

Experiment and Theory

Supercomputer simulations have become essential to making discoveries and advancing science.

EVEN before Lawrence Livermore opened in September 1952, cofounders E. O. Lawrence and Edward Teller recognized the need for a computer and placed an order for one of the first production Univacs. Equipped with 5,600 vacuum tubes, the Univac had impressive calculational power for its time, although much less than that contained in today's \$5 calculator. Computing machines quickly demonstrated to the Livermore staff the ability not only to perform complicated calculations but also to simulate physical processes.

A computer's predictive capabilities were vividly demonstrated in 1957, when the Laboratory received an urgent call from the Pentagon. Livermore had the only U.S. computer able to compute the orbit of Russia's Sputnik I. Researchers accurately predicted the satellite's plunge into the atmosphere in early December of that year. In time, they showed how computer simulations could lend insight into a broad range of physical problems.

"No institution in the world has more consistently invested in new generations of supercomputers than Livermore,"

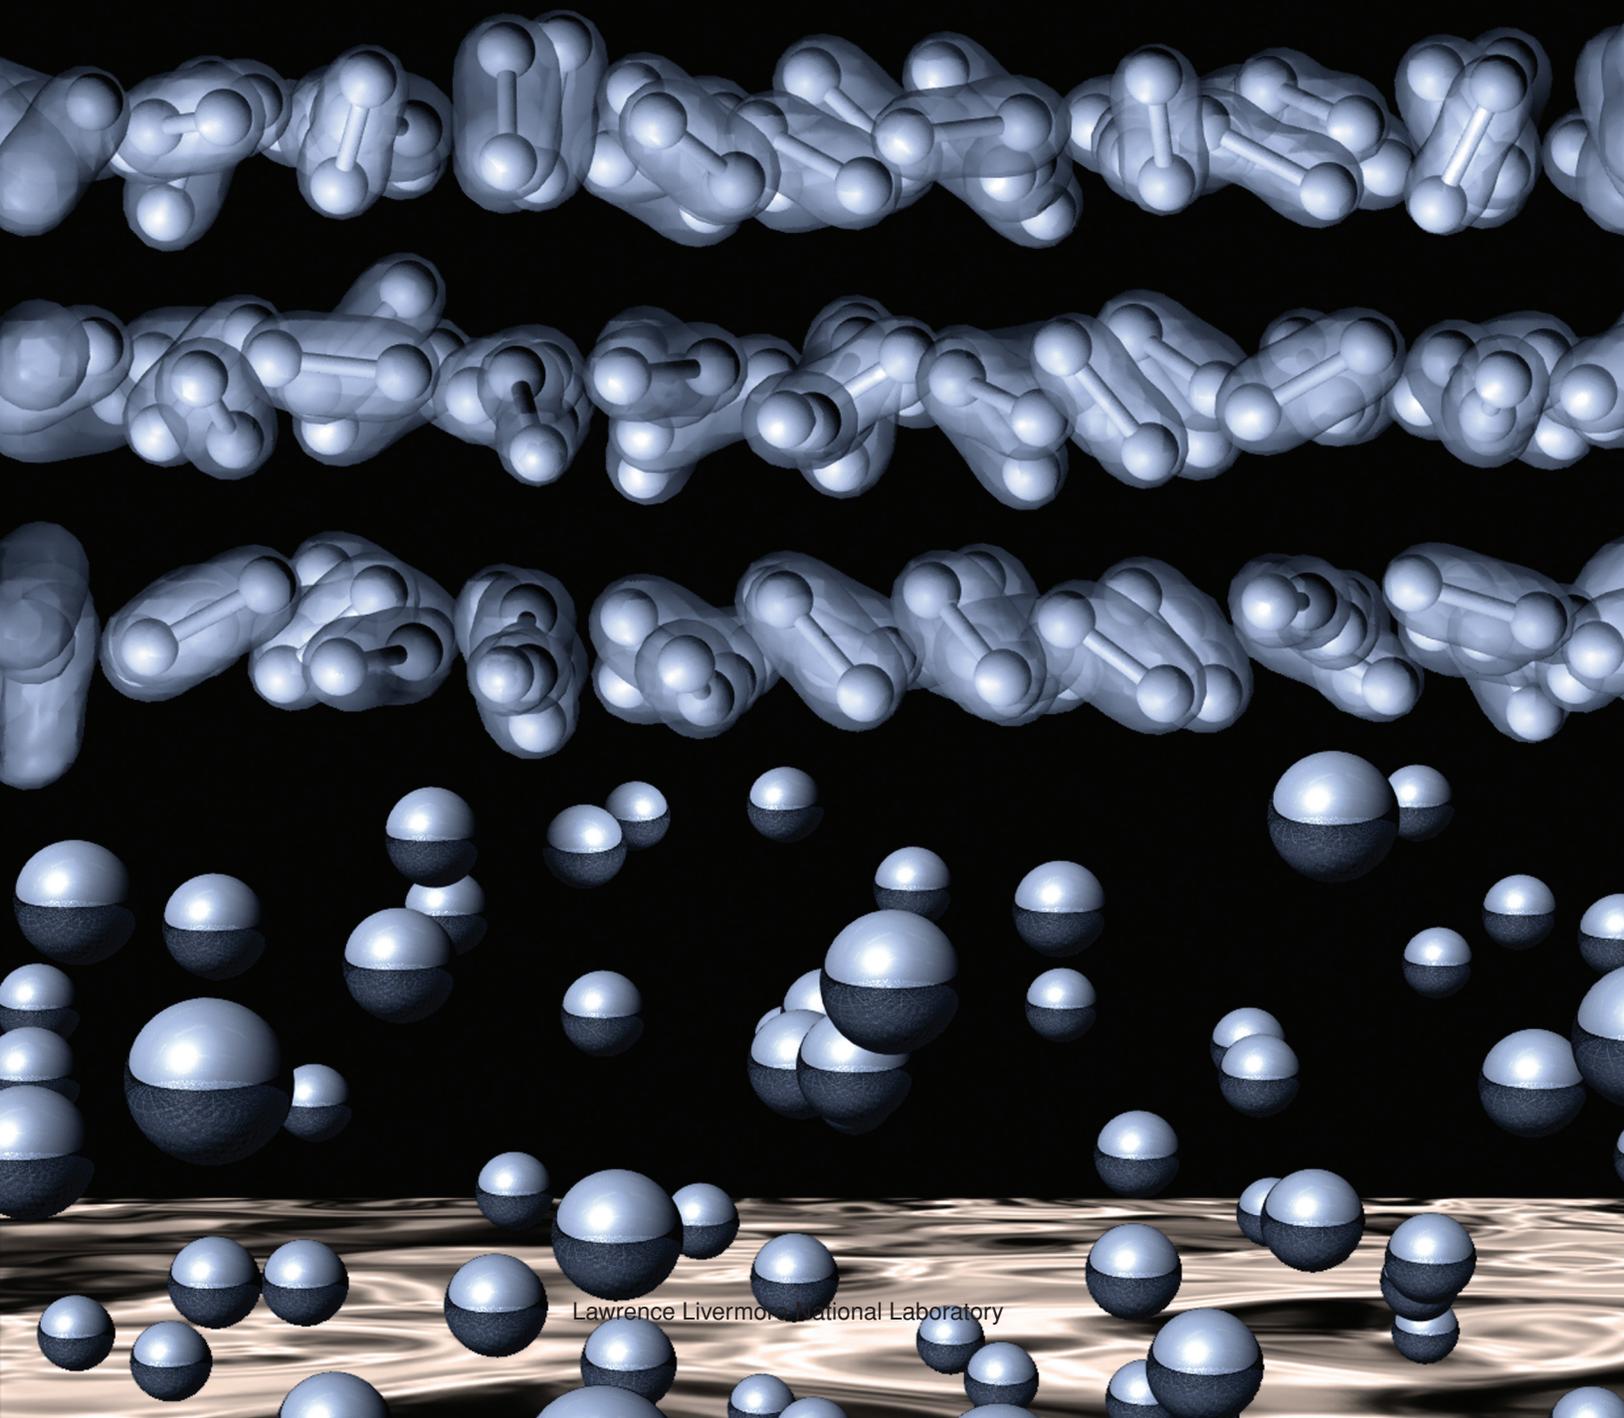
says physicist Tomás Díaz de la Rubia, associate director of Chemistry and Materials Science. During the past five decades, supercomputers have advanced every discipline and have helped to attract some of the brightest minds to the Laboratory. Díaz de la Rubia, for example, was drawn to Livermore in 1989 because of the opportunity to work on the world's most powerful computer (at the time, made by Cray Computer) and with some of the nation's top simulation experts.

Researchers haven't been shy about tapping the increased power of a new machine. "Every new machine has brought new insight," says physicist Francois Gygi of Livermore's Center for Applied Scientific Computing.

"Simulation has changed the way science is done at Livermore," says computer scientist Mark Seager, head of Platforms in the Integrated Computing and Communications (ICC) Department, part of the Computation Directorate. "Today, experiment and simulation are more tightly coupled than ever. At the same time, theory and computation are more tightly coupled than ever."

Laboratory scientists have predicted a new melt curve of hydrogen, resulting in the possible existence of a novel superfluid—a quantum fluid at about 400 gigapascals. This figure illustrates the transition of hydrogen from a molecular solid (top) to a quantum liquid (bottom), which was simulated using Livermore's GP ab initio molecular dynamics code. A "metallic sea" is shown in the background.

Have a New Partner: Simulation



Simulations mimic the physical world down to the interactions of individual atoms. These simulations, conducted on some of the world's most powerful supercomputers, test theories, reveal new physics, guide the setup of new experiments, and help scientists understand past experiments. Many times, the simulations conduct electronic "experiments," replicating scaled models of experiments that would be too difficult or expensive to perform or would raise environmental or safety issues.

ASC Leads the Way

For the past decade, the driving force behind increasingly realistic simulations has been stockpile stewardship, which is the Department of Energy's (DOE's) National Nuclear Security Administration (NNSA) program to ensure the safety

and reliability of the nation's weapons stockpile. A major element of stockpile stewardship is the Advanced Simulation and Computing (ASC) Program, which had an initial 10-year goal to obtain machines that could run simulations at 100 trillion operations per second (teraops). To meet this requirement, ASC spearheaded a transition during the mid-1990s to scalable parallel supercomputers composed of thousands of microprocessors that solve a problem by dividing it into many parts.

Since 1996, proprietary scalable parallel supercomputers running vendor-supplied system software have been used to simulate the physics of nuclear and chemical reactions. ASC's Purple machine will arrive at Livermore in July 2005. Purple will fulfill the goal set in 1996 to achieve 100 teraops by mid-decade for prototype full-system stockpile stewardship simulations.

In addition to ASC scientists, researchers working in almost every other program at Livermore need to run unclassified simulations on ASC-class supercomputers. The Multiprogrammatic and Institutional Computing (M&IC) Initiative has made this class of platform

available to a wide spectrum of scientific investigators at Livermore since the mid-1990s. Most recently, the M&IC platforms have been deployed using Linux cluster technology. (See *S&TR*, June 2003, pp. 4–13.)

M&IC is so named because it serves both mission programs (multiprogrammatic) and individual (institutional) researchers. A mission program can purchase a block of time on existing machines and share in the investment in new equipment. In addition, M&IC grants Livermore scientists engaged in leading-edge research access to computer time, independent of the program to which they belong. Researchers who are funded by Livermore's Laboratory Directed Research and Development (LDRD) Program, an institutional sponsor of individual researchers pushing the state of science in diverse fields, have significant access to M&IC machines.

Seager recalls, "We mounted an effort in the late 1990s to bring large-scale supercomputing to non-ASC scientists by leveraging everything we learned from using ASC machines and developing codes for them." Over the years, he says, the NNSA-funded ASC Program



The White supercomputer is a current "workhorse" of the National Nuclear Security Administration's Advanced Simulation and Computing (ASC) Program, performing 12.3 billion operations per second (nearly 31 billion times faster than Livermore's first supercomputer). Another ASC supercomputer, Purple, is scheduled for demonstration in June 2005 and delivery to Livermore in July 2005. Purple will have a peak performance of 100 trillion operations per second (teraops) and, as with White and the other ASC machines, will be dedicated to research for the nation's nuclear stockpile. (inset) Livermore's first supercomputer, the Remington-Rand Univac-1, was delivered in 1953. It had over 5,600 vacuum tubes and a memory that could store 9 kilobytes of data—a fraction of what today's handheld devices can hold.

and the institutionally funded M&IC Initiative have cooperated and leveraged their cumulative expertise. For example, unclassified simulations for stockpile stewardship-related projects run on M&IC machines.

Bruce Goodwin, associate director for Defense and Nuclear Technologies, notes, "Livermore has pioneered the development of a cost-effective, terascale Linux cluster technology that provides the high-performance computing environment required by the weapons program." The unclassified computers, Goodwin points out, are an essential part of Livermore's strategy to provide computing for both weapons science and weapons simulation. "As Linux cluster technology continues to advance, we expect it to help shoulder our most demanding requirements as well as more routine uses."

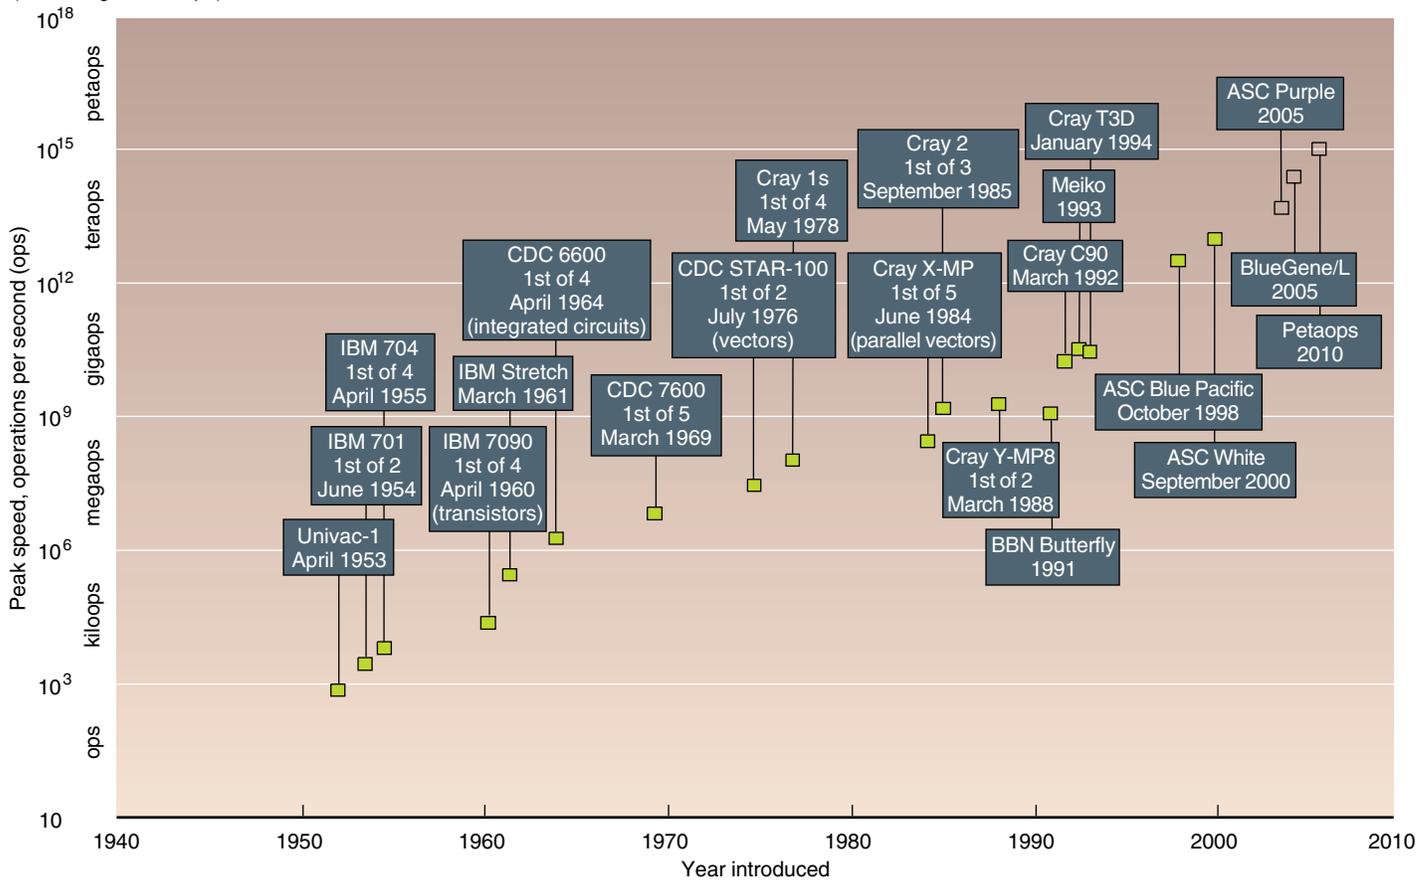
When acquiring new machines for unclassified research, the developers of the M&IC Initiative took a different approach from the ASC Program beginning in 2000. At the time, in what appeared to be a bold gamble to acquire and run supercomputers

at much less cost, Livermore assembled Linux clusters, composed of commercial microprocessors running on the open-source Linux operating system. The results were so impressive that other institutions, from universities to



Thunder, the latest of the Multiprogrammatic and Institutional Computing (M&IC) Initiative's Linux clusters, is composed of commercial microprocessors running on the open-source Linux operating system and cluster software. The 23-teraops Thunder machine runs simulations for scientists from a broad range of scientific disciplines.

(Next range is exaops)



A timeline shows Livermore's key supercomputers and their peak computing power.

corporations, followed Livermore's lead. Today, Linux clusters make up more than one-half of the nation's top-performing supercomputers.

Livermore's Linux clusters range from small platforms, such as the Intel Linux Cluster and Compaq GPS Cluster, that run one-dimensional (1D) codes useful for initial research, to the much more powerful 11-teraops Multiprogrammatic Capability Resource (MCR) and 23-teraops Thunder machines that run 3D codes for ASC-class simulations. MCR and Thunder will be joined later this year by ASC's computational science research machine: BlueGene/L (at 360 teraops).

BlueGene/L is under construction by IBM and recently took over the title of the most powerful supercomputer in the world from the Japanese Earth Simulator. When fully assembled at Livermore in June 2005, the 131,072 microprocessors of BlueGene/L will drive the next generation of simulation codes to advance stockpile stewardship on a path toward petascale (a quadrillion operations per second) computing.

The simulations on MCR and Thunder benefit many scientific disciplines, such

as laser physics, materials science, computational biology, computational physics, and astrophysics. "The breadth and scope of applications are amazing," says Seager.

Making New Science Possible

"The real story is not about the machines but about advancing science," says Brian Carnes, who leads the Services and Development Division in ICC. "Simulations on both ASC and M&IC machines are showing insights and proving new things."

For example, geologist Lew Glenn is using the MCR machine to do fundamental research on damage to large underground structures subjected to shock waves or explosives. "We have developed codes that analyze the behavior and simulate the response of the structures. These codes require the scalable parallel-computing platforms of M&IC," says Glenn.

Large seismic modeling efforts at Livermore include earthquake hazards, oil exploration, nuclear nonproliferation, underground-structure detection, and nuclear test readiness. Geophysicist Shawn Larsen says, "Without M&IC, a significant number of seismic modeling efforts in multiple

directorates would not have been initiated and conducted." In addition M&IC has aided in collaborative research at external institutions. For example, he notes, several University of California graduate students, postdoctoral researchers, and faculty owe much of their research to the computers' availability.

Researchers in Livermore's Biology and Biotechnology Research Program Directorate are using supercomputers to solve biochemical problems related to human health and national security. Projects include designing anticancer drugs, developing detection systems for protein toxins, and investigating the mechanisms of DNA repair and replication. The simulation work includes molecular dynamics software that mimics how individual atoms interact. "We use M&IC computers for simulations that will not fit on our own modest workstations," says biomedical researcher Mike Colvin. "These computers have been crucial to our scientific progress, which has led to dozens of publications in peer-reviewed literature, invited talks, and external collaborations."

Livermore simulations have extraordinarily broad time scales.

Sixteen racks of BlueGene/L (one-quarter of the system) are shown here being assembled at IBM in Rochester, Minnesota. The racks consist of 16,000 nodes with 32,000 microprocessors and have achieved a peak computing speed of 90 teraops. At press time, the racks had been moved to Livermore and were being reassembled in the Laboratory's new Terascale Simulation Facility. BlueGene/L is expected to be fully assembled in June 2005.



For example, physicists often probe the intricacies of nuclear detonations nanosecond by nanosecond. At the other extreme, geologists monitor the slow changes in nuclear waste repositories over centuries. Geologist Bill Glassley says, "M&IC has provided the computational horsepower that enabled us to tackle otherwise intractable simulation." Sponsored by the LDRD Program, Glassley's group conducted the world's only 3D simulations of how a nuclear waste repository would evolve over thousands of years. The group also conducted the world's first thorough simulations of the long-term response of soil water to climate change.

Gygi's simulations model atoms and molecules accurately by using the laws of physics and quantum mechanics. "With these simulations, we can address questions that are difficult to answer even with advanced experiments," he says. In a simulation funded by LDRD, Gygi followed the propagation of a shock front in liquid deuterium. By learning that the propagation of the front is related to the front's electronic excitation, scientists were able to better plan future experiments and understand past experiments. "This was the first time we could describe a shock in a molecular liquid in such detail," says Gygi.

Physicist Giulia Galli says that quantum simulations are playing an increasingly important role in understanding matter at the nanoscale and in predicting the novel and complex properties of nanomaterials. In the next few years, Livermore researchers expect these simulations to acquire a central role in nanoscience and allow them to simulate a variety of alternative nanostructures with specific, targeted properties. In turn, this work will open the possibility of designing optimized materials entirely from first principles.

"Although the full accomplishment of this modeling revolution will be years in the making, its unprecedented benefits are already becoming clear," says Galli. Indeed, simulations based on quantum mechanics are providing key

contributions to the understanding of a rapidly growing body of measurements at the nanoscale. "Quantum simulations provide simultaneous access to numerous physical properties such as structural, electronic, and vibrational, and they allow one to investigate properties that are not yet accessible for experiments," she says. A notable example is represented by microscopic models of the structure of surfaces at the nanoscale, which cannot yet be characterized experimentally with today's imaging techniques. The characterization of nanoscale surfaces and interfaces is important to predicting the function of nanomaterials and eventually their assembly into macroscopic solids.

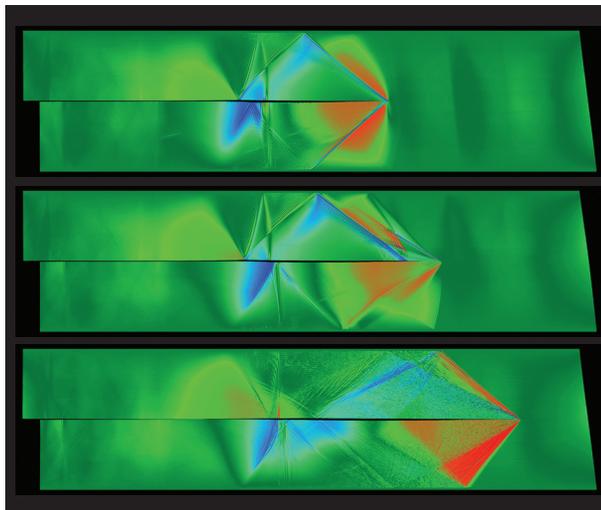
Planning for NIF

Simulations by physicists Bert Still and Steve Langer are playing an essential role in carrying out the first experiments at the National Ignition Facility (NIF) at Livermore. According to Still, "Simulating 1 cubic millimeter of plasma may not seem like a big task, but it requires a supercomputer to track 6.8-billion zones nanosecond by nanosecond." Still and Langer used 3,840 microprocessors of Thunder (94 percent of the machine) to model laser-plasma interactions in the first 4.2 millimeters of a 7-millimeter-long "gas pipe" experiment on NIF. Carbon dioxide gas was ionized in a plastic bag by laser

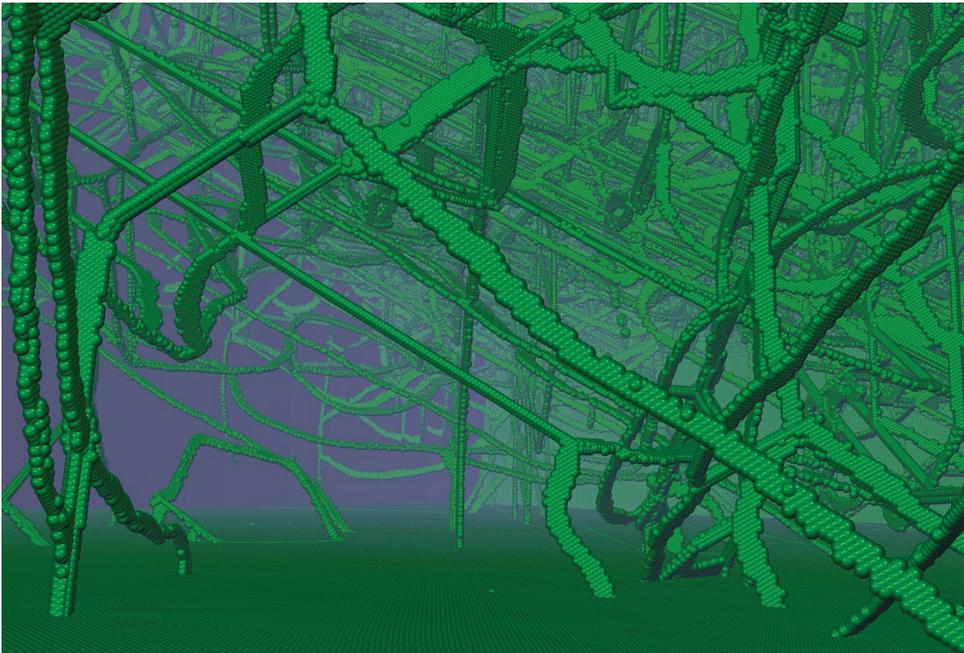
beams, thereby creating plasma similar to that which will be formed by targets in future NIF experiments.

The Thunder simulation unexpectedly showed a delay in the time taken for the laser beam to burn through the plasma at the far end of the gas pipe, when compared to predictions from previous, lower resolution design calculations. One possible explanation for this phenomenon is laser-plasma instabilities. The burn-through-delay issue did not arise in a simulation conducted in August 2003 on MCR. This simulation used 1,600 processors (69 percent of the machine) to model 2.7 millimeters of a similar experiment that used carbon dioxide. A previous simulation, conducted in February 2003, used 1,920 processors of MCR (83 percent of the machine) to model the first 4.5 millimeters of a gas pipe experiment involving neopentane. This simulation turned up another surprise: Scattered light from the neopentane plasma showed strong variability at subpicosecond time scales.

"MCR- and Thunder-class systems make it possible to simulate effects like these for the first time. Earlier systems could not model a NIF-scale plasma," says Still. "These scalable clusters are indispensable. Few platforms exist that can run such simulations, and many of them are located at Livermore, either



A Livermore multimillion-atom simulation to study crack propagation in rapid brittle fracture was performed on the 12.3-teraops ASC White supercomputer to help answer the question, Can crack propagation break the sound barrier? The simulations showed that crack behavior is dominated by local wave speeds, which can be faster than the conventional sound speeds of a solid. The snapshot pictures represent a progression in time (from top to bottom) of a crack traveling in a harmonic wave.



In another study of materials failure, scientists using ASC White examined what happens in ductile failure; that is, what happens when tough materials like metals bend and fail. A close-up snapshot from a simulation of more than one billion atoms shows the true complexity of propagating dislocations and rigid junctions in a crystal structure subjected to ductile failure and work hardening.

as ASC or M&IC machines. Such large calculations help us identify and assess plasma physics issues relevant to NIF experiments and dramatically contribute to achieving ignition.”

Still points out that advanced simulations have 1,000 times more resolution than the detectors that will be used on NIF experiments. As a result, Still and Langer have advised NIF experimenters where best to locate the detectors.

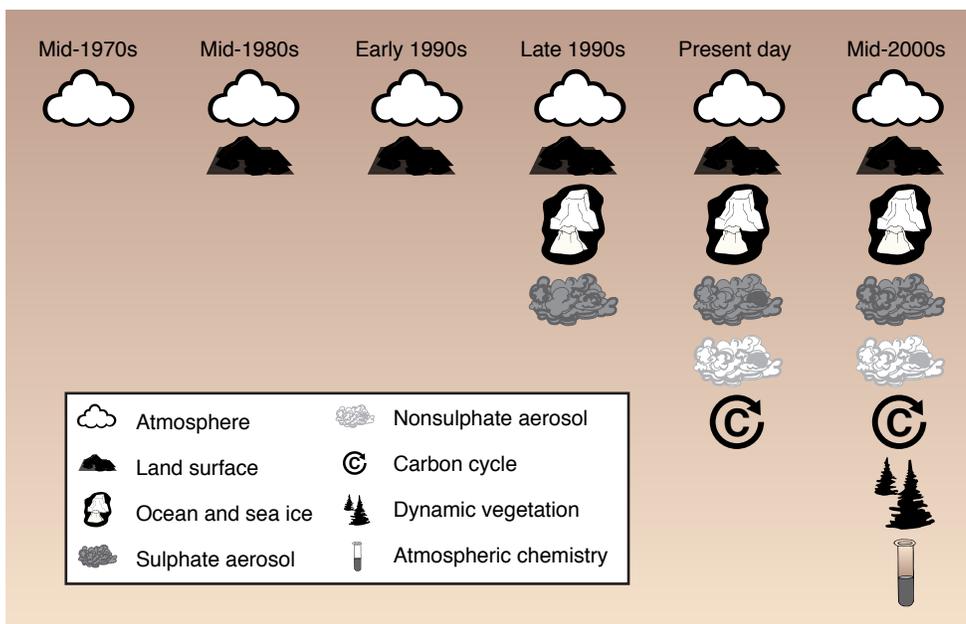
Cracks Break the Sound Barrier

Two landmark simulations, conducted in late 2000, demonstrated the power of advanced simulations to the world’s materials science community. The simulations were performed by Livermore physicist Farid Abraham (at the time working for IBM), computer scientist and visualization expert Mark Duchaineau, Díaz de la Rubia, and Seager.

The simulations, completed on the then newly installed ASC White computer, provided insights into how materials fail. The simulations used molecular dynamics to predict the motion of large numbers of atoms based solely on interatomic forces. “The simulations showed how it is possible to use molecular dynamics to design and perform mechanical tests that complement laboratory experiments,” says Abraham.

The first simulation was a 20-million-atom study of crack propagation in the fracture of brittle materials. The researchers unexpectedly observed the birth of a crack traveling faster than the speed of sound. “Theory stated cracks could move only up to the speed of sound,” says Abraham. He notes that researchers can’t possibly see cracks forming and spreading during experiments. As a result, the simulations serve as a “computational microscope,” in which researchers can see what is happening at the atomic scale.

The second simulation investigated ductile failure by using a record-smashing



Climate models over the past 30 years have become more complex and detailed and can now include more variables as the codes and the machines they run on have become increasingly more powerful.

1 billion atoms. The study simulated the creation and interaction of hundreds of dislocations. In ductile failure, metals bend instead of shatter as a result of plastic deformation, which occurs when rows of atoms slide past one another on slip planes (dislocations).

The simulations, which have been used to produce a 3D movie (www.llnl.gov/largevis/atoms/ductile-failure/index.html), show how a lattice of rigid junctions forms. “We see dislocation moving along, like a little wave, as atomic planes slide over one another,” says Abraham. Dislocations move, interact with one another, and finally become rigid as they stick to one another. This rigidity causes the material to change from ductile to brittle, a phenomenon also called work hardening. The phenomenon had been known for a long time but had never been understood on the atomic level.

Climate Study Needs Simulations

Researchers in climate change have long been avid users of the most powerful supercomputers available. “Supercomputer calculations have altered the direction of climate change research,” says Doug Rotman, head of Livermore’s Carbon Management and Climate Change Program. Rotman says the latest M&IC machines have helped simulation science in three ways.

First, they have increased the number of runs atmospheric scientists can do in a reasonable amount of time. “Climate is a statistical science,” he says. “We have to do ensembles of runs to discover what is happening. Starting with slightly different conditions leads to different weather patterns.”

Second, the machines provide increasing resolution, both horizontally and vertically. He notes that until recently, most climate change research has focused on the global scale at low resolution (100 kilometers at best). The computational power of MCR and

Thunder is permitting researchers to focus for the first time on the regional scale at higher resolution. “We’ll be looking at how certain emissions from a city, for example, affect air quality in a particular region,” says Rotman. “And we’ll be able to see, for the first time, how rain fluctuations affect snow packs in selected mountain ranges and also see the effects in other areas, such as reservoirs.”

Finally, the machines permit codes to incorporate an unprecedented amount of chemical and physical reactions and chemical species. “The computational power of M&IC computers has enabled Livermore to develop the IMPACT chemistry model and to push the scientific edges of atmospheric chemistry modeling,” adds Rotman. IMPACT is the only atmospheric model capable of interactively modeling the combined troposphere and stratosphere, which together make up Earth’s atmosphere. IMPACT can examine the processes that determine the distribution of ozone and other chemicals in the tropopause (the boundary between troposphere and stratosphere).

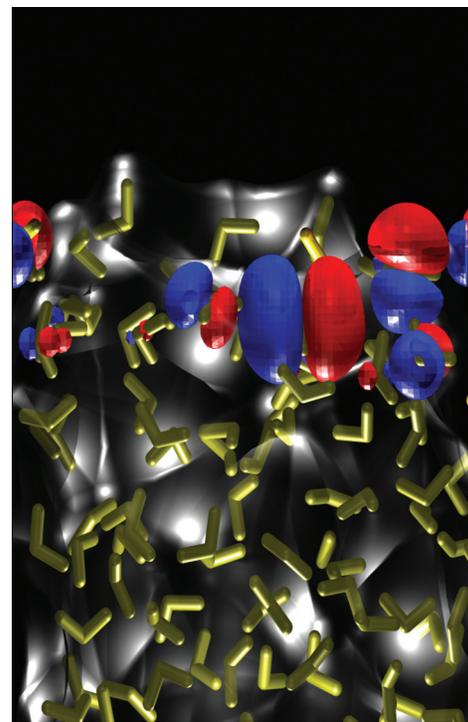
Rotman points out that Livermore is the only U.S. institution that can simulate carbon as it cycles through the planet in different forms, including sequestration in the oceans. “We’ve used every machine and are ramping up to use BlueGene/L,” he says.

Simulating Liquid–Vapor Interfaces

Chemist Chris Mundy and postdoctoral researcher I-Feng Kuo simulated the interface between liquid water and water vapor in a landmark simulation published in *Science* last year. Their results are important to both biologists and atmospheric scientists because knowledge of the reactions on surfaces and interfaces of liquid water are important in both fields. The pioneering study was made possible by using 440 microprocessors on MCR (63 percent of the machine) to simulate, over several picoseconds, 200 water

molecules comprising a film of water 3 nanometers deep.

“Recent experiments are probing the surface of liquids, and Livermore is playing a vital role in providing a microscopic picture using these terascale simulations,” says Mundy. He and colleagues are using Thunder to investigate the liquid–vapor phase interfaces of other species. The potential applications of these techniques include homeland security, such as calculating the physical characteristics



A snapshot is shown of the calculated aqueous liquid–vapor interface, as simulated in an ab initio molecular dynamics study. The individual water monomers are represented by the yellow cylinders. The top isosurface shows an area of the water slab that is reactive to excess protons and electrons. This pioneering study used 1,440 microprocessors on the Multiprogrammable Capability Resource (63 percent of the machine) to simulate, over several picoseconds, 200 water molecules comprising a film of water 3 nanometers deep.

of a possible terrorist chemical weapon without having to test an actual sample.

“When one has access to these kinds of computational resources, one’s first inclination may be to take a normal system and double it, but one doesn’t

often find new physics by taking that approach,” says Mundy. “Livermore’s terascale resources allow us to turn quantity, that is, system size, into a new quality—understanding complex chemistry in different environments. We couldn’t

simulate interfacial systems via first principles without MCR and Thunder. These resources enable us to answer many important scientific and programmatic questions from first principles. It’s very exciting!

Making Breakthrough Science Possible

Breakthrough computer simulations that lend new insight or reveal new physics require the most powerful machines and advanced codes available. For Livermore scientists running unclassified simulations, those machines have recently become Linux clusters.

The previous generation of scalable parallel supercomputers rely on vendor-integrated systems using high-performance proprietary microprocessors (the type found in expensive servers), proprietary interconnects and vendor (Unix-based) system software. This approach has matured to the point where the cost per teraops (trillion operations per second) improves only slowly.

In the late 1990s, Livermore experts turned to Linux cluster technology, in which large groups of commodity microprocessors (the type found in inexpensive desktop personal computers) are combined with third-party interconnects and open-source (Linux-based) system software. The result has been extraordinary advances in cost-effective computing, with Livermore taking the lead in harnessing this cluster technology for programmatic needs. Computer scientist Mark Seager says, “By using the economic force of nature known as a commodity ecosystem, we’ve changed the way scalable systems are architected, procured, and used. We invented the methods and the partnerships to scale these machines way beyond the previous state of the art.”

The first cluster to gain national recognition was the Multiprogrammatic Capability Resource (MCR) scalable parallel supercomputer, with 2,304 processors capable of 11.2 teraops. This system began running large-scale applications in December 2002. In June 2003, MCR was ranked as the world’s third fastest computer, the first time that a computer based on Linux-cluster technology was one of the top 10 computers in the world. The machine dramatically increased Livermore’s unclassified computing capability. MCR nearly matched the Advanced Simulation and Computing (ASC) Program’s 12.3-teraops White machine in power, but at \$1.2 million per teraops, it was 10 times less expensive than White. Seager notes that MCR would not have been possible without the ASC Program’s investments in technology developed at Livermore.

In 2004, as the competition for access to MCR increased, Livermore procured a more powerful system called Thunder. This machine features a peak speed of 22.9 teraops and uses 1,024 high-

performance nodes each with four Intel Itanium2 processors (4,096 processors total) and 8 gigabytes of memory. The machine debuted as the world’s second fastest computer. Most simulations on Thunder—and MCR—support projects in biology, materials science, lasers, and atmospheric science as well as unclassified simulations for stockpile stewardship.

A new generation of supercomputers uses system-on-a-chip technology and low-cost, low-power embedded microprocessors. This technology is embodied in BlueGene/L, which is scheduled to arrive at Livermore this summer. The machine will have 131,072 advanced microprocessors and a peak computational rate of 367 teraops.

Physicist Steve Langer says, “Clusters seem to be the way to go for massively parallel codes. There are a few problems that we can still cram into a single SMP (a computer architecture in which memory and other components are shared), but we’re not going to put in billions of zones and have enough memory and computing capacity to get a simulation done anytime soon. The performance-to-cost ratio of clusters is higher than we’ve had on any other machine.”

The software running on Multiprogrammatic and Institutional Computing (M&IC) Initiative computers consists of three major components: an operating system, a parallel file system, and a resource management system. The Clustered High Availability Operating System, developed at Livermore, augments the standard Linux system with modifications for high-performance networks, cluster management and monitoring, and access control.

LUSTRE is an open-source parallel file-sharing system developed in part through a collaboration between Livermore, ClusterFile Systems, Hewlett-Packard, Intel, and the ASC Program. Langer says, “LUSTRE ties everything together with a unified file system: multiple Linux clusters, visualization resources, and the archive. Before we had to copy data multiple times and then keep track of it.”

Simple Linux Utility for Resource Management (SLURM) is a tool developed in a collaboration between Livermore, Linux Networx, Hewlett-Packard, and others to manage a queue of pending work, allocate access to nodes, and launch and manage parallel jobs. SLURM has proven to be reliable and highly scalable. As a result of its success on Linux-based systems, it is being deployed on ASC platforms.

“To do these problems, you need to be at a place like Livermore,” says Mundy. “It takes not only the machines but also the staff of experts who can run the machines and write new software.” He notes that the newest generation of machines allows scientists to simulate a new class of physical reactions, which replace homogeneous systems with heterogeneous systems. “Most things in life are heterogeneous.”

Hydrogen Meltdown Uncovered

The October 7, 2004, cover of *Nature* reported a new melt curve of hydrogen at extremely high pressures predicted by Livermore scientists using the ab initio molecular dynamics code GP. This new curve—the result of nearly two years of work on an LDRD-funded project by physicists Stanimir Bonev, Eric Schwegler, Tadashi Ogitsu, and Galli—presents the melting point of hydrogen at pressures from 50 to 200 gigapascals at temperatures from 600 to 1,000 kelvins.

At about 80 gigapascals, the melt line hits a maximum and goes from a positive to a negative slope. This maximum, the scientists say, relate to a softening of the intermolecular interactions and to the fluid and solid becoming similar in structure and energy at high pressure. Melting point maximums are unusual but are also found in water and graphite. In these materials, liquid is denser than solid when they coexist.

The extremely complex simulations were run on several Livermore machines including MCR, Frost, and other Linux clusters. The GP code calculated the interactions of 720 atoms over time spans from 5 to 10 picoseconds. Thanks to the power of the code and machines, scientists were able to perform numerous simulations under varying thermodynamic conditions.

Results from these first-principles simulations provided strong evidence

of the existence of a low-temperature quantum fluid in hydrogen, notes Bonev. These findings led the team to propose new experimental measurements that could help verify the existence of a maximum melting temperature and the transformation of solid molecular hydrogen to a metallic liquid at pressures close to 400 gigapascals.

The success of the project, Galli notes, is due to a combination of codes, machines, and expertise accumulated over the years. “Results such as these do not just happen,” she emphasizes. “You need all three elements—people, codes, and machines—in place.”

Amazing Progress in 10 Years

Seager summarizes that simulation science at Livermore has gone from about 50 gigaops in 1995 to 12.3 teraops on ASC White in late 2000 to 100 teraops on ASC Purple. Later this year, BlueGene/L, with 360 teraops of processing power, will begin its shakedown. Seager says, “It’s been an amazing ramp-up from the first clusters, and all those systems have had direct and measurable impact on programs at this Laboratory.”

Mike McCoy, head of ICC and deputy associate director for Computation, calls today’s simulations “science of scale” because they are predictive. “The computing is performed at a resolution and degree of complexity, with inclusion of sufficient physics, that scientists have confidence they can predict the outcome of an experiment. This is an exciting time to be at Livermore,” he says.

These examples show over and over, in multiple physical disciplines and programs at the leading edge of scientific discovery, the tight coupling between experiment and simulation at the Laboratory. Simulation is essential in the design of modern experiments. In addition, simulations are now of sufficient resolution and size and

contain enough physics that their results can be directly compared with experimental results. Hence, simulations are essential to understand the physical phenomena involved in experiment.

Even with all this progress, Díaz de la Rubia notes that much work still needs to be done. Many fields, such as computational biology, are still in their infancy in taking advantage of the simulation power of the latest machines. “Supercomputers are helping us tremendously to solve problems we couldn’t attack otherwise,” he says, “but we can’t claim victory yet.” He sees a need for improved models and for coupling experiments with simulations more tightly.

With ASC Purple and BlueGene/L coming on line this year, Livermore computer experts are preparing for the newest generation of scalable parallel supercomputers. In addition, these systems position the ASC Program and the Laboratory for the next leap forward to petascale computing. For Livermore researchers, another generation of these systems means an even more powerful tool for understanding—and predicting—the physical universe.

—Arnie Heller and Ann Parker

Key Words: Advanced Simulation and Computing (ASC) Program, atmospheric chemistry modeling, BlueGene/L, climate change, crack propagation, ductile failure, GP code, hydrogen melt curve, IMPACT code, laser–plasma interactions, Linux clusters, liquid–vapor interfaces, molecular dynamics, Multiprogrammatic and Institutional Computer (M&IC) Initiative, Multiprogrammatic Capability Resource (MCR), National Ignition Facility (NIF), Purple, seismic wave analysis, supercomputers, Thunder, White.

For further information contact Mark Seager (925) 423-3141 (seager1@llnl.gov).