Where has your career at the Lab taken you?

Webster: I started in fusion, and now I'm in weapons. I went from fusion to submarine detection to radar to lasers (doing work related to strategic defense) to oil and gas. At the Lab, you can always find someone who will help mentor and take you through that next transition. It's like being in a library, where you can go to the shelf and find any information you want, but here you pick out a person instead of a book.

When you go around the Laboratory you find that we have enormous bench strength in scientific capability. You can find somebody who can work on or who knows something about almost every different problem. That's an amazing thing.

NSS: Recently, there have been some concerns about the Lab's intellectual integrity because of the current "for profit" business model. As an institution are we compromising our intellectual scientific integrity to a business model?

Webster: That one is actually easy to answer: No. The Laboratory is full of highly educated people who spent years in colleges, which are often fairly liberal places. At those institutions, debate was valued—the open exchange of ideas. We hire people who are selected from that background and training. We have so many scientists here that it would be virtually impossible to compromise the Laboratory's intellectual values. The intellectual, scientific culture here wouldn't allow it.

For a manager, the Los Alamos culture can sometimes be frustrating; managing scientists can be a lot like herding cats. But the upside is that because we're herding cats, it's almost impossible to compromise our intellectual integrity. The staff will speak up; they're not afraid to speak up. That's a huge power here.

Yes, folks will throw arrows at the management. I've been here since 1984, when I first came here as a student. As a scientist, I grew up here. I didn't stay here to get rich, and my value system didn't suddenly change when the contract model changed. That's equally true for all of my colleagues at the Lab. They're all doing this because they believe they're making a difference, with science, for society.

You can find somebody who can work on or who knows something about almost every different problem.

It's frustrating for people like me in management right now, people who have put our lives into doing this. We are getting outside pokes from people who know nothing about the Laboratory but still say, "Oh, you're just doing it for a profit."

That stings. Sure, we're occasionally going to be accused of compromising our intellectual integrity, but that just doesn't happen, not as an institution. Not at Los Alamos.

REFLECTIONS ~ FAREWELL TO TOM D'AGOSTINO

As I look back at Tom D'Agostino's leadership of the National Nuclear Security Administration (NNSA)—and his unique relationship with Los Alamos—there is certainly no shortage of contributions.

His career has spanned most of the history of the Stockpile Stewardship Program. He led the NNSA during a period of transition for the nuclear weapons complex— guiding us with a steady hand through challenges and changes at Los Alamos and Livermore national laboratories. The Greek philosopher Heraclitus noted, "Nothing endures but change." But change needs good leadership if it is going to become positive change.

The laboratories are thankful that Tom has been at NNSA to help guide us through these years of change.

It was on his watch that Los Alamos achieved success testing with the second axis of our Dual Axis Radiographic Hydrodynamic Testing (DARHT) facility. DARHT is a flagship facility at Los Alamos for doing research for the Stockpile Stewardship Program.

I could talk about his unique ability to make budget work—and his effectiveness working with Congress. He provided deep budget insights during his time as NNSA's administrator.

I know these things from working with Tom for the past decade and a half.

I'd like to share a story our readers might not have heard about—one that dramatically illustrates Tom's integrity, his deep sense of service to the nation, and his straightforward way of "taking care of business."

On June 26, 2011—only 26 days after I became director of Los Alamos National Laboratory—a wildfire started burning in the remote areas west of Los Alamos. That evening and night, this fire grew into a monster: from a few thousand acres to more than 46,000 acres! It eventually grew to more than 156,000 acres. The Los Conchas wildfire became, at that moment, the largest wildfire in New Mexico history.

The Laboratory and surrounding communities faced a grave danger. That's when I picked up the phone and I called Tom. I will remember that phone call forever. Within minutes of that phone call, the Los Conchas wildfire moved to the top of the national priority list.

Federal assets quickly descended on Los Alamos, including two of the most elite firefighting units in the nation. Thanks in no small measure to that quick response, a disaster was averted and a 70-year-old national treasure was spared. The men and women of Los Alamos owe Tom a deep debt of gratitude.

And so does the nation. He is a patriot, with a profound sense of duty to country and mission—and a deep faith.

On behalf of the Laboratory, I presented Tom a plaque displaying a completely used-up target wheel from the second axis of DARHT. This target wheel is one of the unique technologies that allows us to get *four* radiographic images: a stunning achievement, one that's needed to continue to improve our program in Stockpile Stewardship. We are now "addicted" to getting four images from the second axis.

We will miss Tom, and we wish him well. Congratulations, Tom. Fair winds and following seas!

~ Charlie McMillan, Laboratory Director

National Security Science
Mail Stop D416
Los Alamos National Laboratory
Los Alamos, NM 87545
Email: NSSPub@lanl.gov
Tel: 505-667-7814
www.lanl.gov/science/NSS

Presorted Standard
U.S. Postage Paid
Albuquerque, NM
Permit No. 532

The Laboratory's first administration building, now demolished, as it looked 50 years ago, is shown in this black and white photograph. The new administration building, the National Security Sciences Building, is pictured to the upper right.

April 2013 • LALP-13-005

National Security Science highlights work in the weapons and other national security programs at Los Alamos National Laboratory. Current and archived issues of *NSS* are available at www.lanl.gov/science/NSS/. *NSS* is unclassified and funded by the Weapons Program Directorate.

Los Alamos NATIONAL LABORATORY

70 YEARS OF CREATING TOMORROW

Managing Editor • Clay Dillingham

Writers/Editors • Necia Grant Cooper, Lisa Inkret, Eileen Patterson,

Designers/Illustrators • Barbara Maes, Kelly Parker, Leslie Sandoval

Art Director • Donald Montoya

Photographer • Ethan Frogget

Editorial Advisor • Jonathan Ventura

Technical Advisor • Sieg Shalles

Laboratory Historian • Alan Carr

Printing Coordinator • David Van Etten

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-ACC52-06NA25396. A U.S. Department of Energy Laboratory.