

Science & Technology

REVIEW

April 1998

Lawrence
Livermore
National
Laboratory



Meeting the
ASCI Challenge

Also in this issue:

- **Argus Advances in Security Technology**
- **Supersensitive Superconductor Detectors**

About the Cover

The computer simulation of compressible instability and turbulent mixing on the cover was produced at Lawrence Livermore on the Accelerated Strategic Computing Initiative (ASCI) Blue Pacific computer system using advanced codes developed at the Laboratory. The Department of Energy's ASCI program is a fundamental component of DOE's Stockpile Stewardship Program. A cooperative effort involving the three national weapons laboratories, private industry, and academia, ASCI is rapidly fulfilling its mission of substantially increasing the nation's computational power so that numeric simulations can contribute to the assurance of the safety and reliability of our nuclear deterrent without testing. Turn to p. 4 for a report on ASCI's progress in meeting its stockpile stewardship responsibilities.



What Do You Think?

We want to know what you think of our publication. Please use the enclosed survey form to give us your feedback.

Electronic Access

S&TR is available on the Internet at <http://www.llnl.gov/str>. As references become available on the Internet, they will be interactively linked to the footnote references at the end of each article. If you desire more detailed information about an article, click on any reference that is in color at the end of the article, and you will connect automatically with the reference.

About the Review

Lawrence Livermore National Laboratory is operated by the University of California for the Department of Energy. At Livermore, we focus science and technology on assuring our nation's security. We also apply that expertise to solve other important national problems in energy, bioscience, and the environment. *Science & Technology Review* is published ten times a year to communicate, to a broad audience, the Laboratory's scientific and technological accomplishments in fulfilling its primary missions. The publication's goal is to help readers understand these accomplishments and appreciate their value to the individual citizen, the nation, and the world.

Please address any correspondence (including name and address changes) to *S&TR*, Mail Stop L-664, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94551, or telephone (925) 422-8961. Our electronic mail address is hunter6@llnl.gov.



Prepared by LLNL under contract
No. W-7405-Eng-48

SCIENTIFIC EDITOR

J. Smart

MANAGING EDITOR

Sam Hunter

PUBLICATION EDITOR

Dean Wheatcraft

WRITERS

Sam Hunter and Gloria Wilt

ART DIRECTOR AND DESIGNER

Kitty Tinsley

INTERNET DESIGNER

Kitty Tinsley

COMPOSITOR

Louisa Cardoza

PROOFREADER

Al Miguel

S&TR is a Director's Office publication, produced by the Technical Information Department, under the direction of the Office of Policy, Planning, and Special Studies.

Printed in the United States of America

Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161

UCRL-52000-98-4
Distribution Category UC-700
April 1998

2 The Laboratory in the News

**3 Commentary by Michael Anastasio and David Cooper
The Accelerated Strategic Computing Initiative: A Daunting
Technical Challenge**

Feature

4 Computer Simulations in Support of National Security

The Accelerated Strategic Computing Initiative is multiplying computing power and vastly improving simulation technology as part of DOE's Stockpile Stewardship Program, which seeks to support assessment and certification of the nation's nuclear deterrent in the absence of nuclear testing.


Research Highlights

13 Argus Oversees and Protects All

16 Of Josephson Junctions, Quasi-particles, and Cooper Pairs

19 Patents and Awards

20 Abstract



Page 4



Page 13



Page 16

Laboratory spotlights movie special effects

In a presentation open to the public at Lawrence Livermore, three computer graphic artists from Industrial Light and Magic, a special-effects company based in San Rafael, California, explained how they work movie magic behind the scenes on some of today's hottest films. Their presentation, "The Digital Creature Feature: Putting It All Together," walked the audience through a typical sequence of events, from the director's first phone call to the computer modeling work to assembling all the elements that make up the final shots as we see them on the screen.

The free presentation was part of the Laboratory's "Science on Saturday" lecture series, a nine-week series of talks geared to middle and high school students, and their teachers, parents, and chaperones. Topics are selected from the forefront of scientific research, covering a variety of disciplines.

Contact: Dolores Doyle (925) 422-5056 (doyle5@llnl.gov).

Laboratory produces largest optical crystal

Scientists at Livermore have succeeded in using a "rapid-growth" method to produce the world's largest single-crystal optical elements. The pyramid-shaped KDP (potassium dihydrogen phosphate) crystal, measuring about 3 feet tall and over 20 inches wide at the base and weighing nearly 500 pounds, was grown in a 6-foot-high tank filled with nearly a ton of supersaturated solution.

The fast-growth method was pioneered in Russia by Natalia Zaitseva at Moscow State University and perfected at Livermore by Zaitseva and Laboratory scientists over the past few years. It allowed scientists to grow the record crystal in six weeks. Previous methods would have required a growing period of 12 to 24 months to achieve the same result. Slices of the KDP crystals will be critical components of the world's largest laser, the National Ignition Facility, currently under construction at Livermore.

Contact: Natalia Zaitseva (925) 423-1505 (zaitseva1@llnl.gov).

Scientists fine-tune mine detection

Lawrence Livermore scientists are working with Ukrainian scientists to evaluate a new device to detect land mines.

"This may be a solution to a small part of a large problem," says Arthur Toor, physics project leader for Lawrence Livermore. An estimated 100 million land mines in 70 countries cause 26,000 casualties each year.

Toor and his Livermore colleague David Eimerl are working with Ukrainian scientists Gregory Pekarsky and Vitaly Bystritski to evaluate a new detection device that is similar to one the Russian military is using. The device is used like a metal detector, but its technology involves detecting large amounts of hydrogen. When it detects a large amount of hydrogen, its gauge rises, indicating that a land mine may be present. Then the spot is tested for certain high levels of nitrogen that indicate a land mine is present.

In January, the scientists demonstrated Pekarsky's technology at the Buried Objects Detection Facility at the Nevada Test Site, a secure area with 296 defused mines.

Contact: Arthur Toor (925) 422-0953 (toor1@llnl.gov).

UC-managed labs rated excellent

The three national laboratories operated by the University of California for the Department of Energy received overall ratings of excellent from DOE for fiscal year 1997, based on annual assessments of the laboratories by the University as well as the Department's own reviews.

The findings, which a University spokesman said are consistent with those of previous years, stem from a five-year contract signed by DOE and the University in 1992, which pioneered the concept of performance-based management for nonprofit operators of DOE laboratories. The UC President's Council on National Laboratories characterized the science and technology activities at Los Alamos and Livermore as outstanding, the highest ratings available. Lawrence Berkeley National Laboratory received an excellent rating. The UC Laboratory Administration Office concluded that the administration and operations performance of each of the three laboratories exceeds expectations, which the University said is equivalent to an excellent rating. Based on these University assessments, DOE gave overall ratings of excellent to all three laboratories.

Contact: Public Affairs Office (925) 422-4599 (garberson1@llnl.gov).

NIF mammoth didn't live alone

The ancient mammoth excavated from Lawrence Livermore's National Ignition Facility construction site had loads of company, including at least one other mammal that apparently made a meal out of the mammoth, a scientist says.

Paleontologist C. Bruce Hanson found tooth marks on an upper leg bone of the elephant-like animal, which died roughly 10,000 years ago. "There are a few carnivores living then that could have been responsible," Hanson said. "I have to take a look and see if there's any chance of matching up the tooth marks with the known suspects."

The most likely culprit would have been a dire wolf similar in size to modern wolves, he said. But saber-toothed tigers and a type of lion that lived in the ancient, stream-crossed valley might have killed the mammoth or simply gnawed on its carcass.

Two other clusters of bones lie on the superlaser site, Hanson said. One includes the skull of an ancient horse, probably from the Pleistocene era, and the other cluster includes several ribs and a shoulder blade that appear to have come from a giant ground sloth, which probably stood 10 feet tall on its hind legs.

Contact: Public Affairs Office (925) 422-4599 (garberson1@llnl.gov).



The Accelerated Strategic Computing Initiative: A Daunting Technical Challenge

THE Department of Energy's Stockpile Stewardship Program is designed to maintain high confidence in the safety, reliability, and performance of the U.S. nuclear weapons stockpile into the indefinite future under a Comprehensive Test Ban Treaty. This objective requires a fundamental understanding of nuclear weapons science and technology through an integrated program incorporating archival nuclear test data, high-fidelity nonnuclear experiments, and advanced computer simulations. The Accelerated Strategic Computing Initiative (ASCI) was developed to provide the simulation capabilities required to meet the stockpile assessment and certification requirements in this uncharted environment. A key element of this initiative is an aggressive plan to acquire computing systems that are substantially beyond those that can be expected simply from market forces. It also represents systems that are (1) at least an order of magnitude greater in speed and memory than any currently in use and (2) three orders of magnitude greater in capability than servers used by most industry and commercial firms and some government laboratories. Unprecedented advances will also be necessary in several areas of computer science and technology to integrate the data transmission, networking, storage, and visualization needs of the ASCI simulation codes.

To meet this tremendous challenge requires much more than just developing a hardware environment. New simulation codes must be developed that incorporate more fundamental science,


a more accurate database describing the properties of materials under nuclear weapon conditions, and new algorithms that take advantage of the massively parallel architectures of the ASCI computers. And the ASCI program's extremely short deadline amplifies the challenge—the target for assembling this high-fidelity three-dimensional simulation capability is the year 2004.

The ASCI approach to accelerating the nation's computer simulation capability reflects the fact that expertise in the many disciplines required to do this job does not necessarily reside within any single laboratory. Many of the important contributions to our nation's technological future must be found within the wider circle of scientific expertise represented in a variety of ASCI partnerships—involving the Sandia, Los Alamos, and Lawrence Livermore national laboratories and featuring collaborations with the commercial computer industry and some of the nation's leading universities.

Here at Livermore, participation in the ASCI program extends across two directorates: the Defense and Nuclear Technologies Directorate and the Computation Directorate. This issue's feature article, which begins on [p. 4](#), describes ASCI, a critical element of DOE's Stockpile Stewardship Program.

-
- Michael Anastasio is Associate Director, Defense and Nuclear Technologies.
 - David Cooper is Associate Director, Computation, and the Laboratory's Chief Information Officer.

Computer Simulations in Support of National Security



The Accelerated Strategic Computing Initiative is making significant progress toward meeting its major challenge: dramatically increasing the nation's computing power as a necessary contribution to the assurance by scientists of the safety and reliability of our nuclear deterrent in the absence of testing.

TIME is running out on the U.S. nuclear weapons stockpile. As the weapons age beyond their design lifetimes, important questions arise: Are the weapons still safe? Will they still perform reliably? How long will they continue to be reliable? What maintenance and retrofitting should be prescribed to extend their working life? These questions must be answered with confidence as long as nuclear deterrence remains an essential part of U.S. national security policy.

With the U.S. commitment to the Comprehensive Test Ban Treaty, the viability of the U.S. nuclear arsenal can no longer be determined through underground nuclear testing. Thus, new approaches are being taken to maintain and preserve the U.S. nuclear deterrent through DOE's Stockpile Stewardship Program.

One key component of the multifaceted Stockpile Stewardship Program is the Accelerated Strategic Computing Initiative (ASCI), an effort to push computational power far beyond present capabilities so scientists can simulate the aging of U.S. nuclear weapons and predict their performance. To calculate in precise detail all the complex events of a thermonuclear explosion requires computational power that does not yet exist, nor would it exist any time soon without the ASCI push, even at computer development speeds predicted by Moore's Law (that computer power doubles about every two years). ASCI's goal is to put such a high-fidelity simulation capability in place in the near future. To do that, the American computer industry must dramatically speed up the pace of computational development. Currently, computing's top

speed is 1.8 teraflops, that is, 1.8 trillion floating-point (arithmetic) operations per second. This speed must increase to at least 100 teraflops by 2004, growth that must be coordinated with a host of accomplishments in code development and networking.

Why is this accelerated schedule necessary? Not only are weapons aging, so are the nuclear weapons experts with experience in designing and testing them. The Stockpile Stewardship Program must have this high-fidelity, three-dimensional simulation capability in place before that expertise is gone. "It's a tremendously ambitious goal, especially under such a short schedule," says Randy Christensen, ASCI's deputy program leader at Lawrence Livermore National Laboratory. Christensen describes the work as something akin to "trying to get a computer code to run in a few days a simulation that would have taken so long with current capability that it would not have been attempted."

Orchestrating Integration

ASCI is reaching for computational powers in the hundreds of teraflops, but the ASCI challenge demands more than hardware. Meeting it will require careful integration of the major elements of a national effort: platform development, applications development, problem-solving environment, and strategic alliances—coordinated work conducted at three national laboratories in partnership with the commercial supercomputer industry and the nation's great universities (Figure 1).

Figure 1. Meeting the challenge of ASCI requires careful integration of the major elements of the program across three national laboratories and ASCI's industrial and academic partners.

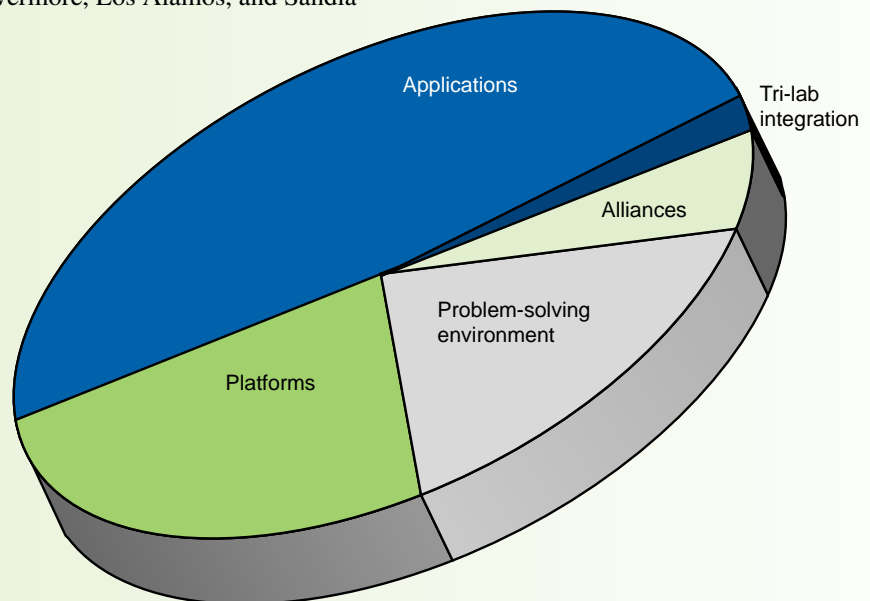
To ensure this balanced development, ASCI planning began with a "one program—three laboratories" approach. Project leaders at each laboratory, guided by the DOE's Office of the Assistant Secretary for Defense Programs, are implementing this collaboration and extending it to ASCI's industrial and academic partners. The overriding challenge for the ASCI scope of work is to synchronize the various technological developments with each other. For example, sufficient platform power must be delivered in time to run new advanced codes, and networking capabilities must enable the various parts of the system to behave as if they were one. The success of ASCI depends on this integration as much as it depends on the success of ASCI's individual elements.

Developing the Platform

ASCI's computer hardware is being developed by a consortium of three national laboratories and a select group of industrial partners in a prime example of government—industry cooperation. The national laboratories—Lawrence Livermore, Los Alamos, and Sandia—

are each teamed with a major commercial computer manufacturer—IBM, Silicon Graphics—Cray, and Intel, respectively—to design and build parallel, supercomputing platforms capable of teraflops speeds.

The development of infrastructure technologies seeks to tap all available resources to make these computer platforms perform the kind of high-fidelity simulation that stockpile stewardship requires. ASCI has a PathForward component, a program that invites computer companies to collaborate in developing required technologies. For instance, the program's first PathForward contracts, announced on February 3, 1998, awarded more than \$50 million over four years to four major U.S. computer companies to develop and engineer high-bandwidth and low-latency technologies for the interconnection of 10,000 commodity processors that are needed to build the 30-teraflops computer. (See box, p. 6.) As a result of this effort, subsequent collaborations involving other agencies, academia, and industry are expected.



Developing Applications

ASCI is an applications-driven program. Unprecedented computer power with a first-rate computing environment is required to do ASCI's stockpile stewardship job, which is to run new computer codes programmed with all of the accumulated scientific knowledge necessary to simulate the long-term viability of our weapons systems. The new generation of advanced simulation codes being developed in the ASCI program must cover a wide range of events and describe many complex physical phenomena. They must address the weapon systems' normal performance from high-explosive initiation through

final nuclear yield and the effects of changes introduced by remanufacturing (perhaps using different materials and fabrication methods) or defects brought on by aging. In addition, they must simulate weapon behavior in a wide variety of abnormal conditions to examine weapon safety issues in any conceivable accident scenario. If this weren't difficult enough, the new codes must provide a level of fidelity to the actual behavior of weapons that is much higher than their predecessors provided.

The major challenges facing the developers of these advanced simulation codes are to base them on rigorous, first-principles physics and eliminate many

of the numerical approximations and simplified physics that limit the fidelity of current codes; make them run efficiently on emerging high-performance computer architectures; validate their usefulness by means of nonnuclear experiments and archival nuclear test data; and do all of these in time to meet stockpile needs.

Meeting these challenges requires the coordinated efforts of over a hundred physicists, engineers, and computer scientists organized into many teams. Some teams create the advanced weapon simulation codes, writing and integrating hundreds of smaller programs that treat individual aspects of weapon behavior into a single, powerful simulation engine that can model an entire weapon. Other teams are devoted to developing the advanced numerical algorithms that will allow these codes to run quickly on machines consisting of thousands of individual processors—a feat never before achieved with programs this complex. Still others are developing much improved models for the physics of nuclear weapon operation or for the behavior of weapon materials under the extreme conditions of a nuclear explosion. Both the scale (the largest teams have about 20 people) and the degree of integration demanded by this complex effort have required a much greater level of planning and coordination than was needed in the past.

One example of the advanced simulation capabilities being developed in the ASCI program is its material modeling program. Enormously powerful ASCI computers are being used to carry out very accurate, first-principles calculations of material behavior at the atomic and molecular level. This information is then used to create accurate and detailed models of material behavior at larger and larger length scales until we have a model that can

PathForward Contracts Awarded February 1998

Industrial Partner	PathForward Project
Digital Equipment Corporation (DEC) Maynard, Massachusetts	Develop and demonstrate a processor interconnect capable of tying together 256 Digital UNIX-based AlphaServer symmetric multiprocessing (SMP) nodes.
International Business Machines (IBM) Poughkeepsie, New York	Develop future high-speed, low-latency, scalable switching technology to support systems that scale to 100 teraflops.
Silicon Graphics–Cray Research (SGI/Cray) Chippewa Falls, Wisconsin	Develop and evaluate advanced signaling and interconnect techniques. The technology will be used in future routers, switches, communication lines, channels, and interconnects.
Sun Microsystems (SUN) Chelmsford, Massachusetts	Perform hardware and software viability assessments by constructing interconnect fabric and verifying scalability and correctness of the interconnect monitoring facilities, resource management, and message-passing interface (MPI) capabilities.

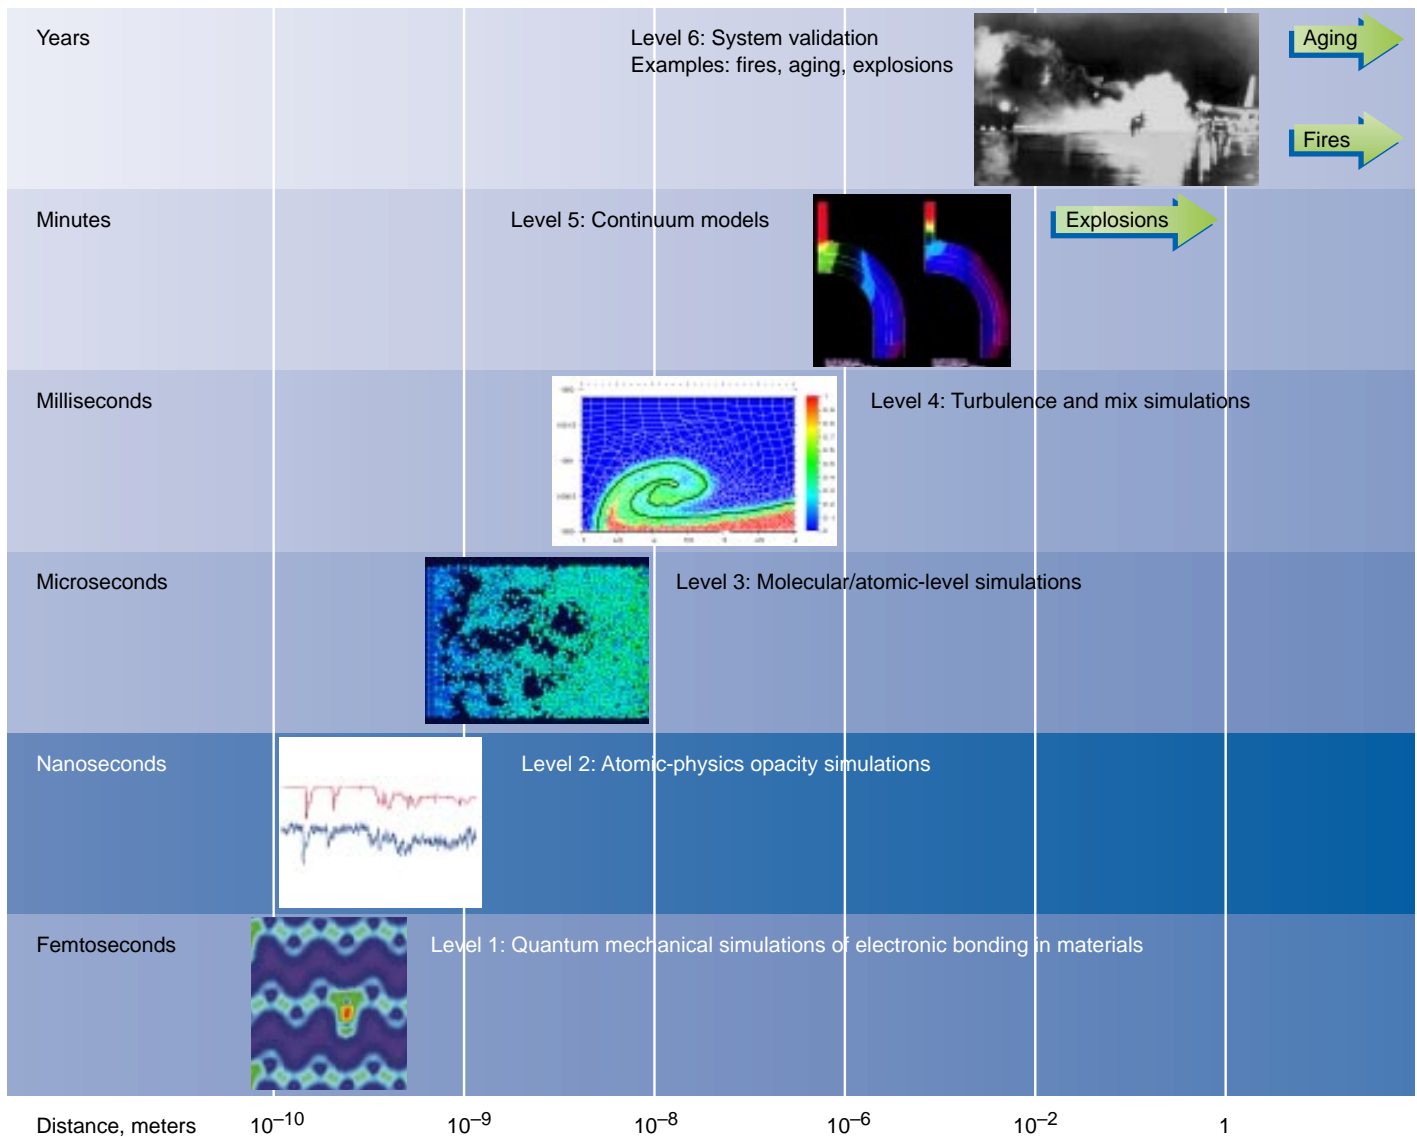
be used directly in the weapon simulation codes (Figure 2). This computational approach to material modeling has already produced a much better understanding of the phase changes in actinides (the chemical family of plutonium and uranium). The new approach is expected to be applied to many weapons materials, ranging from plutonium to high explosives. When fully developed, it will become a powerful tool for understanding and

predicting the behavior of any material (for example, alloys used in airplane construction, steel in bridges), not just those used in nuclear weapons.

Developing the Infrastructure

In addition to platform and applications development, ASCI is also developing a powerful computer infrastructure. A high-performance problem-solving environment must be available to support and manage the

Figure 2. ASCI is providing DOE's Stockpile Stewardship Program with a hierarchy of models and modeling methods to enable predictive capability for all processes relevant to weapon performance.



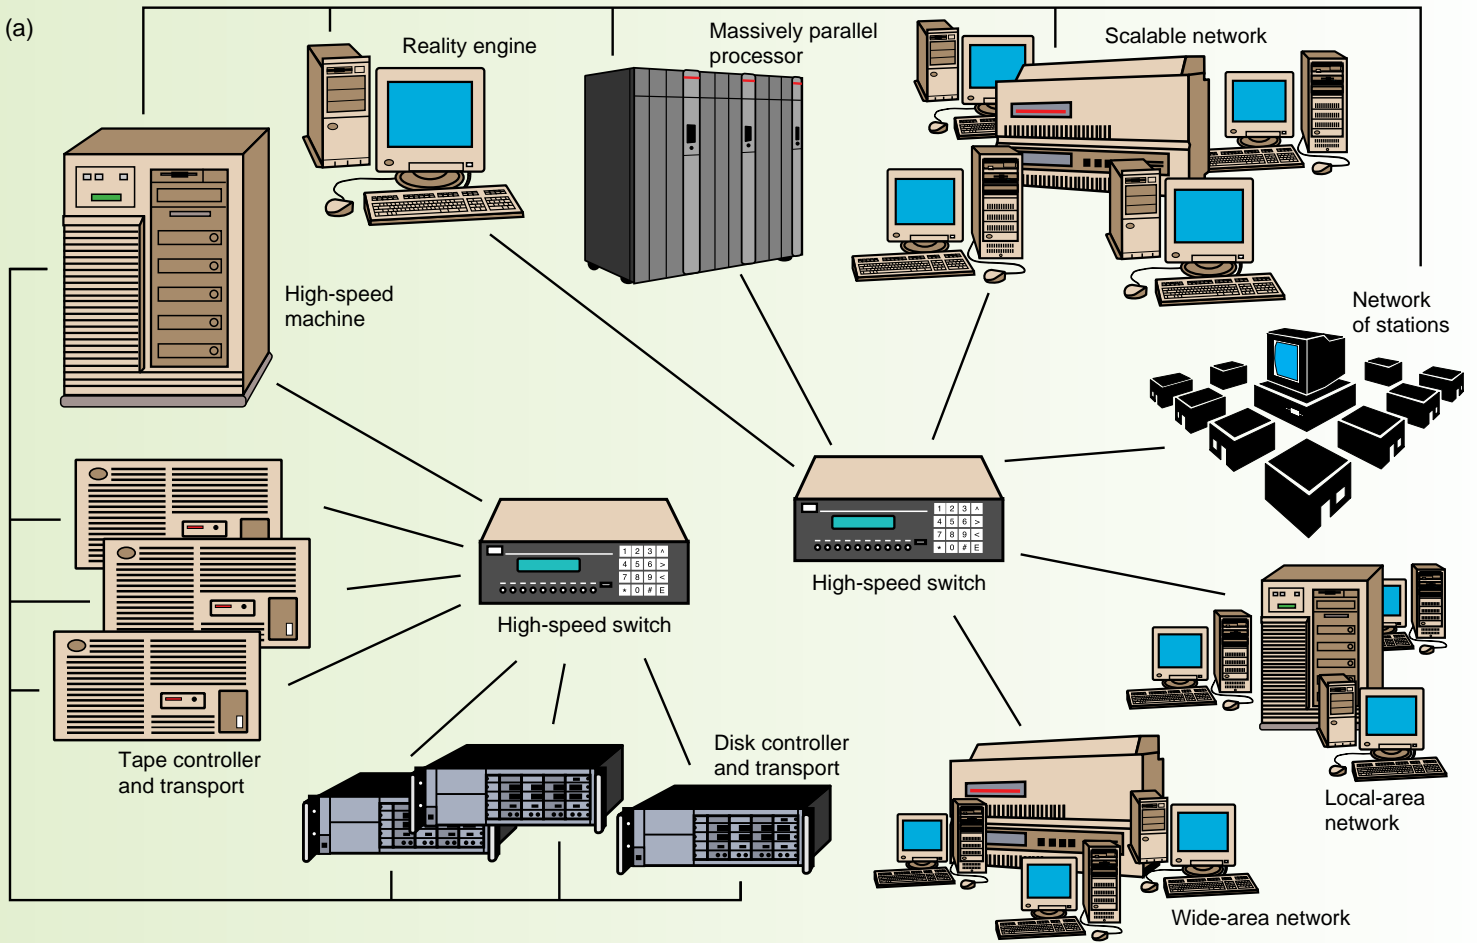



Figure 3. (a) A high-performance problem-solving environment manages the workflow and communications among all ASCI computers. (b) The result is ASCI's ability to bring three-dimensional images resulting from calculations to scientists on their desktops regardless of the physical location of the processors doing the work. This distance-computing option features remote caching of simulation data.



workflow and the communications between all the ASCI machines. At any time, over 700 classified and unclassified code developers and testers may be accessing ASCI computers, either from within the national laboratories or via the Internet. A scalable network architecture, in which individual computers are connected by very high-speed switches into one system, makes this high-demand access possible. With such a configuration, the network is, in effect, the computer (Figure 3).

Allowing large numbers of computers to communicate over a network as if they were a single system requires sophisticated new tools to perform scientific data management, resource allocation, parallel input and output, ultrahigh-speed and high-capacity intelligent storage, and visualization. These capabilities must be layered into the computer architecture, between user and hardware, so that the two can interact effectively and transparently. The applications integrate the computing environment and allow users, for example, to access a file at any of the three national laboratory sites as if it were a local file or to share a local file with collaborators at any ASCI site.

At Livermore, ASCI staff are performing numerous projects to develop this integrated computing environment. One team is working on a science-data management tool that organizes, retrieves, and shares data. An important objective of this tool is to reduce the amount of data needed for browsing terascale data sets. Another team is developing data storage that will offer a vast storage repository for keeping data available and safe 24 hours a day. The repository will store petabytes (quadrillions of bytes) of information, equivalent to one hundred times the contents of the Library of Congress. The storage device will also rapidly deliver information to users, at a rate

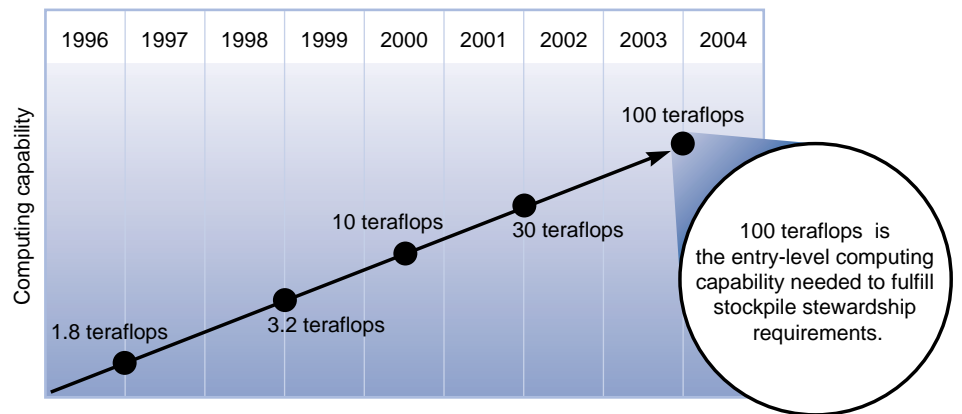


Figure 4. The ASCI goal is to achieve the 100-teraflops (trillion-floating-point-operations-per-second) threshold by 2004.

of billions of bytes per second. Another ASCI team is writing scalable input and output software to move data from computer to computer and reduce congestion between computers and storage. The changes resulting from its improvements will be tantamount to moving busloads of data, as compared to carloads—a sort of mass transit for data.

Weapons scientists will be confronted with analyzing and understanding overwhelmingly large amounts of data derived from three-dimensional numerical models. To help them, ASCI is developing advanced tools and techniques for computer visualization, wherein stored data sets are read into a computer, processed into smaller data sets, and then rendered into images. The development of visualization tools for use across three national laboratories will require close collaboration with regard to programming language, organization, and data-formatting standards. The Livermore team is focusing on how to reduce data sets for visualization—because they surely will become larger and larger—through the use of such techniques as resampling,

multiresolution representation, feature extraction, pattern recognition, subsetting, and probing.

While the fast, powerful machines and complex computer codes garner most of the headlines, this problem-solving-environment effort is fundamental to fulfilling the ASCI challenge. As we come to understand that “the network is the computer,” the significance of this element of the ASCI program comes sharply into focus.

In Pursuit of 100 Teraflops

The 100-teraflops milestone, the entry-level computing capability needed to fulfill stockpile stewardship requirements, is ASCI’s goal for 2004 (Figure 4). Fulfilling it will require enough computational power to run calculations distributed over 10,000 processors, which is just enough to conduct three-dimensional weapons simulations at a level of complexity that matches the current understanding of weapons physics. While this computing capacity is not the final goal, it is already 100,000 times more than the computing power used

(a)



(b)



Figure 5. (a) The IBM Blue Pacific computer arrived at Livermore on September 20, 1996, just two months after the IBM/Livermore partnership was announced by the White House and about six weeks after the contract was signed. (b) The initial-delivery system has already begun significant calculations in important areas of stockpile stewardship such as three-dimensional modeling of material properties, turbulence, and weapon effects. Upgrades will bring system power to 3.28 teraflops (trillion floating-point operations per second) by 1999, with an option to upgrade to 10 teraflops in fiscal year 2000.

by weapons scientists today, represented by Livermore's J-90 Cray computer. At 100 teraflops, all of the calculations used to develop the U.S. nuclear stockpile from the beginning could be completed in less than two minutes.

ASCI's approach to the 100-teraflops goal has been to use off-the-shelf, mass-market components in innovative ways. It aggregates the processors developed for use in desktop computers and workstations to scaleup computing power. It is this approach that makes ASCI development cost-effective; and leveraging of commercially available components will encourage technology development in the commercial sector. The mass-market approach will take advanced modeling and simulations into the computational mainstream for universal PC use.

Improvements to ASCI power will occur over five generations of high-performance computers. To ensure success, multiple-platform development approaches are being attempted. This strategy will reduce risk, allow faster progress, and result in greater breadth of computing capability. For example, the Sandia/Intel Red machine, which was put on line in August 1995, has achieved 1.8-teraflops speed (currently the world's fastest) and is now being used for both code development and simulation. The Lawrence Livermore/IBM Blue Pacific and the Los Alamos/Silicon Graphics-Cray Blue Mountain systems, which resulted from technical bids awarded in late 1996, are already running calculations.

Blue Pacific was delivered to Lawrence Livermore on September 20, 1996, with a thousand times more power than Livermore's existing Cray YMP supercomputer (Figure 5a). The Lawrence Livermore/IBM team installed and powered up the system and had it running calculations within two weeks. Already, it is conducting

some of the most detailed code simulations to date.

The Blue Pacific initial-delivery system, which arrived in 340 refrigerator-sized crates, takes up a significant portion of Livermore's computing machine room space, operates at 136 gigaflops, and has 67 gigabytes of memory and 2.5 terabytes of storage (Figure 5b). Initially, each of its 512 nodes contained one processor. During March 1998, these nodes were replaced with four-way symmetrical multiprocessors, quadrupling the number of processors. A further improvement will endow it with thousands of significantly improved processor nodes for the ASCI production model. These reduced-instruction-set computing microprocessors operate at a peak of 800 megaflops and, in this configuration, will bring the system to a total of 3.28 teraflops.

In that three-teraflops configuration, the Blue Pacific's "Sustained Stewardship Teraflops" system alone would more than fill up all the space in Livermore's current machine room. For that reason, construction crews are now building and wiring new space to accommodate it. In new, larger quarters, workers have been installing electric power, replacing air handlers and coolers, and hooking up new fans as part of necessary building upgrades. The numbers are impressive: 12,000 square feet of building extension, 5.65 megawatts of power, 11 tons of air conditioning, 16 air handlers that replace the air four times per minute, and controllers that keep the temperature between 52° and 72°F at all times. This machine is scheduled to be installed in March or April of 1999 (Figure 6).

Involving Academia

Although work on weapons physics is classified, work on the methods and techniques for predictive materials models encompasses unclassified research activities. ASCI thus can pursue

a strategy of scientific exchange with academic institutions that will more rapidly establish the viability of large-scale computational simulation and advance simulation technology. This strategy is embodied in the Academic Strategic Alliances Program. The program invites the nation's best scientists and engineers to help develop the computational tools needed to apply numerical simulation to real-world problems. In this way, a broader scientific expertise is at work making the case for simulation; simulation algorithms are tested over a broad range of problems; and the independently produced simulations provide a peer review that helps validate stockpile stewardship simulations (see box below).

Computers Changed It All

In the short span of time since computers came into general use, the

nature of problem-solving has changed, by first becoming reliant on computers, and then becoming constrained by the limits of computer power. ASCI will develop technologies that will make computational capability no longer the limiting factor in solving huge problems. Just as important, ASCI will change the fundamental way scientists and engineers solve problems, moving toward full integration of numerical simulation with scientific understanding garnered over decades of experimentation.

In the stockpile stewardship arena, the ASCI effort will support high-confidence assessments and stockpile certification through higher fidelity simulations. Throughout American science and industry, new products and technologies can be developed at reduced risk and cost. Advanced simulation technologies will allow scientists and engineers to do such things as study the workings of

The Academic Strategic Alliances Program

In July 1997, the Academic Strategic Alliances Program awarded Level I funds to five universities to perform scientific modeling to establish and validate modeling and simulation as viable scientific methodologies.

- Stanford University will develop simulation technology for power generation and for designing gas turbine engines that are used in aircraft, locomotives, and boats. This technology is applicable to simulating high-explosive detonation and ignition.
- At their computational facility for simulating the dynamic response of materials, the California Institute of Technology will investigate the effect of shock waves induced by high explosives on various materials in different phases.
- The University of Chicago will simulate and analyze astrophysical thermonuclear flashes.
- The University of Utah at Salt Lake will provide a set of tools to simulate accidental fires and explosions.
- The University of Illinois at Urbana/Champaign will focus on detailed, whole-system simulation of solid-propellant rockets. This effort will increase the understanding of shock physics and the quantum chemistry of energetic materials, as well as the effects of aging and other deterioration.

These Level I projects are part of a 10-year program, in which projects can be renewed after five years. Also under the Alliances program, smaller research projects are being funded at universities across the country as Level II and III collaborations.

Figure 6. Rendering of the 3.28-teraflops IBM Blue Pacific “Sustained Stewardship Teraflops” system in its new home being constructed at Lawrence Livermore. The machine is scheduled for installation in early 1999.



disease molecules, so they can design drugs that combat the disease; observe the effects of car crashes without an actual crash; and model global weather to determine how human activities might be affecting it. The uses are limitless, and their benefits would more than justify this investment in high-end computing, even beyond the benefits of ASCI’s principal national-security objective.

—Gloria Wilt

Key Words: Academic Strategic Alliances Program, Accelerated Strategic Computing Initiative (ASCI), computer infrastructure, computer platform, parallel computing, PathForward, problem-solving environment, Stockpile Stewardship Program, simulation, teraflops, weapons codes.

For further information contact

*Randy B. Christensen (925) 423-3054
(christensen5@llnl.gov).*

About the Scientist



RANDY CHRISTENSEN is Deputy Program Leader of the Department of Energy’s Accelerated Strategic Computing Initiative (ASCI). He has broad management responsibilities within the ASCI program as well as specific responsibility for applications development. He holds a B.S. in physics from Utah State University and an M.S. and Ph.D. in physics from the University of Illinois. Following a postdoctoral fellowship at the

Joint Institute for Laboratory Astrophysics (1978–1981), he joined Lawrence Livermore National Laboratory as a code physicist in the Defense and Nuclear Technologies Directorate. He held a number of leadership positions in that directorate before becoming Deputy Associate Director of the Computation Directorate in 1992, where his responsibilities included management of the Livermore Computing Center.

Argus Oversees and Protects All

FROM outside Lawrence Livermore National Laboratory, the public sees a site protected by chain link fences and guards at entry gates. But this Department of Energy national laboratory, home to a variety of classified research, requires much higher level security measures. Therefore, it is guarded as well by a sophisticated, computerized security system called Argus. Argus was designed, engineered, and installed at Livermore and is continually being upgraded and enhanced. It is also available to other Department of Energy and Department of Defense facilities.

Although named for the hundred-eyed monster of Greek myth, Argus security comprises much more than visual capabilities. A highly interconnected network engineered with comprehensive security features, Argus lives up to such stringent security requirements that DOE's Office of Safeguards and Security has cited it as the standard for physical security systems protecting facilities where the consequences of intrusion are significant. In addition to Lawrence Livermore, the Argus system has been installed at three other DOE sites and at one DOD site to protect top-priority assets or nuclear material.

As it monitors and controls entry into the Laboratory's high-security buildings, Argus is simultaneously monitoring the entire site for security threats and can alert and direct security forces to those threats. Argus security is all-encompassing and omnipresent, but it is surprisingly noninvasive. Employees of Lawrence Livermore enter and move about the Laboratory campus with relative ease. Yet, the Laboratory's Top Secret documents, materials, and facilities are thoroughly protected, intruders can be detected in real time, and intrusions and emergencies get instantaneous response from police and investigative personnel. The Laboratory is provided with maximum security 24 hours a day, 7 days a week.




Figure 1. Remote-access panels are the primary interface between users and the Argus system and are the means by which badged employees enter controlled buildings and areas.

This security results from a software system that comprises some 1.5 million lines of code, offering a wide range of security features. Extensive features are necessary, because Argus must accommodate many different configurations of security rules within one security complex, and sometimes one complex may have multiple geographical locations (for example, Livermore's Argus system controls the main site and the nearby Site 300 high-explosives testing facility). Moreover, Argus must be reconfigurable at any time. Extensive features also translate into flexibility and simplicity for end users. That's important because every authorized person in a high-security site accesses and interfaces with the Argus system. To ensure that designers, operators, and users understand Argus, DOE's Central Training Academy in Albuquerque, New Mexico, has 14 classes available, ranging from one hour to one week, that cover the complete set of Argus features.

While protecting a security complex, Argus also protects itself. A high degree of redundancy has been incorporated to prevent system failure, and tamper-indicating devices and data encryption have been used throughout to protect

surveillance equipment and data from intruders and thieves. Insider threats to weaken the system have been addressed with a comprehensive set of system-enforced and procedural measures, including consistency checking, captive accounts, and a rule prohibiting people from working alone.

How Users Work with Argus

Argus is implemented through four integrated computer subsystems. One subsystem controls access into buildings and areas. Another monitors alarms and sensors installed throughout the site. A third integrates and displays security data so security personnel can assess and control incidents. The fourth provides central computing and data storage to support the overall system configuration and databases. These four elements provide what Greg Davis, the manager of Livermore's Argus program, calls a "God's eye view" of the site. They connect into a real-time, interactive security assessment and response system.

User interactions with the Argus network are made possible through two hardware components: a remote access panel (RAP) and the Argus field processor (AFP).

The RAP (Figure 1) is a microprocessor-based, programmable input-output device connected to an AFP (Figure 2). It is the primary user interface to the Argus system. When Livermore



Figure 2. Argus field processors are microprocessor-based devices that verify badge information transmitted from remote access panels against locally stored encrypted access authorization databases.

employees are "badged," they are enrolled into the Argus system and can then use RAPs to gain entry into controlled buildings and areas on site. They swipe their badges, which have been coded with a unique identification number and a decryption key. The RAP communicates the badge information to the AFP, another microprocessor-based device, which verifies it against locally stored encrypted access authorization databases.

Argus software allows access based on credentials (determined by a badge), a user's identity (determined by personal identification number and biometrics), clearance level, and privilege. Although access rules can be very restrictive, the access system provides flexibility by being able to make fine distinctions within those rules. Thus, it might allow a person into a high-security building within a classified area but prevent that person from entering an even higher level security vault within that building. The access system also allows changes in user privileges, within rule confines; for example, regular users can be enrolled to escort visitors through high-security areas. The system eliminates the need for labor-intensive badge checking, and it monitors, tracks, and logs all badge usage.

In addition to controlling and monitoring the RAP access controls, AFPs also control and monitor the networks of thousands of electronic sensors and other surveillance equipment that comprise the alarm stations of a security complex.

The AFP determines the status of security in the alarm station by polling its sensors, controls station operating mode (that is, whether the station is open or secured, in maintenance, etc.), and provides entry authorization via the RAP interface. Alarm station caretakers can also use the RAP to modify access lists, change the rules of the alarm station, and authorize maintenance on the station.

Alarm stations are of many types—outdoor perimeter exclusion zones, normal interior rooms, vaults of concrete or steel, or even entire buildings. They can have sensors and surveillance equipment installed on walls, floors, and ceilings. Because as many AFP modules can be installed as necessary to monitor alarm stations, site security is scalable. At the same time, its modularity restricts problems and makes maintenance and diagnostic work easier.

Real-Time Command and Control

Occasionally at the Laboratory, police cars with flashing lights and howling sirens speed through the streets in response to an alarm or other security incident. They have been dispatched by security personnel who monitor site security 24 hours a day from Argus consoles (Figure 3). The consoles



Figure 3. Security personnel monitor site security 24 hours a day from Argus consoles like this one at Lawrence Livermore.

integrate and display graphical data from controlled entryways and alarm stations, and they are linked to telephone, radio, and intercom systems. They provide Livermore security staff with a real-time command-and-control capability.

At each console workstation, an operator controls two high-resolution, color display screens that show maps of security areas and the security equipment contained in them (sensors, entry control devices, cameras). The system display lists any security anomalies that are occurring and indicates the security status of surveillance equipment by color code. Green, for example, indicates normal or secure, while red indicates a potential security threat, an alarm, or a failure. When security anomalies occur, an operator is alerted by the lists and can view them on the screens; the views can be enlarged or adjusted for seeing additional details.

Operators may also be able to zoom in on the anomaly. Consoles can be linked to closed circuit televisions. Console video subsystems have computer-controlled switches capable of delivering signals from any linked television camera simultaneously to any display monitor and to all recording devices. Video options also include pan-tilt-zoom cameras and video motion detectors.

The consoles are ergonomically designed, providing comfort and ease of use to operators. The number of consoles in operation depends on site requirements and operator workloads; Argus can support any number of workstations without degradation.

Continuing Improvements, Ever More Uses

The installation of Argus at a major DOE nuclear weapons storage and dismantlement site is nearing completion. There,

Argus was modified to accommodate access authorization procedures that require observation of the two-person rule for entry and exit. In addition to RAPs, the entry portals have devices that read stored hand-geometry data, and booths may have special detectors to monitor the transport of sensitive materials. To serve this site and other users, Argus program staff are developing a 24-hour help line.

They are also moving ahead to evolve Argus to the next technological level, with such features as topology-independent network-based sensors and capability to simulate intrusions and attacks. In the first, Argus staff are in the midst of developing a neuron chip that can be embedded into sensors, adding the capability to communicate with sensors instead of merely receiving signals from them. This feature will enhance AFP line supervision of alarm stations, enhance sensor security, and dramatically reduce installation costs. In the second, Argus staff are beginning research and development to endow Argus with simulation capabilities that can be used in conjunction with conflict simulation exercises. Argus console operators will soon be able to detect simulated attacks and send virtual security dispatches to contain and control them. Such simulation would hone a site's emergency response tactics and provide realistic training to console operators.

—Gloria Wilt

Key Words: Argus, Argus field processor (AFP), remote access panel (RAP), security technology.

For further information contact
Gregory Davis (925) 422-4028 (davis19@llnl.gov)

Of Josephson Junctions, Quasi-particles, and Cooper Pairs

If Josephson junction brings to mind an intersection of two small back roads, it's time to change gears and think science. This term, along with quasi-particle and Cooper pair, is part of the large area of superconductors.

Simon Labov and his colleagues in Lawrence Livermore's Physics and Space Technology Directorate say these concepts and discoveries show great promise for applications in areas such as wireless communication, energy storage, and medical diagnostics. Labov and his fellow researchers are using superconductors to create a new generation of supersensitive detectors for nondestructive evaluation and astrophysics.

When ordinary metal conducts electricity, the electrons carrying the current collide with imperfections in the metal, thereby creating resistance. But when a superconducting material is cooled to its critical temperature, electrons pair off into Cooper pairs, named for Leon Cooper, one of the scientists who won a 1972 Nobel Prize in physics for explaining the now widely accepted theory. Any movement of one electron is matched by equal and opposite movement of the other. As a result, they don't hit the imperfections, no electrical resistance is generated, and electrons flow freely, without the addition of more energy.

But to put these theories to practical use in detectors requires a Josephson junction. Named for Brian Josephson, who described the theory when he was a graduate student at Cambridge University in 1962, a Josephson junction is two pieces of superconducting material linked by a weak insulating barrier. When an x ray hits a Josephson junction, the Cooper pairs break up, and quasi-particles are created. These quasi-particles, which are electronlike or holelike excitations in the

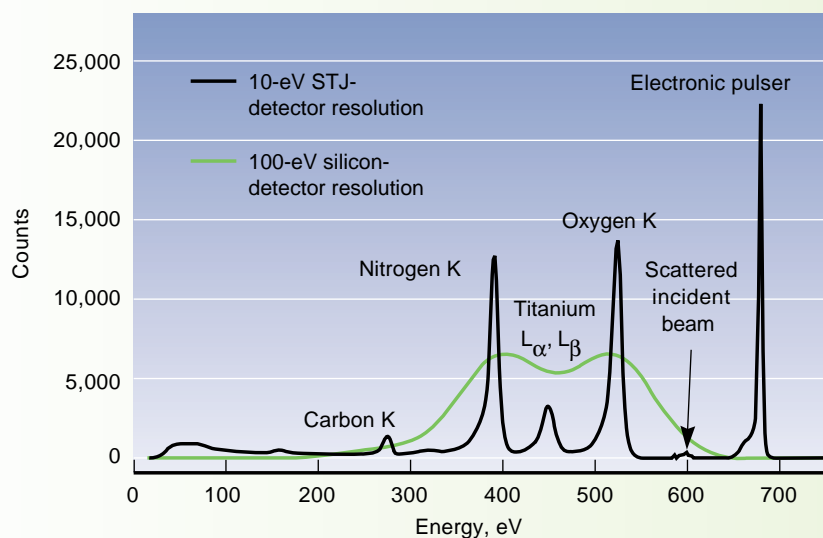


Stephan Friedrich inserts a sample in the Lawrence Livermore superconducting-tunnel-junction (STJ) detector cryostat at the Stanford Synchrotron Radiation Laboratory. X rays from the synchrotron enter the cryostat through the beamline on the left and strike the sample, producing x-ray fluorescence, which is measured by the STJ detector.

superconductor, can tunnel through the weak insulating barrier of the Josephson junction, producing a pulse of electrical current. By measuring the number of Cooper pairs that are broken, scientists can determine the energy of the x ray up to ten times better than with conventional technology and can identify the material that emitted the x ray. These superconducting-tunnel-junction (STJ) detectors also work with optical, ultraviolet gamma-ray photons and large biomolecules. Labov and his team are working to use this new technology in applications for analyzing all of these particles.

Measuring Large, Slow Molecules

The Livermore group has, for example, teamed with scientists at Lawrence Berkeley National Laboratory and a commercial firm, Conductus Inc., of Sunnyvale, California,



X-ray fluorescence measured with a superconducting-tunnel-junction (STJ) detector shows ten times better energy resolution than can be achieved with a conventional silicon detector. A titanium nitride sample was irradiated by 600-electron-volt (eV) x-ray photons at the Stanford Synchrotron Radiation Laboratory. The STJ detector cleanly separates the titanium L lines from the nitrogen K line and the background oxygen K line.

to measure massive, slow-moving macromolecules in DNA research. In a typical time-of-flight mass spectrometer using a microchannel-plate (MCP) ion detector, large ions move too slowly to be efficiently detected. Using an STJ detector, the team found that they could achieve close to 100% detection efficiency for all ions, including the slow, massive macromolecules. “A comparison of count rates obtained with both detectors indicated a hundred to a thousand times higher detection efficiency per unit area for the STJ detector at 66,000 atomic-mass units,” Labov says. “For higher molecular masses, we expect an even higher relative efficiency for cryogenic detectors because MCPs show a rapid decline in detection efficiency as ion mass increases.”

Even more exciting, STJ detectors can measure independently the mass and charge of the molecule. Current MCP detector technology cannot measure the charge of the molecule, and this inability often causes confusion in interpreting mass spectrometer data. According to Labov, if nonfragmenting ionization techniques can be perfected, cryogenic detectors could make possible the rapid analysis of large DNA molecules for the Human Genome Project and might be used to analyze intact microorganisms to identify viruses or biological weapons materials.

High Resolution for Soft X Rays

In another experiment using an STJ, Labov again teamed with Conductus and seven other Lawrence Livermore scientists to study energy resolution for soft x rays with energies between 70 and 700 electron volts. The results showed that STJ detectors can operate at count rates approaching those of semiconductor detectors while still

providing significant improvement in energy resolution for soft x rays. “In this region, the STJ detector provides about ten times better resolution,” Labov adds.

Astronomers also are looking to STJs as single-photon detectors of both x rays and visible wavelengths. In the visible band, silicon-based, charge-coupled devices cannot measure a photon’s energy, but STJs can. One photon, depending on its energy, can generate thousands of quasi-particles. By measuring the photon’s energy, STJ detectors will allow astronomers to study galaxies and stars that are barely bright enough to be seen with the largest telescopes.

Detecting Impurities as Semiconductors Shrink

As semiconductor devices continue to shrink, the industry needs to detect and identify small amounts of contamination on the devices. Microanalysis systems with conventional energy-dispersive spectrometers “excite” contamination on chips with fairly high (10-kiloelectron-volt) energy, which results in the surrounding material also being excited. When the surrounding material is excited, a flood of unwanted signals or noise is created, making it impossible to detect the contamination. But STJ detectors can operate with excitation energies of less than 2 kiloelectron volts, which produce signals from the contamination only, allowing the imperfections to be detected.

Helping to Enforce Nonproliferation

One of the Laboratory’s important missions is to help guard against the proliferation of nuclear weapons. Labov and his team are conducting a research and development project that involves using a superconducting tantalum detector to improve gamma-ray resolution. This technology provides better

diagnostic capability, particularly when there are large amounts of one isotope and small amounts of another. For example, when small quantities of nuclear materials are present, most of the gamma rays detected will be from background sources. Conventional detectors aren't sensitive enough to distinguish clearly between gamma radiation from the background source and from the nuclear material.

The team's high-resolution, superconducting spectrometer can detect special nuclear materials by isolating emissions from different radioisotopes. For example, if an inspector suspected that a heavily shielded barrel of spent plutonium from a reactor plant also contained weapons-grade plutonium, the superconducting spectrometer can measure the composition of the materials in the barrel much more accurately than a conventional detector. The technology also holds promise in environmental monitoring for the analysis of trace contaminants because it can detect levels that conventional detectors would miss.

Looking toward the Future

Traditional energy-dispersive and wavelength-dispersive spectrometers are fully developed technologies, leaving little

room for significant performance improvements. Cryogenic detectors are still in a developmental stage, with significant progress having been made over the past few years. STJ detectors, although an "old" concept, are now better able to resolve low-energy x rays without sacrificing count-rate capability, and the x-ray collection efficiency of these detectors can be increased by orders of magnitude with focusing x-ray optics, which concentrate the x rays on the detector. These developments could greatly increase the use of these detectors in a wide range of applications.

—*Sam Hunter*

Key Words: atomic spectroscopy, Cooper pairs, detectors, Josephson junction, mass spectrometry, quasi-particles.

For further information contact

Simon Labov (925) 423-3818 (labov1@llnl.gov).

Each month in this space we report on the patents issued to and/or the awards received by Laboratory employees. Our goal is to showcase the distinguished scientific and technical achievements of our employees as well as to indicate the scale and scope of the work done at the Laboratory.

Patents

Patent issued to	Patent title, number, and date of issue	Summary of disclosure
Bernard T. Merritt Gary R. Dreifuert Richard R. Post	Halbach Array DC Motor/Generator U.S. Patent 5,705,902 January 6, 1998	This direct-current (DC) motor/generator based on a Halbach array of permanent magnets does not use ferrous materials; thus, the only losses are winding losses and losses due to bearings and windage. The rotor is on the outside of the machine, and the stator forms the inside of the machine. The rotor contains an array of permanent magnets that provide a uniform field. The windings of the motor are placed in or on the stator. The stator windings are switched or commuted to provide a DC motor/generator.
George E. Vogtlin Bernard T. Merritt Mark C. Hsiao P. Henrik Wallman Bernardino M. Penetrante	Plasma-Assisted Catalytic Reduction System U.S. Patent 5,711,147 January 27, 1998	This inexpensive and reliable system for reducing NO _x emissions, particularly in gasoline-powered vehicles and engines with oxygen-rich exhaust, combines nonthermal plasma gas treatment with selective catalytic reduction. A first reactor converts NO to NO ₂ , and a second reactor reduces the output of the first to produce an exhaust with reduced NO _x emissions.

Awards

David Leary, head of the Laboratory's Business Services Department, received a **Hammer Award** for his efforts as part of a Department of Energy team to streamline property management regulations and performance evaluation, so that the latter is not regulation-based but outcome-based as it is in private industry. The award was given to a joint DOE/contractor value-based self-assessment team of which Leary was a member. The Hammer awards were created by Vice President Al Gore to recognize special achievements in the efforts to reinvent government by cutting red tape and making government more efficient. They are given to individuals and teams that have made significant contributions in support of the President's National Performance Review principles.

Stan Trost has been named an **Executive Fellow of the Institute of Electrical and Electronics Engineers** and will serve one year with the Federal Communications Commission (FCC) in Washington, D.C. Trost will work in the FCC's Office of Plans and Policy on a broad range of deregulation issues resulting from the Telecommunications Act of 1996. Among the deregulation issues the FCC faces are the Year 2000 problem and critical communication infrastructure, reliability, openness, and service issues related to the Internet. Currently, Trost is a program manager in the Nonproliferation, Arms Control, and International Security Directorate. Previously, he was deputy associate director for electronics engineering and has held a variety of development and leadership positions focusing on information technology.

Computer Simulations in Support of National Security

The Accelerated Strategic Computing Initiative (ASCI) is an ambitious component of the DOE Stockpile Stewardship Program to maintain the safety and reliability of the nation's nuclear weapons. ASCI's goal is to provide the computing power and simulation technology that will allow weapons scientists to support assessment and certification of the nation's nuclear deterrent from their computer desktops. To provide this capability, the ASCI program is developing computer platforms powerful enough for teraflops calculations, writing three-dimensional weapons codes that reflect our understanding of the physics of the entire weapons system, and developing computer infrastructure that can run the codes on the platforms. In the long term, ASCI's achievements, developed to support our national security interests, will change the fundamental way scientists and engineers—and eventually, all of us—solve problems. Because it is a scientific grand challenge with a short schedule, ASCI work must be well planned and well coordinated throughout the three participant national laboratories—Livermore, Los Alamos, and Sandia. It must also take advantage of the world-class expertise of America's universities and ASCI's industrial partners.

Contact:

Randy B. Christensen (925) 423-3054 (christensen5@llnl.gov).

© 1998. The Regents of the University of California. All rights reserved. This document has been authored by the The Regents of the University of California under Contract No. W-7405-Eng-48 with the U.S. Government. To request permission to use any material contained in this document, please submit your request in writing to the Technical Information Department, Publication Services Group, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94551, or to our electronic mail address report-orders@llnl.gov.

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California and shall not be used for advertising or product endorsement purposes.

U.S. Government Printing Office: 1998/683-051-60042

Science & Technology Review
Lawrence Livermore National Laboratory
P.O. Box 808, L-664
Livermore, California 94551



Printed on recycled paper.

Nonprofit Org.
U. S. Postage
PAID
Livermore, CA
Permit No. 154