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-

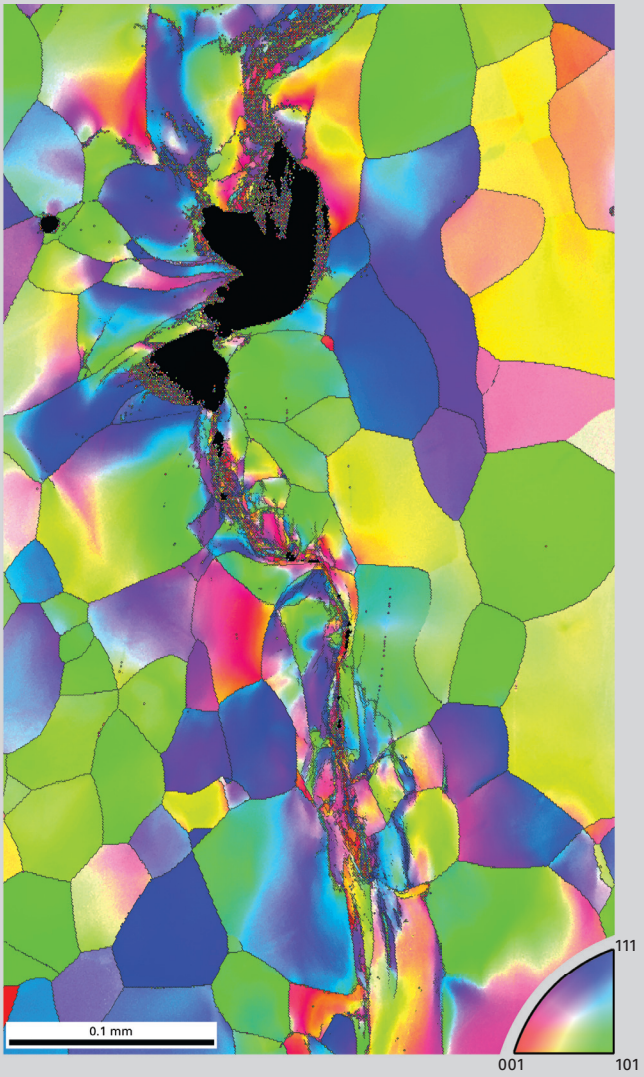

Microstructure

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.

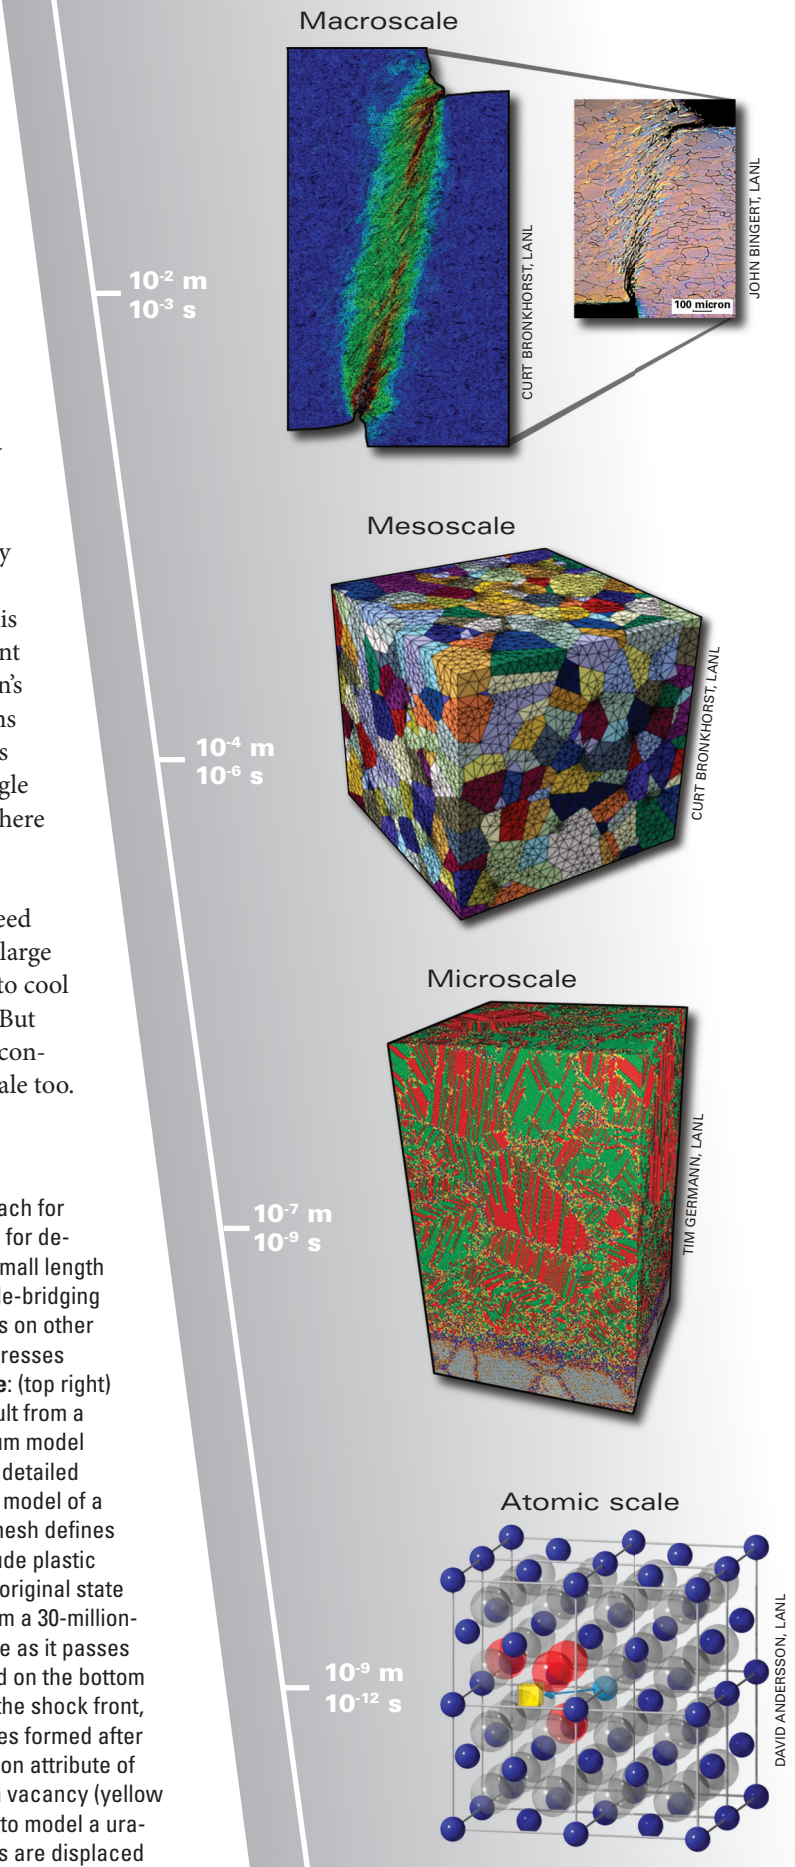

Simulations often use a multiscale, multiphysics approach for investigating material behaviors. Models are optimized for describing phenomena on one scale—with changes on small length scales taking place within small time scales—and scale-bridging algorithms pass the information to ongoing calculations 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