LOS ALAMOS SCIENCE AND TECHNOLOGY MAGAZINE

JUNE 2012

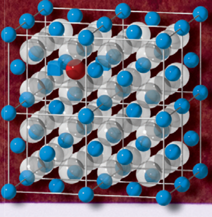

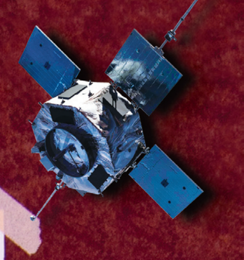

1663

Stormy Space Weather
Modeling Extreme Materials
Antibody Protein Pipeline
Killing Killer Asteroids

Making Sense of Sight

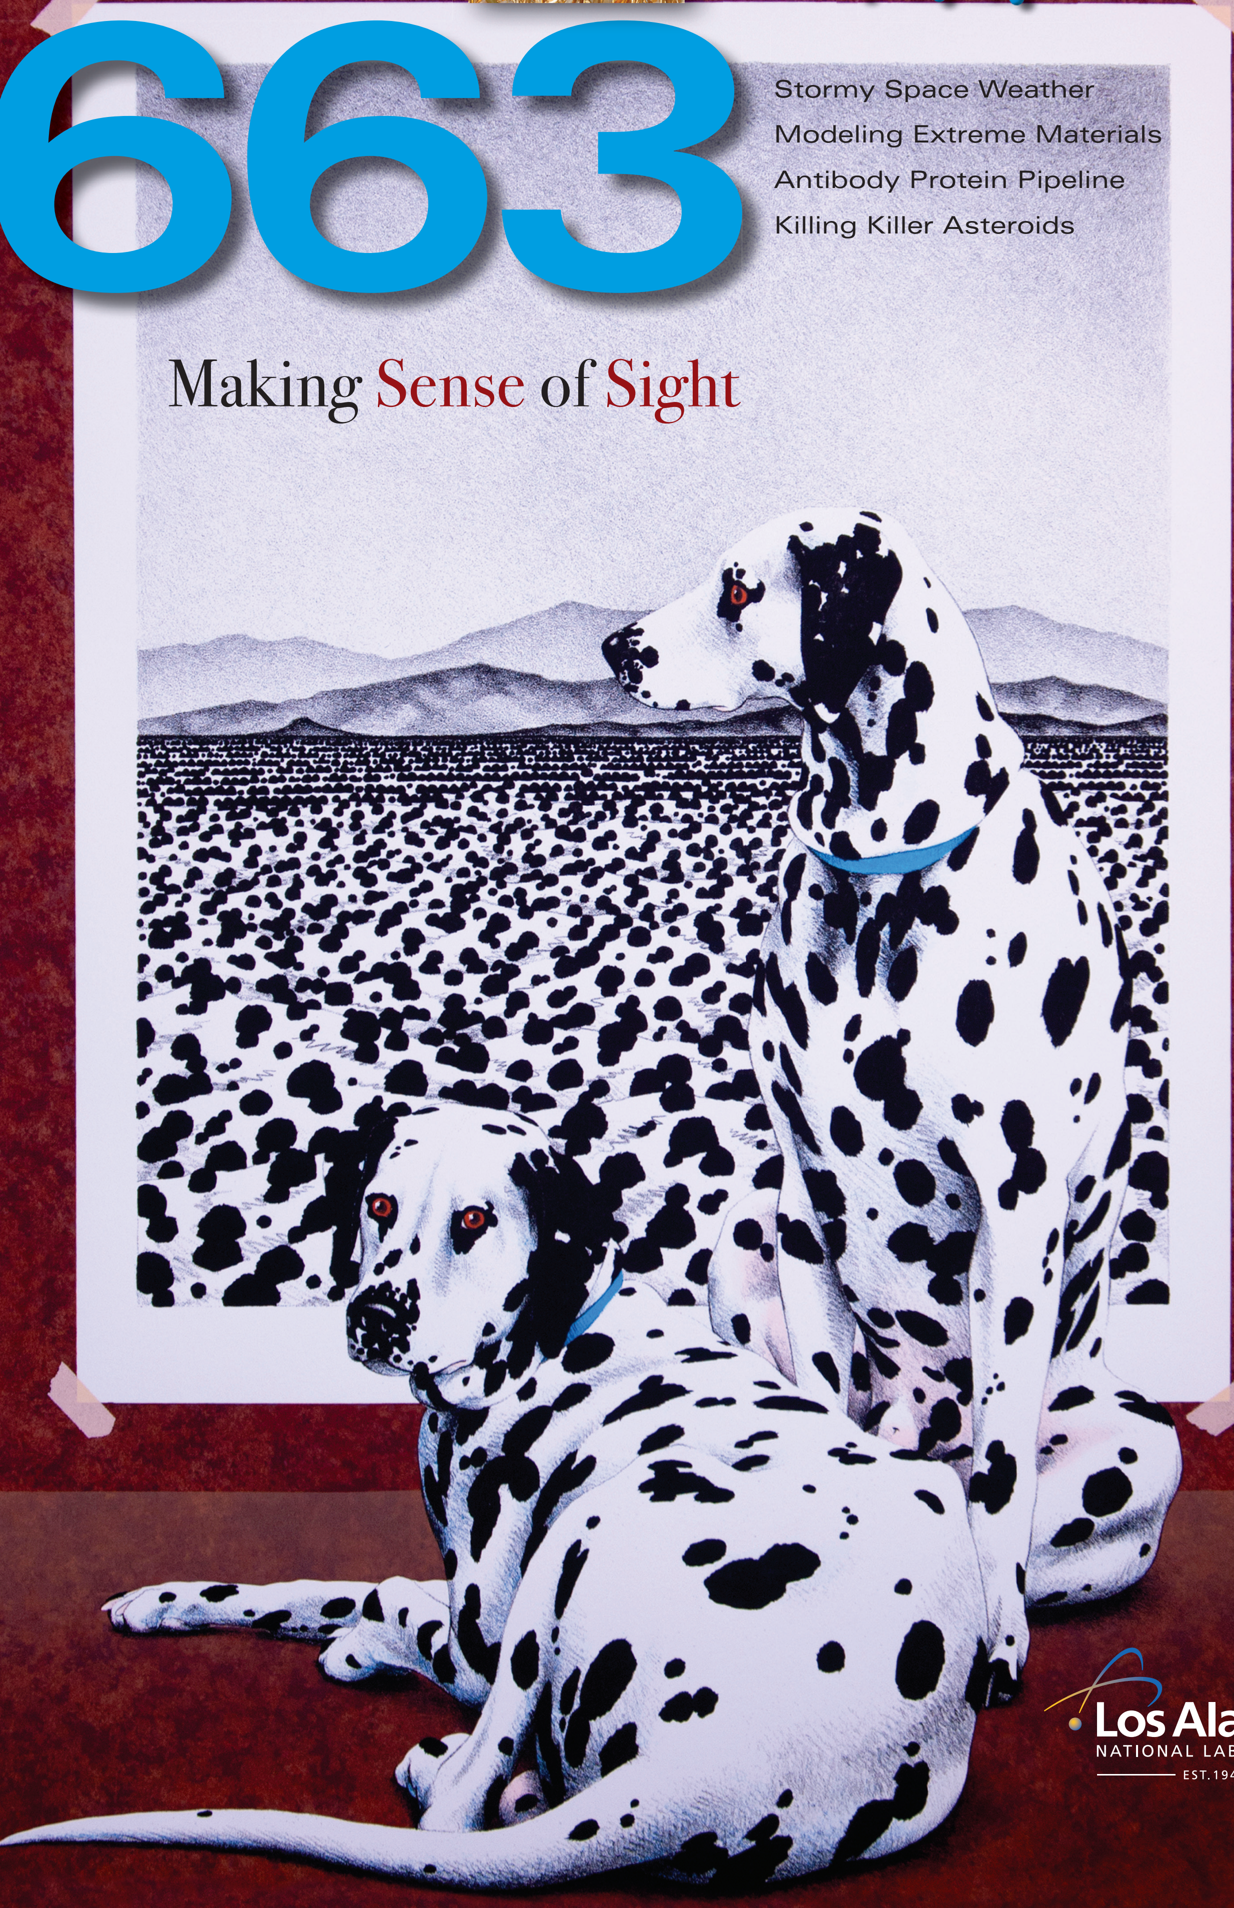

Los Alamos
NATIONAL LABORATORY
EST. 1943

1663

LOS ALAMOS SCIENCE AND TECHNOLOGY MAGAZINE

About Our Name: During World War II, all that the outside world knew of Los Alamos and its top-secret laboratory was the mailing address—P. O. Box 1663, Santa Fe, New Mexico. That box number, still part of our address, symbolizes our historic role in the nation's service.

About the LDRD Logo: Laboratory Directed Research and Development (LDRD) is a competitive, internal program by which Los Alamos National Laboratory is authorized by Congress to invest in research and development that is both highly innovative and vital to our national interests. Whenever *1663* reports on research that received support from LDRD, this logo appears at the end of the article.

About the Cover: The human visual cortex is capable of interpreting complex and confusing scenes in a small fraction of a second. Few people would fail to distinguish the two dogs from the background in the cover artwork, whereas few computers would be up to the task. Los Alamos scientists are developing new methods to understand how the brain processes its visual inputs—an important step toward developing artificial (robotic) eyesight and understanding how the brain performs any of the remarkably complex tasks it seems to handle effortlessly every day. *Two Blue Collars*, a serigraph by Dick Mason, is reprinted here with permission from the Windsor Betts gallery in Santa Fe, NM.

Los Alamos Firsts

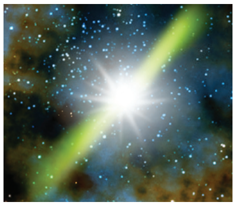

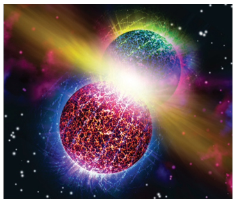

Discovery of Gamma-Ray Bursts

For decades, Los Alamos National Laboratory has developed satellite technologies for space exploration and national security purposes. The Laboratory's first foray into satellite development started in 1959, less than two years after the Soviet Union's launch of Sputnik. The project was dubbed Vela, which, when translated from the Spanish word *velador*, refers to someone who vigilantly watches over something. Indeed, the first Vela satellites were launched in late 1963 to watch over the Earth and ensure compliance with the Limited Test Ban Treaty against above-ground nuclear tests. Pictured below, instrument scientist Richard Belian installs protective covers on sensitive surfaces just prior to launch.

Despite their non-astronomical mission, in 1967, Vela satellites recorded data that would reveal an important astronomical discovery—the gamma-ray burst (GRB). For decades, the origin of GRBs remained a mystery, and some peculiar GRBs still puzzle scientists today (see page 23). However, most GRBs can be assigned into one of two categories: One is the death of an extremely massive, fast-spinning star; as its core collapses to a black hole, a jet of energetic particles produces the burst (left top). The other is a collision between a neutron star and either another neutron star or a black hole (left bottom).

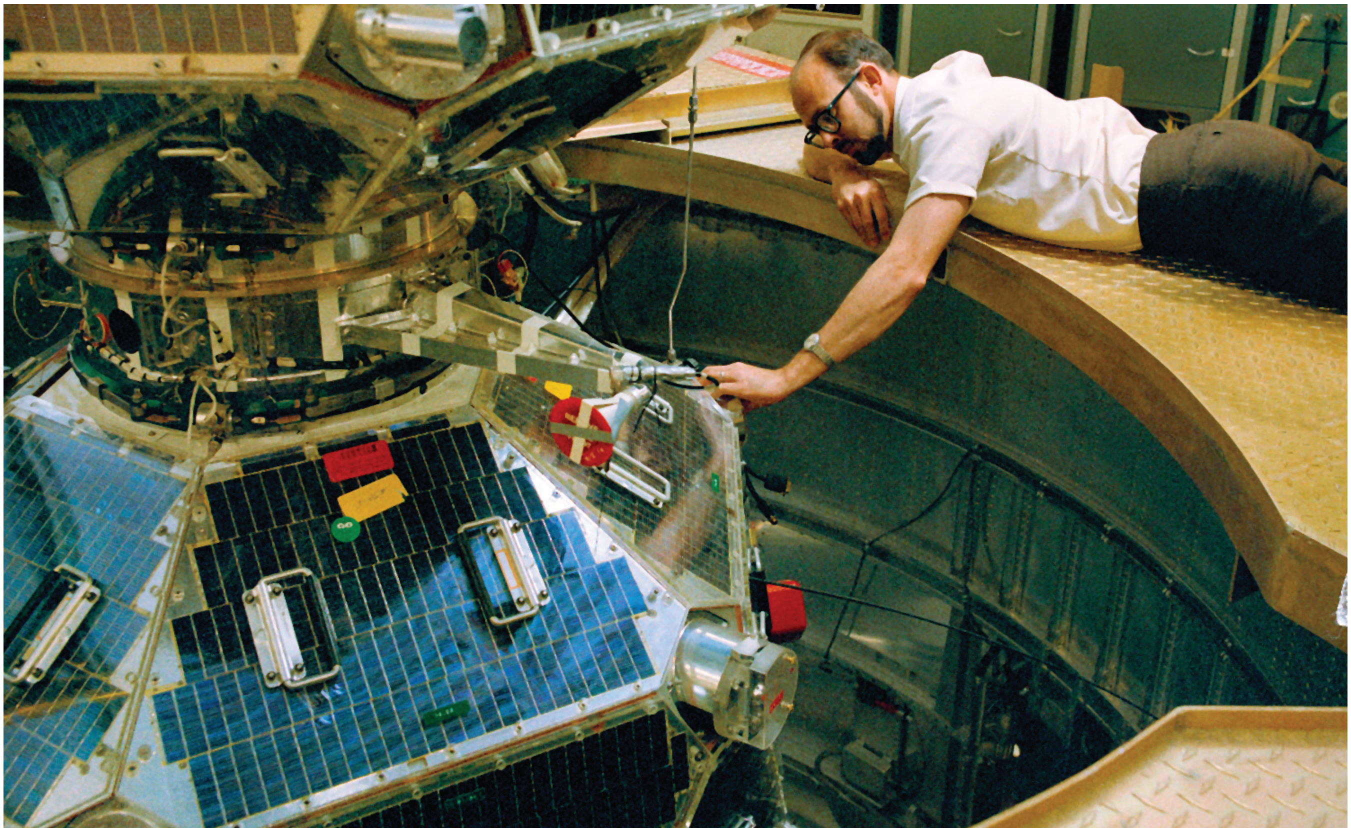

IN THIS ISSUE

FEATURES

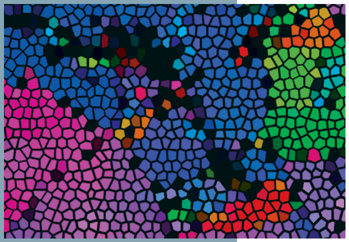

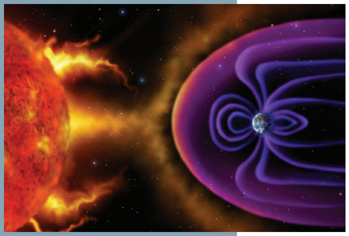

SPOTLIGHT

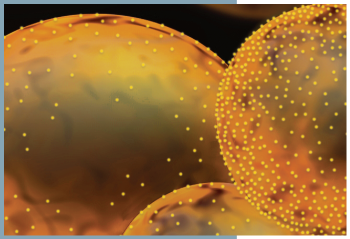

The Mind's Eye

Los Alamos researchers are at the forefront of a revolution in experimental brain science

A real brainstorm: In order to examine a visual scene, the brain parses signals from the eyes into components, such as edges, shapes, colors, sizes, locations, recognized objects, and motion. It accomplishes this analysis, even for complex scenes like those shown on the following page, with far greater speed and accuracy than a computer. How exactly the brain does this is not yet known, but Los Alamos scientists are pioneering the experimental techniques that may provide the answers.

TWO BLUE COLLARS BY DICK MASON REPRINTED WITH PERMISSION FROM THE WINDSOR BETTS GALLERY

Consider the human brain as a computer. It is an electrical signaling system capable of carrying out mathematical and logical operations. It has short-term and long-term memory. It exchanges inputs and outputs with external devices, like ears and arms. Estimates of human brain performance vary widely because no direct method of comparison to a computer is known, but based on the brain's hardware and architecture, some experts peg its computing power roughly on par with the world's fastest supercomputer.

Yet it is clear that the brain is not like a computer. For one thing, humans are notoriously bad at arithmetic. Even humans who excel at arithmetic are bad at arithmetic when compared to even the most limited calculators and computers. But while humans lose every "mathletic" contest hands-down, they utterly obliterate the electronic competition when it comes to more sophisticated tasks, such as recognizing other people—even when seen from different angles or illuminated by different light sources—and reading their emotional states from the subtle variations of their facial muscles. And while a supercomputer might store more bytes or achieve more operations per second, it takes up an entire room and consumes enough electricity to power thousands of homes. The human brain, on the other hand, fits neatly between the ears and runs on chicken and broccoli.

Researchers at Los Alamos National Laboratory and elsewhere have been striving to program a computer to accomplish complex tasks as well as a human being, with only limited success. One reason for the difficulty stems from the significant technical differences between a brain and a computer. For human beings, there is no rigid distinction between processors and memory chips (the same neurons are both), nor is there even a simple distinction between hardware and software. In addition, the basic processing unit of the brain, the synapse, is substantially more complex than a computer chip's transistor. The brain does amazing things, but it's not yet clear how its organization contributes to its success.




Teaching a computer to make sense of complex visual scenes like these may depend on first learning how humans do it.

One such amazing thing is the brain's ability to understand what the eyes see. In a fraction of a second, a person can recognize any of the tens of thousands of objects he or she frequently sees in the world, regardless of how each object appears in the scene. It could be a sycamore tree in the fog or a particular type of pen in a messy supply closet. "Humans doing object recognition are essentially flawless," says cognitive psychologist Amy Guthormsen, part of a team of researchers at Los Alamos led by John George and Garrett Kenyon trying to reveal how the human visual system works. The team's ingeniously programmed, state-of-the-art computer model for human vision, she says, "scores a B+ at best."

Taking In the View

What little is currently understood about human vision goes like this: The rods and cones in the retina respond to the intensities and colors of light entering the eye from each direction within the field of view. This information is then transmitted, like a video stream, to the thalamus near the center of the brain. The retina and the thalamus (in that order) each perform some minimal processing of the visual data; for example, the retina identifies regions of contrast and relative amounts of color within the scene. After acquiring such preliminary information, the data stream is relayed to the visual cortex at the rear of the brain.

The visual cortex is organized into several component regions believed to process visual data in a hierarchical fashion—with more complex information extracted at higher levels of the visual cortex (see upper figure on page 5). The lowest level, called V1, seems to extract some basic information about edges, orientations, and motion. Up a level to V2, you get some other edge feature and color information. At V3 and V5 you get more insight about motion, and at V4 you get simple shape recognition. It is clear that the human visual system shows compartmentalization: different regions do somewhat different tasks, generally in order of increasing complexity.

Interestingly, what seems to be missing entirely from the brain's image processing system is an actual image. "There's no JPEG file in the brain," explains Michael Ham, a Los Alamos physicist who studies computer vision. "It's not as though the brain forms an image and shows it to some kind of mental processor for analysis; the brain extracts different pieces of information from the visual data without ever assembling it into an image." Indeed, this may be why it has been so difficult to design an artificial (robotic) visual system to mimic the human one: With computer image processing, you start with a still image and try to identify its components. But with human visual processing, you begin dissecting and reorganizing the data stream before it even leaves the eye.

Zhengping Ji, of the Laboratory's applied mathematics group, also works on computer modeling of human vision. Using results from human-subject experiments run by Guthormsen and others, Ji structures a computer model to process information in the complicated manner seemingly employed by the human visual cortex. This has allowed the model to outperform earlier models, but it still can't compete with an actual human being. The problem is partly a lack of knowledge about exactly how (and how often) the various regions like V1 and V2 communicate with one another and within themselves. For example, some activity in V1 has been observed to occur after other activity in V2, implying feedback. But it's not clear what exactly is accomplished by this feedback or how each region contributes.

Another poorly understood aspect of human vision is the higher-level processing that takes place further up the hierarchy of the visual cortex. Starting after V1 at the back of the brain, visual signal processing splits along two main pathways through the brain. The dorsal pathway runs forward along the top of the brain's surface (broadly called the cortex), while the ventral pathway runs forward along the bottom of the cortex. The upper pathway terminates at the posterior parietal (PP) lobe, while the lower pathway terminates at the inferior temporal (IT) lobe. Each pathway appears to serve a different purpose. The PP seems to determine the position and size of objects in the field of view. The IT, on the other hand, identifies the objects themselves. The PP might say, for example, "On top of the lamppost," and the IT would say, "is a black crow." (PP: "A big one, too.") How the PP and IT do their jobs is not yet known, nor is it known how they communicate with the lower levels of the visual cortex. But Los Alamos scientists might know how to find out.

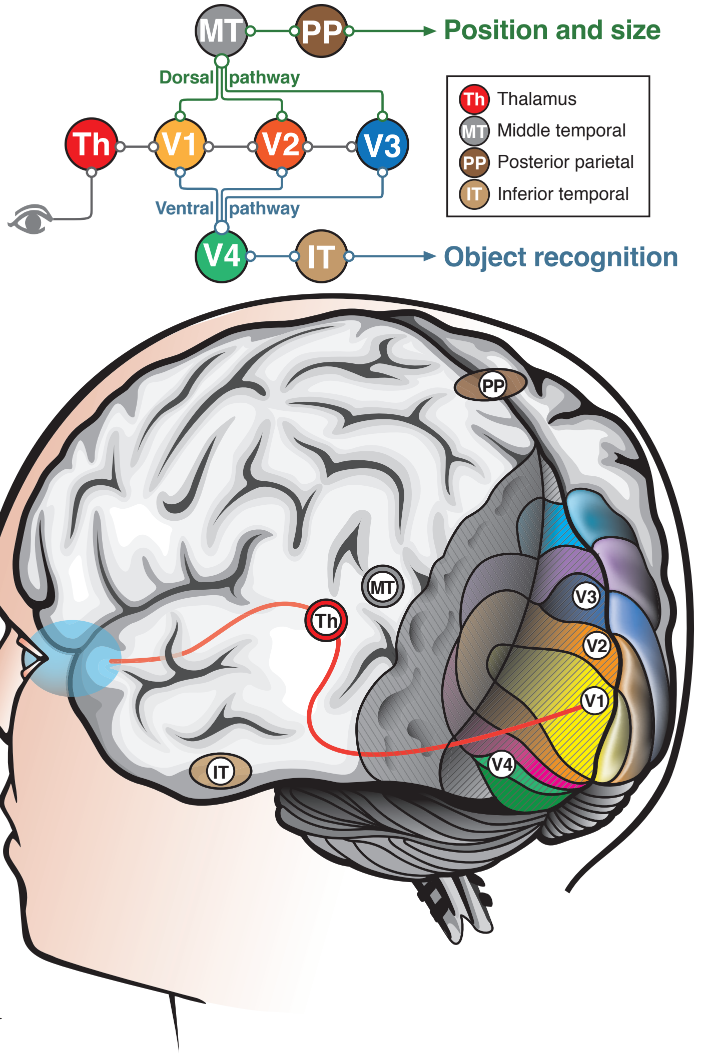

Once a visual signal leaves the eyes and is relayed by the thalamus to the visual cortex, its progression divides into two major pathways. Along the dorsal (upper) pathway, successively higher-level processing leads to the posterior parietal (PP) lobe, from which emerges detailed size and position information about objects in the visual scene. Along the ventral (lower) pathway, the inferior temporal (IT) lobe ultimately identifies what the objects in the scene are. A variety of intermediate steps also help dissect the scene—its edges, colors, and motion, for example.

EEG, MEG, MRI, Oh My!

The most tried-and-true method for observing the brain in action is functional magnetic resonance imaging (MRI). This type of brain scan uses strong magnetic fields to probe for oxygenated hemoglobin. The logic goes, when part of the brain is in use, it requires more oxygenated blood, causing the MRI scan to take notice. The scan can resolve the location of the extra oxygenated blood in the brain within a few millimeters in any direction.

But while the MRI's spatial resolution ("where") is excellent, its temporal resolution ("when") is poor. Because it takes 2–5 seconds for the body to supply the extra oxygenated blood once some part of the brain has "requested" it, the MRI can only tell what parts of the brain were in use 2–5 seconds ago, and therefore it is most useful for studying brain tasks that last at least that long. Visual processing, however, happens much faster than that. Complex object recognition takes less than half a second, and more "primitive" tasks, such as triggering the response to duck when something not-yet-identified is coming toward your head, are virtually instantaneous.

Fortunately, there are other types of brain scans with better temporal resolution than the MRI: electroencephalography (EEG) and magnetoencephalography (MEG). When a current flows in a circuit—or charged ions flow in a neuron—an electromagnetic signal is produced, and EEG and MEG scans pick up different parts of that signal. The strength of the signal depends on the angle between the detector and the actual "wire" (neurons are long and thin, like wires). Due to the geometry of the cortex, neurons on the smooth part of the brain's surface show up better with EEG, while the neurons located within the brain's folds show up better with MEG.

Both EEG and MEG have the advantage of capturing actual electrical activity in the brain, rather than using oxygenated blood as a proxy for it. Both employ an arrangement of sensors on the head, and both have excellent temporal resolution, allowing them to determine the timing of field changes down to about a millisecond. A drawback of EEG is

Los Alamos's computer model for visual object recognition has as much processing power as the human visual cortex and knows how to recognize certain objects. But the model underperforms a human being in both speed and accuracy. Boxes in this image indicate where the model successfully identified the features of a vehicle.

that materials in the human head block electric fields to varying degrees, and any effort to compensate requires an approximation. The same is not true for magnetic fields, making MEG signals less ambiguous, but the equipment necessary for performing MEG is sufficiently specialized and expensive that its use is restricted to larger laboratories, including Los Alamos.

Although both EEG and MEG have sufficient temporal resolution to study what happens when, they both suffer from the same serious flaw: they can't tell where the signals originate. The same EEG or MEG signal can be produced at a particular sensor location by many different combinations of neurons firing all over the brain, making it impossible to uniquely identify the region or regions of the brain responsible for the combined signal. What's needed is a way to obtain the spatial resolution of the MRI with the temporal resolution of an EEG or MEG (or both).

The brute force approach would be to measure both at the same time. Unfortunately, the MRI uses such powerful magnetic fields that it completely drowns out the subtle signals measured by EEG or MEG. Nonetheless, a team led by Michelle Espy of the Lab's Applied Modern Physics group is close to taking simultaneous readings from an MEG during an MRI scan. If Espy succeeds, it should be possible to sample what the different parts of the brain are doing in real time, in the visual cortex or anywhere else. Then it will be necessary to develop a technique for blending the two different types of data into a coherent picture of how signals are shuffled around in the brain.

"We don't have to wait for simultaneous location and timing data," Guthormsen points out. "As long as we know in detail what parts of the brain carry out some particular mental operation, we can blend MEG and MRI data today." That's where the choice to study vision comes in. The retina sends specific parts of a visual scene to specific parts of the visual cortex. Thus, if Guthormsen shows a test subject a series of lights appearing in different parts of the subject's field of view during an MRI examination (allowing 2–5 seconds each time), the MRI will determine where in the brain those lights are processed. This allows the construction a retinotopic map, which shows the parts of the field of view that are directed from the retina to specific parts of the V1 cortex (see figure at right). And while one can't work backward from an EEG or MEG signal to locate the neurons involved in producing that signal, one can use the combination of neurons identified in a retinotopic map to project the EEG or MEG signal that should result. Guthormsen needs only show her test subjects visual stimuli in particular parts of the visual field.

This methodology solves a core physics problem associated with localizing MEG or EEG data: that it is not possible to isolate the unique set of firing neurons responsible for generating a measured MEG or EEG signal. However, it is possible to attribute that signal, arising in a combination of sensors around the head, to the visual processing that produces it if the neurons involved have already been identified—in this case, by a retinotopic map.

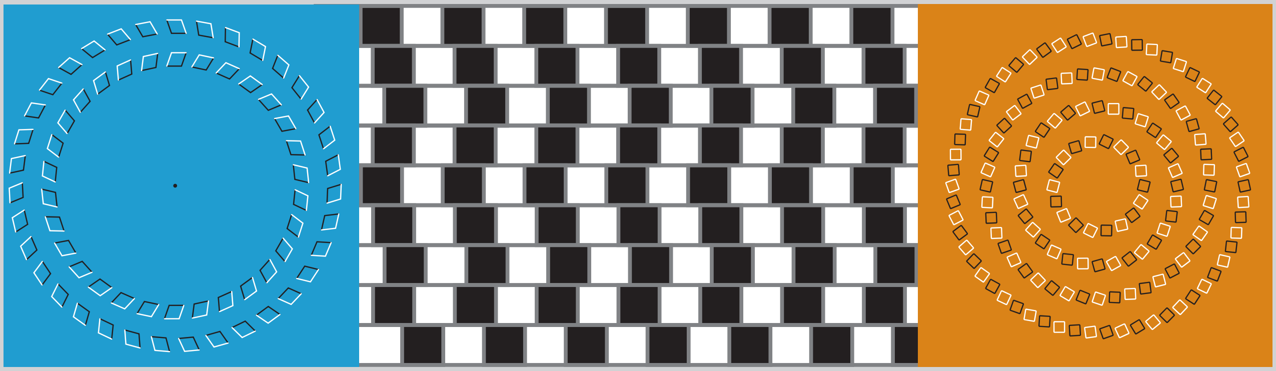

The human visual system builds a model of the world—a best guess as to what real-world scene could have given rise to an observed pattern of data. This is an interpretive rather than algorithmic process. Each of these optical illusions plays upon this distinction by inducing the brain to falsely identify visual elements, such as those associated with 3-D perspective or motion. Left: Watch the central dot as you move your head toward and away from the page. You will see the circular patterns appear to spin. Middle: The horizontal lines appear to alternately converge and diverge, but they are in fact parallel. Right: What appears to be a set of spirals is really just a set of concentric circles.

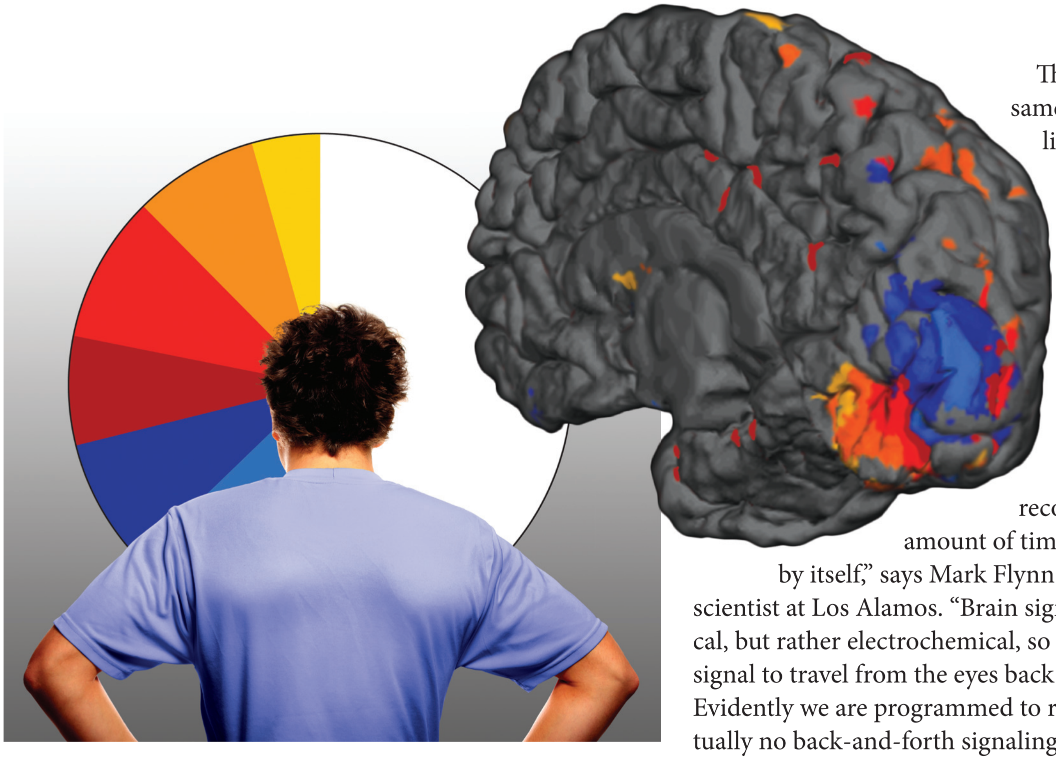

This retinotopic map shows where in the visual cortex at the rear of the brain different parts of a test subject's field of view are processed. The pinwheel pattern shows angular position on the left side of a subject's view. (Radial position, which would appear as concentric circles, is also mapped but is not shown here.) From both eyes, signals pertaining to this left-side view are relayed to the right hemisphere of the brain, shown here. For example, the dark red color indicates that objects seen at eye-level on the subject's left side are being analyzed, in part, at the center of the V1 area of the visual cortex on the right hemisphere.

The Truth about Cats and Dogs

Los Alamos scientists are working to pioneer the merging of MRI (retinotopic map) and electromagnetic (EEG or MEG) data from the visual cortex. If successful, they hope to identify how the brain coordinates information between various parts of the visual cortex to comprehend a scene. But they have already glimpsed the kind of surprising results such an approach can provide in a related experiment.

The team showed a test subject a series of photographs; for each, the subject was instructed to indicate quickly whether he or she saw a "target" animal in the photograph. In some experiments the target was a cat and in others it was a dog, but the photographs included pictures of cats, dogs, and various inanimate objects. An EEG apparatus recorded electrical activity over time at various locations, allowing the construction of waveforms (see figure on page 8) that can be compared to identify at what point the brain's electrical activity begins to differ when processing different images.

The waveforms look the same for the first 150 or so milliseconds (ms); at that point, a waveform associated with observing an animal begins to diverge from one associated with observing an inanimate object. While it's not clear what exactly the brain is doing differently in the two cases, it is clear that some degree of recognition must occur in that amount of time. "That's quite amazing all by itself," says Mark Flynn, a biologist and computer scientist at Los Alamos. "Brain signals are not purely electrical, but rather electrochemical, so it takes about 150 ms for a signal to travel from the eyes back through the visual cortex. Evidently we are programmed to recognize animals with virtually no back-and-forth signaling within the cortex because 150 ms just isn't enough time for it."

On the other hand, the waveform for a cat image doesn't begin to differ from that for a dog image until about 350 ms have elapsed. This suggests that there may be substantial crosstalk needed within the visual cortex to recognize the difference between similar objects (animals in this case). The researchers believe that these kinds of results can help them discriminate between competing theories of how the brain understands vision. Prevailing theories, for instance, have held that the visual cortex processes signals upward along the hierarchy only—from V1 to V2 and straight up each pathway to the PP and IT. That may be adequate, and perhaps necessary, when comparing puppies with inanimate objects. But for more sophisticated comparisons, the extra 200 ms may imply the need for signaling back and forth across and within levels. For example, to distinguish a cat from a dog, the IT may (somehow) request more detail from V2, say, to see if there are any whiskers and from V4, perhaps, to determine the shape of the eyes. The IT may then coordinate the results. For example, finding both whiskers and vertical-sliver eyes, the IT concludes the animal is a cat.

"At this point, we can only speculate about how the visual cortex actually functions," Guthormsen admits. In order to rigorously demonstrate how the cortex communicates across levels and within them, the team needs to combine temporally sensitive waveforms with spatially sensitive MRI data. When MEG sensors obtain signals consistent with neuron activity in the regions prescribed by retinotopic

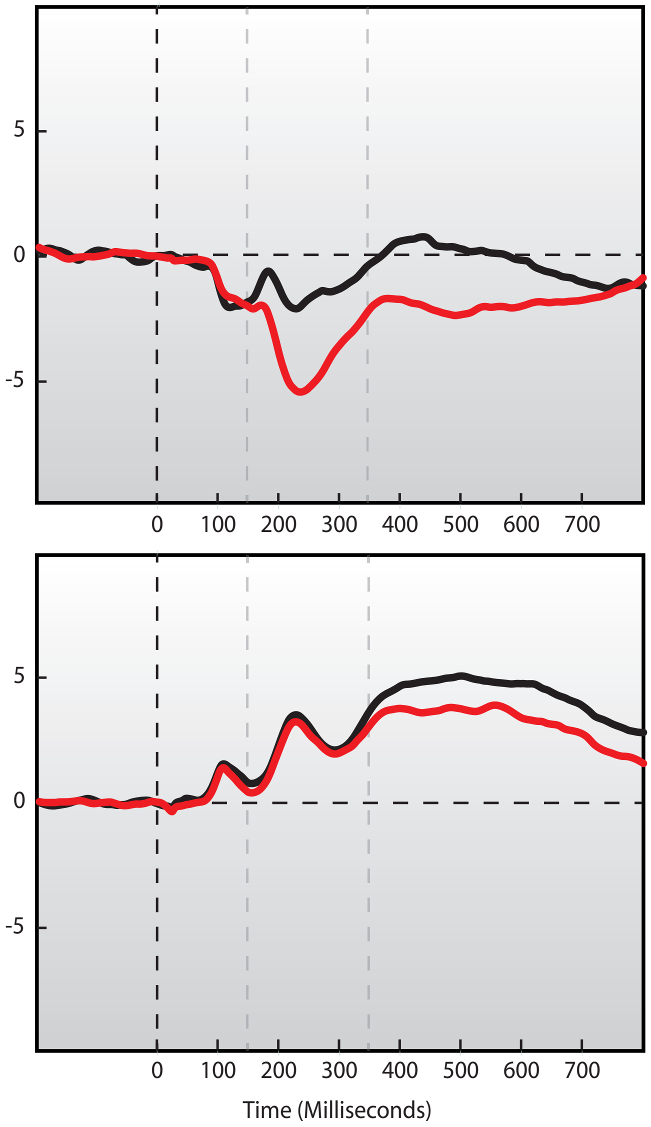

These EEG waveforms show differences in brain activity when people recognize different objects. In each case, the test subject is instructed to distinguish a specified target object (e.g., a dog) from a nontarget object. In the upper frame, the black line resulted from showing a test subject a picture of the target (a dog), and the red line resulted from showing an inanimate object. The two waveforms track one another for about 150 milliseconds—approximately the time needed for a brain signal to travel directly from the eyes into the visual cortex—at which point some recognition that the target and nontarget differ evidently kicks in. In the lower frame, the red line shows the waveform obtained when the test subject is shown a cat as the nontarget object. Due to the similarities between dogs and cats, it takes longer for the brain to recognize the difference and the waveforms track together for about 350 milliseconds, allowing enough time for different parts of the visual cortex to signal back and forth to "compare notes."

mapping, they will have the data they need. "Of course," she says, "obtaining the data is one thing; making sense of it will take some time. But at least now we've shown it can be done."

Doing What Comes Naturally

Presumably, the complicated manner in which the human brain processes visual information is an evolutionary optimization. Some other animals, for example, are particularly good at identifying the shape and motion patterns of their predators and prey; their brains may have organized to maximize these abilities. Because human evolution followed a particular path—walking upright and using arms and hands to manipulate objects, socializing for cooperative benefit and protection, choosing mates based on various visual clues to their genetic quality—it stands to reason that human brains are organized to succeed at these tasks. The visual cortex needs to understand objects and people seen from different angles and in different contexts. Evolution, therefore, needed

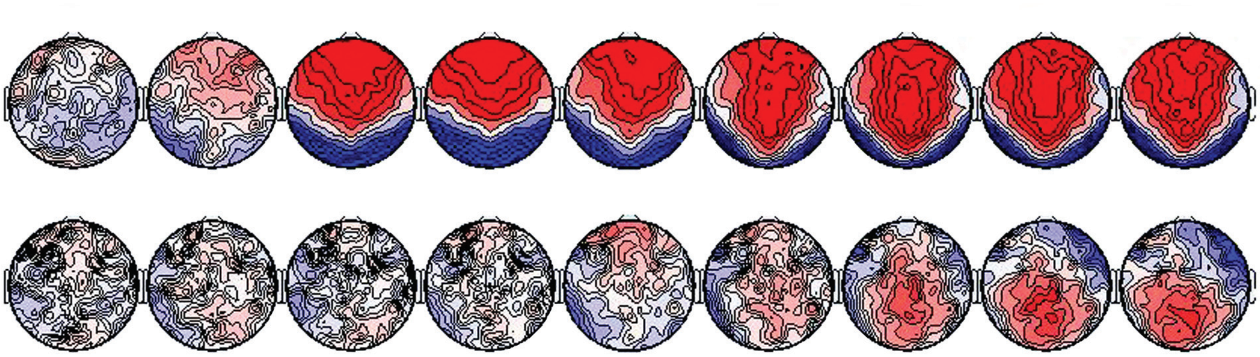

These "scalp plots" show a top view of a test subject's head and indicate the variation in EEG data at different locations (ears and noses are shown for orientation). Left to right, the top row shows the progression of EEG signal differences every 50 milliseconds after showing an image of a target animal versus an inanimate object. Red and blue both indicate brain locations where there are different EEG results for the two cases, while white indicates identical EEG data. The bottom row is the same for target animals versus nontarget animals (dogs vs. cats). Greater spatial detail will be needed to constrain theories of what the brain is actually doing in each case.

Left to right: Michael Ham, Amy Guthormsen, and Mark Flynn pose to demonstrate just how capable the visual cortex is.

to find every possible trick to enhance these abilities because it couldn't house (or supply adequate power to) an entire supercomputer inside a primate's head. The result is the clever, capable, and very complicated brain.

So far, scientists and engineers have been unable to construct an artificial intelligence to match the capability of the brain, so they study the brain in the hope of duplicating its methods. One advantage evolution had over today's researchers, however, is time. Humans spend years in early childhood accumulating information and learning how to understand what their eyes see, while artificial systems are generally expected to function right out of the gate. It may be more fruitful to invent a robot-computer-camera system that can acquire visual sense over years of experience, just as human children do. In the meantime, the Los Alamos team and others in the field think it's wise to try to understand the brain as designed by nature.

If they succeed, the results could be world-changing, allowing robotic systems to attain human-level object recognition capability. This could allow automation of many tasks currently carried out by human labor, and it could lead to new technologies for assisting people with vision disabilities—perhaps eventually including computer and camera elements that link to the brain. And if Los Alamos succeeds in blending MRI and MEG systems to obtain data simultaneously, the benefits need not be limited to artificial vision. Without the need for a retinotopic map to provide the spatial detail, brain researchers could uncover the tricks behind human processing of language, emotion, humor, and so on.

But far from trying to speed the world along toward a cyber-science-fiction future, Guthormsen and her colleagues obtain their daily thrill in the pure science of studying how the brain works. "We find challenge and reward enough," she says, "just trying to uncover how people do the incredibly difficult things they do everyday without appearing to make the slightest effort." ❖ LDRD

—Craig Tyler

A QUINTIL MATTERS

Can Exascale Computing Help Us Understand Extreme Materials?

Some things are difficult to understand—higher math, relationships, the appeal of reality TV—whereas other things are understood to be difficult—brain surgery, two-year olds, learning to speak Finnish. Then there's the response of a material hit by a shock wave, which is not only difficult to understand, but trying to simulate it, even using the world's most powerful computers, is sufficiently difficult that it currently can't be done.

A shock wave is an extremely energetic disturbance that moves through matter at supersonic speeds. Like a flash flood tearing through a slot canyon, it arrives without warning. Matter suddenly finds itself immersed in the wild pressure and temperature maelstrom that trails the wall-like shock front. As the shock propagates through, say, a solid, it generates enormous mechanical stresses that can deform, crack, even shatter the material. Even if there is no structural damage, will the material properties be the same as they were before?

Only select groups of people—demolition experts, makers of body armor, certain types of physicists—know that the answer to that question is "We don't know" and are frustrated by it. But the much larger materials-science community is similarly frustrated by a related problem: the inability to produce the next generation of so-called extreme materials that can survive and function in extreme environments. The core of an advanced nuclear reactor is an extreme environment. So is the radiation-filled vacuum of near-Earth space or any environment where a shock wave comes to visit.

Extreme materials deserve our attention because if researchers could create polymers that withstand high temperatures and pressures, alloys that resist corrosion, or Earth-friendly materials that can tolerate excessive exposure to chemicals, radiation, or electromagnetism, then a bevy of already-thought-of advanced technologies could come off the drawing boards and possibly turn our world into the sustain-

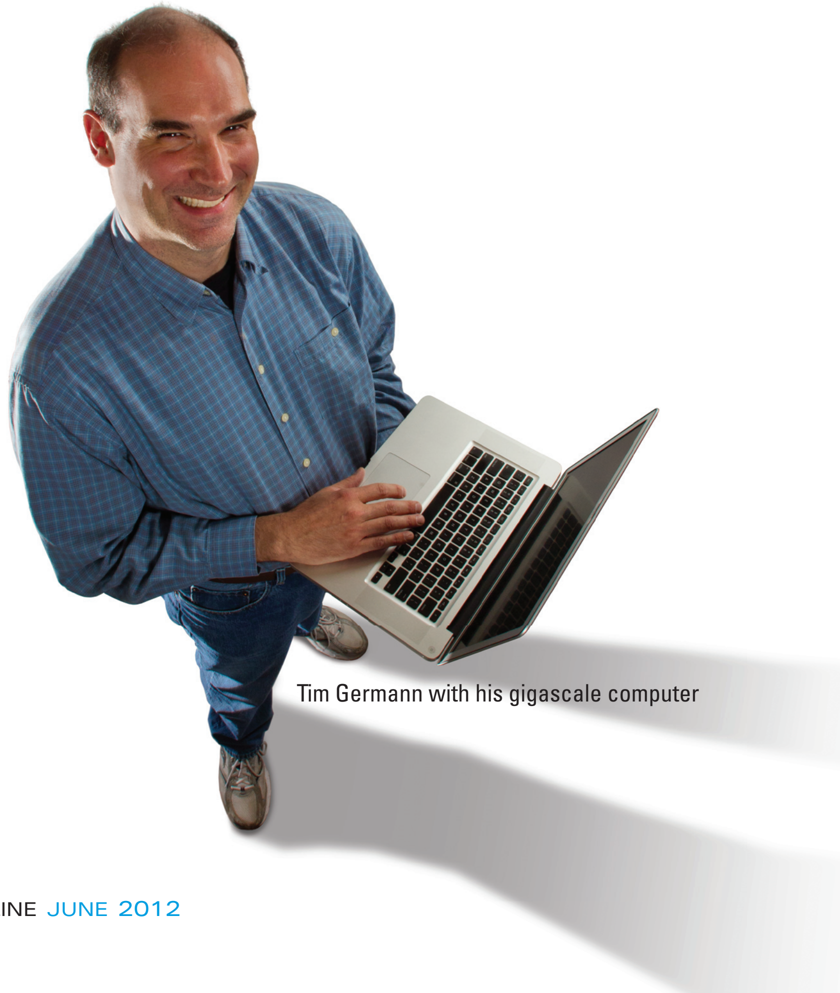

Tim Germann with his gigascale computer

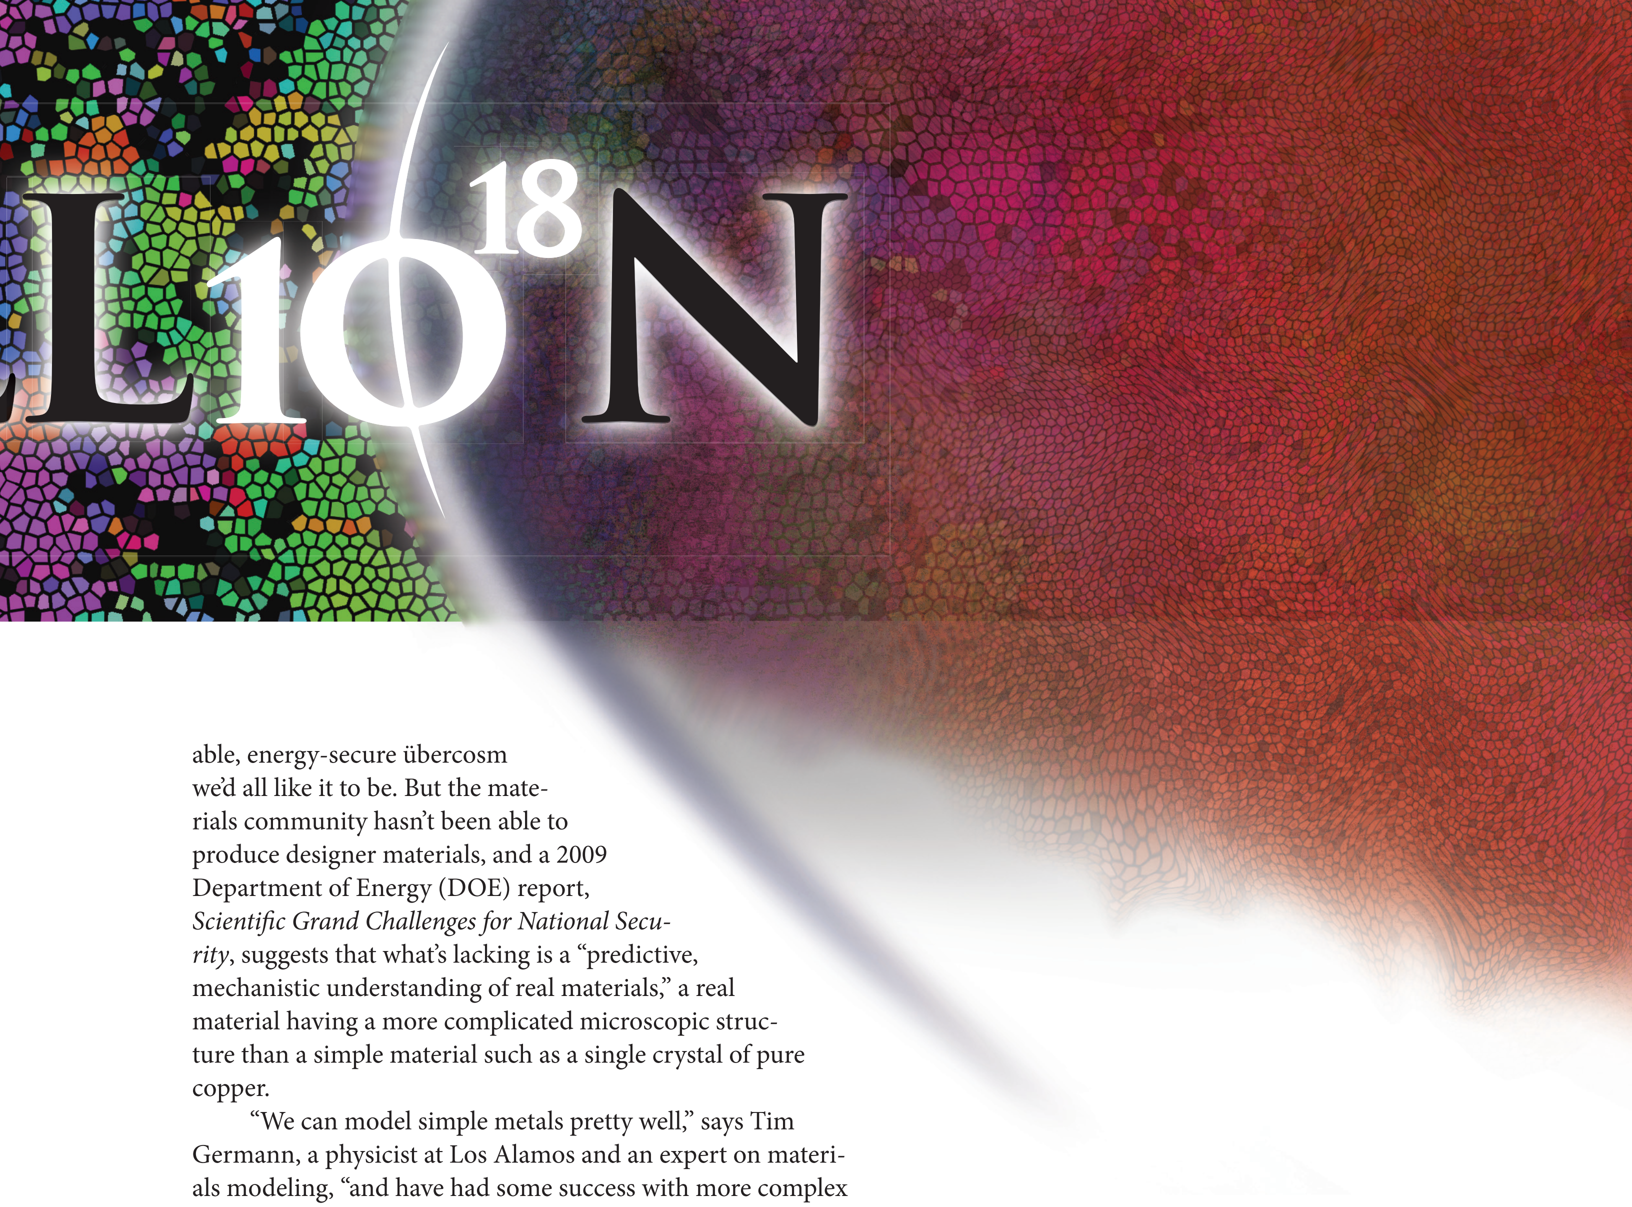

able, energy-secure übercosm we'd all like it to be. But the materials community hasn't been able to produce designer materials, and a 2009 Department of Energy (DOE) report, *Scientific Grand Challenges for National Security*, suggests that what's lacking is a "predictive, mechanistic understanding of real materials," a real material having a more complicated microscopic structure than a simple material such as a single crystal of pure copper.

"We can model simple metals pretty well," says Tim Germann, a physicist at Los Alamos and an expert on materials modeling, "and have had some success with more complex materials. But our ability to predict the properties of real, engineering-scale materials in extreme environments is close to nil."

Extreme materials and shocked matter are of particular interest to scientists at Los Alamos National Laboratory because one of the Laboratory's missions is to ensure the continued safety, reliability, and performance of our nation's nuclear deterrent. It so happens that the performance of a nuclear weapon depends intimately on how its components fare when hit by the shock waves generated inside the detonated device.

After five decades of nuclear tests followed by another two decades of laboratory experiments, computer simulations, and hands-on inspections, weapons scientists know how the weapons in the nuclear arsenal work and how to keep them safe. They know the weapons will perform as expected when triggered properly and won't perform at all when not—devices will not go nuclear if dropped or jarred.

But in the absence of any future nuclear tests, how long can such certainty be maintained? The interior of a nuclear weapon is an extreme environment. The radioactive decay of the nuclear materials produces radiation that changes the internal structure of the weapons components, atom by atom. All of the weapons in the stockpile were originally fielded decades ago, so at what point does the sum of many individually insignificant changes become significant? The answer is not known to any acceptable degree of accuracy, and gaining such knowledge will require the ability to simulate chunks of matter containing perhaps a billion billion atoms, simulations so challenging that they will take an ultra-supercomputer operating at phenomenal speed to do them. That means moving on up to the exascale.

Exa-Size

Exa- is a numerical prefix meaning 10^{18} , as in, "Gee, I'd love to hang with you, Ted, but my to-do list is exalines long." And while modern living has familiarized us with the large (gigahertz, or 10^{9} cycles per second) and the very large (terabytes, or 10^{12} bytes), a factor of 10^{18} (a quintillion, or 1,000,000,000,000,000,000) is unlike anything we have previously encountered. A quintillion M&M candies laid end-to-end would form a line a light-year long; laying that line down, one candy per second, would take nearly 32 billion years—more than twice the age of the universe—and so on.

An exascale computer would execute an astonishing quintillion floating-point operations per second (10^{18} flops or 1 exaflops, with flops being the standard unit for measuring computing prowess). That would make it about a thousand times more powerful than Los Alamos's Roadrunner supercomputer, which is currently the 10th most powerful supercomputer in the world. The huge thousand-fold upgrade in computational power might be enough to make predictive simulations a reality, but achieving that upgrade won't be easy.

An exaflops can't be reached by simply adding more parallel computing branches to a Roadrunner-like supercomputer. Roadrunner has available about 122,000 processor cores. An exascale computer might have roughly a billion. Its system software, which is responsible for ferrying data between processors and memory and for coordinating pro-

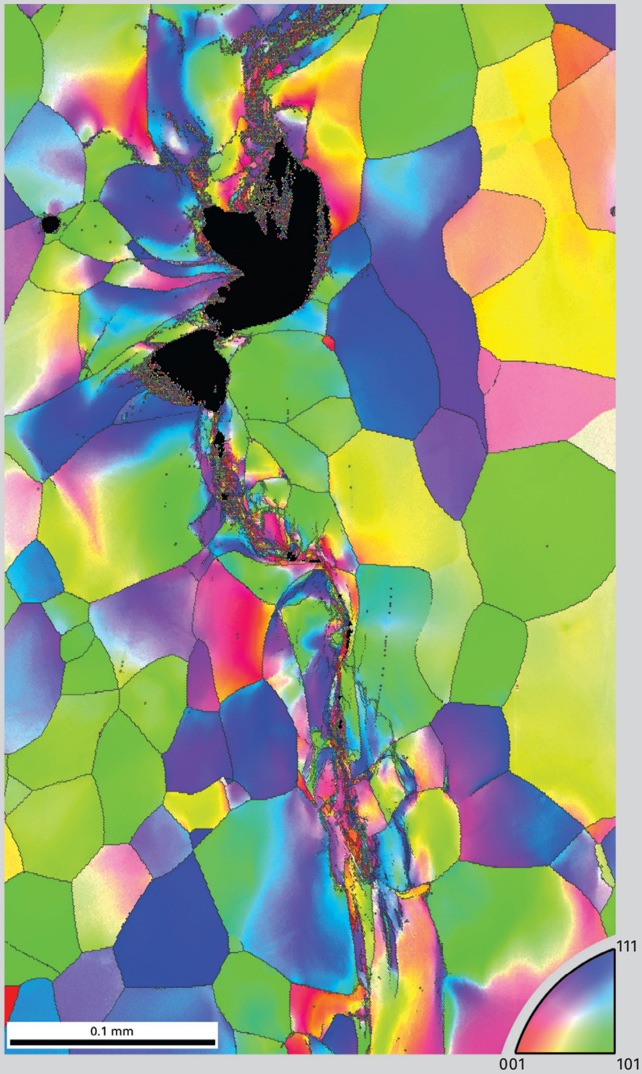

Microstructure

While a material's properties ultimately derive from the interactions of its constituent atoms, many properties are better understood in terms of large groups of atoms, or structures.

For metals, the basic structure is the crystal grain—typically between a thousandth and a millionth of a meter long. Atoms in a grain sit at precise locations within a three-dimensional lattice. How a material responds to external forces depends largely on how each grain responds, which in turn depends on the grain's composition and lattice structure.

Grain properties are therefore sensitive to lattice defects, including missing atoms (vacancies), different atoms (substitutions), and dislocations—line defects where atoms are miss aligned in a different lattice. (You can create a so-called line defect at home simply by mis-buttoning a checkered shirt. The checkered lattice becomes misaligned along a line.)

The interface where grains meet also affects material properties; for example, the atoms at the edge of each grain either line up with each other (so the two grains can stick tightly together) or they don't (so the grains more easily slip apart). Interfaces are an important structural element, as are voids (the absence of material), gaps between grains, or cracks that run between larger-scale domains.

Crystal grains, defects, interfaces, voids, and cracks are collectively known as the material's microstructure. To make predictions about material properties, one needs to know not only what its atoms are doing but also how the material's microstructure influences those properties. That's difficult, mainly because the various components of the microstructure can differ by four or five orders of magnitude in size, and their influences are poorly understood.

Crystal grains (large, colored shapes) dominate the landscape in this microscopic view of a piece of tantalum metal. The angle at which electrons scatter from the prepared surface depends on the orientation of the crystal lattice, thus the colors indicate the grains' orientations. This sample was shocked, causing tiny voids to form and coalesce into larger ones (black regions) and leaving a trail of highly deformed regions.

CREDIT: VERONICA LIVESCU, LANL

cessor activity, would need to be something unworldly, since it would have to integrate 10,000 Roadrunner-sized supercomputers into one machine. (Imagine where your bags would end up if your airline added 10,000 new routes for every existing route in its schedule.)

Another area that doesn't scale well is "resiliency," or the ability of a computer to continue carrying out calculations in the face of system glitches. As inevitable as tax season, the glitches range from almost benign soft errors—mild "hiccups," such as the errant flip of a single bit in memory, which occur frequently but are routinely handled by the aptly named "error-correcting memory"—to hard errors, such as the death of a processing node. The passing of a node is relatively rare, but it requires substitution or replacement.

The one saving grace of hard and soft errors is that the system notices them. It's the so-called silent errors that bring an involuntary pause to Germann's breathing. Silent errors corrupt data or instructions without leaving any indication, and while the odds are very long that such a fate should befall any single processor core, the odds become alarming when there are a billion cores.

Simple scaling of the power needs of today's supercomputers implies an exacomputer would need about a gigawatt of electric power, the output of a large nuclear reactor. There is also the question of how to cool a billion processors and what to do with the heat. But these issues may all be secondary to one looming concern: the price. Metaphorically, that's on the exascale too.

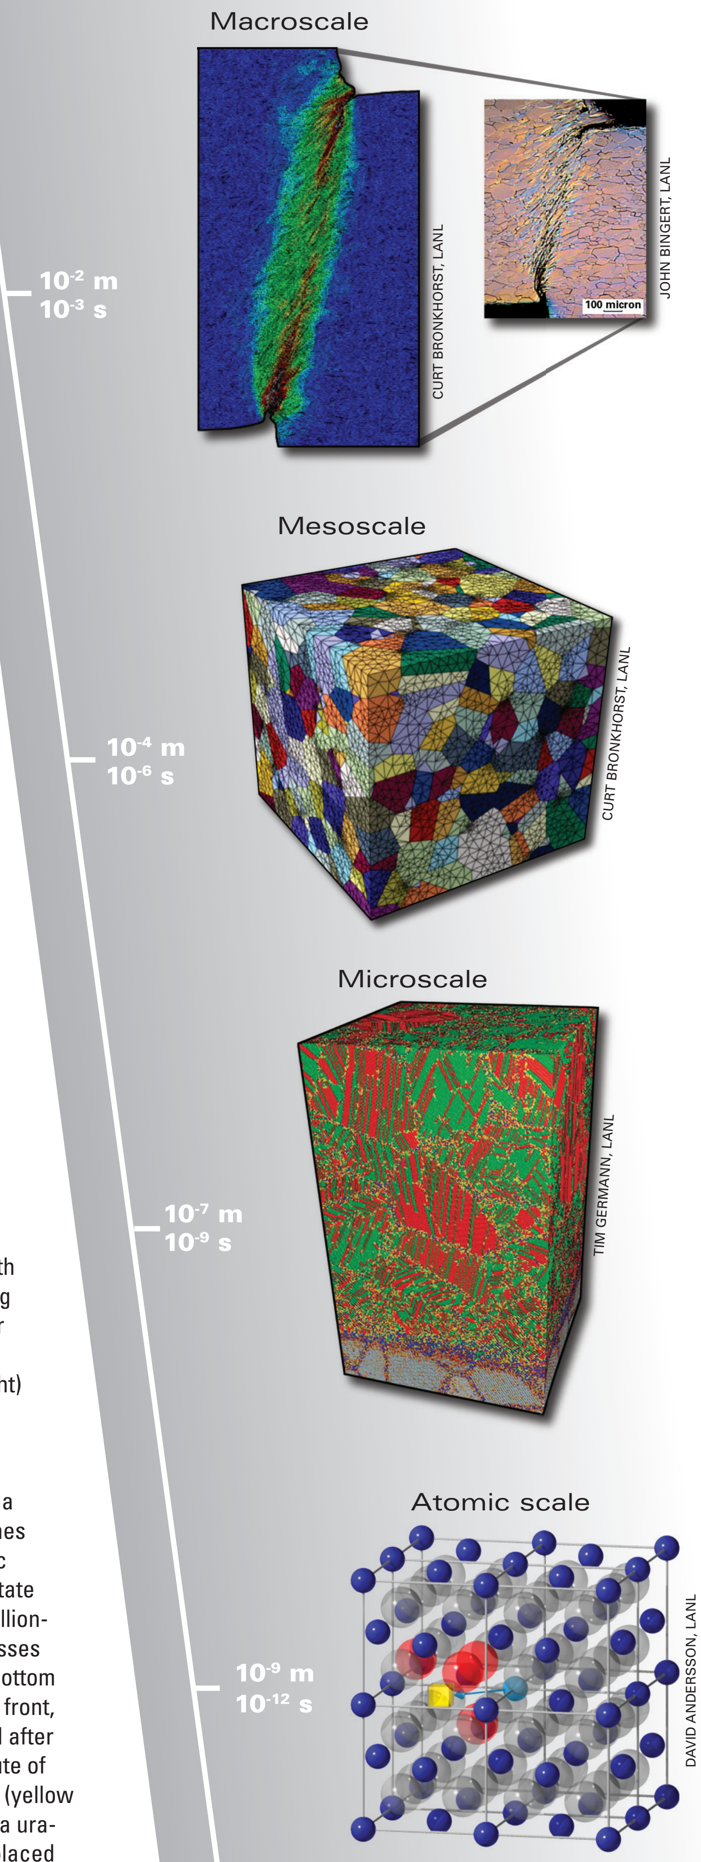

CURT BRONKHORST, LANL

JOHN BINGERT, LANL

CURT BRONKHORST, LANL

TIM GERMANN, LANL

DAVID ANDERSSON, LANL

Simulations often use a multiscale, multiphysics approach for investigating material behaviors. Models are optimized for describing phenomena on one scale—with changes on small length scales taking place within small time scales—and scale-bridging algorithms pass the information to ongoing calculations on other scales. From top to bottom, the illustration at right progresses from macroscopic to atomic length scales. **Macroscale**: (top right) a piece of tantalum metal, sheared, and (top left) a result from a simulation. The simulation used a macroscale continuum model to describe the response outside the shear zone and a detailed polycrystal model within the shear zone. **Mesoscale**: a model of a polycrystalline material. The variably sized triangular mesh defines calculational cells. Mesoscale phenomena would include plastic deformation, wherein the material doesn't return to its original state once the stress is removed. **Microscale**: the output from a 30-million-atom simulation showing the aftermath of a shock wave as it passes through a nanometer-sized piece of iron. The gray band on the bottom is un-shocked matter, the narrow band just above it is the shock front, and the red and green regions are new lattice structures formed after the shock wave passed. Elastic deformation is a common attribute of this scale. **Atomic scale**: uranium oxide with a uranium vacancy (yellow square). Density functional theory methods were used to model a uranium ion migrating into the vacant site. The red spheres are displaced oxygen ions.

Exascale Simulations

Despite the difficulty, many scientists feel that pursuing such a computer will be worth it for economic, social, national competitiveness, and, of course, scientific reasons. The base-level scientific argument revolves around the size of atoms—even large ones are only about 0.3 nanometer across—and the fact that the largest materials simulation to date could only handle about 10 billion of them, equivalent to simulating a cubic chunk of matter barely 300 nanometers on a side.

That's not big enough. To a large extent, a material's properties depend on the details of its internal structure (see "Microstructure" on page 12). The structural features span many length scales, from single-atom vacancies in the material (sub-nanometer scale) to cracks that run through the entire bulk (macroscale). Predicting the material response to external forces means understanding how structures of one length scale respond to these forces and what effect that has on all other scales.

To capture the full range of behaviors of an extreme material, scientists feel they will need to simulate a chunk of matter at least 0.1 millimeter on a side—about the size of a grain of salt—which has a billion times more atoms and would require a similarly large increase in the size of the simulation. Neither Roadrunner nor any other supercomputer has anywhere near the computational resources or memory to handle it. We have to move to the exascale.

Not surprisingly, scaling up a simulation in size is accompanied by a severe increase in its complexity and sophistication. As a simple example, physicists like Germann will construct various models to account for the material response on each length scale. The simulation then uses scale-bridging algorithms to let the different responses influence each other. If a continuum-level constitutive model is used to determine the bulk response, but that model clearly breaks down when the material is severely stressed, the simulation will automatically look to a finer scale and begin to use, say, a model based on the detailed microstructure, or maybe one that uses an atom-by-atom description, to generate a more realistic response.

ExMatEx

With a common mission of settling the exa-frontier, the DOE's Office of Advanced Scientific Computing Research (ASCR) and the National Nuclear Security Administration's Advanced Simulation and Computing (ASC) program are coordinating the United States' effort to achieve a Hulk-like leap in computing power and capability to the exascale. The DOE wants to do this in less than a decade, making it one whopper of a mission.

The DOE's strategy has been to establish numerous "co-design centers," where everyone who has anything to do with solving the problem sits at the table: scientists who develop the physics models, programmers who translate those models into algorithms and who construct a simulation that will run on the exascale machine, computer architects who design the hardware, systems people who establish the infrastructure and network capabilities, experimentalists who gather data, plus data analysts, managers, accountants—anyone needed to make it happen.

Three co-design centers have already been established by ASCR. The Center for Exascale Simulation of Advanced Reactors (CESAR) is headed by Robert Rosner of Argonne National Laboratory. Another is the Center for Exascale Simulation of Combustion in Turbulence (ExaCT) led by Jacqueline Chen of Sandia National Laboratories. The Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx) is headquartered at Los Alamos. Like the other centers, ExMatEx partners with universities, such as Stanford and Caltech, and other national laboratories, including Livermore, Oak Ridge, and Sandia. Headed by Germann, its goal is to create a robust and cost-effective exascale computing environment that would enable research into extreme materials, with an emphasis on understanding shocked materials.

A lot of brilliant people have journeyed into that area without finding a way out, fueling the notion that extreme materials are a scientist's version of a perfect storm: they don't yet have the right physics models, don't have enough data to help guide model development, and they're still limited by computing resources. An exascale computer will do much to quiet that storm.

What will come of the exascale effort waits to be seen. But history consistently shows that with each new material development—think iron or aluminum, Styrofoam or silicon—society advances, sometimes by a little, sometimes by a lot. Here's hoping for a lot. ❖ LDRD

—Jay Schecker

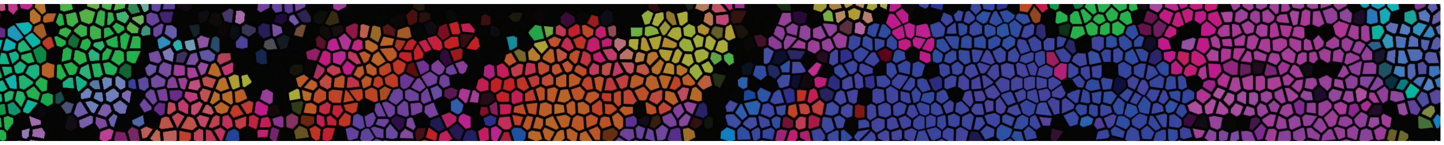

The Stuff That DREAM Is Made Of

Los Alamos physicist Geoffrey Reeves knows that the space just outside of Earth's protective atmosphere is a tempestuous, radiation-filled environment that can knock an orbiting satellite dead. So Reeves and a small team developed DREAM: software that gives satellite operators a heads up about the conditions surrounding their spacecraft and, thus, a chance to prepare for the worst of space-stormy weather.

In 1958, as the United States and the Soviet Union jockeyed for the lead in the newly inaugurated space race, a simple Geiger counter on the first U.S. satellite, Explorer 1, revealed that a belt of high-energy electrons and ions surround the Earth. Trapped by the Earth's magnetic field, the particles girdle the planet in a broad, donut-shaped cloud—the Van Allen radiation belt—and as it happens, nearly every commercial or military satellite flying today orbits either completely or partially within this radiation zone.

That's a problem. A fraction of the charged particles in the belt are relativistic (moving at an appreciable fraction of the speed of light), and relativistic electrons are a potent form of ionizing radiation. The electrons will blaze a trail of ionized atoms within almost any satellite material before losing enough of their energy to effectively come to a stop. They are notorious disrupters of computers and flippers of computer memory bits. The electrons can accumulate within a material, especially a dielectric, until they discharge as a spark. Furthermore, the electrons emit x-rays as they slow down. The x-rays fan out in all directions and effectively widen a single electron's sphere of potential damage to include the entire spacecraft.

But people are clever, and satellite engineers are generally able to counter the effects of an electron assault. They use radiation-hardened computer chips, shield critical electronics, rely on error-correcting software to repair data corrupted by radiation, and install redundant circuitry to compensate for hardware failures. These measures work well under normal circumstances. But then there are the storms.

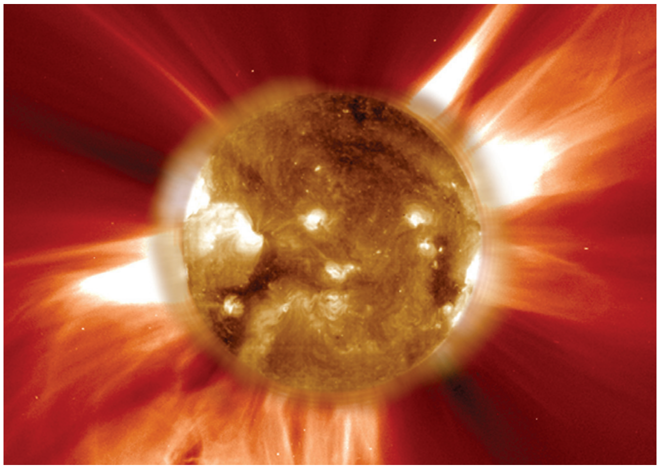

The Sun ejects several billion tons of magnetized plasma into space on close to a daily basis, and when the plasma hits the Earth's magnetic field, it can result in any number of dramatic events. For example, an eruption in 1989 caused a complete blackout of the Quebec province power grid in Canada, while one in January 1997 was the likely cause of the catastrophic failure of the Telstar 401 communications satellite. The photo shows the ejection of a large solar mass. (An ultraviolet image of the Sun is superimposed over the solar disk.)

CREDIT: SOHO (ESA & NASA)

Magnetic Storms

The Sun is ultimately the source of all weather within the solar system, but the Sun's influence on the Earth's space environment is largely conveyed through the solar wind—a gusty flow of particles that stream outward from the Sun at about a million miles per hour. The wind transfers solar energy into the Earth's magnetic field (the geomagnetic field), and during periods of intense solar activity, when the solar wind turns into a fierce gale, the amount of energy transferred gets proportionately larger. The transfer process is disruptive and catastrophic, somewhat analogous to the way a stretched rubber band snaps and transfers energy to your fingers, only in this case the result is an energized and distorted geomagnetic field. The enhanced geomagnetic activity is referred to as a magnetic storm.

Magnetic storms can last anywhere from hours to days, during which the intensity of relativistic electrons and the rate at which they pepper a satellite can increase several thousand-fold. Television, telephone, or radio reception can be disrupted, and the electrons can wreak havoc on satellite-transmitted data or, in what amounts to extraordinary bad luck, knock a satellite unconscious—forever.

If they know their "bird" is in for nasty weather, satellite operators can re-route communications or take other measures to protect the data streams. But usually they don't know because most satellites lack sensors to monitor the local space weather. And though there are satellites that monitor what's going on around them as they orbit, it's not a simple matter to use that data to infer the weather conditions along a different orbit. It's a tricky business.

That's where Reeves and his DREAM team can help.

DREAM

The whole point of the Dynamic Radiation Environment Assimilation Model, or DREAM, is to provide a snapshot of the belt's *global* electron environment, despite having sparse data that provides only a sample of the *local* environment surrounding a few satellites. If the Van Allen belt were relatively static and uniform, those local conditions would allow satellite operators to estimate the conditions facing their satellites. But conditions within the belt are too variable because the glue that holds the belt together is the dynamic and stormy geomagnetic field.

DREAM can provide that global picture. It takes whatever data is available and, in effect, combines it with data created from theoretical models of the geomagnetic field and the radiation belt. The technique called data assimilation then produces a more complete data set—an optimized solution that best represents the true electron environment surrounding the Earth out to about seven Earth radii (a distance of approximately 42,000 kilometers, or a little more than 10 percent of the distance to the moon).

It's a remarkable process, somewhat akin to using traffic conditions on the beltway around Washington, D.C., to describe traffic in downtown Chicago or along the tumbleweed-lined margins of Interstate 40 through New Mexico. Once DREAM has produced an output, an informed satellite operator can take actions as needed.

Says Reeves, "Even with partially complete physical models for the radiation belt dynamics and for the geomagnetic field, DREAM will do a surprisingly good job of calculating the local environment anywhere in near-Earth space."

A lite version of DREAM set up as a demonstration works so well that DREAM is already being regarded as a major asset for situational awareness. That's the government euphemism for everyone keeping their heads up to better anticipate emergency or hostile situations. And it may turn out that DREAM's ability to pinpoint hostile weather conditions will help everyone *keep* their heads should a critical satellite suddenly go *poof!*

Increased space-situational awareness has certainly been a motivating factor in DREAM's development. But its software framework was designed to be modular and highly flexible, so it's a simple matter to test different theories about, say, electron diffusion in the Van Allen belt by swapping out different electron-diffusion modules. DREAM is as much a research tool for understanding near-Earth space as it is an aid to the satellite community. And there is indeed much to learn about that space.

Surprise! Surprise!

Because the amount of energy transferred to the Earth's field fluctuates wildly, and because there are many ways to distribute that energy, the geomagnetic field is remarkably dynamic and complex.

"It's a rotating, asymmetric field strongly coupled to a highly variable magnetized plasma [the solar wind]," says physicist and DREAM team member Mike Henderson. "It supports the Van Allen radiation belt, but the belt has its own dynamics—its particles gain energy, diffuse in and out. It inflates in size, even disappears sometimes."

The Radiation Belt Storm Probes, a matched pair of satellites designed to gather the data needed to understand the dynamics of the radiation belts, are scheduled to be launched in the fall of 2012. Los Alamos was heavily involved in justifying and defining their mission and in designing a suite of instruments carried on each probe, including building one of the instruments (the HOPE spectrometer). This is an artist's rendering of the probes in orbit.

CREDIT: NASA/JOHNS HOPKINS UNIVERSITY APPLIED PHYSICS LABORATORY

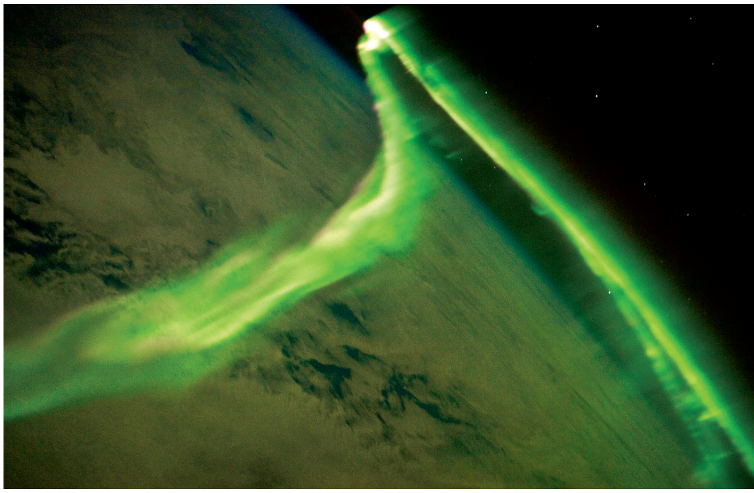

Particularly during magnetic storms, particles from the radiation belt can follow the Earth's field lines right into the Earth's upper atmosphere, where they collide with oxygen and other atoms. The atoms absorb the particles' energy and subsequently release it in the form of light, which stargazers observe as an aurora (the aurora borealis, or northern lights, in the northern hemisphere, and the aurora australis, or southern lights, in the southern hemisphere). The upper figure shows the aurora australis as seen from space; the lower shows the aurora borealis from the ground in Canada.

UPPER FIGURE CREDIT: ISS EXPEDITION 23 CREW, ISAL, NASA

Sweet **DREAM** Is Made of These

DREAM fills in the gaps of a sparse data set and produces a representation of the charged-particle environment surrounding the Earth. Some of the components that go into DREAM are illustrated below. To learn more, visit the DREAM website at dream.lanl.gov.

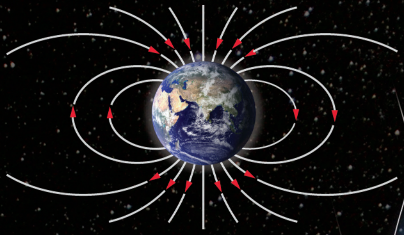

(Left) **The Geomagnetic Field** is often modeled as a simple dipole field, with field lines—invisible lines that run in the direction of the field—symmetric about the magnetic axis.

(Right) A more realistic model of the geomagnetic field includes the distortions caused by the solar wind, which compresses the daytime side of the field and drags the nighttime field far downwind. There is no symmetry about the magnetic axis. By adding more physical processes, such as a ring current that runs around the planet's midsection, the model can be made more complete, at the price of requiring more data to constrain the increased parameter space and significantly more time to perform the algorithms.

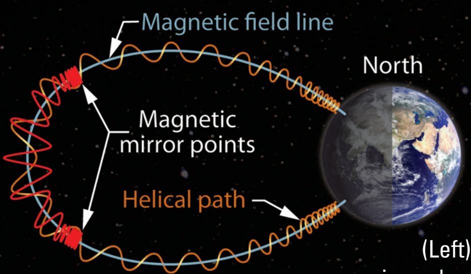

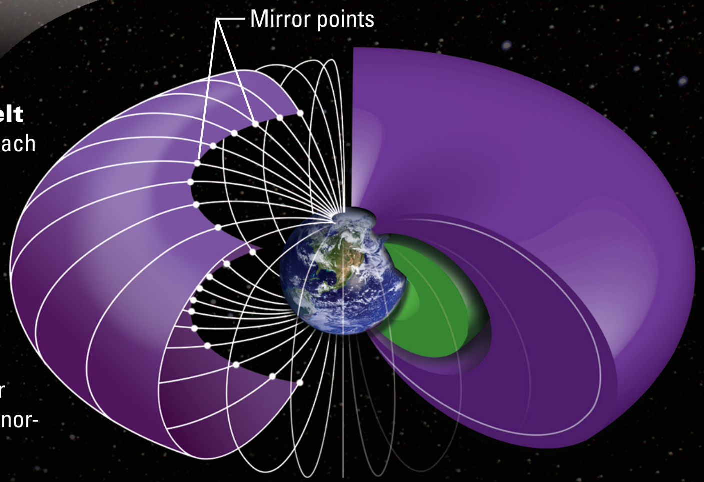

(Left) **The Van Allen Radiation Belt** is made up of energetic ions and electrons. Each particle follows a helical path that's centered around a geomagnetic field line. At a so-called mirror point, determined by the field strength and the particle's pitch angle, the particle reverses direction until it "bounces" off a similar mirror point in the other hemisphere. (The pitch angle determines how far a particle goes along a field line for one turn of the helix.)

(Right) Particles also undergo a slow circular "drift" around the magnetic axis, so that particles of a given energy and pitch angle are trapped—constrained to occupy a mostly circular "drift shell" that's bounded by the north and south mirror points. With a distribution of particle energies and pitch angles, the result is the enormous, donut-shaped Van Allen radiation belt.

Evidently, there are aspects of the belt that are simply not understood.

For example, it's natural to suppose some correlation between the solar wind and radiation-belt electrons, and 30 years ago, the data indicated a nice, linear relationship between electron intensities and the solar-wind velocity. But after 30 more years of data collection, much of it gathered with Los Alamos instruments, the relationship is clearly nonlinear, with high electron intensities occurring for a wide range of velocities. This is not understood.

Then there is the mystery of the magnetic storms. It turns out to be only partially true that an influx of solar energy leads to higher particle energies. Reeves' team looked at several hundred storms to see how each affected the energy and population of electrons in the belt. The expectation was that more intense storms would result in more particles with higher energies. But often the storm had no effect, and 19 percent of the storms—nearly one out of five—actually depleted the belt. It's like a thunderstorm comes through and leaves your backyard drier than it was.

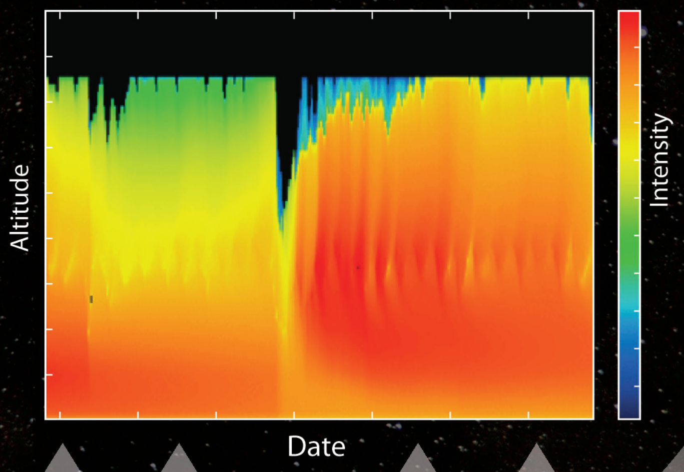

DREAM Output is the result of assimilating the data with results obtained from models of the geomagnetic field and the Van Allen radiation belt. It specifies the global particle environment, with electron flux given as a function of time and altitude.

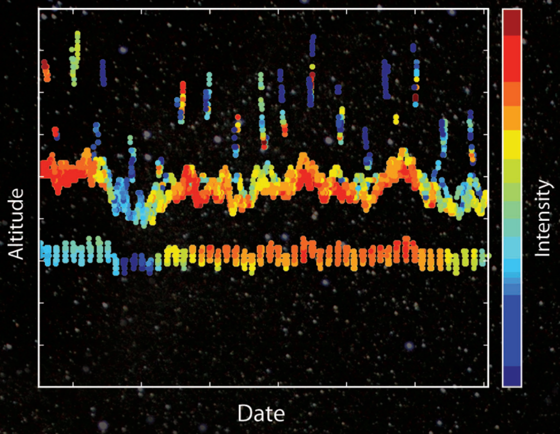

Data Used by DREAM can come from a variety of instruments from any spacecraft, often in the form of electron flux (number of electrons per unit area, time, and energy) versus altitude and time. This data set comes from the GOES-13 satellite, a space weather data-gathering station in a geostationary orbit. (It moves at the same rotational speed as the Earth so it always has the same view of the planet.)

Manmade Radiation Belts can result from the injection and initial trapping of radiation from high-altitude nuclear explosions (HANE). DREAM contains a module, used for national security applications, for estimating the effects of the artificial belts. The picture shows a successful atmospheric nuclear test conducted by the United States in 1962. Known as Starfish, the 1.4-megaton device was detonated over the Pacific Ocean at an altitude of 400 kilometers. In the months following the test, seven orbiting satellites failed, presumably due to the addition of HANE radiation, which persisted for about five years.

A Prediction of Prediction

New light will be cast upon many of these poorly understood aspects of the belt come the fall of 2012 when NASA plans to launch the Radiation Belt Storm Probes (RBSP). The matched pair of satellites will probe all regions of the belt, gathering data that should revolutionize our understanding of the dynamic charged cloud. And once the belt and Earth's magnetic field are better understood, DREAM may transition from being a real-time specifier of the local space environment to a real-time forecaster of space weather. Then it may help prevent satellites from going *poof!* in the first place.

Whether DREAM gains predictive capability or not, Reeves and the team are in an ideal situation. "It's estimated that 99 percent of the visible universe is plasma," says Reeves, "and plasma interactions are the same whether they occur in the geomagnetic field or in jets shooting out from a supermassive black hole. It's a little difficult to probe the dynamics around a black hole, but we can launch probes into near-Earth space and watch those interactions as they are happening."

And they can do that while pursuing their DREAM. ❖ LDRD

—Jay Schecker

Taming the Tempest

Some celebrities achieve such fame that they are known by a single name—Elvis, Oprah—and that name will forever be attached to them and requires no explanation. Katrina is such a name.

In late August 2005, an African tropical wave moved across the Atlantic, intensified to a hurricane, and made landfall in Florida. This hurricane, named Katrina according to an alphabetical system, sidled away and regained power, more than doubling in size to strike again, then again, and again, along much of the Gulf Coast. The hurricane produced 175-miles-per-hour winds, spawned 143 tornadoes, caused 1,836 human deaths (more than 700 are still missing), and ultimately amassed $110 billion in damage. Seven years later, swaths of land remain unusable due to catastrophic damage and pollution, and regional economies have not recovered. Normally, hurricane names are reused by the naming organization, but Katrina's name has been stricken from the list, and the name has plummeted in popularity among baby names. And she wasn't even the deadliest: 8,000 people died in 1900 when a hurricane struck Texas. Fortunately for baby name books, that storm went unnamed.

Although predictions about a hurricane's path have improved substantially in recent years, forecasting its strength is difficult because of unknowns concealed in the depths of the storm. Now Los Alamos atmospheric scientist Jon Reisner and his team are using lightning as a predictor of a storm's strength.

Lightning does not accompany all tropical storms, but during the record-setting hurricane season of 2005, three of the most powerful storms—Rita, Katrina, and Emily—did have lightning and lots of it. According to NASA, hurricanes are most likely to produce lightning when they're making landfall. In broad terms, lightning activity within and surrounding a hurricane has been known to indicate when and where a hurricane may intensify. Research also shows that if the electricity moves from the eye (or center) to the periphery of the cloud, the energy is probably dissipating. Los Alamos researchers are trying to flesh out that connection between lightning activity and hurricane intensity.

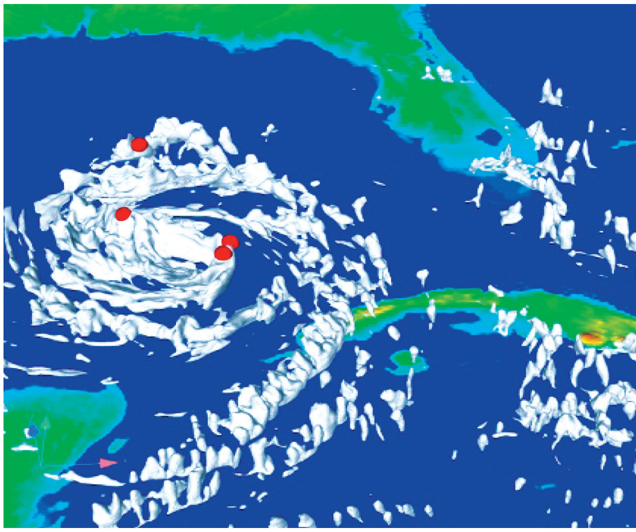

In this simulation of Hurricane Rita, red dots denote areas of active lightning within cloud regions, which correlate well with observed lightning locations during a period of rapid storm intensification.

Lightning produces electromagnetic waves with high and low frequencies during a strike. Los Alamos sensors detect these frequencies, allowing researchers to characterize the charge associated with the flash. The detectors and instruments are mounted aboard planes and flown into the hurricane's eye. Reisner brings the lightning data into a simulation—the first to incorporate a three-dimensional model of the lightning activity in the hurricane.

Reisner also discovered that by looking at individual water particles, he was able to construct a more realistic representation of the cloud structure within a hurricane. Traditional hurricane models, his team found, improperly express cloud boundaries, structuring clouds as continuous objects rather than collections of particles. When a storm develops as warm ocean water evaporates, winds force humid air to rise until the vapor condenses back into tiny liquid particles. During condensation, energetic water particles release heat as they collide and condense, fueling the storm and forming the hurricane's eye wall. These particles interact or change at miniscule scales—between

More than 1.2 million people along the northern Gulf coast were ordered to evacuate to escape Hurricane Katrina's floodwaters in 2005, and this New Orleans neighborhood is still devastated. Reisner's research may help citizens prepare for such hurricane emergencies.

10 and 100 nanometers (small enough to slip through a surgical mask)—challenging researchers because of the difficulty integrating these small spatial scales. However, once Reisner's team accounted for these "nano-droplets," they found that the previous structural mischaracterization of clouds caused a false rendering of temperature conditions. The discovery led to a vital correction to the temperatures in the hurricane's eye wall, the area in the storm where the most damaging winds and rainfall are located.

Reisner's two developments—proper cloud representation and lightning predictive models—together help demystify hurricane intensity, which may lead to more accurate predictions that help the public prepare for potential devastation. ❖ LDRD

—Kirsten Fox

Los Alamos scientists are improving methods to convert agricultural waste into fuels and other valuable chemicals.

Agricultural Alchemy

Susan Hanson wants to convert grass into gold. Stumps or stalks or weeds in a field can be burned to generate energy—emitting harmful byproducts—but Hanson and others in her field have a more elegant solution in mind: turning agricultural and forestry waste into valuable chemicals and fuel. The sources are abundant, and this form of chemical production would not compete with demands for food, fertilizer, or water since it uses only waste products rather than requiring additional agricultural production. But the nation's energy problems are not easily solved, and where the enormous hurdle of efficiently extracting the energy from plant matter has stumped others, this Los Alamos chemist is getting to the root of the matter: the lignin.

Plant cellular walls are primarily composed of cellulose, hemicellulose, and lignin. Comprising nearly 30 percent of the biomass on Earth, lignin conducts water in stems, provides the mechanical support for the plant, and strengthens cell walls. One gram of lignin contains about 2.27 kilojoules of energy—comparable to coal and 30 percent more than cellulose alone. But just as this woody material "glues" the cell wall together to protect plants from pests and pathogens, it protects them from researchers as well. An efficient method is needed by which researchers can break apart the complex polymer's strong bonds and degrade the plant into energy-rich simple sugars.

Most research has focused on pretreating lignin with environmentally and economically unfriendly solvents, heat, or high pressure to rupture the lignin bonds. But Hanson and Los Alamos colleagues Pete Silks and Ruilian Wu found a "green" catalyst, which enables a desired chemical reaction without being consumed by it, to break down lignin into smaller chemical components. These components could potentially be used to produce alcohols, waxes, surfactants (for detergents and other applications), and fuels.

Historically, precious metals such as platinum have been the basis for most catalysts, but Hanson focuses on vanadium. Found in many minerals and marine organisms, and often collocated with iron ores and petroleum, vanadium is an earth-abundant metal that is not toxic in the small amounts used in catalysis. The reaction proceeds in air at atmospheric pressure with only mild heat, making vanadium easier to use than most other metals, which are sensitive to air and damaged by oxidation. Hanson's team designed and synthesized its catalysts by combining vanadium with other components.

The Los Alamos work demonstrates that a lignin model compound may be broken down selectively into useful components using a vanadium catalyst. The strong carbon-carbon and carbon-oxygen bonds are cleaved via oxidation that breaks the lignin into smaller, usable pieces. And the only byproduct—besides the desirable, energy-rich sugars—is water. ❖ LDRD

—Kirsten Fox

What if the Sky Is Falling?

One doesn't have to be paranoid to believe the sky is falling: asteroids and comets do rain down on Earth. These rocky or icy chunks usually orbit the sun between Mars and Jupiter or beyond Neptune, but gravity and collisions occasionally redirect them inward, sometimes on a collision course with Earth.

What happens next depends on the size of the asteroids. Some are too small to survive the passage through the atmosphere, where they burn up. Larger objects, however, do sometimes crash into Earth and create craters or worse—and they don't need to be terribly large to be devastating. The rock that likely exploded over Siberia in 1908 and knocked down trees for hundreds of miles in all directions from its shockwave alone, equal to 10–15 megatons (million tons) of TNT, is thought to have been less than 60 meters wide. An object only a few times that size could not only cause tremendous regional damage on land but could also spawn deadly tsunamis from an impact at sea—potentially reaching shorelines far from the impact. Moreover, the vapor and debris produced by a substantially larger collision could block out the Sun, resulting in massive climate change and numerous extinctions.

Near-Earth objects (NEOs) are asteroids and comets in orbits that allow them to enter Earth's neighborhood as they orbit the Sun. NEOs whiz through Earth's vicinity on a regular basis, leading the U.S. government to launch a neighborhood watch. NASA detects, tracks, and characterizes NEOs using ground- and space-based telescopes. According to NASA, 8,871 near-Earth objects have been discovered, of which 1,300 have been classified as potentially hazardous. If the ship-sized asteroid that passed by Earth last year (closer than the Moon) had slightly altered course, its impact would have been equivalent to 10 million of the earliest atomic bombs. This threat has led some to consider using a nuclear energy source to disrupt an NEO in space before it can unleash that level of destruction down here.

Los Alamos astrophysicist Robert Weaver has already demonstrated this is possible. To simulate the phenomenon in 3-D, Weaver ran a detailed simulation on Cielo, one of the Lab's supercomputers, to reveal how a potentially hazardous object could be disrupted. In January, Weaver released a video (available online at www.youtube.com/user/LosAlamosNationalLab) revealing how a one-megaton nuclear energy source might affect a half-kilometer-long asteroid capable of destroying a continent. The simulation shows how a single explosion directed at the asteroid would send a shock wave through it, impacting the individual granite chunks comprising it and shattering the entire object. Weaver pronounced the hazard "fully mitigated." Score one for happy humans, zero for killer asteroid.

"Prior to my calculations, it was merely speculation that a nuclear source might be a good option to deter an asteroid. We provided the physics and hydrodynamics—definitive scientific results about how a nuclear burst could do the job," says Weaver. The simulation revealed that the high velocities of the asteroid debris fragments prevented them from reassembling afterward and, in fact, caused them to disperse well beyond their original Earth-crossing trajectory. "It's highly unlikely for them to pose a secondary threat."

Nonnuclear alternatives have been explored previously, including conventional explosives and some other creative ways to steer an Earth-crossing object off course, but according to a NASA study delivered to Congress in 2007, a nuclear device remains the most effective option. Nonetheless,

How could the world be saved from a killer asteroid?

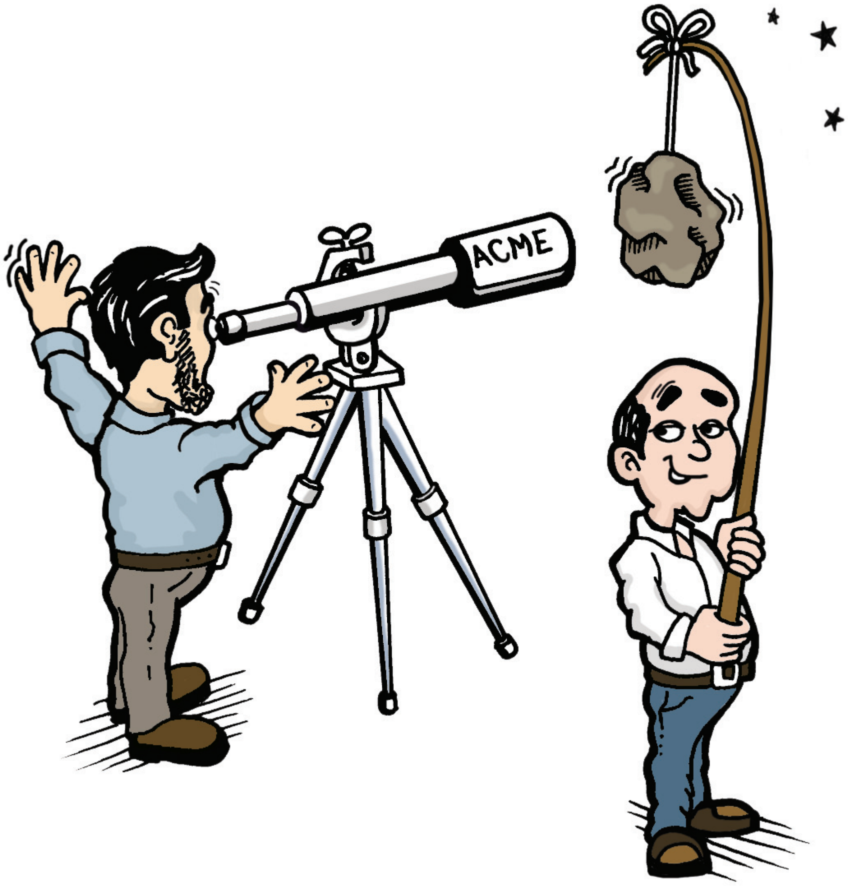

Weaver is sensitive to concerns about using a nuclear device. He proposes that it would be detonated in deep space, where, according to his simulation, neither the explosion nor the radioactive fallout would pose any threat to Earth.

The real issue, according to Weaver, is how much advance notice of a nearby asteroid is obtained and how long it would take to execute a mission; deflecting an asteroid could require a two-year warning. Yet even with an appeal to the nuclear option, there is currently no feasible defense against asteroids more than a mile wide. Fortunately, they rarely impact Earth. The last occurrence was the celestial celebrity that brought about the end of the dinosaur age 65 million years ago.

—Kirsten Fox

Ghost of Christmas Past

On Christmas day, 2010, the Burst Alert Telescope onboard NASA's Swift satellite, running software developed at Los Alamos, detected a new type of gamma-ray burst (GRB). GRBs are exactly what they sound like—quick bursts of gamma rays—often followed by less energetic radiation. Although the gamma-ray component typically lasts less than a minute, this particular burst lasted a half hour before being followed by an x-ray afterglow. Its emission spectrum, too, contained a blend of familiar and unfamiliar features: One part resembled an energetic jet of matter and another part resembled a supernova; both of these are frequently associated with stellar explosions and the expanding shock waves that accompany them. But this particular GRB had one feature that could not be so neatly explained as originating with an exploding star: shortly after the shock wave broke out of the star, it appeared to run into an unexpected outer shell of material, emitting a burst of ultraviolet, visible, and infrared light. Evidently the GRB, now known as the Christmas Burst, represents a rare astronomical event.

The cause of the Christmas Burst is the subject of lively debate in astronomy circles. One leading hypothesis is that the GRB originated in a binary star system in which a red giant star with a core made of helium closely orbited a neutron star. Eventually the stars spiraled into each other, producing the outburst. This scenario was originally proposed in 1997 by Los Alamos computational physicist Chris Fryer, who has continued to pioneer theoretical efforts to better study its characteristics. If Fryer's scenario is indeed the cause of Christmas Burst, then based on the burst's brightness, the collision must have taken place in a distant galaxy. An alternate proposal involves a small, comet-like object falling onto a neutron star and producing a much dimmer burst, in which case it must have occurred closer, within our Galaxy. Normally this could be resolved by simply looking for a galaxy in the location where the GRB was observed, but only an inconclusive hint of a glow was found. Indeed, Fryer and his team submitted a proposal to take a longer-exposure image of that region using the Hubble Space Telescope, while still pursuing other ways to identify the source of the unusual GRB.

During the year following the burst, Fryer and Los Alamos colleague Wesley Even put the distant stellar collision hypothesis to the test, using a sophisticated computer simulation to calculate the emission spectrum produced by such a collision. Working with an international team of researchers running a variety of simulations, they were able to confirm that the observed Christmas Burst matched the neutron star-red giant collision model studied by Fryer and his team.

For most of its "life," a star produces energy by nuclear fusion, with hydrogen nuclei in its core fusing together. But eventually, the hydrogen in the core fuses into helium, which is unable to fuse with itself to generate additional heat and pressure. Without that source of pressure, the helium core begins to collapse under its own weight while the outer layers of the star, spurred by nuclear reactions outside the core, expand outward. The result is known as a red giant. For massive stars, the helium core eventually gets hot enough to ignite further nuclear fusion and thereby produce

This artist's conception of the Christmas Burst astronomical event shows the merger of a neutron star with a red giant star, with the resulting body collapsing to a black hole (central dot). The collapse produces an extremely energetic jet of matter that interacts first with the collapsing stellar core (central sphere) and then with the outer layers of the red giant star (surrounding red swirl) that were previously cast outward by the approach of the neutron star. Los Alamos scientists recently performed a computer simulation to demonstrate how this event produced the complex pattern of radiation observed.

CREDIT: AURORE SIMONNET, SONOMA STATE UNIVERSITY, AND NASA

heavier and heavier elements in the core. When that progression of elements reaches iron, which is incapable of releasing energy through nuclear reactions, a supernova explosion blasts away the outer layers of the star and, in most cases, leaves behind an ultra-dense object known as a neutron star.

In Fryer's model, the more massive star in a binary star system evolves through its life, ultimately collapsing to form a neutron star. As the second star evolves into a red giant in the final stages of its life, it expands and envelops the neutron star, causing the neutron star to spiral into its helium core. Meanwhile, the outer layers of the red giant are flung outward by the neutron star's passage, creating a shell of material surrounding the colliding pair. When the collision takes place, the helium drives the neutron star over its maximum stable mass, causing it to collapse into a black hole. The energy released by that collapse powers an explosion and a jet of matter slams into the surrounding shell.

This sequence of events—black hole formation, explosion with a jet, and interaction with the ejected shell—led to the complicated spectrum of radiation observed in the Christmas Burst. As the jet progressed through the helium core, it deposited energy into the core, causing the core to emit the strong x-ray afterglow that followed the initial GRB. And when the weakened jet subsequently passed through the surrounding shell, the shell produced ultraviolet, visible, and infrared light, also exactly as observed.

"What's great is that we're starting to be able to identify the rare types of GRB that used to be considered too weird to explain," says Fryer. "And often it's the weird ones that have the most to teach you." ❖ LDRD

—Craig Tyler

Deciphering DNA and Disease

More than a decade ago, the field of medicine waited with bated breath while the international Human Genome Project was completed. This project's goal was to reveal and provide understanding of the genetic makeup of humans, no small task. The sequence of the 3 billion chemical base pairs that make up human DNA was determined, potentially opening the door to finding the genetic roots of disease and developing treatments. A tremendous amount of data was compiled, but the question remained, how to decipher it?

As Los Alamos molecular biologist Csaba Kiss describes, "We have the book, now we have to translate it." Scientists still do not know the function of more than half of the discovered genes. Los Alamos researchers are attacking that problem by trying to link the genes to the corresponding proteins they specify in order to determine their role in cellular activity—a lofty goal with the potential for many applications, including drug design.

Kiss is a member of a Los Alamos research team led by renowned molecular biologist Andrew Bradbury. A former physician, Bradbury thinks that one key to future medical advances is the use of antibodies, proteins naturally produced by the immune system that help the body fight infectious diseases. Antibodies are also necessary for basic research purposes on a bench-top scale in order to identify key protein interactions without using animals as surrogates. In fact the ability to selectively target proteins with engineered antibodies is a major goal in biotechnology. Success offers a route to attack human diseases such as cancer, a means to stay ahead of evolving infectious bacteria such as multi-drug-resistant tuberculosis, and a countermeasure for potential bioagents used in weapons of mass destruction.

Antibodies are Y-shaped proteins that each respond to a specific antigen, or *anti*body *gen*erator. Usually a protein, an antigen is any substance—either formed within the body such as a cancer cell, or in the external environment such as pollen or a virus—that causes the immune system to produce antibodies against it. Each antibody has a special section (at the tip of each branch of the Y) that binds to an antigen, as a lock to a key. Antibodies act as bodyguards by blocking entrance to human cells, disabling the antigen's cellular function, or triggering other parts of the immune system to attack the antigen.

Bradbury aims to generate antibodies that react to human proteins, using the proteins as antigens. What is the benefit of finding an antibody that attacks a protein from your own body? Consider how it could fight breast cancer, using antibodies made from human cells to destroy these cancer cells that kill nearly 40,000 American women every year. In fact, the anti-breast cancer drug Herceptin is an antibody that recognizes a specific human protein found frequently on some breast cancer cells. More broadly, antibodies against human proteins can be used as treatments to fight many types of cancers and autoimmune disorders, such as multiple sclerosis, Crohn's disease, and rheumatoid arthritis.

Bradbury starts with the approximately 30,000 genes identified in the human genome encoding for the proteins that will serve as the antigens. His goal is to identify an antibody that matches and binds tightly to each specific antigen. Los Alamos colleague Geoff Waldo uses a fluorescence technique to help Bradbury select the best proteins, or parts of proteins, to use as antigens. The technique allows for rapid detection of the antibody-antigen interaction.

Selection involves testing one antigen at a time against many different antibodies. To begin the antibody selection process, Bradbury and his team insert the genes for billions of different antibodies into a bacteriophage, a virus that infects bacteria. The phage infects an E. coli cell. Each bacteriophage carries a different antibody gene, and the collection of the billions of bacteriophage, each carrying different antibodies, is known as an "antibody library." When the E. coli reproduces, it also produces more phage. Each phage displays only one copy of the antibody on its surface and is mixed with the target antigen previously anchored to a solid support. After allowing matching pairs to bind to each other, the "non-binders" are washed away. The remaining phage-antibody is released from the antigen using a very basic solution, and the gene inside the phage is isolated and analyzed to determine the precise structure of the antibody that bound to the specific antigen.

A limitation of this approach is that the extremely small size of the phage makes it impossible to directly see one phage-antibody-antigen interaction, which invariably leads to false positives. However, this process narrows the prospective antibodies to a manageable number for further analysis.

To perfect his selection process, Bradbury introduces a larger medium—baker's yeast—that provides better control and more sophisticated screening via a technique known as flow cytometry. The prospective antibody is introduced to yeast, which uses its own natural cellular machinery to pop the antibody onto its surface. Unlike the bacteriophage, which displays only a single copy, each yeast has tens to hundreds of thousands of copies of identical antibodies on its surface. Next, Bradbury takes the target antigen (labeled by Waldo's fluorescent marker) and mixes it with the yeast. A flow cytometer forces the yeast solution into a narrow stream so cells can be analyzed one by one and hits them with a beam of light that causes the antigen to glow. Yeasts that have antibodies on their surface that recognize the target antigen will fluoresce. The cytometer sorts the mixture, tossing yeasts that do not glow, and collecting those that do, allowing the viable antibodies to be further validated and tested. Once the number of antibodies is pared down to a handful of possibilities, they will be tested individually in animals in order to determine if the antibody does recognize the specific antigen, a validation that may lead to a revolutionary approach for antibody discovery against killer diseases where no efficient therapies currently exist.

The antibody sequences will be made available online, free of charge to anyone to aid research in immunology and human protein function. "Now we can understand what proteins do," Kiss says. "For a decade, scientists have wondered what to do with the human genome findings—this is the next step. We're finally starting to translate the book. It's cutting-edge."

—Kirsten Fox

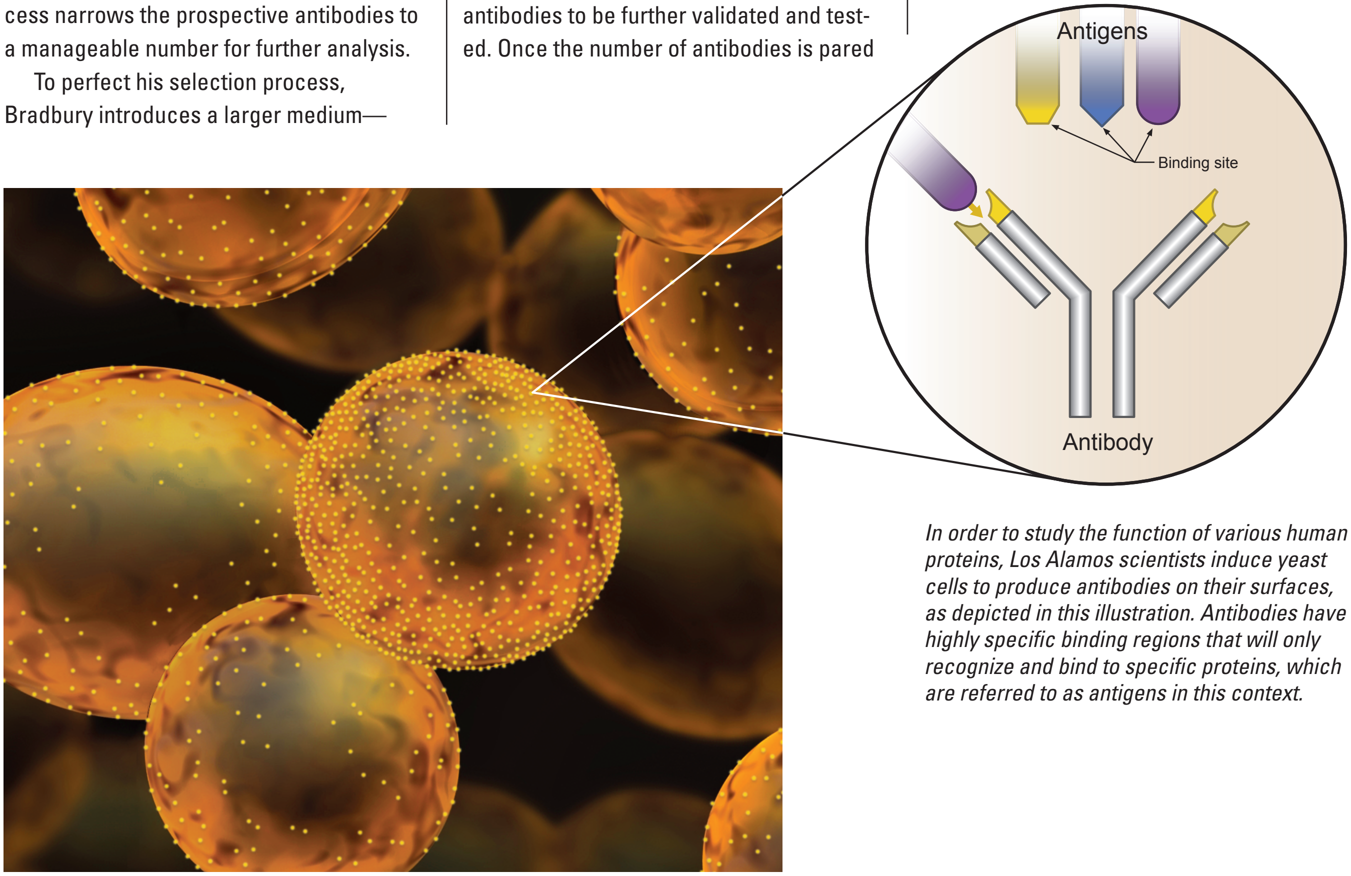

In order to study the function of various human proteins, Los Alamos scientists induce yeast cells to produce antibodies on their surfaces, as depicted in this illustration. Antibodies have highly specific binding regions that will only recognize and bind to specific proteins, which are referred to as antigens in this context.

ISSN: 1942-6631

Address mail to:
1663
Mail Stop M711
Los Alamos National Laboratory
Los Alamos, NM 87545
email: 1663magazine@lanl.gov
Fax: 505-665-4408
www.lanl.gov/science/1663/

PANORAMAS FROM NEW MEXICO

Principal Associate Director of Science, Technology, and Engineering—Alan Bishop

Editor-in-Chief—Craig Tyler

Science Editor—Jay Schecker

Managing Editor—Kirsten Fox

Art Director—Donald Montoya

Graphic Designer—Gloria Sharp

Illustrator—Leslie Sandoval

Photographers—Sandra Valdez, LeRoy N. Sanchez

Laboratory Historian—Alan Carr

Advisors—Terry Wallace, Duncan McBranch, and Kim Thomas

Los Alamos
NATIONAL LABORATORY
EST. 1943