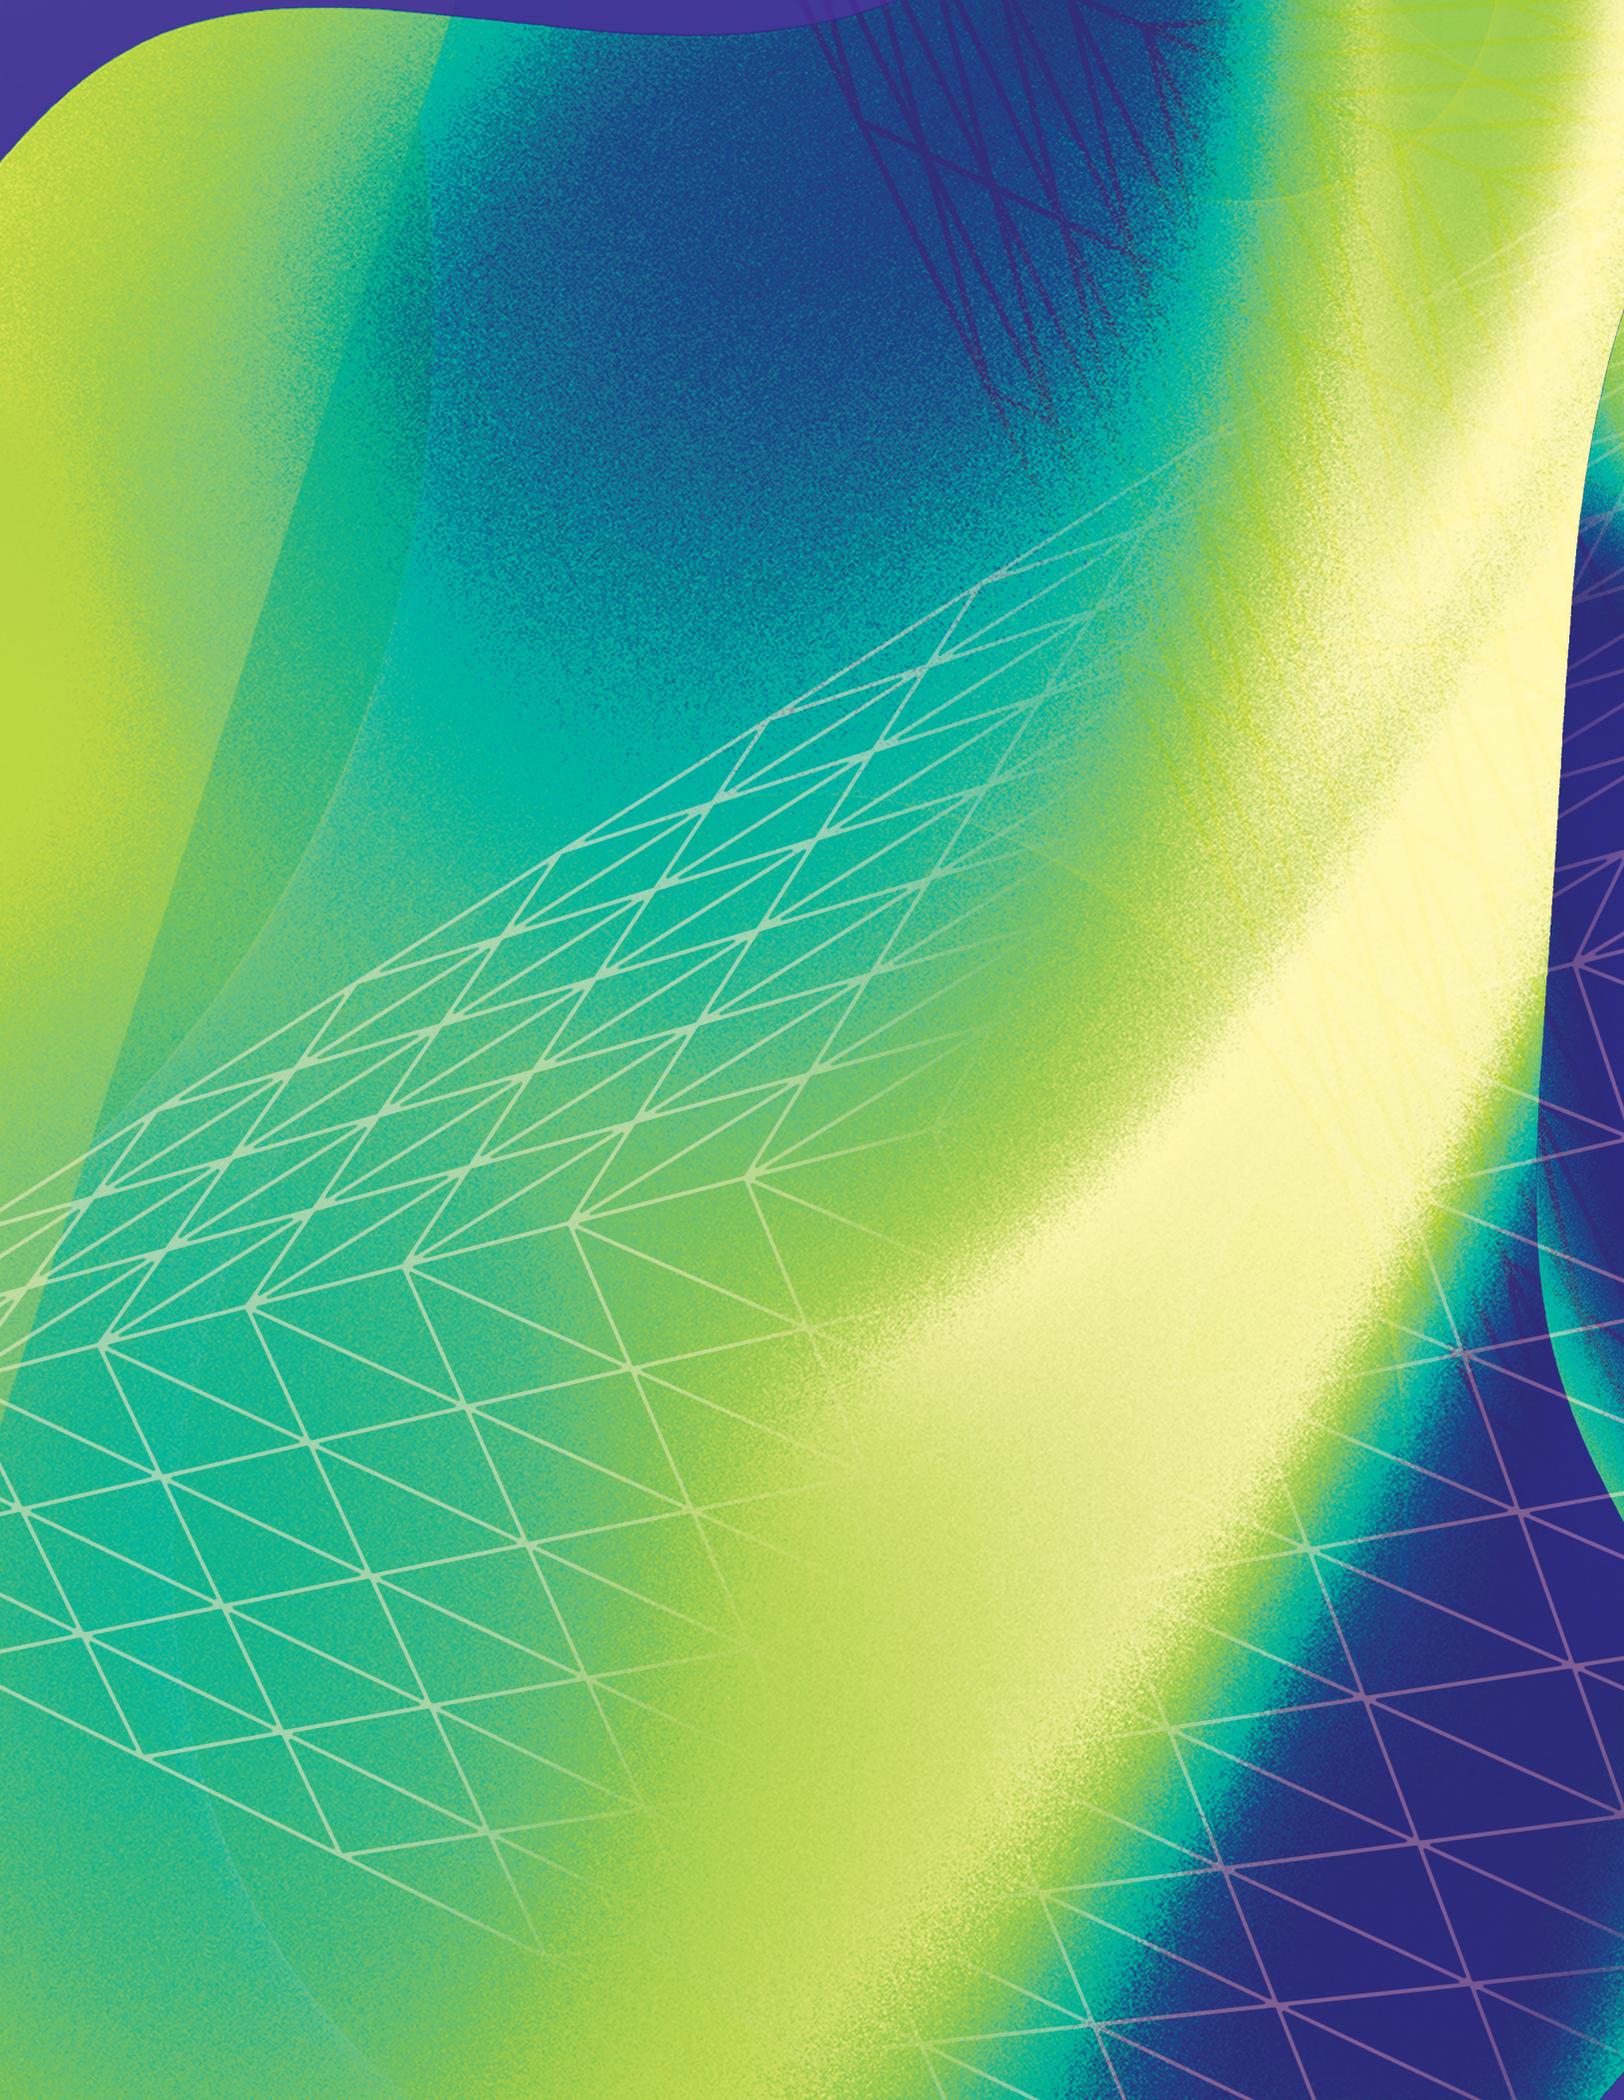
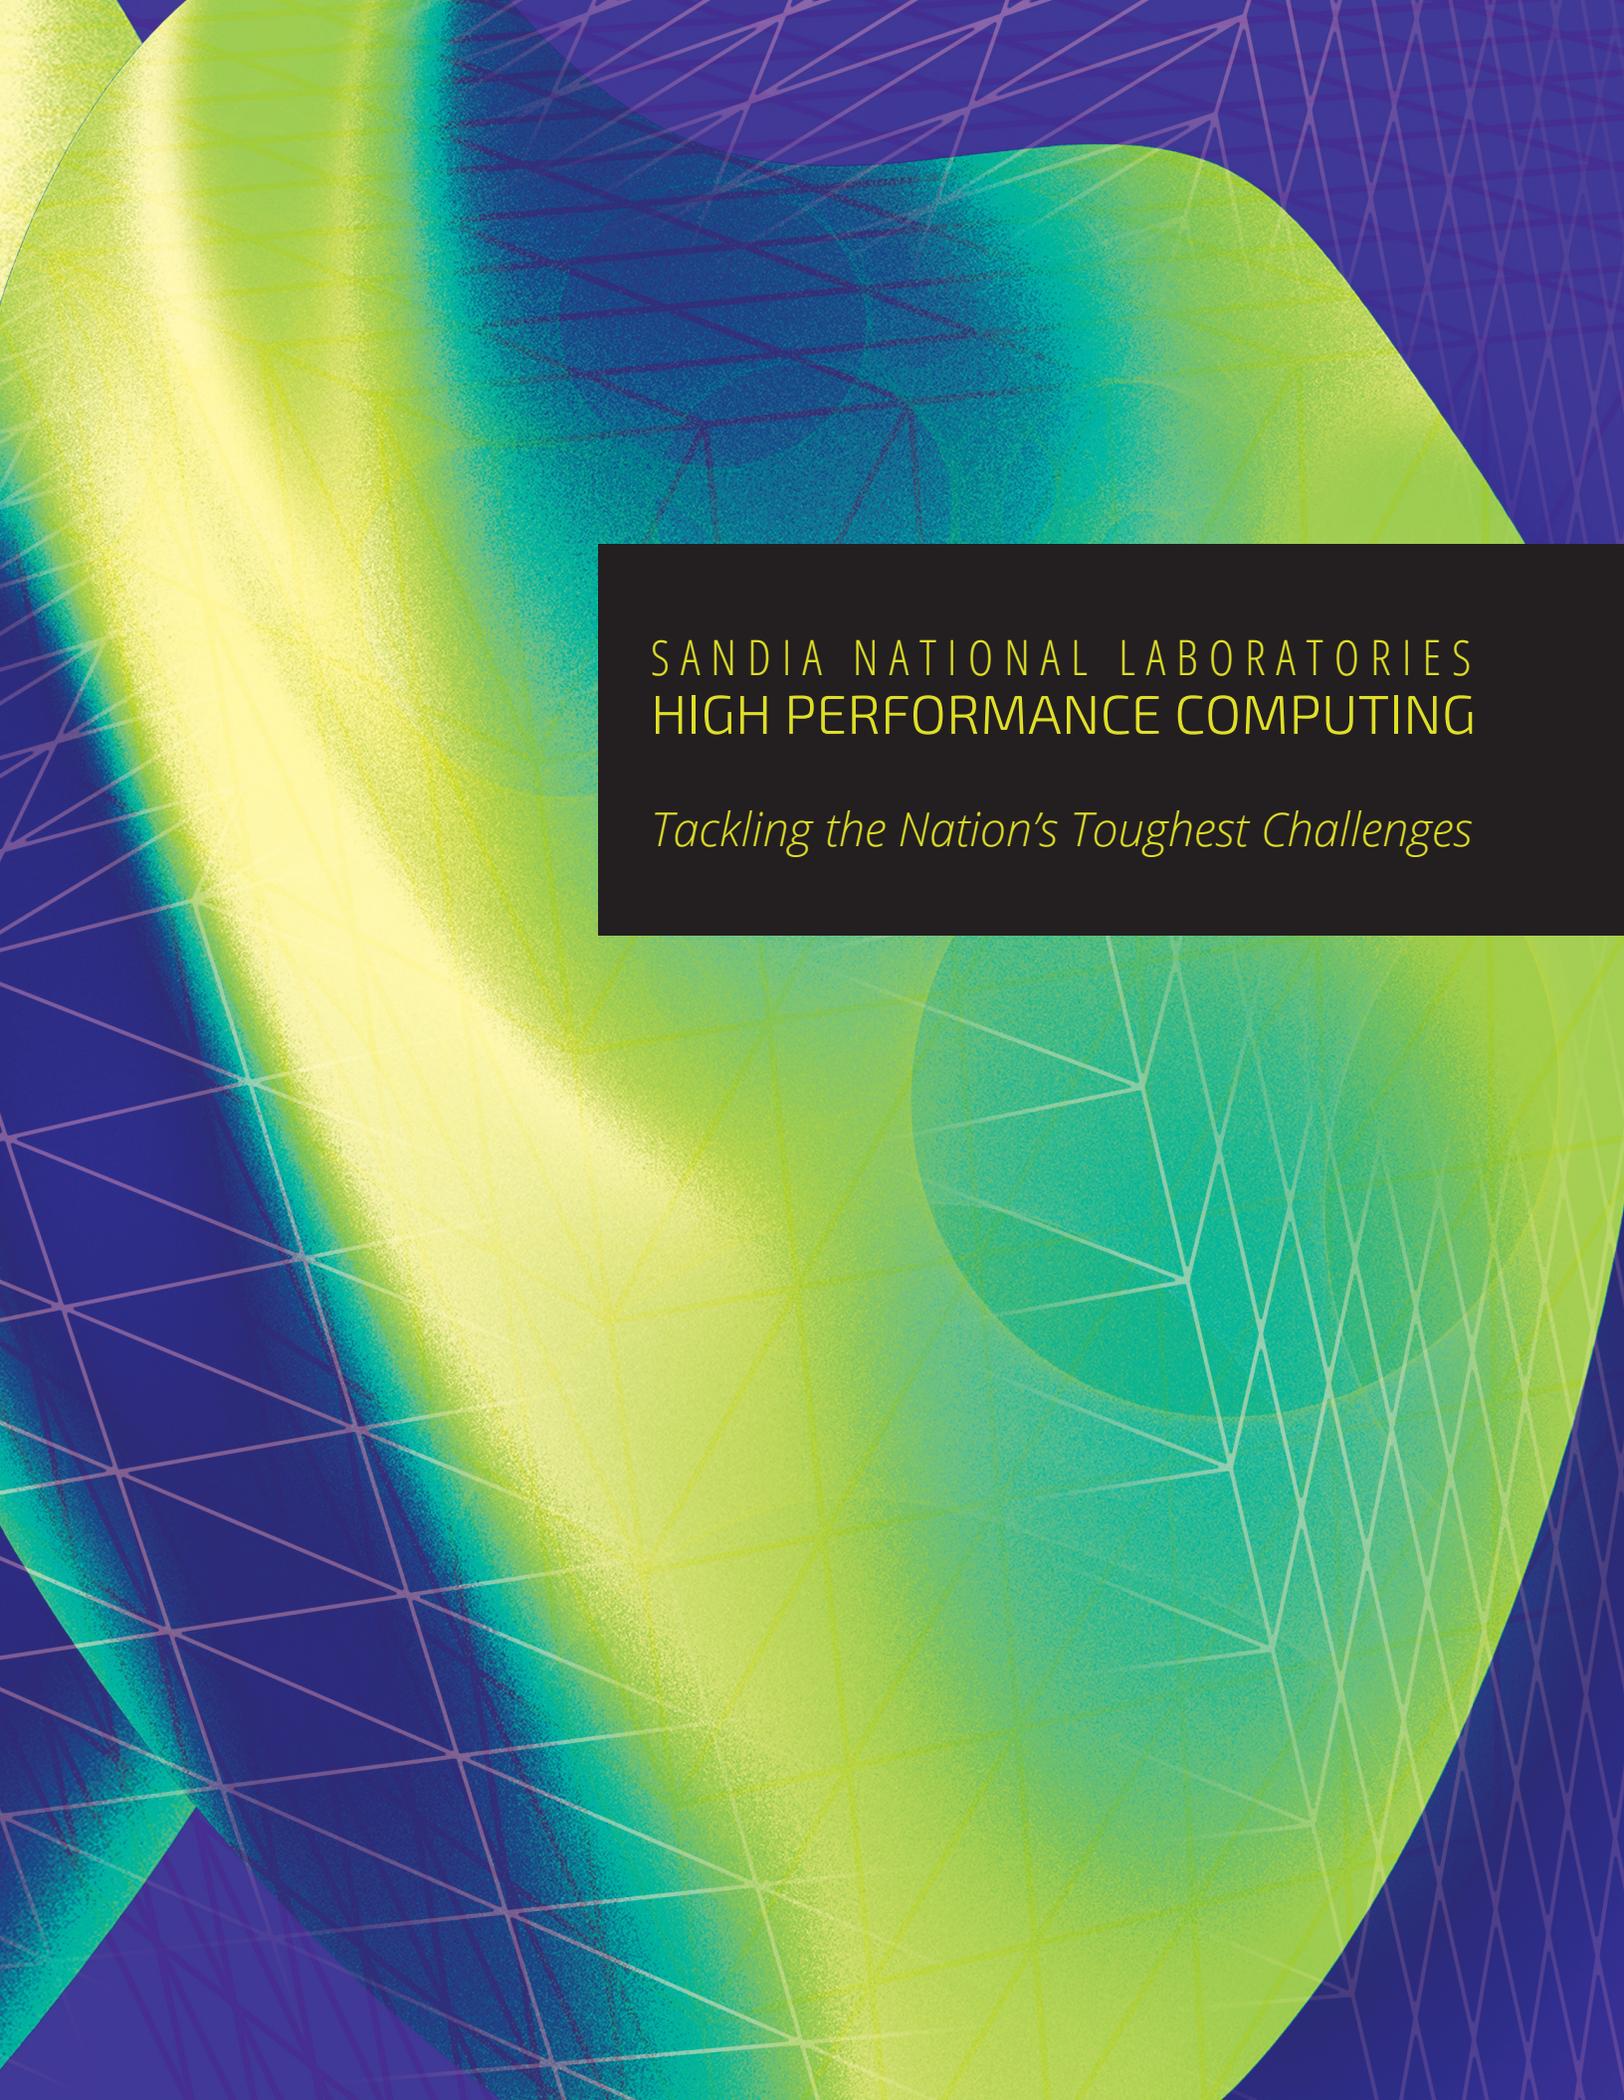


2022  HPC

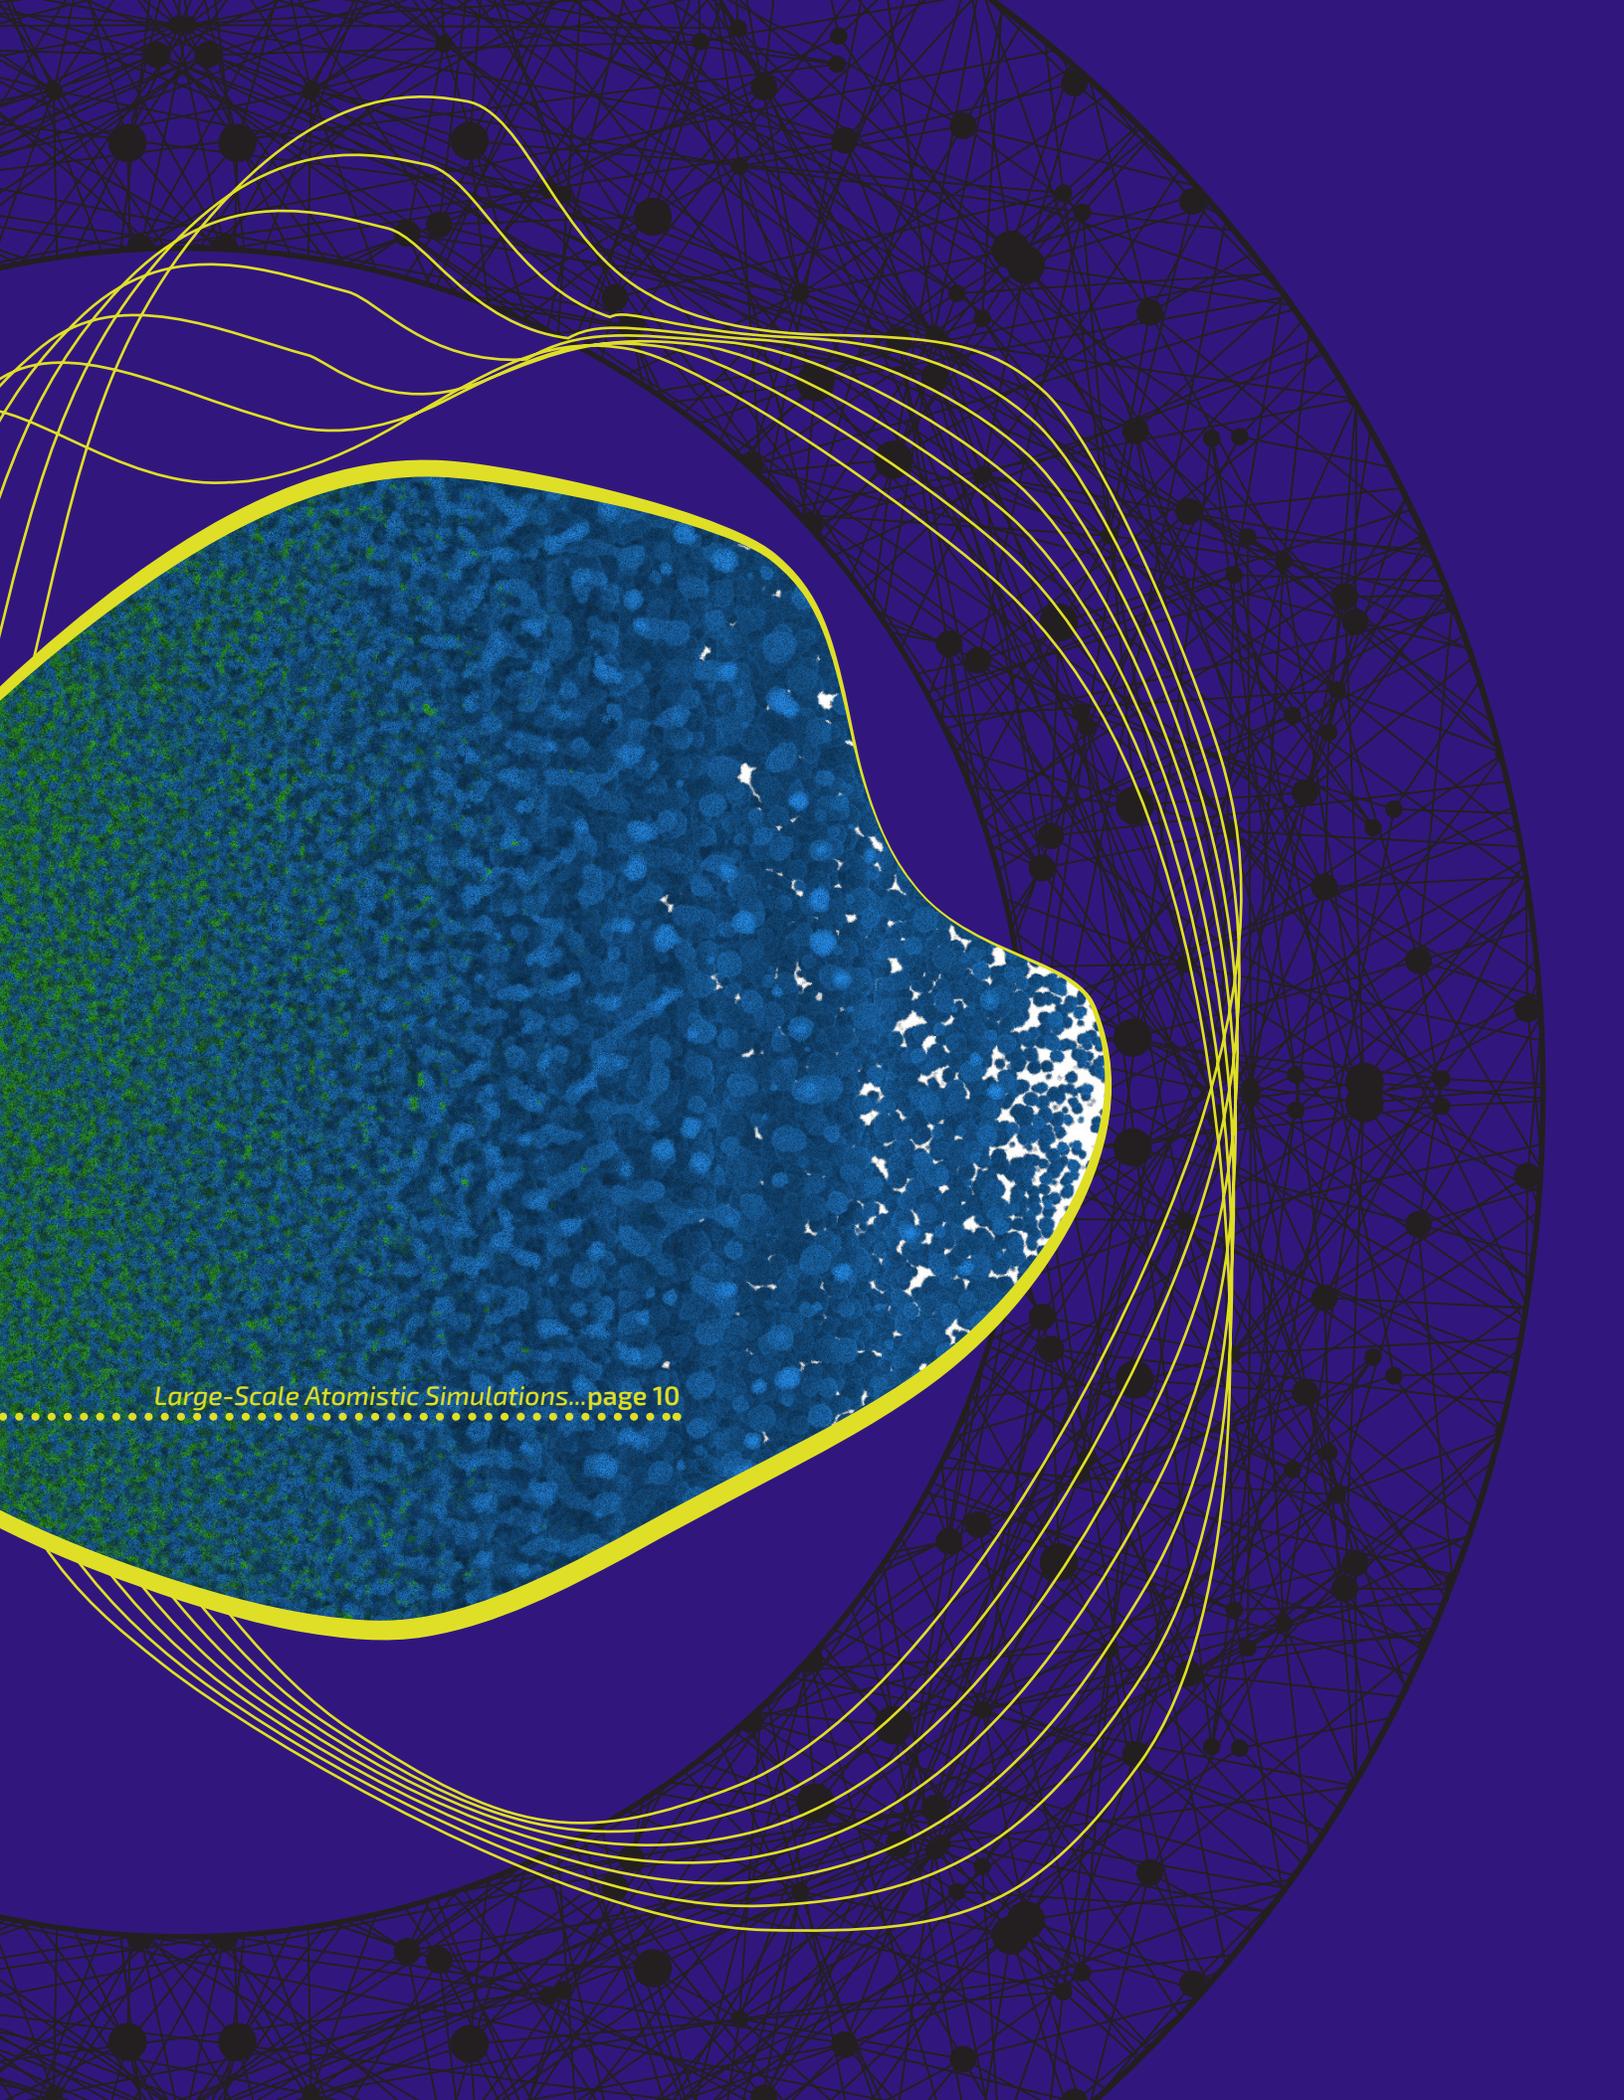
SANDIA NATIONAL LABORATORIES
HIGH PERFORMANCE COMPUTING





SANDIA NATIONAL LABORATORIES
HIGH PERFORMANCE COMPUTING

Tackling the Nation's Toughest Challenges



Large-Scale Atomistic Simulations...page 10

Contents

3 Director's Message

MISSION FOUNDATIONAL

6 Credibility Framework for Intrinsic Credibility Evidence in Analysis Workflows

10 Large-Scale Atomistic Simulations Investigate Expansion of Molten Metal

14 Neural Mini-Apps: Lighter, Faster Diagnostic Tests for Neuromorphic Computing

16 Neuromorphic Object Recognition at the Speed of Light

20 Uniform Mesh Refinement Enables Ultra-High-Resolution Simulations

22 Simulating Radiation-Induced Shocks in Thermal Spray Coating Materials Could Help Protect Diagnostic Equipment at Pulsed Power Facilities

EMERGENT COMPUTING

28 HPC-Derived Models That Smooth Re-Entry for Hypersonic Vehicles

32 Electromagnetic Environment Modeling for Nuclear Weapons

36 LEONIDAS: Leveraging the International Space Station to Demonstrate Advanced Sensor

40 Next-Generation Concentrating Solar Power Particle Receivers Enable Inexpensive Electricity Generation with Energy Storage

44 Valhalla: Designing Custom Satellites for National Security Missions with High-Performance Computing in Weeks, Not Years



VALHALLA...page 44

DIRECTOR'S MESSAGE



Labs Director,
Dr. James S. Peery

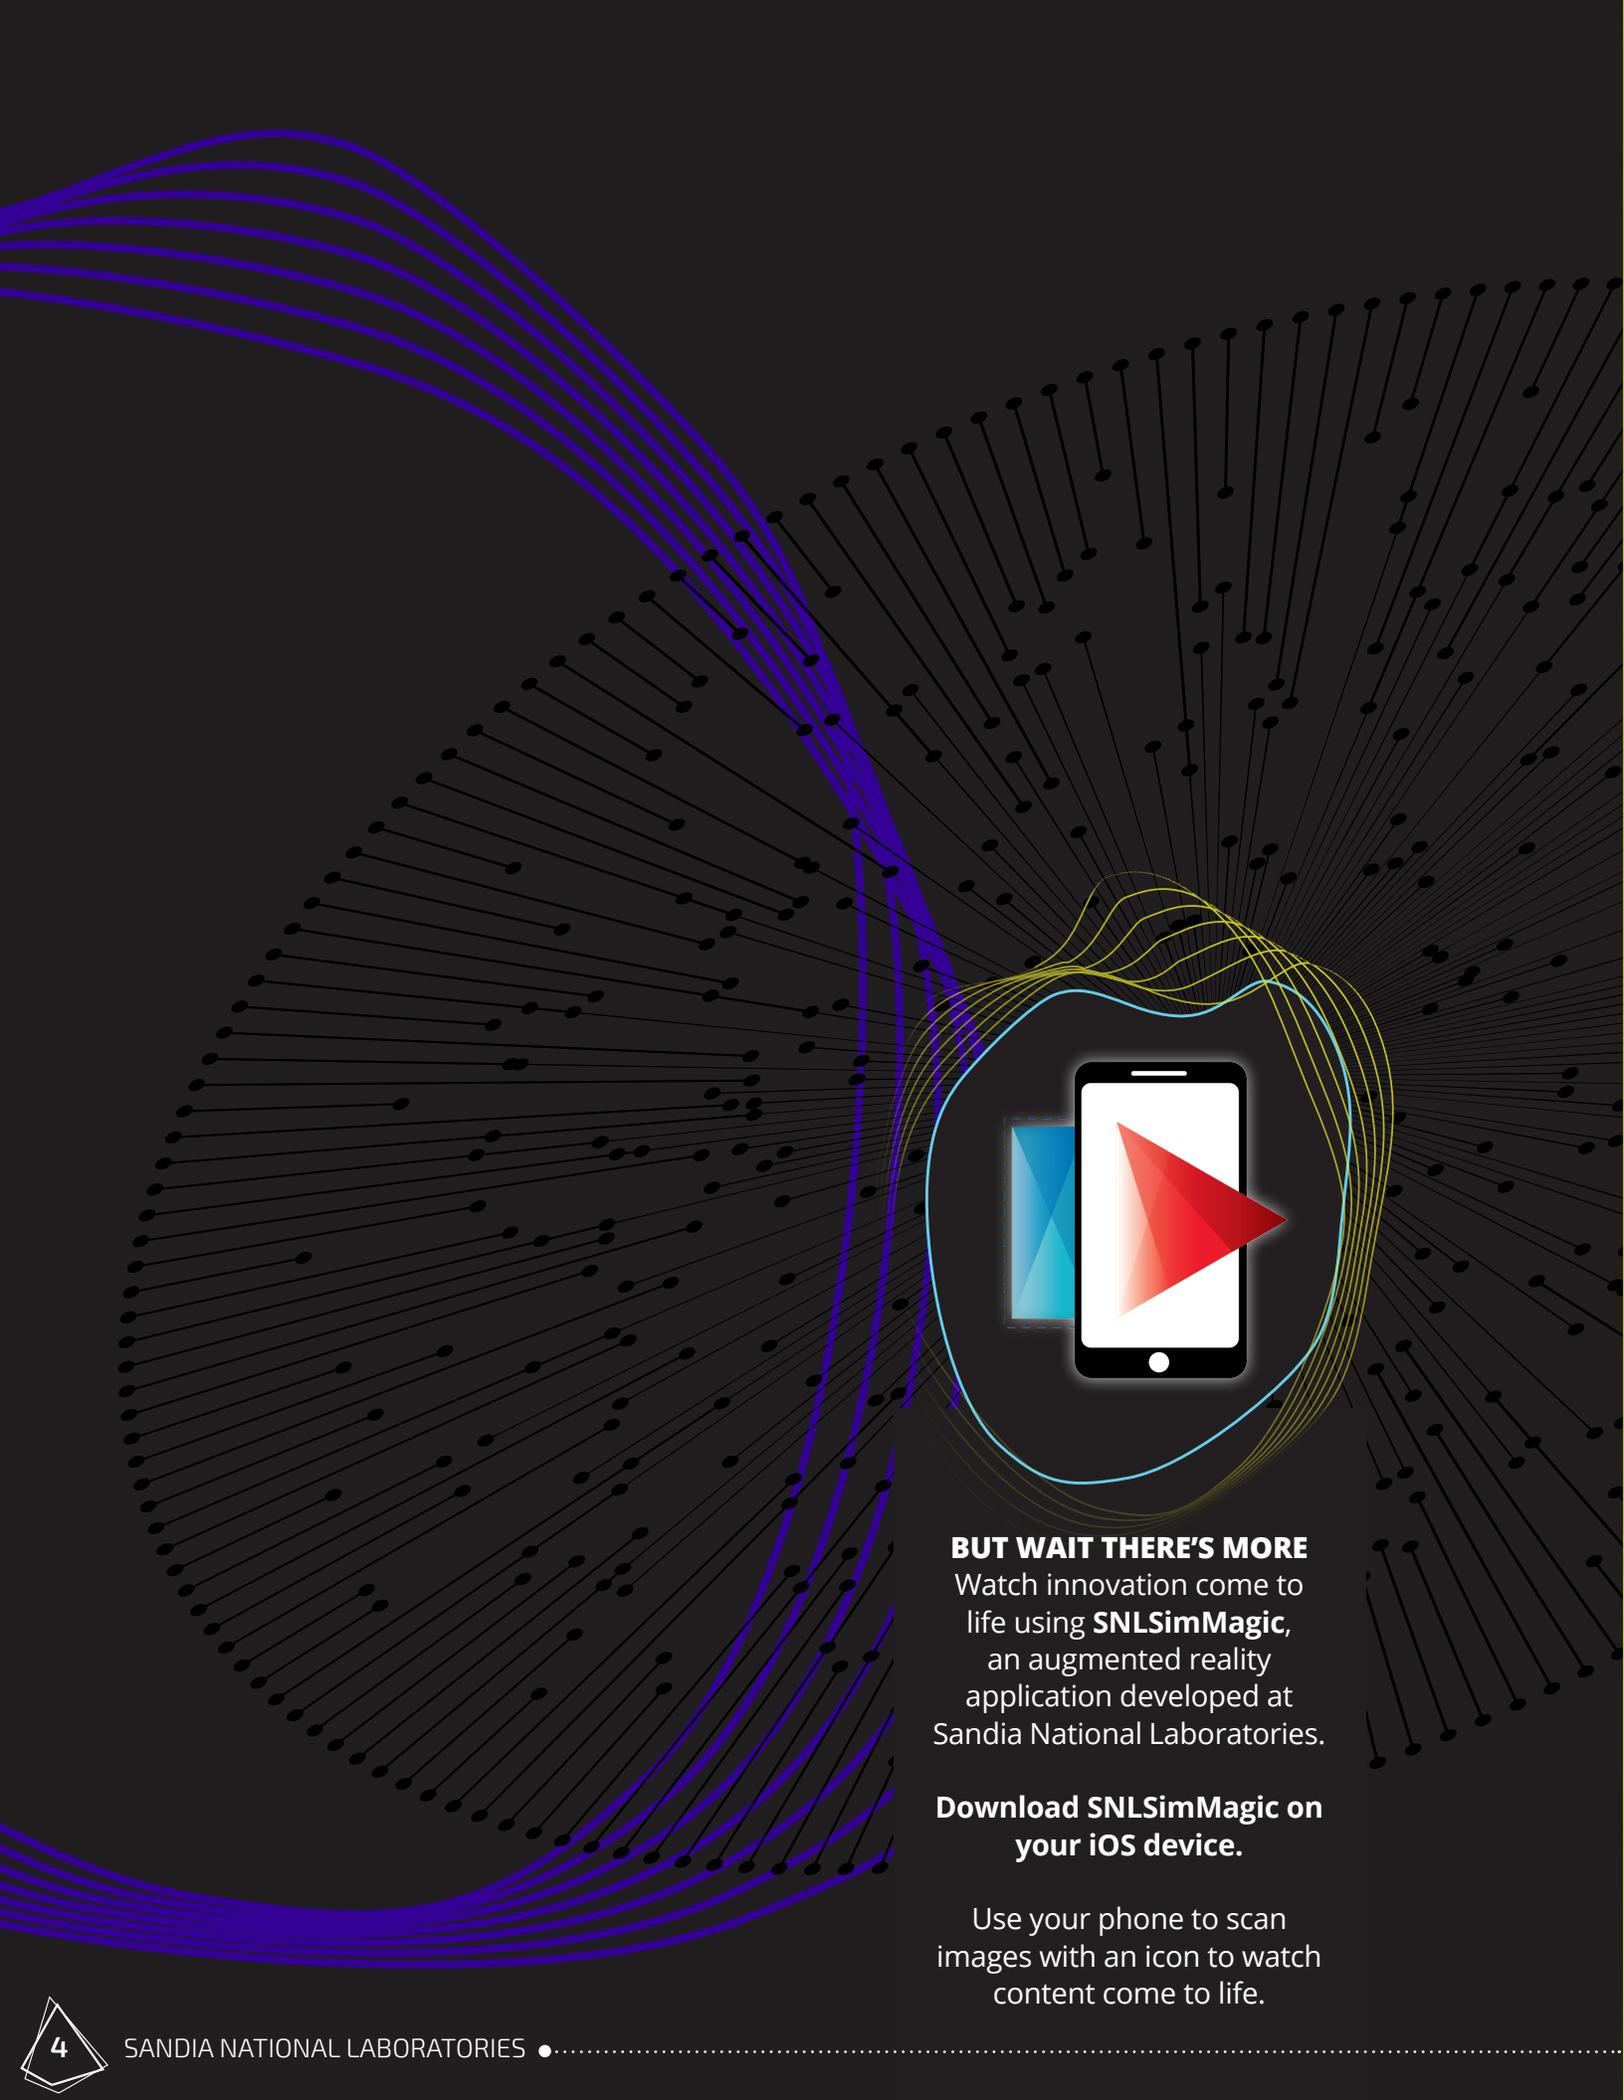
High-performance computing (HPC) is integral to the success of Sandia National Laboratories, now and in the future. With the arrival of the tri-lab “El Capitan” computer located at Lawrence Livermore National Laboratory next year, the National Nuclear Security Administration (NNSA) will enter the Exascale era with computers that make one quintillion calculations per second and let researchers solve problems that once seemed beyond reach.

Sandia is a world leader in HPC, developing solutions to the most sweeping and critical challenges our country faces. In this HPC Annual Report, you’ll see our researchers using supercomputers in mission areas including nuclear stockpile stewardship, hypersonics, machine learning, energy, material design, and more. HPC is instrumental in Sandia’s nuclear deterrence work. It helps us understand how materials respond to harsh, rapidly changing environments and ensures that weapons will perform as intended, which is essential to national security.

The forces behind our HPC breakthroughs are advanced physics and engineering simulation codes developed at Sandia that let our scientists and engineers deliver accurate, high-fidelity results. We’re developing new algorithms and codes to ensure the potential of next-generation supercomputers is realized.

In this report you will see our advances in foundational meshing software that opened the door to ultra-high-resolution simulations in aerospace, national security, energy, and medicine. You’ll see many other examples of how HPC helps our country. Sandia’s Aerosciences Department used HPC developed under NNSA’s Advanced Simulation and Computing Program to run large-scale calculations of the turbulent flow past re-entry flight vehicles, generating data that will be used in future designs. Our HPC resources contributed to a concentrating solar power receiver using particle technology for low-cost energy generation with thermal energy storage. As we advance our national security concepts to product realization, these tools are an integral part to prototyping through engineering development to qualification and finally sustainment.

With new and emerging challenges, it is critical to continue expanding our computing capabilities and resources to remain on the cutting edge of innovation. Sandia will be a driving force in the decades to come, turning big ideas into reality as we provide exceptional service in the national interest. 



BUT WAIT THERE'S MORE
Watch innovation come to
life using **SNLSimMagic**,
an augmented reality
application developed at
Sandia National Laboratories.

**Download SNLSimMagic on
your iOS device.**

Use your phone to scan
images with an icon to watch
content come to life.



A woman with blonde hair, wearing a light-colored short-sleeved shirt, is kneeling in a server room. She is holding a server tray filled with various components. The server racks in the background are filled with equipment. The entire scene is overlaid with a yellow-green tint and several blue lines that curve across the top and left sides, suggesting data flow or network connections.

MISSION FOUNDATIONAL



CREDIBILITY FRAMEWORK FOR INTRINSIC CREDIBILITY EVIDENCE IN ANALYSIS WORKFLOWS

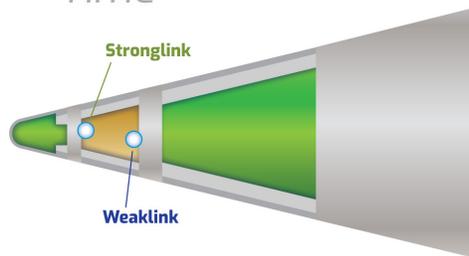
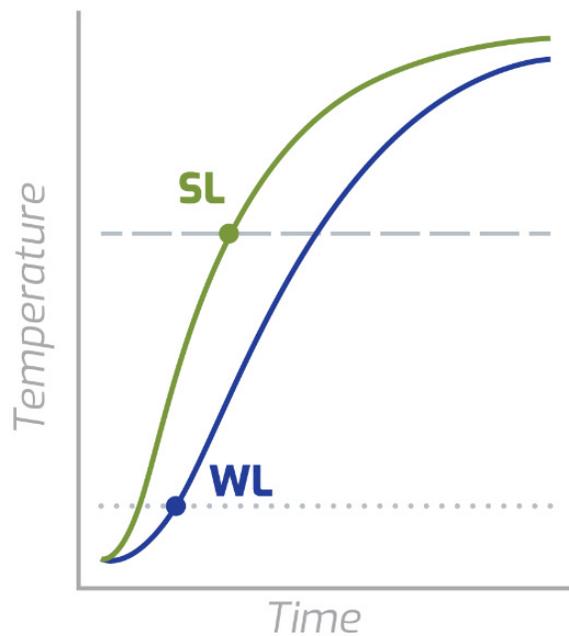
*Team: George Orient, Aubrey Eckert
Contributing Writer: Steve Scott*

*How confident are we that
a simulation is accurate?*

Scientists and engineers at Sandia and elsewhere rely on computational simulation (CompSim) to help address a multitude of highly complex problems, from predicting weather patterns or the spread of diseases to understanding the physical forces at work in a weapons system.

One of the central challenges in developing simulations of this type revolves around a crucial question: how confident are we that the simulation is accurate?

Answering that question involves the significant task of gathering and presenting “credibility evidence,” key indicators from the simulation’s results that show fidelity to established benchmarks and expectations.



The Credibility Framework, a recently released, open-source application, streamlines that work by providing an integrated suite of tools for planning and carrying out credibility processes.

At Sandia, a methodology for credibility processes has been in place for the nuclear enterprise for more than a decade. Key components of the methodology seek to identify and rank computational capability gaps and to focus the CompSim team's attention on crucial facets of credibility: code verification, physics and material model fidelity, representation of geometric fidelity, solution verification, uncertainty quantification and validation.

Gathering evidence for several of these credibility "sub-elements" requires large-scale model assemblies that can only be executed on HPC platforms. For example, an uncertainty quantification study on the B61-12 weapon system took one week on Los Alamos National Laboratory's Trinity using 120 samples each running on 1200 processors for several hours. To record and present the needed data, development teams have previously relied on an ad hoc array of spreadsheets and other electronic documents.

The Credibility Framework (CF) application guides developers in planning for, gathering, organizing and presenting the large volumes of evidence needed to support a simulation's reliability. The application provides direct links to analysis and ensemble workflows, and to evidence recorded by the CompSim team in CF, establishing a hierarchical, organized gateway into the detailed technical basis of credibility. The CF app also supports communication between team members, peer reviewers and customers.

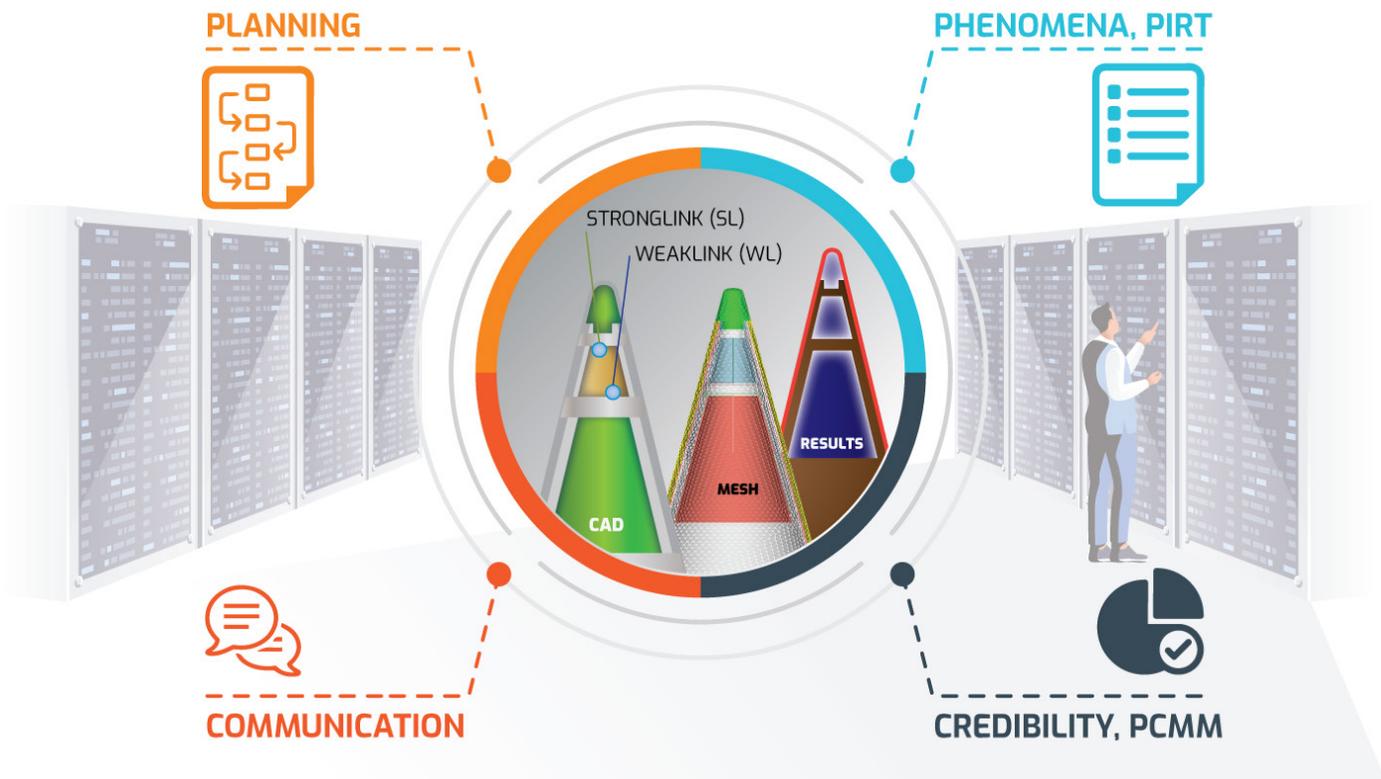


Figure 1: The Credibility Framework guides and organizes CompSim credibility process

The CF application represents three years' effort by a team of Sandia engineers and software developers working with small-business partner NexGen Analytics. The application comprises a suite of tools and a standardized workflow for implementing them in doing the following:

- Planning the simulation itself. The application provides procedures and dashboards for identifying the simulation's requirements, inputs and outputs, and an inventory of model uncertainties.
- Planning the collection of credibility evidence.
- Developing and carrying out activities that build credibility evidence.
- Gathering the credibility evidence.
- Reporting the credibility evidence. The application presents the top-level engineering information needed by reviewers when assessing the credibility of a computational model.

By organizing, reporting and archiving development of credibility evidence, CF supports CompSim model readiness and other forms of peer reviews. Before CF's public release, in a joint ASC V&V and IC level 2 milestone project, developers, analyst partners and UI/UX contributors successfully demonstrated use of the application on three nuclear deterrent (ND)-related exemplar

simulations developed specifically to showcase a CompSim model development approach that produces and records credibility evidence as integral part of the modeling process.

A range of differing mission goals and credibility evidence needs were represented by these exemplars:

- A thermal fluids simulation related to an unclassified reentry vehicle (URV).
- A structural dynamics simulation W87-1 motivated mid bulkhead assembly.
- A structural mechanics Joint Technology Demonstrator (JTD) bolted joint.

The resulting credibility evidence packages and linked CompSim analysis artifacts are archived in the Sandia Analysis Workbench (SAW). These will assist with training new analysts and V&V practitioners, regression testing. They will also serve as realistic test beds for strategic initiatives such as the Engineering Common Modeling Framework (ECMF) and, by proxy, Accelerated Digital Engineering (ADE).

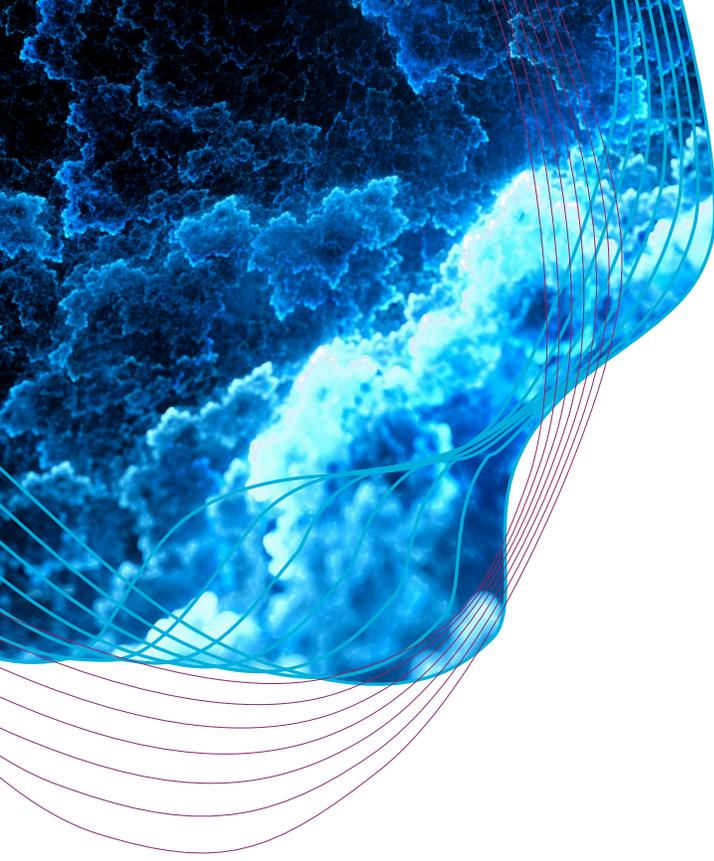
The CF application itself is now available to any organization performing high-consequence CompSim work requiring credibility evidence. As an Eclipse plug-in, it can be integrated with SAW or operated within a generic Eclipse instance.

On-going development will provide team collaboration capabilities as well as enhancements to the reporting subsystem that generates LaTeX/PDF and MS Word report from the content of a CF project.

In conjunction with an SDM (Simulation Data Management) system such as the one implemented in SAW, the information cataloged within CF is also useful for training and knowledge archival purposes as well as model sustainment efforts.

CF supports current ND programs such as W80-4 and W87-1 as well as on-boarding engineers to support training them in Sandia's credibility processes. Additionally, credibility evidence from CF is attached to analysis packages entered in the ECMF to inform long-term application of these models by archiving model credibility. [▲](#)





LARGE-SCALE ATOMISTIC SIMULATIONS INVESTIGATE EXPANSION OF MOLTEN METAL

*Team: Stan Moore, Mitchell Wood, Kyle Cochran,
John Carpenter, Aidan Thompson
Contributing Writer: Neal Singer*

*Can we visualize the
moment between liquid and
vapor at the atomic level?*

As computing power grows, materials science researchers are developing higher fidelity models to capture smaller length and briefer time scales — where the action takes place — of materials under stress.

Ultimate accuracy in modeling at the atomic level comes from quantum chemistry calculations such as density functional theory (DFT), which solve Schrödinger's equation and simulate interactions between electrons. However, these methods don't scale well with large number of atoms; quantum system sizes are typically limited to less than a thousand atoms.

Instead, a preferred method called molecular dynamics (MD) uses Newton's equations of motion to describe atom behavior. MD tracks the motion of billions of atoms by repeatedly calculating the atomic forces over very many, exceedingly tiny, intervals of time. MD is much faster and scales better than DFT, but it does not explicitly include electrons, a weakness. Instead, MD typically uses simple empirical models to relate atom positions to forces. The model parameters can be adjusted to approximately mimic specific materials, but these models can not accurately match DFT results in detail.

MACHINE LEARNING TO THE RESCUE

Recently, however, empirical models are being replaced with machine learning methods, which can approach DFT accuracy at a fraction of the cost, and also scale much better than DFT with respect to the number of atoms included in the target sample. MD simulations then can bridge the gap between expensive quantum chemistry calculations and lower fidelity continuum methods that omit atomic level details. (Continuum methods resolve the behavior of 3D “chunks” of continuous material, while MD resolves the motion of individual atoms.)

MD length scales are still extremely small compared to a grain of sand or a glass of water, and time scales are similarly small. However, with state-of-the-art computing power, MD simulations of billions to trillions of atoms can approach the length and time scales necessary for real-world experiments in some cases. Additionally, information from MD simulations can be used improve continuum methods in a multi-scale workflow.

Better understanding of the kinetics of the liquid to vapor phase transition is essential in building models that more accurately represent what is happening and thereby reduce the difference between code prediction and experimental measurement.

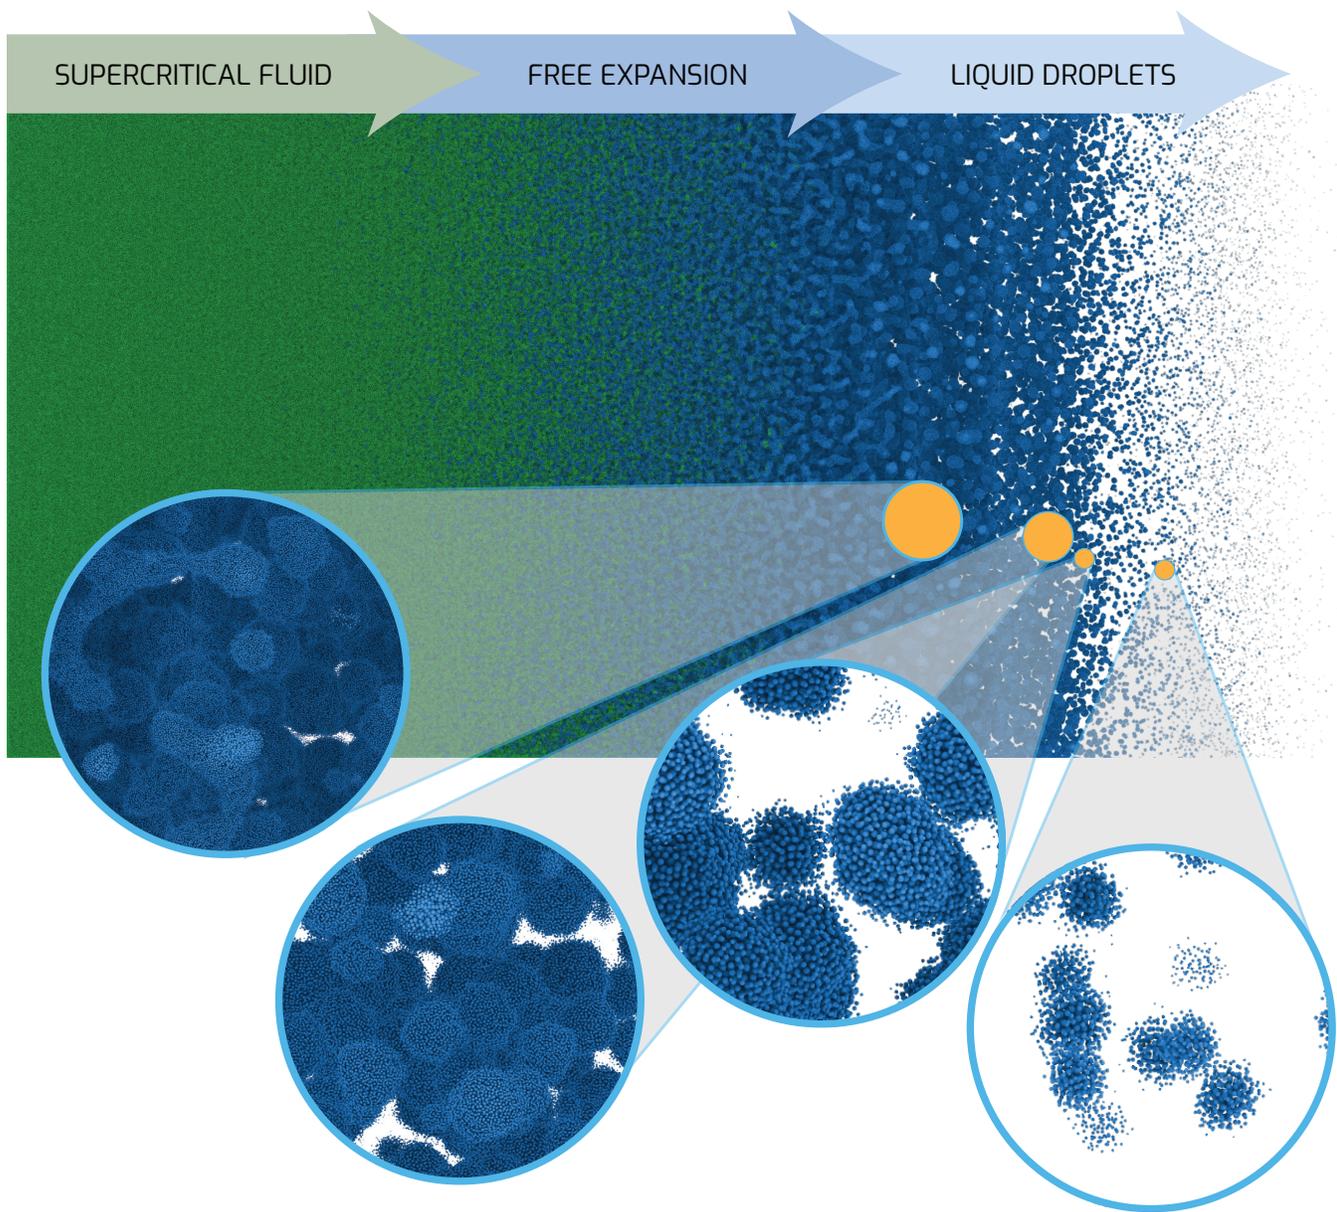
One such application is the free expansion of molten metal, for example when an electrical pulse rapidly heats and melts a thin wire up to a supercritical temperature. Supercritical means the molten metal is so hot that there is no longer a clear distinction between the liquid and vapor phases. However, when the supercritical fluid expands, the temperature drops below the critical temperature, where the fluid rapidly separates into liquid droplets and vapor bubbles.

...with state-of-the-art computing power, MD simulations of billions to trillions of atoms can approach the length and time scales necessary for real-world experiments in some cases.

Figure 1: Molecular dynamics simulation of supercritical fluid (green) expanding into vapor/liquid two-phase region (blue) resulting in spinodal decomposition and coalescence of liquid droplets. The LAMMPS simulation used ~1.5 billion Lennard-Jones atoms running on 1024 GPUs on the LLNL Lassen supercomputer. Atom diameter is visualized proportional to local density.



View video with SimMagic app



BRIDGING PRACTICE AND THEORY OF FREE EXPANSION AT SANDIA'S Z MACHINE

For some flyers and wire vaporization experiments on Sandia's Z machine, as well as other platforms, the expanding material enters this mixed phase (liquid/vapor) regime. Continuum simulation codes generally assume an equilibrium material approximation in the mixed phase and apply some type of averaging to account for how much vapor vs. liquid is present.

For a system in equilibrium, this is a valid approximation. However, a rapidly expanding material is not in equilibrium. When the leading edge of the expanding material reaches a low enough density below the critical temperature, it should transform from a liquid to a vapor at a material dependent kinetic rate, or even experience spinodal (spontaneous) decomposition. The material further inside will continue to expand based on how much back pressure is at and near the leading edge, and such information will continue to propagate through the rarefaction wave.

If the transformation rate from liquid to vapor is wrong due to incorrect kinetics approximations, then the expansion velocity that simulations predict will also be wrong after some time. This information propagates through the material until, in the case of flyers on the Z machine or Thor, the measured velocity at the front of a flyer and the simulated flyer velocity will diverge. Similar effects would occur on other platforms and has caused severe problems in wire initiation simulations where the energy to vaporize thicker wires is predicted to be less than seen in Z experiments. New computational work using MD is intended to lessen this disparity.

Researchers at Sandia started out using an efficient, empirical potential called the Lennard-Jones (LJ) model to examine free expansion with MD. One of LJ's nice features is that it can be written as a single equation with only two free parameters. Even though it is simple, the LJ model captures many relevant physics phenomena such as the transition from supercritical fluid to subcritical phase separation into liquid and vapor phases seen in free expansion. Huge (multi-billion atom) LJ simulations have been run on the Lawrence Livermore National Laboratory Sierra supercomputer to investigate the physics of free expansion.

The key advantage of MD in this case is that it avoids making any explicit assumptions about the material behavior (droplet formation, coalescence, break-up, surface tension, heat transfer) commonly needed for continuum models. Researchers are now transitioning to a more realistic — but computationally expensive — model for aluminum metal using the SNAP (Spectral Neighbor Analysis Potential) machine-learning method fit to DFT data. SNAP, developed at Sandia, can be used to scale up models to much larger length and time scales than DFT while preserving high accuracy. Thus, SNAP atomistic simulations will offer unprecedented insight into phase change kinetics and fluid microstructure evolution directly relevant to experimental studies of free expansion. ▴

Figure 3:
Sandia's Z-Machine

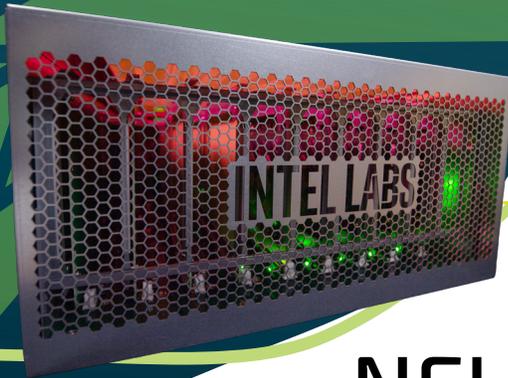


Figure 1: Neural Mini-Apps help researchers advance neuromorphic systems, such as the one seen close-up here.

NEURAL MINI-APPS: LIGHTER, FASTER DIAGNOSTIC TESTS FOR NEUROMORPHIC COMPUTING

*Team: James B. Aimone, Suma Cardwell, Frances Chance, Srideep Musuvathy, Fred Rothganger, William M. Severa, J. Darby Smith, Corinne Teeter, Craig M. Vineyard, and Felix Wang
Contributing Writer: Troy Rumlner*

Is it possible to compute more with less energy?

How can advanced science and engineering redesign hardware and software for supercomputers, satellites, drones and ground sensors to do more with less energy? World-leading research at Sandia looks to the brain — a remarkably energy-efficient computer compared to conventional computing systems — for inspiration.

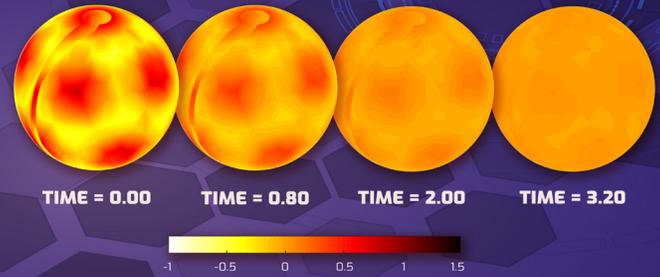
But while research has shown neuromorphic, or brain-inspired, computing can save energy, scientists need new tools to understand whether it can outperform conventional computing because traditional measures don't always apply.

Sandia is addressing this challenge with a suite of Neural Mini-Apps developed to enable researchers to understand, predict — and ultimately improve — the performance of neuromorphic computing technologies.

Mini-Apps are simple programs that stand in for more complex computations. Neural Mini-Apps are similar to mini-apps that have successfully advanced conventional HPC research over the past decade but have been tailored for neuromorphic research. The new apps run different tasks, operate on different hardware and sometimes diagnose different kinds of problems.

Each of Sandia's initial three programs is a task commonly run on neuromorphic systems and has been expressly developed to make it easy for researchers to collect data about what is happening in each part of the system and predict how small-scale technologies will perform in a larger system.

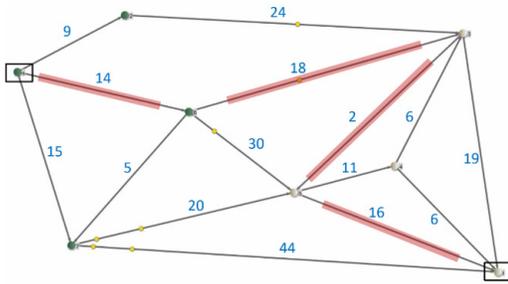
Figure 2: The Neural Random Walk mini-app, developed at Sandia, can produce this heat diffusion map on a neuromorphic computer while allowing researchers to probe how well their system worked. Image courtesy of Craig Vineyard.



View video with SimMagic app

These are:

- **Neuromorphic Random Walk.** This mini-app performs the same calculation many times with random inputs, the same way meteorologists model a storm's possible trajectories. Research has shown that small neuromorphic computers are especially efficient at this task.
- **Neural Sparse Coding.** Sparse coding is a way of modeling data. It breaks a data set into small pieces, simplifies each piece and reconstructs the set, reducing the complexity of computations.
- **Neural Graph Analysis.** Graphs represent relationships in data. They show how things are or are not connected. Graph analysis examines graphs to determine properties about them. If you think of a map as a graph, a property of the graph could be the shortest path between two points.



View video with SimMagic app



Figure 3: The Neural Graph Analysis mini-app, developed at Sandia, can identify shortest paths on a neuromorphic computer while allowing researchers to probe how well their system worked. Image courtesy of Craig Vineyard.

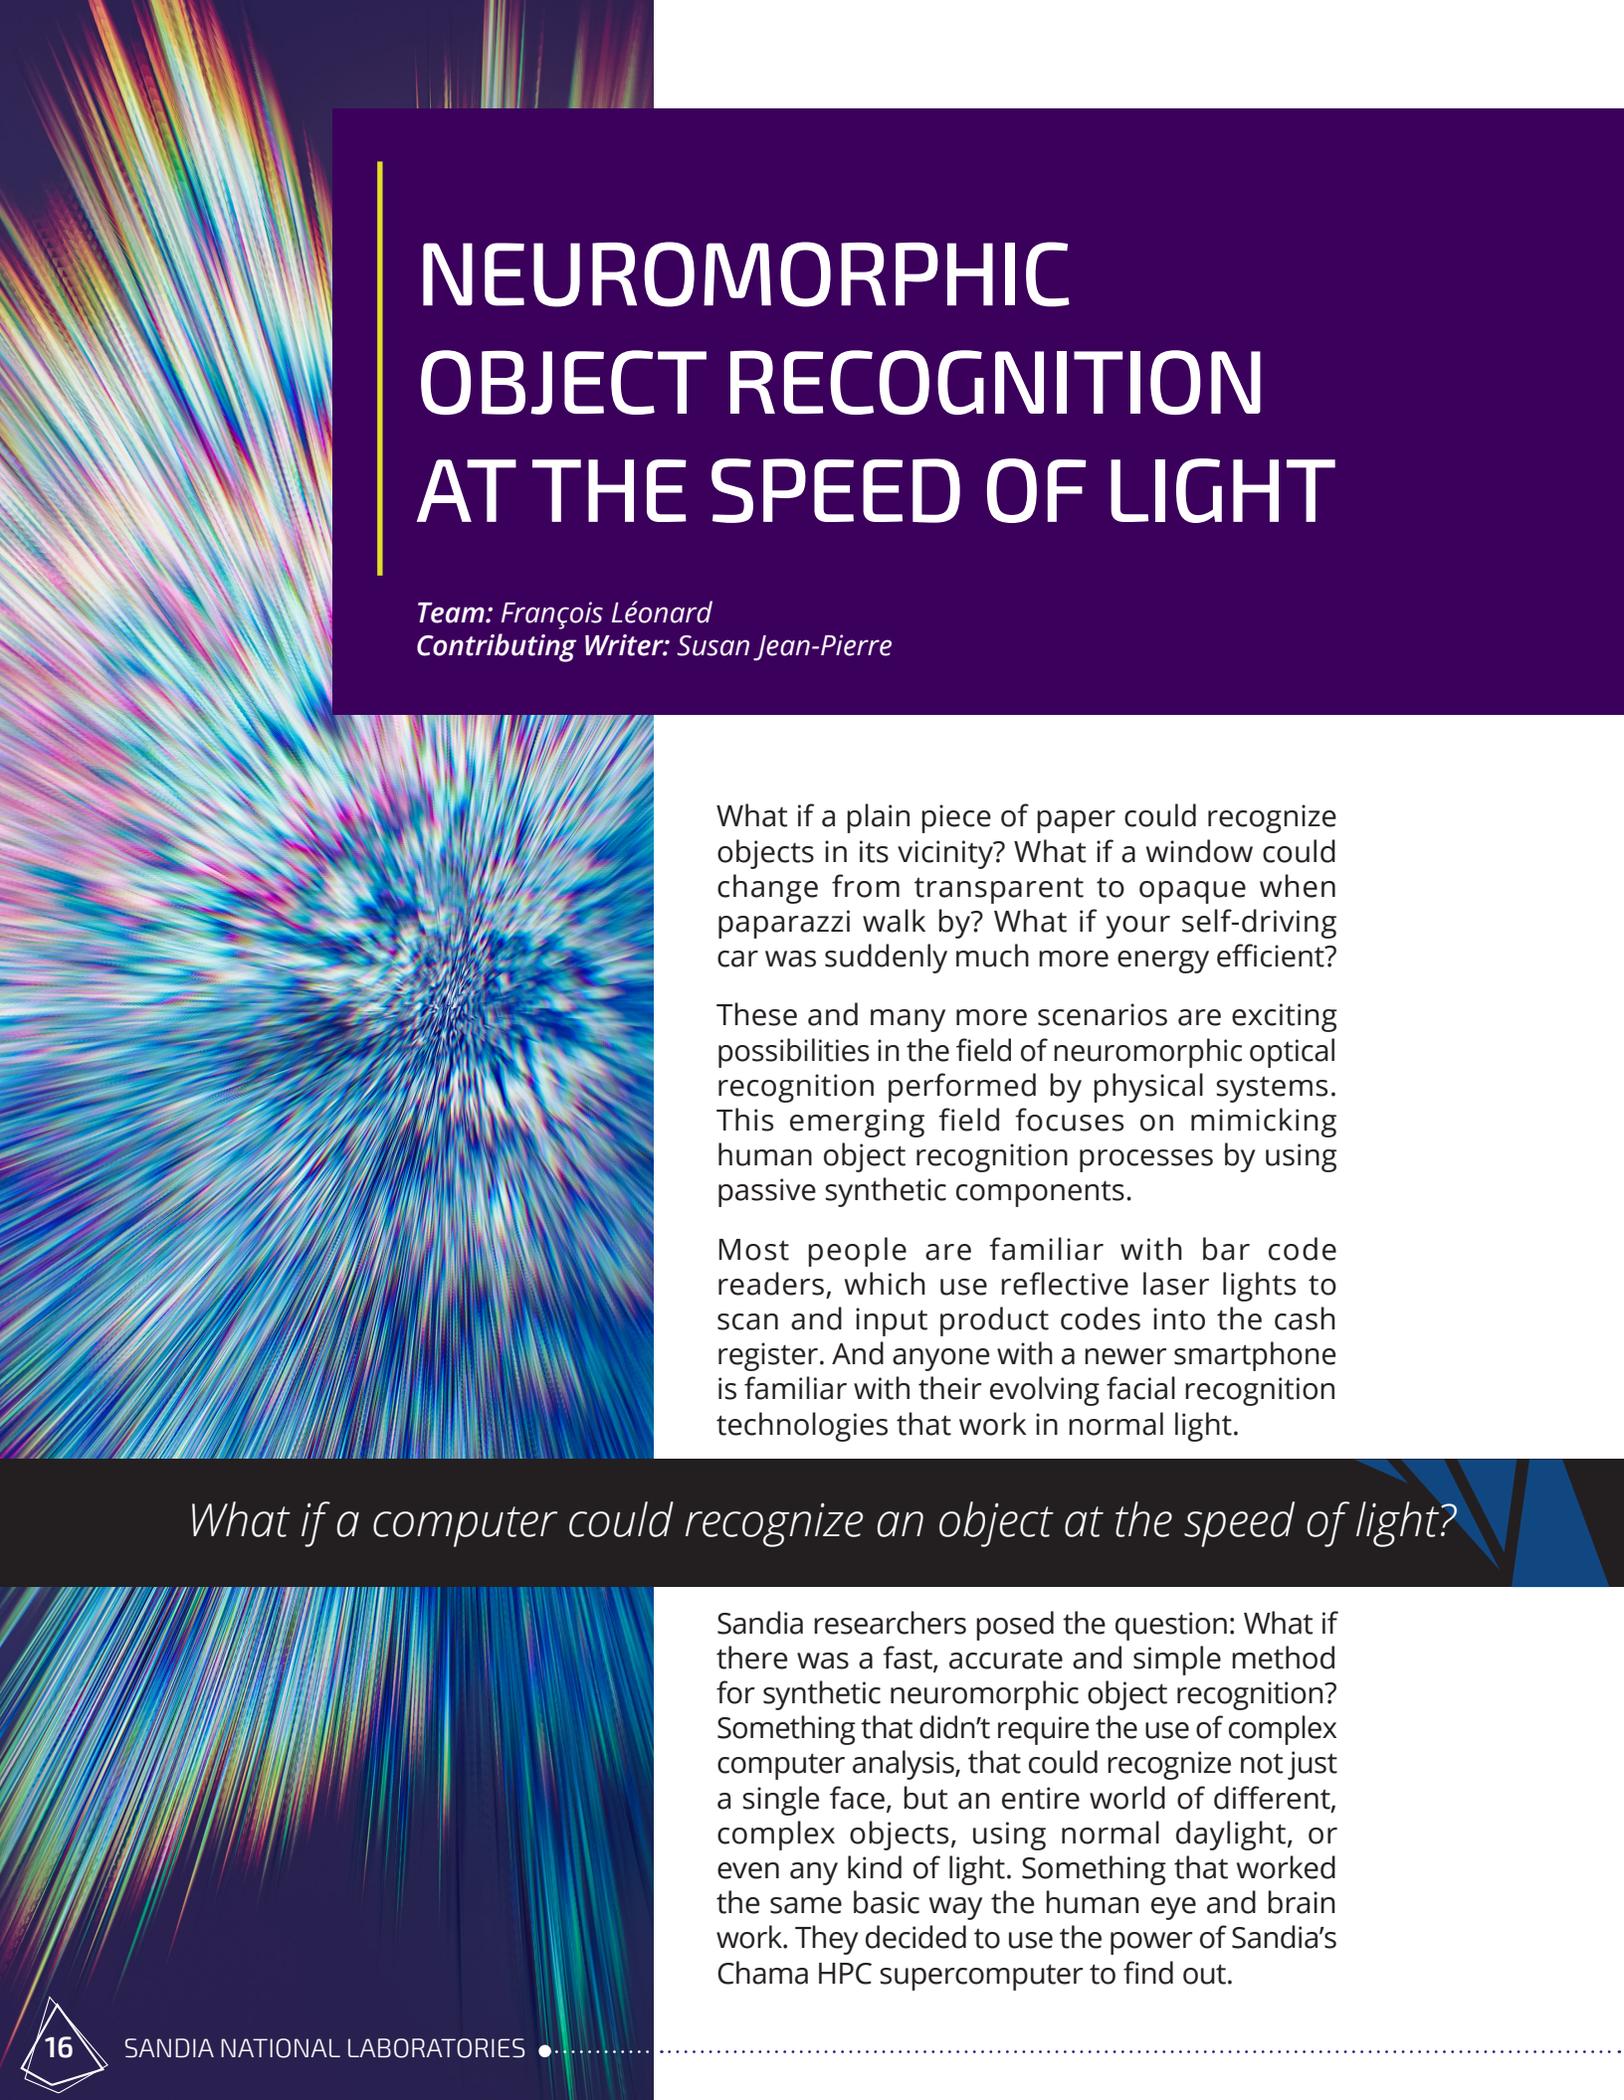
Neural Mini-Apps are similar to mini-apps that have successfully advanced conventional HPC research over the past decade but have been tailored for neuromorphic research.



Figure 4: Sandia developed Fugu, logo pictured here, as a software tool for building neuromorphic computing applications. Image courtesy of Craig Vineyard.

Sandia also created a companion app-building tool called Fugu. With Fugu, the worldwide neuromorphic research community will be able to create, customize and exchange neural applications, including mini-apps, advancing the frontiers of the technology.

Inside Fugu, apps are built out of so-called bricks, pre-designed neural circuits that configure artificial neurons to interact in certain ways. App designers arrange bricks to build complex operations from simpler ones. Sandia designed these bricks to be compatible with several common neuromorphic architectures. ▲



NEUROMORPHIC OBJECT RECOGNITION AT THE SPEED OF LIGHT

*Team: François Léonard
Contributing Writer: Susan Jean-Pierre*

What if a plain piece of paper could recognize objects in its vicinity? What if a window could change from transparent to opaque when paparazzi walk by? What if your self-driving car was suddenly much more energy efficient?

These and many more scenarios are exciting possibilities in the field of neuromorphic optical recognition performed by physical systems. This emerging field focuses on mimicking human object recognition processes by using passive synthetic components.

Most people are familiar with bar code readers, which use reflective laser lights to scan and input product codes into the cash register. And anyone with a newer smartphone is familiar with their evolving facial recognition technologies that work in normal light.

What if a computer could recognize an object at the speed of light?

Sandia researchers posed the question: What if there was a fast, accurate and simple method for synthetic neuromorphic object recognition? Something that didn't require the use of complex computer analysis, that could recognize not just a single face, but an entire world of different, complex objects, using normal daylight, or even any kind of light. Something that worked the same basic way the human eye and brain work. They decided to use the power of Sandia's Chama HPC supercomputer to find out.

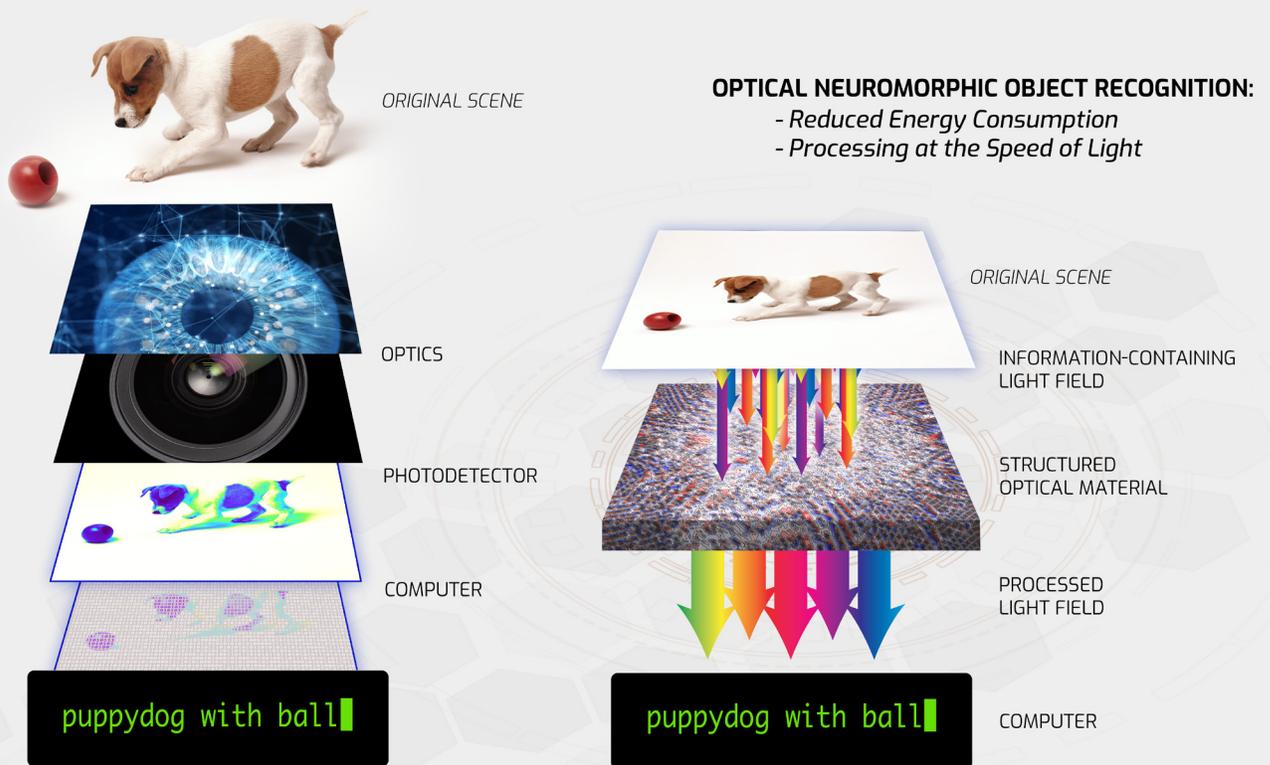


Figure 1: In biological and synthetic object recognition (left) the light is captured by optical elements (eye lens or optical lenses) and directed onto a photodetector (rods and cones, or a camera). This last process converts photons into electrons, which then are analyzed by an algorithm (in the brain or in a computer) to reach a conclusion ("what a cute puppy!"). The new approach studied here is shown at right, where a passive material replaces several steps in the process, enabling low energy consumption and fast processing.

CAN IT BE DONE?

The most complex part of object recognition is the computation stage. Figure 1 depicts the mental analysis in a biological system and the algorithmic analysis in a synthetic system. Both systems require the optical data to be converted to electrical data so it can be received by the brain or a computer chip for analysis. In synthetic systems, the conversion and computation stages are resource intensive, which means they cost both energy and time. In the human brain, the same task is performed quickly and easily because the neurons and synapses are highly interconnected, and the brain can generalize after seeing only a few examples.

The Sandia team decided to determine whether structured optical materials could be taught to act more like the human brain and recognize classes of objects after learning by example. Rather than design, manufacture, and test experimental hardware components, they developed complex algorithms written in Fortran to determine the feasibility of their design, train a virtual "smart chip" and then optimize their design.

Instead of using the more traditional approach of converting optical data into electrical signals and then analyzing those signals with complex algorithms, the team decided to combine the entire process into a single trainable structured optical material, requiring no photon-to-electron conversion and no algorithmic analysis.

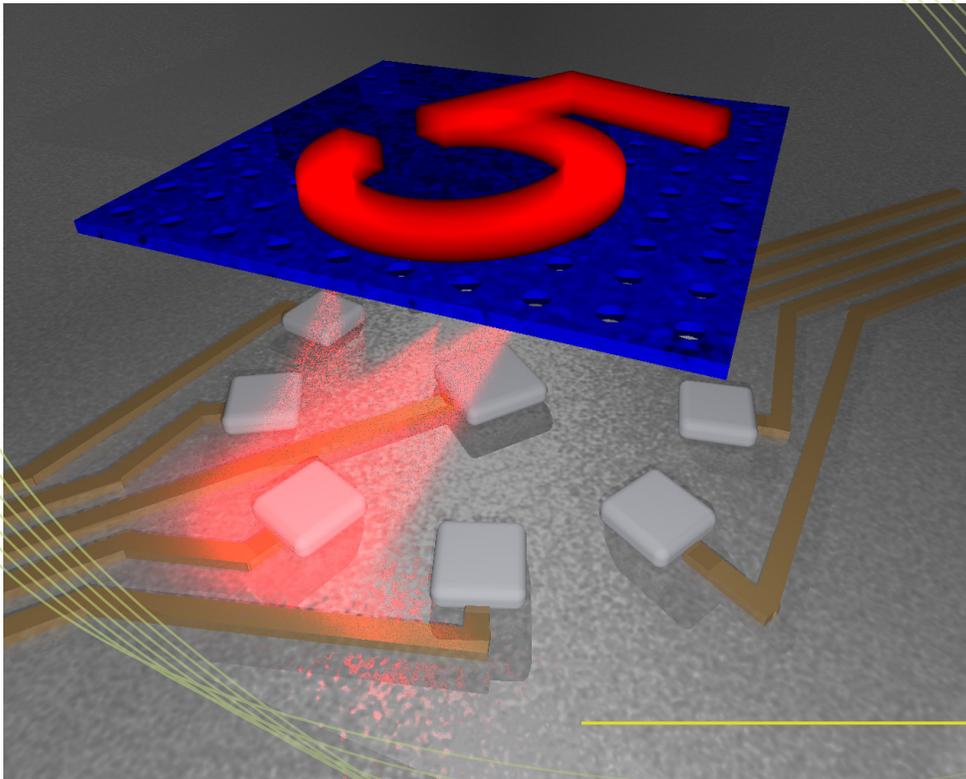


Figure 2: Three-layer design showing the input light field (the digit "5"), the structured material (blue plane) and the output plane containing single-pixel detectors for each class of interest.

DESIGN AND SIMULATION APPROACH

In the virtual world of supercomputer simulations, the team used Fortran to design a three-part system: An input light field corresponding to the scene of interest, a materials layer containing several thousand apertures that were filled to different thicknesses, and a detector layer consisting of 10 sensors, each sensor designated to recognize a class of objects (see Figure 2).

The team then trained the materials layer to recognize an object type, either a specific Arabic numeral (all of the numerals from 0-9) or a particular clothing type (all of 10 possible types), by illuminating the materials layer with thousands of example images from one of two standard image databases (MNIST for numerals, Fashion-MNIST for clothing types). As the images were projected onto the materials layer, the structure of each aperture was adjusted until it "learned" the structure needed to classify the images. Once the materials structure was fine-tuned in this process, the materials layer could perform inference, i.e., accurately recognize previously unseen images with high accuracy. The example shown in Figure 3a illustrates a light field in the shape of the handwritten digit "0" impinging on the trained material, leading to a light pattern on the sensor plane. The maximum intensity on the sensor plane is on the sensor designated to recognize a "0". As seen in Figure 3b, this optical system far surpasses the performance of several well-known electronic linear classifiers.

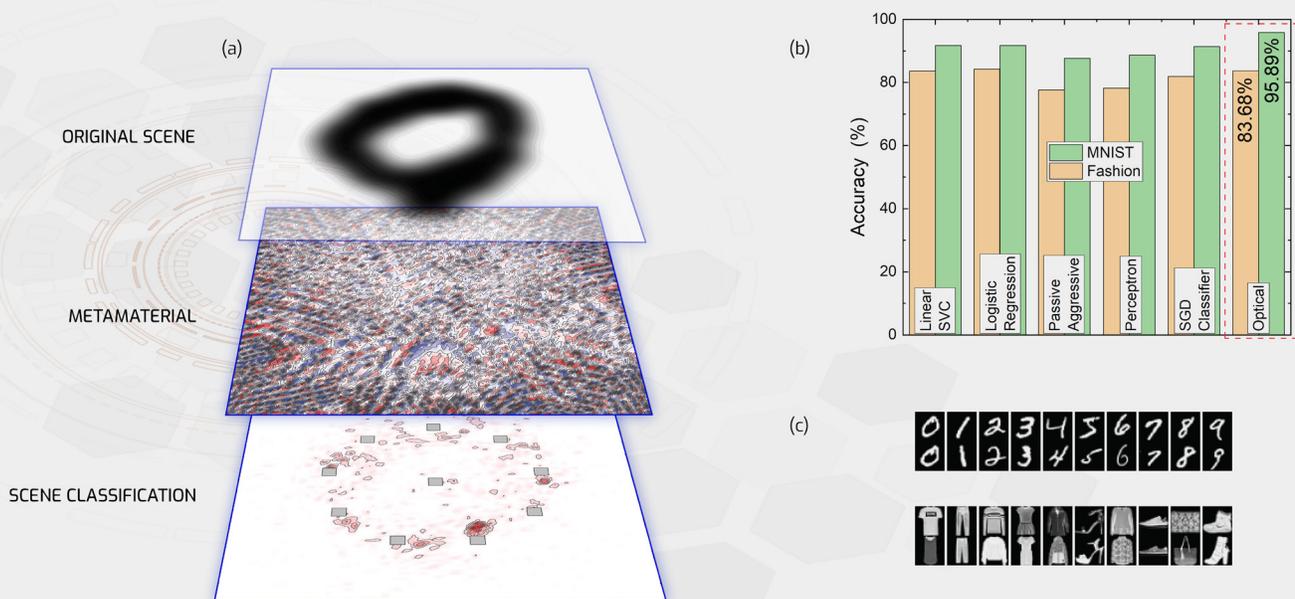
Detailed analysis suggests that these systems could use orders of magnitude less energy while performing inference orders of magnitude faster. In contrast to other approaches for on-chip optical computing, which use waveguides that reduce the speed of light propagation, this system uses light propagation in free space: it works at the speed of light. Also, only a few low-energy sensors need to be monitored, compared to the millions of pixels in a camera.

To demonstrate the concept, the team relied heavily on Sandia's HPC resources. They ran hundreds of single processor simulations to probe the complex design space for these systems. Furthermore, their recent work on incoherent light sources used about 1000 times more computational resources, requiring the development of a new multi-processor simulation approach. Because the researchers are directly including the physics of light propagation in their simulations, they did not rely on traditional software for machine learning but developed from scratch a Fortran simulation approach that allows them to fine-tune all aspects of the neuromorphic system.

FUTURE WORK

The team is currently testing their predictions with an experimental system they recently designed and built. This work opens up new possibilities for materials that can passively perform neuromorphic computing functions. While the current work focused on light fields, it could be broadened to analyze other fields such as thermal signatures, radio frequency signals, mechanical variations and chemical reactions. [↗](#)

Detailed analysis suggests that these systems could use orders of magnitude less energy while performing inference orders of magnitude faster.



View video with SimMagic app

Figure 3: (a) Example of an input light field in the form of the handwritten digit "0", the trained material, and the light intensity on the output plane containing the 10 sensors. (b) Results on the MNIST and Fashion datasets show high performance exceeding that of electronic linear classifiers. (c) Examples from the MNIST and Fashion-MNIST image databases.

UNIFORM MESH REFINEMENT ENABLES ULTRA-HIGH-RESOLUTION SIMULATIONS

*Team: Brian Carnes, Kevin Copps, Byron Hanks, Roshan Quadros
Contributing Writer: Troy Rummier*

*Can variations be predicted precisely
in billions of details at a time?*

Exascale computers, these are the first machines with enough processing power to run calculations on meshes made of tens or hundreds of billions of elements. Without a mesh, though, the potential of these computers would go untapped.

For an anxious air traveler, it's important to know the bumpy seats are in the back. For scientists and engineers, understanding how conditions like vibration, pressure and temperature change throughout an object is important, too — whether it's on the wing of an airplane or in the center of a nuclear reactor.

Using what is called a computational mesh, scientists have been able to divide digital objects into millions of polygons, or elements, to predict these kinds of variations in HPC simulations with impressive precision. Now, Sandia is enabling scientists to create meshes with billions of elements, far past the limit of standard meshing software, opening the door for ultra-high-resolution simulations in aerospace, national security, energy and medicine.

The new code, called Uniform Mesh Refinement (UMR), is especially important as a new generation of supercomputers is emerging. Called, exascale computers, these are the first machines with enough processing power to run calculations on meshes made of tens or hundreds of billions of elements. Without a mesh, though, the potential of these computers would go untapped.

Refining a mesh means subdividing pieces of a coarse, blocky one, while smoothing rough edges until it reaches the desired resolution. Creating a basic mesh is complex enough, so building one with billions of elements on conventional software takes too much time and creates a file that's too large to be practical. That's where UMR takes over and finishes the job with a more efficient algorithm.

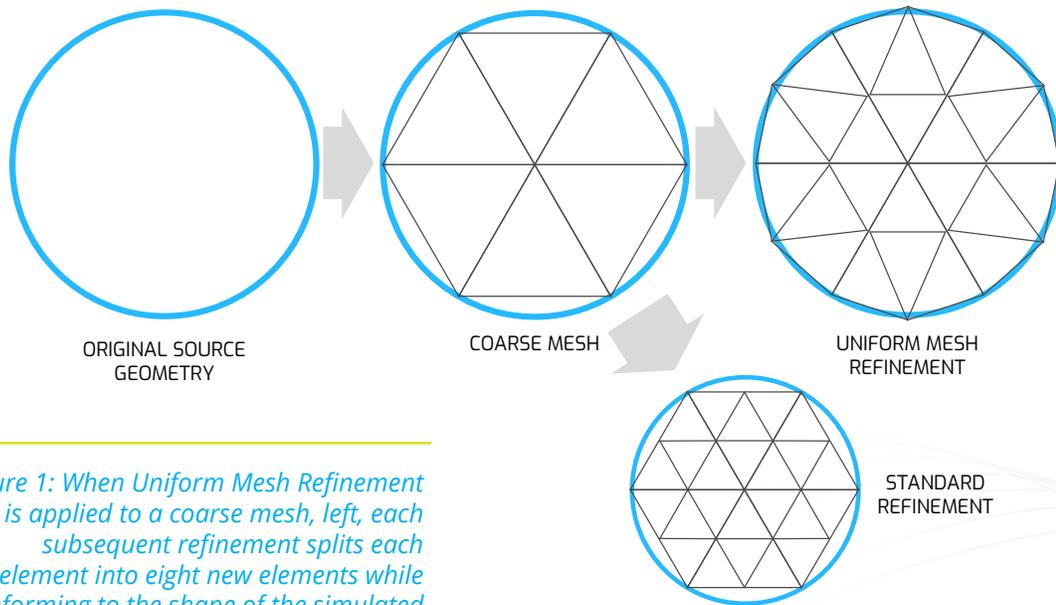
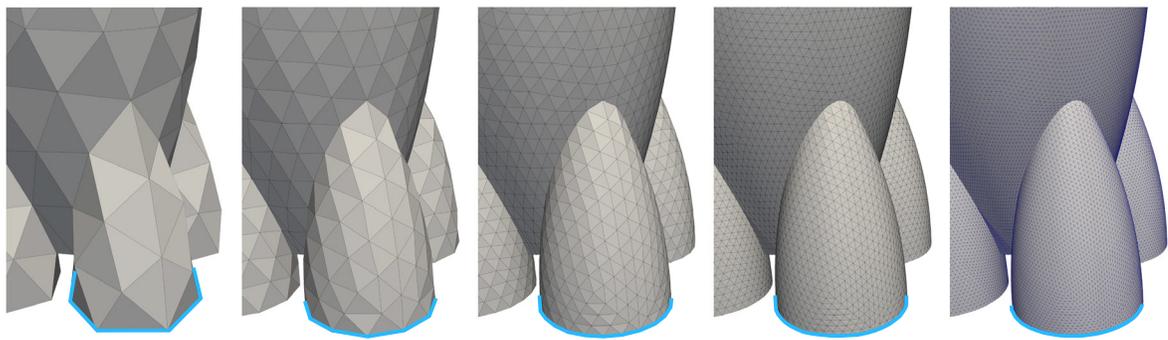


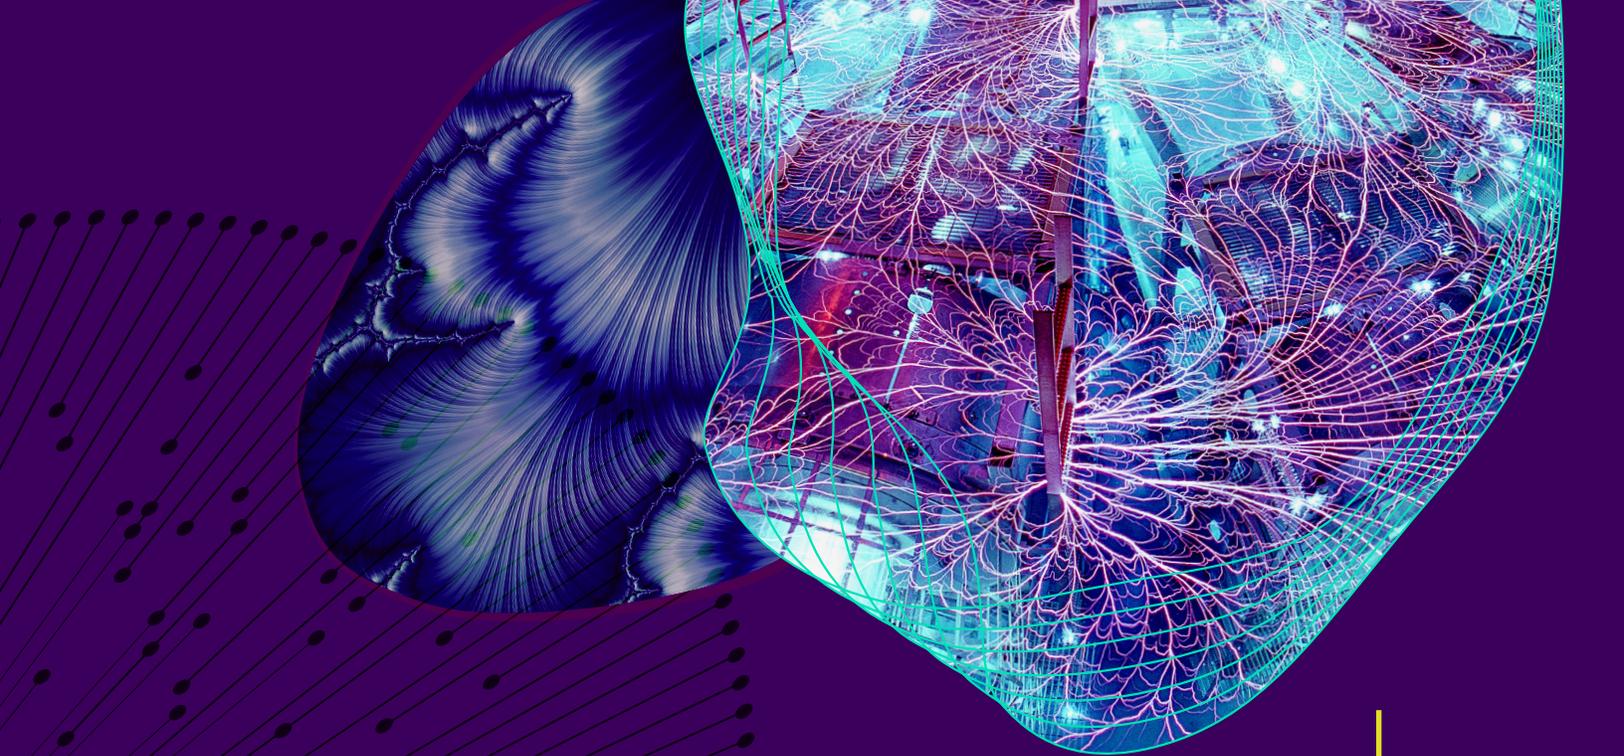
Figure 1: When Uniform Mesh Refinement is applied to a coarse mesh, left, each subsequent refinement splits each element into eight new elements while conforming to the shape of the simulated object. Image courtesy of Brian Carnes.

After scientists use available software to build a crude mesh, they can copy that file onto an HPC system. There, the new Sandia code takes advantage of parallel processing in ways conventional meshing software can't, so it runs much faster — generating over 100 billion elements in minutes. It can also operate entirely in the computer's temporary memory, so it doesn't generate massive files.

Originally developed for the code EMPIRE, a plasma physics simulation that will be used to advance aerospace, fusion and astrophysics research, Uniform Mesh Refinement is compatible with a commonly used, open-source mesh interface called SEACAS, which means many application codes at Sandia, other Department of Energy labs and other research institutions can use it.

The Sandia Geometry and Meshing team also created a version of the tool for laptops and small workstations. This version generates up to four billion elements in less than a minute to help researchers gauge quickly how fine a mesh they need to get the results they're after. It benefits individuals who do not necessarily need exascale computing power or have limited access to high performance computers and can't afford to run a simulation only to find they had used the wrong resolution.

This version of UMR is distributed along with CUBIT™ to about 2,000 users across about 165 sites within the United States for government purposes. ▲



SIMULATING RADIATION-INDUCED SHOCKS IN THERMAL SPRAY COATING MATERIALS COULD HELP PROTECT DIAGNOSTIC EQUIPMENT AT PULSED POWER FACILITIES

Can a simulation depict materials in motion under stress in high fidelity?

*Team: Jimmy Carleton, Theron Rodgers, Mikhail Mesh, Corbett Battaile, Bill Scherzinger, Jason Sanchez, Dan Bolintineanu, Carianne Martinez, Don Bruss, Aaron Olson, Brian Franke, Shawn Pautz, Nathan Moore
Contributing Writer: Kristen Meub*

Figure 1: In the thermal spray process, metal droplets impact, solidify and deform on a substrate to form a porous, layered structure.

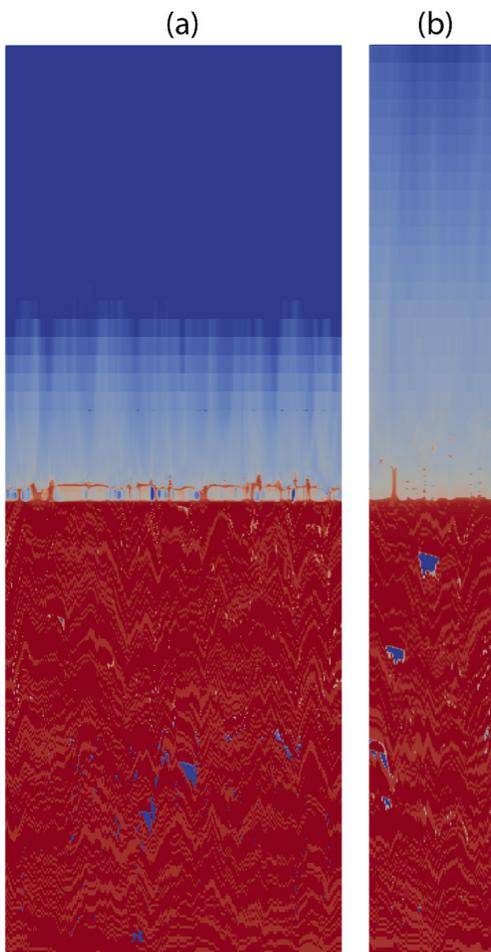
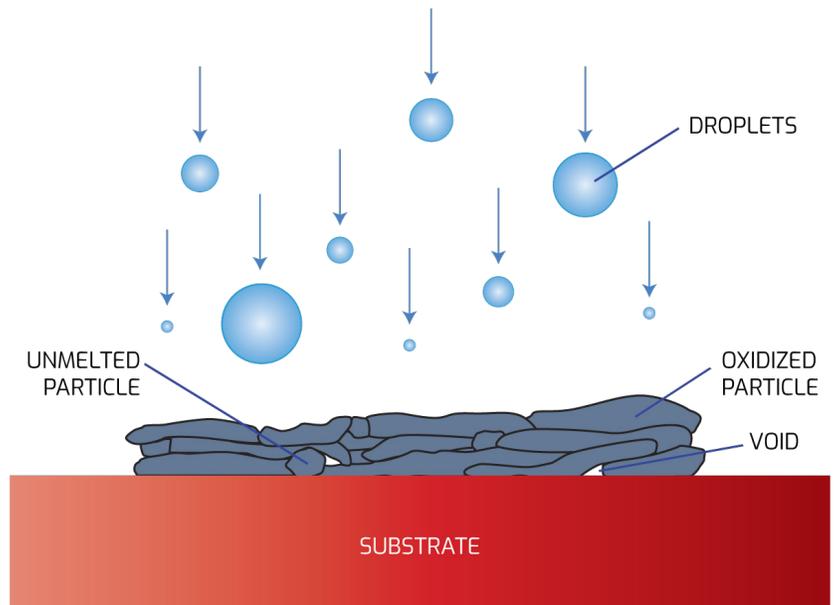


Figure 2: Density (log scale) from the simulation after the surface has been irradiated by (a) photons at the NIF and (b) electrons at Saturn. The metal forms high-density ejecta near the surface and a low-density plume. Pore compaction is seen near the surface.

Thermal spray coating materials could help protect sensitive equipment and sensors at pulsed power, high energy-density facilities like Sandia's Z machine from intense radiation, extreme heat and the flying pieces of debris generated during a shot. The ability to simulate how radiation and other thermally induced shocks travel through these spray-formed coating materials is important for understanding how the materials will perform and helps researchers analyze and design new coatings for next-generation pulsed power, inertial confinement fusion and other high energy density physics facilities.

Until now, previous quasistatic radiation transport simulations modeled just a subset of the physics occurring during a shot, showing researchers how much energy would be absorbed at any one point in space in the material, but unable to account for all aspects of the radiation environment. These old simulations also treated the spray coating material as a solid block of homogeneous material that didn't fully capture the details of its structure. When the coatings will be sprayed on equipment, metal droplets will impact, solidify and deform on a substrate and form a porous, layered surface comprised of blobs of metal and small pockets of air, not a solid homogeneous material.

Using Sandia's multiphysics code ALEGRA, researchers have developed fully coupled radiation hydrodynamic simulations that model radiation-induced shocks traveling through the

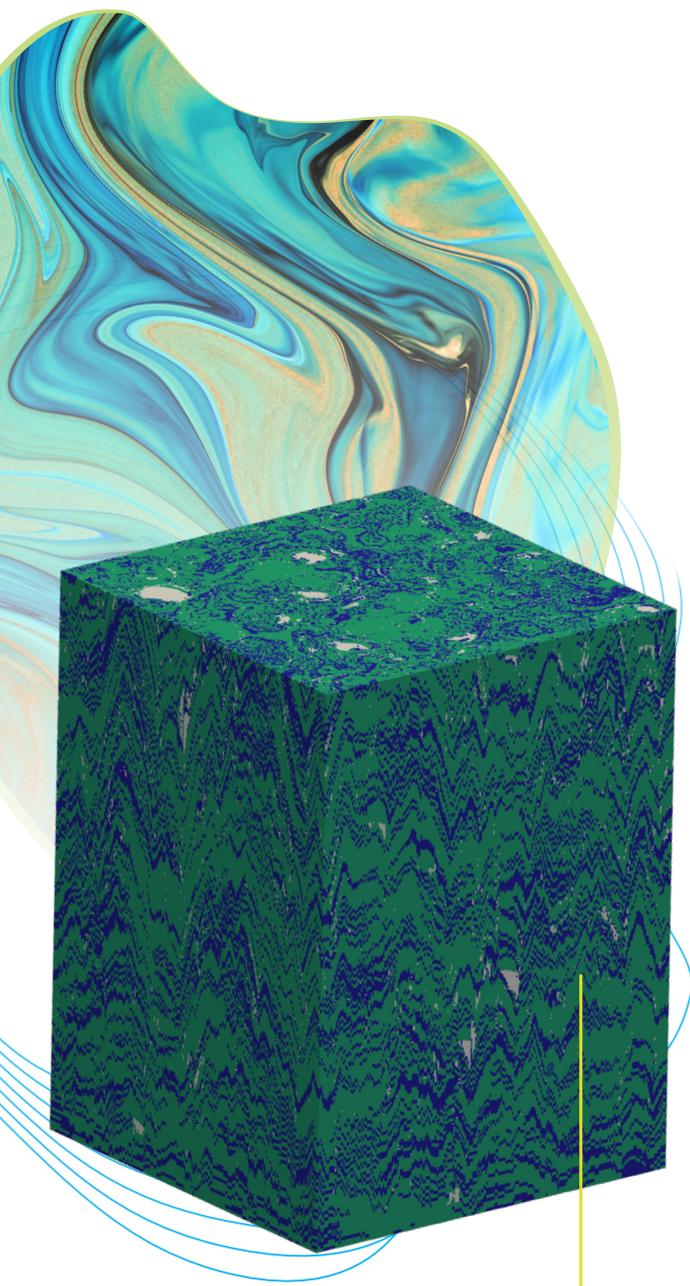


Figure 3: This simulated mesoscale geometry is a roughly 0.7 x 0.7 x 1.0 mm representative volume element containing metal “splats” (blue and green) and air-filled pores (light gray).

heterogeneous mesoscale structure of spray-formed materials. These HPC simulations solve equations that include the physics of both radiation transport and solid dynamics to predict how the material would evolve in harsh environments.

Because these high-fidelity physics simulations now include mechanics – materials in motion and stress – the team can see, at any point in time and at any point in space, what the state of the material is, what the velocity of the material is, what the stress of that material is and what the energy is.

When a pulsed power facility like Sandia’s Z-machine fires a shot, the radiation that is produced and hits the equipment is so intense that it sends a stress wave through the material. These new advanced simulations show how the mechanical stress wave moves through the material, which is something the previous simulations couldn’t do. Researchers can predict stress wave profiles and mechanical surface phenomena when the system is placed in different facilities. This information can help researchers design a material that not only survives the intense environment but also reduces the stress created when the same radiation pulse is applied in the future.

To better model the spray-coating materials, the team generated representative volume elements of the mesoscale geometry of the coatings by using Sandia’s SPPARKS Kinetic Monte Carlo Simulator with an algorithm that mimics real structures generated by the spray process (Figure 1 and Figure 2). ALEGRA couples with the Kull IMC and SCEPTRE radiation transport codes to do radiation hydrodynamics simulations of the stress waves generated when the spray-coated material is subjected to photon and electron radiation environments in pulsed power facilities such as Z, the National Ignition Facility and Saturn. The insights from these simulations and experimental observations could help researchers design effective debris and thermal protection systems for these facilities.

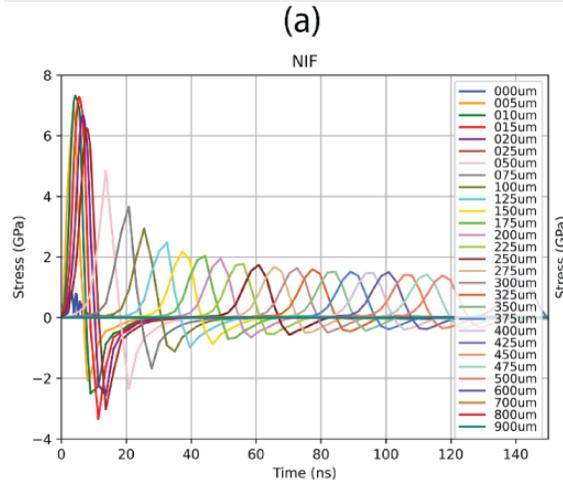
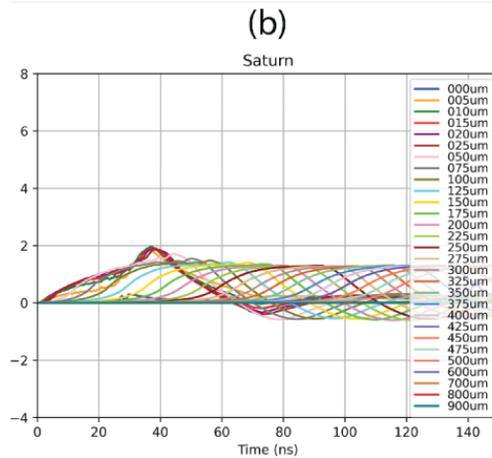


Figure 4:
Axial stress
(compressive
positive)
vs time at
multiple
depths below
the irradiated
surface
from the
simulation for
(a) NIF and
(b) Saturn.



In addition to their use in pulsed power facilities, refractory metals could potentially be used in inertial or magnetic confinement fusion applications, such as tokamak reactors, owing to their high temperature and corrosion resistance. The advent of spray-coating technologies and the ability to better simulate and analyze them could lead to these materials being used on reusable or even disposable machine components.

For future improvements, the team is working on improving the fidelity of the models to capture more of the relevant physics, including the interaction of stress waves with microcracks and better material model calibration for the constituent materials, whose material properties can differ from their non-sprayed counterparts.

This work was funded by NNSA's Engineering and Technology Maturation program and Sandia's Laboratory Directed Research and Development program. 

These new advanced simulations show how the mechanical stress wave moves through the material, which is something the previous simulations couldn't do.





EMERGING COMPUTING

HPC-DERIVED MODELS THAT SMOOTH RE-ENTRY FOR HYPERSONIC VEHICLES

*Team: Matthew Barone, Cory Stack
Contributing Writer: Johann Snyder*

How can HPC make re-entry smoother for hypersonic vehicles?

Re-entry vehicles face an immensely harsh environment at hypersonic speeds not only from aerodynamic heating, but also from severe, unsteady, turbulence-induced pressure fluctuations on the vehicle body, termed the “random vibration” environment. The amplitude of these pressure fluctuations may cause a reaction in the vehicle body (it may start to shake, rattle and roll) that can impact the vehicle’s internal components, thereby raising survivability concerns for all structural components of the vehicle.

Predicting and characterizing the turbulence physics that produce these pressure fluctuations is very complicated due to the chaotic, nonlinear behavior of the equations governing fluid motion at hypersonic conditions. The turbulent events responsible for these pressure fluctuations have a wide range of space and time scales, and thus the body of the vehicle is subjected to a complex, turbulence-induced forcing function.

Historically, researchers utilized a combination of experiments and simulations in simplified configurations (i.e., a vehicle at 0-degree angle-of-attack) to characterize the random-vibration environment

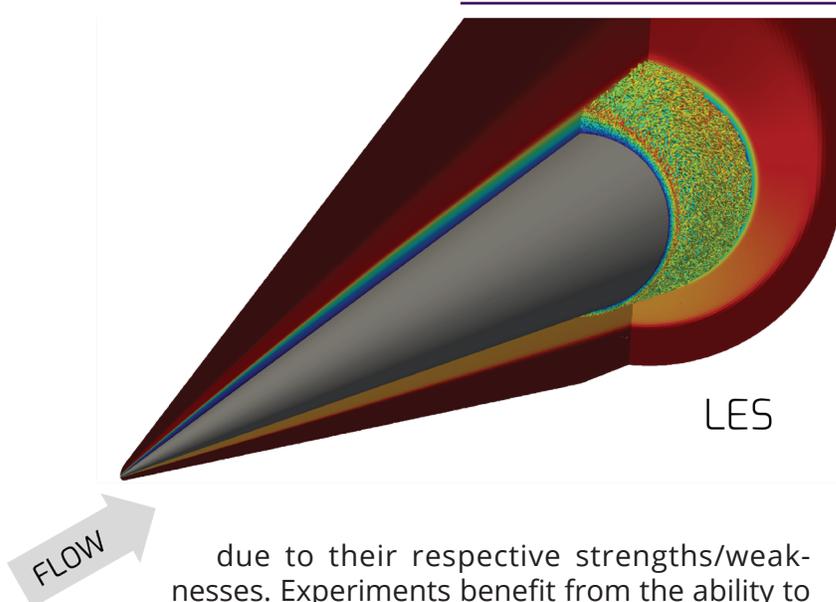


Figure 1: Visualization of the flow over a cone at a 6-degree angle-of-attack. The cone body is colored gray, and the colors depict the fluid velocity; with red as high and blue as low. Near the back of the cone an arrow points to the turbulent structures resolved by the large-eddy simulation.

due to their respective strengths/weaknesses. Experiments benefit from the ability to rapidly test numerous models/flight conditions but are hindered by the spatio-temporal resolution of the sensors, a difficulty that is exacerbated at hypersonic conditions due to the small turbulence length and time scales. Conversely, high-fidelity simulations can resolve all relevant spatio-temporal turbulence scales, but these simulations require vast amounts of computing resources; thus, only a small subset of conditions can be studied.

Due to the limited number of experiments and high-fidelity simulations that can be performed, researchers must devise a simplified model to describe the random vibration environment across a breadth of flight conditions using the data from the limited subset of available test cases. Often, this reduced-order model (ROM) is generated by leveraging unique expertise and advanced physical insight. This ROM is a tool for vehicle designers to quickly assess the random vibration environment of a re-entry vehicle, which is inherently coupled to the vehicle's survivability. However, the tool's validity only falls within the bounds (and potentially the very near-vicinity) of the subset of flight conditions that it was constructed upon.

Within the context of this body of work, researchers expanded the validity of the random vibration ROM into more complex, realistic environments by generating new datasets to train the ROM. Namely, high-fidelity Large-Eddy Simulations (LES) of the turbulent flow over representative re-entry vehicles (see Figure 1) were performed over a range of angles-of-attack and flight speeds using the Sandia HPC Parallel Aerodynamics and Re-Entry Code (SPARC). To expedite the time-to-solution for the simulations, researchers leveraged the computational power of the NNSA's 94.6 petaflop Lawrence Livermore National Laboratory's Sierra ATS-2 platform.

Generating new high-fidelity datasets via large scale computational fluid dynamics simulations made it possible to identify the deficiencies of previous ROMs for random vibration environments...

CONE AT 0° AoA

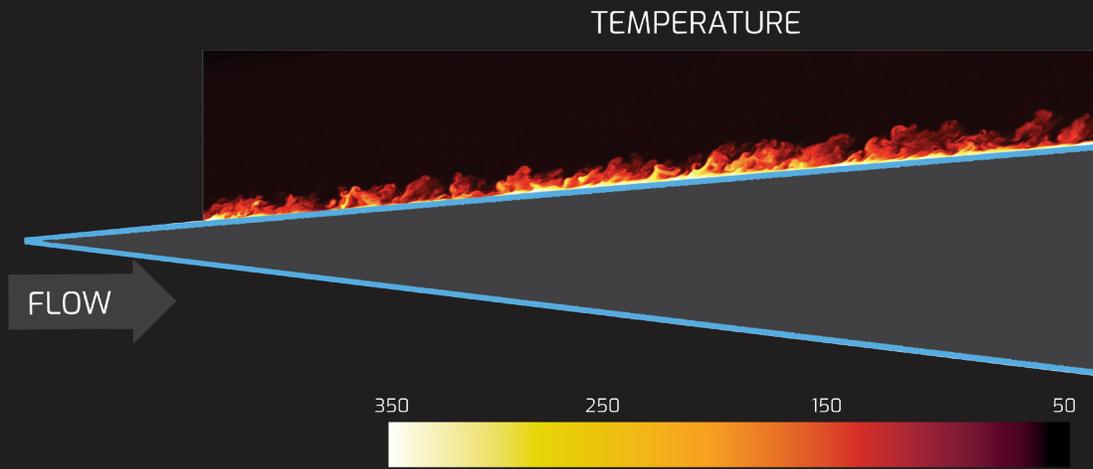


Figure 2: Visualization of the temperature field for the flow over a cone at a 0-degree angle-of-attack. The surface of the cone is represented by blue lines.



View video with SimMagic app

CONE AT 6° AoA

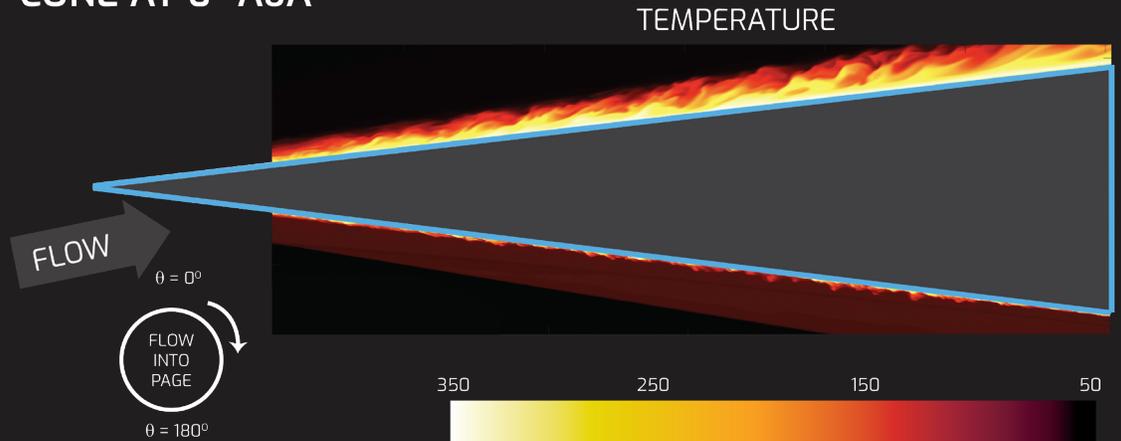


Figure 3: Visualization of the temperature field for the flow over a cone at a 6-degree angle-of-attack. The surface of the cone is represented by blue lines.



View video with SimMagic app

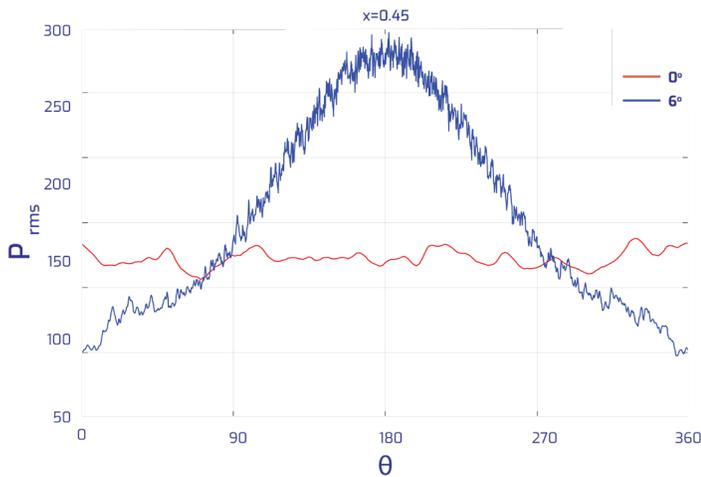


Figure 4: Comparison of the root-mean-square of the surface pressure fluctuations between a cone at a 0-degree angle-of-attack and a cone at a 6-degree angle-of-attack. See Figure 3 for a reference of the azimuthal coordinate system(θ).

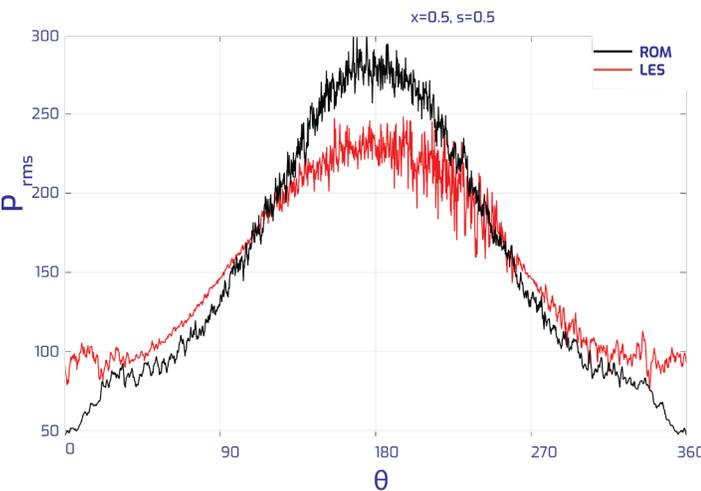
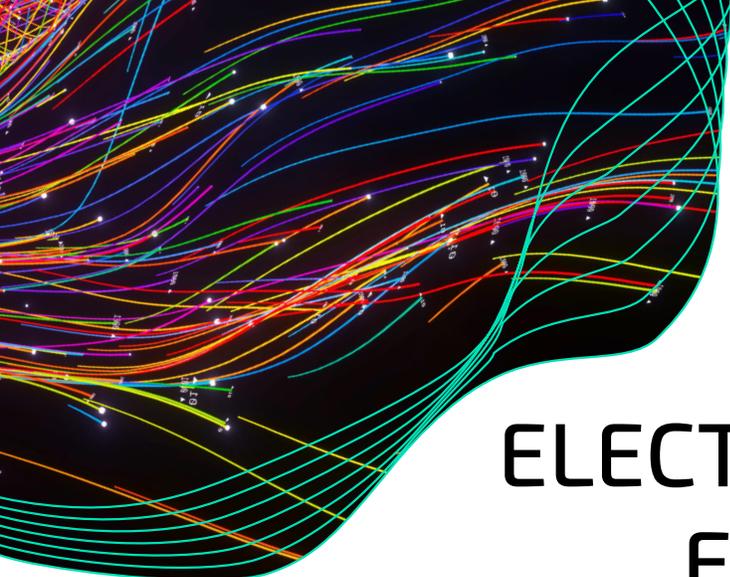


Figure 5: Comparison of the root-mean-square of the surface pressure fluctuations between a cone at a 6-degree angle-of-attack and the previous ROM.

Sample results from the flow over a conical flight body at a 0-degree angle-of-attack are shown in Figure 2, while results from the flow over a cone at a 6-degree angle-of-attack are shown in Figure 3. Visual observation makes it clear that the near-wall turbulence is drastically different for the 6-degree case compared to the 0-degree case. On the wind-side (180-degrees) the turbulence is confined to a much smaller region near the wall, while on the lee-side (0-degrees) the turbulence occupies a larger spatial region. These differences in the turbulence dynamics in turn affect the surface pressure fluctuations (Figure 4), where the 0-degree case and 6-degree cases clearly have different amplitudes of wall-pressure fluctuations. Comparing the results from the 6-degree case to the previous ROM highlights, noticeable deficiencies are evident (Figure 5). Specifically, the previous ROM under-predicts the amplitude of the pressure fluctuations on the wind-side (180-degrees), while also over-predicting the pressure fluctuation amplitude on the lee-side (0-degrees). As a result of the LES simulation campaign, these inherently three-dimensional features of the flow around the re-entry vehicle are now incorporated into the new ROM, thereby increasing its validity across a broader range of flight conditions.

Generating new high-fidelity datasets via large scale computational fluid dynamics simulations made it possible to identify the deficiencies of previous ROMs for random vibration environments, and it allowed the expansion of the ROM to a broader range of flight

conditions, thereby mitigating the use of current ROMs outside of their physically applicable scope. Future work will continue to expand the physical scope of the ROM by generating additional training datasets. This improved ROM gives designers an invaluable tool to assess the random vibration environment in more complex scenarios than ever before, permitting more efficient and reliable designs in the development of advanced re-entry systems. [▲](#)



ELECTROMAGNETIC ENVIRONMENT MODELING FOR NUCLEAR WEAPONS

Team: J D Kotulski, Vinh Dang, William Langston, Robert Pfeiffer, Alden Pack, Aaron Krueger, William Johnson, Larry Warne, Jonathan Robey, Luis San Martin, Alexander Moreno, Jeffrey Osterberg, Ian Timmins, Javier Rivero
Contributing Writer: Susan Jean-Pierre

The testing and analysis of these systems for such environments is a vital part of Sandia's "Always/ Never" nuclear deterrence responsibility

How can you test Nuclear Weapons in extreme electro magnetic environments?

Nuclear weapon design is complex and demands that different aspects of the system are tested and analyzed to ensure their safety and reliability in different environments. One of these aspects is the system's response to and hardening for the many potential electromagnetic environments. Additionally, the qualification of weapons in the current stockpile must be re-examined against the evolving potential electromagnetic environments in today's world.

The testing and analysis of these systems for such environments is a vital part of Sandia's "Always/ Never" nuclear deterrence responsibility: Ensure that a nuclear weapon will always work as intended when needed and never work under any other circumstances. This assures the safety, control, and survivability of existing nuclear weapons as well as new weapons designs.

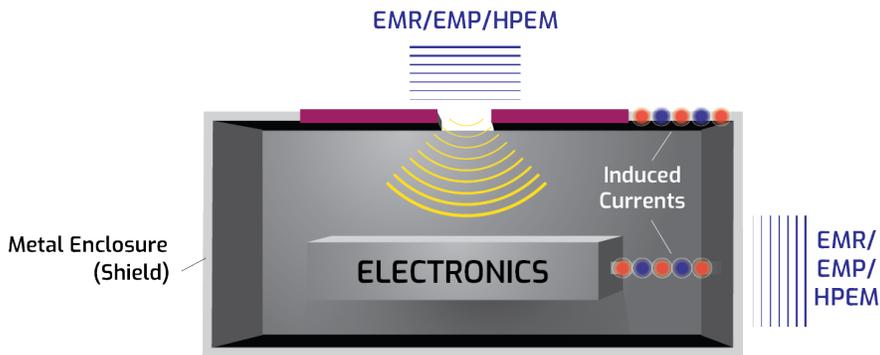


Figure 1: The coupling problem: Electromagnetic radiation (EMR), electromagnetic pulses (EMP) and high power electromagnetics (HPEM) can impact the exterior and interior of an enclosure

Penetration into enclosure through apertures

The Sandia Advanced Simulation and Computing (ASC) electromagnetics team specifically has considered this aspect of nuclear weapons stockpile qualification while developing and maintaining the EIGER and GEMMA electromagnetic ASC codes. EIGER is the legacy code, written in Fortran. GEMMA is the next generation electromagnetic simulator, written in C++ and containing additional advanced electromagnetic algorithms. Both codes have been designed to perform robust, complex assessment and analysis of nuclear weapons design and weapons qualification at Sandia while focusing on electromagnetic environments.

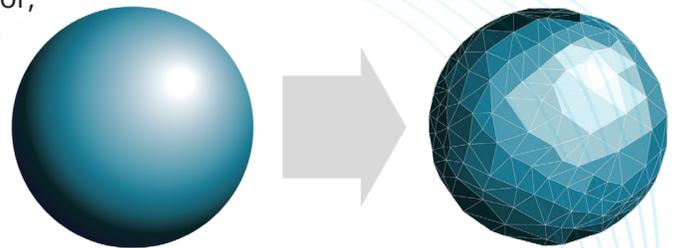


Figure 2: Meshing a sphere

The team works actively with other Sandia organizations, particularly experimental electromagnetic groups. These groups provide validation data to inform code credibility and helps the team improve the code, while the code simulations can also provide feedback to experimentalists to help them plan and understand tests.

The EIGER and GEMMA codes are designed to predict the electromagnetic fields inside and outside complex nuclear weapons structures that have potential entry points. This naturally requires the code to be complex too. The team has worked for years to optimize the code so it can perform fast, high-fidelity analyses of these systems. Nevertheless, EIGER and GEMMA require the computing power of Sandia’s HPC resources as well as Los Alamos National Laboratory’s Trinity and Lawrence Livermore Laboratory’s Sierra systems.

More specifically, EIGER and GEMMA simulate the major electromagnetic interference or coupling issue illustrated in Figure 1. Since electromagnetic penetration can affect internal components, it is vital for designers and analysts to understand how each specific test object is affected by such interference.

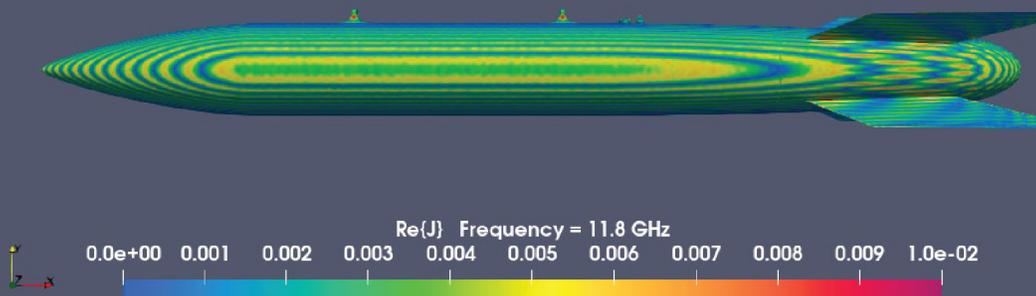


Figure 3: Simulation of an 11.8 GHz current on the exterior of a virtual mesh object (the entry points and internal components are omitted from this simulation).

Sandia's HPC as well as other advanced next generation computing platforms are a key part of our "Always/Never" responsibility.

HOW EIGER AND GEMMA WORK

First, a three-dimensional CAD model is prepared for meshing, creating the electromagnetically-significant features and neglecting the unimportant features. Then the model is meshed with surface elements. A simple example is shown in Figure 2.

Next, a linear system is constructed by applying appropriate boundary conditions and a series of complex equations are solved for the meshed object, which include the exterior, interior and entry points. See Figure 3 for an example showing exterior currents on a B61-12.

Once the algorithms are run, the simulation results are processed to provide detailed feedback to the nuclear weapons designers and the weapons qualification engineers.

Although the above steps may seem simple, the calculations are complex and resource intensive, yet must be performed quickly and with extremely high fidelity. The complexity of these simulations depends on the object and the frequencies of the electromagnetic fields. Since more complex objects and higher frequencies require orders of magnitude more computing resources to run the EIGER and GEMMA simulations, Sandia's HPC as well as other advanced next generation computing platforms are a key part of our "Always/Never" responsibility.

OTHER POTENTIAL APPLICATIONS

Although nuclear weapon qualification is the primary mission of EIGER and GEMMA, the codes can be used in other areas, such as: electromagnetic scattering from a system by a radar, radiation from an antenna located on a system of interest, or electromagnetic effects on structures as well as components. After meshing an aircraft, the code simulations are run on an observation plane for the mesh model (see Figure 5).

After the simulation is run and analyzed, the results are mapped to show location and magnitude of the electric field on the aircraft exterior (see Figure 6).

Although EIGER and GEMMA were born from a need to qualify nuclear weapons, they can be used to understand and solve numerous other issues surrounding electromagnetic effects. GEMMA code development continues, focusing on advanced algorithms and solver technology to expand its solution capability and enhance Sandia's response to problems of national importance. Current efforts include analyzing GEMMA's response in combined environments, such as electrical and mechanical, and streamlining the GEMMA workflow. [△](#)

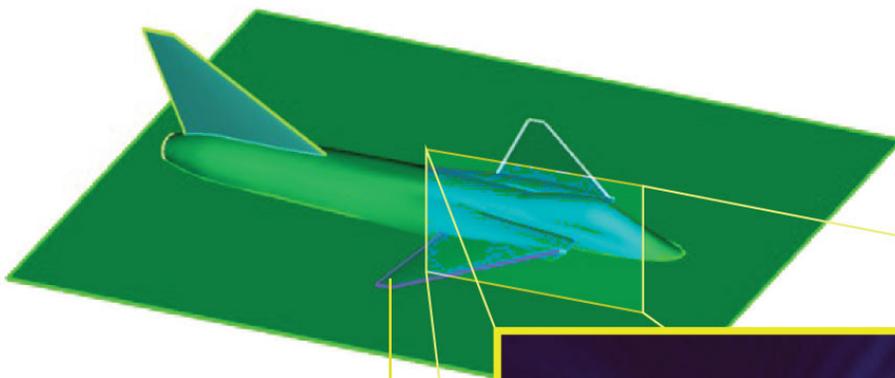


Figure 4: Aircraft meshed model and its observation plane for the simulation of electric fields

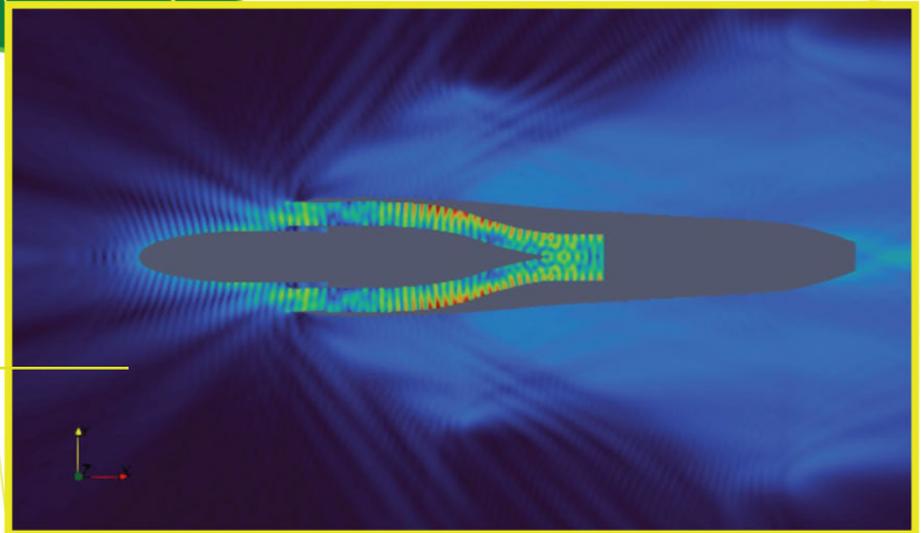


Figure 5: Magnitude of the scattered electric field on the observation plane of Figure 5.



LEONIDAS: LEVERAGING THE INTERNATIONAL SPACE STATION TO DEMONSTRATE ADVANCED SENSOR

*Team: James Meub, Greg Anders, Thomas Bradshaw, Kurt Brenning, Hunter Earnest, Jeff Hunt, Rich Hunt, Sirsedrick Kendrick, Lena Lopez, Willie Marchetto, Andrew McCants, Jeff McCasland, Joseph Reneker, Matt Russell, Ted Winrow
Contributing Writer: Kristen Meub*

Can we pave the way for space-ready high-performance computing systems?

Sandia's LEONIDAS project team is sending a payload to the International Space Station (ISS) to study the effects of radiation from the space environment on advanced computing systems.

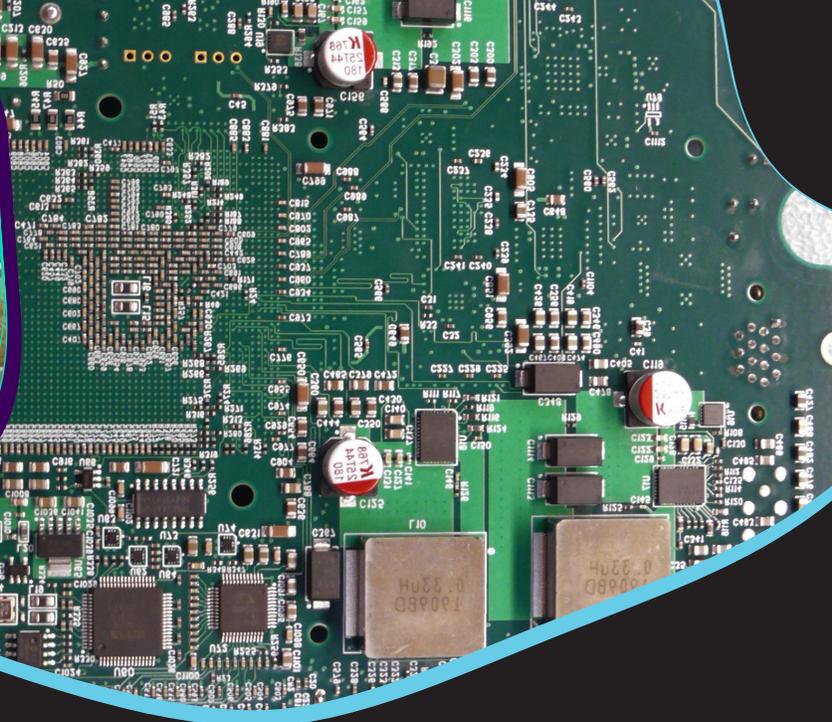
The electronics and sensors that make up an optical payload can be vulnerable to damage from naturally occurring radiation in space. As a result, engineering teams have typically relied on radiation-hardened electronics, but these computing systems can be expensive, challenging to integrate and draw a lot of power.

At the same time, some organizations and agencies flying satellites are looking to collect and send more data than before, but there are physics-based limits to how much data can be sent from space to the ground.



Figure 1: NREP Placement for External Installation

Figure 2: Xilinx Veral Test Board



With the LEONIDAS project, the team integrated a commercial computer chip — the Xilinx Versal system-on-chip — that is more resilient against radiation, on a single board computer designed for space flight. LEONIDAS also hosts two imaging sensors, a focal plane array, which captures visible light, and an event-based sensor, an imager that captures pixel-by-pixel changes in intensity. Data from these sensors is collected, processed and packetized for storage and downlink in the Xilinx Versal on a module and will be communicated through the ISS's communication systems.

Taken together, LEONIDAS demonstrates a payload with event-based sensors co-collecting with traditional sensors feeding real-world in-situ data into cutting edge digital electronics running novel algorithms while path-finding a process to operate in space. The team hopes this project will be a precursor to a high-performance computing space system that enables organizations to rapidly process data on the edge, a requirement to get more relevant data, faster, from their satellites.

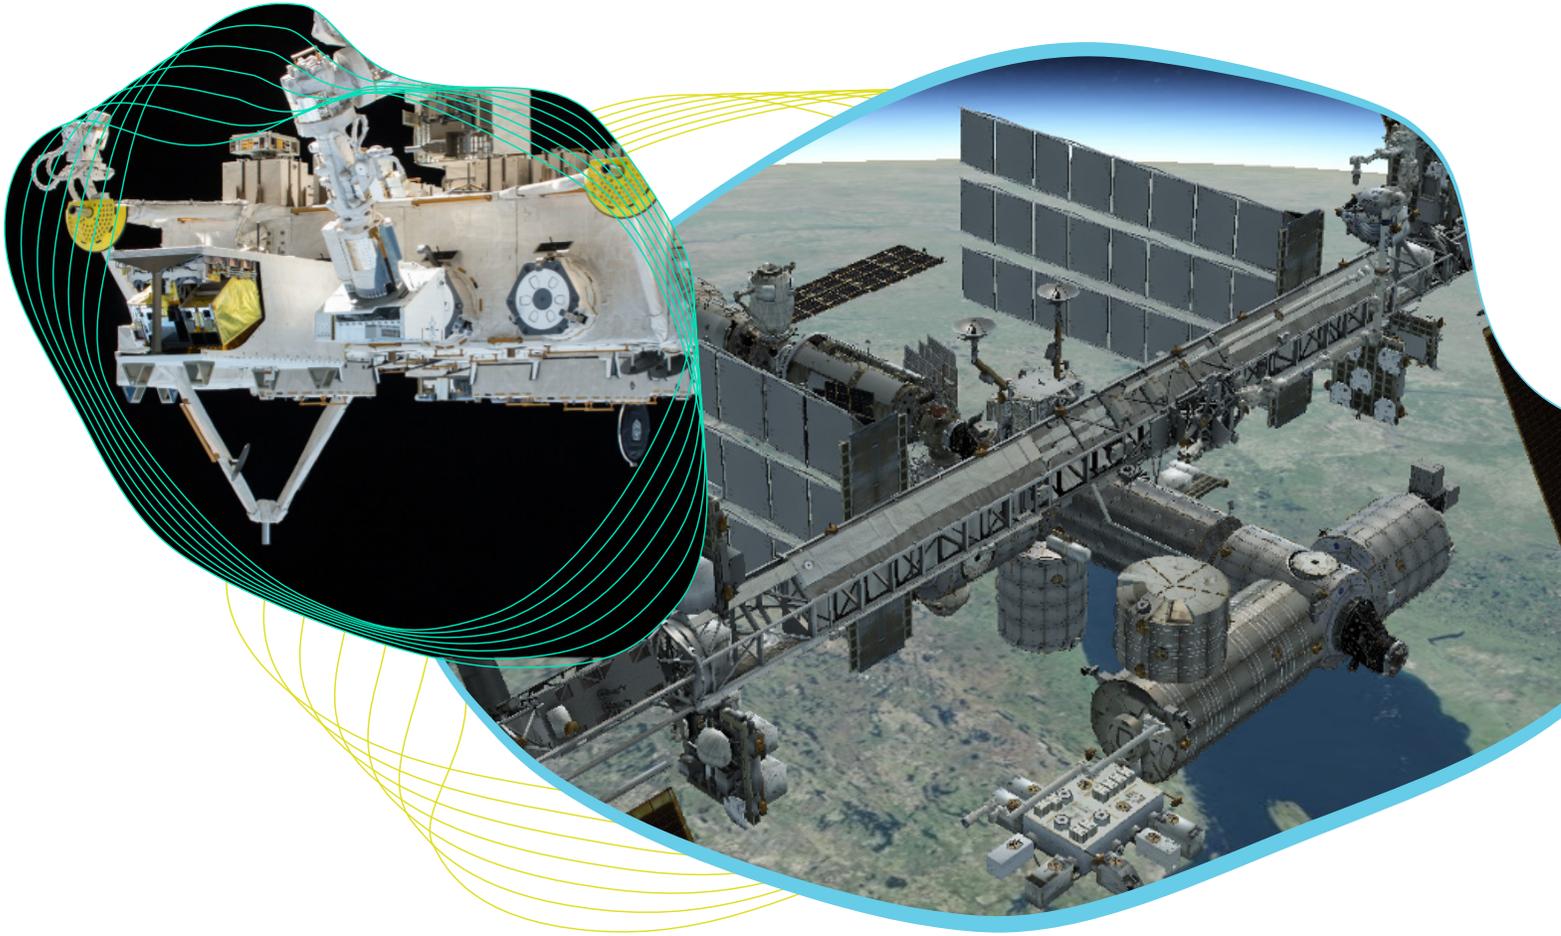
The LEONIDAS payload is expected to launch on a Northrup Grumman resupply mission to the ISS in August 2022. Launch coordination, astronaut collaborations and physical interfacing with the ISS will be provided by NanoRacks, a commercial vendor that maintains and operates payload bays on the ISS. The astronauts will install the payload on the Japanese Experiment Module platform, where it will operate and gather data for its remote sensing mission for four months.

For high consequence national security missions hoping to leverage novel technologies on orbit, like the Xilinx Versal, a rigorous qualification program that replicates the space environment is required. This includes:

- Component screening.
- Radiation testing across a variety of environments, often requiring testing at multiple facilities.
- Thermal cycling, vibration, shock and acoustic testing.
- Electromagnetic interference testing.
- Thermal vacuum testing.



...a precursor to a high-performance computing space system that enables organizations to rapidly process data on the edge, a requirement to get more relevant data, faster, from their satellites.



These tests are done on the ground because they have historically been cheaper and enable inspection, rework and re-testing of the devices. The LEONIDAS project is demonstrating a faster, less expensive approach that tests novel technologies in the actual space environment through the use of commercial space flight operators and the ISS National Laboratory. Another benefit is that the team will get the payload back after its time in space. This will allow for post-flight inspection.

This approach will give payloads space flight heritage and increase technology readiness levels. Beyond achieving data collection and technology demonstration, the goal of LEONIDAS is to streamline government and commercial partnerships for rapid insertion space flight missions.

LEONIDAS achieves several “firsts,” including:

- First event-based sensor on orbit co-collecting with traditional sensors.
- First flight of the Xilinx Versal component.
- One of the fastest to flight operations for a space payload, conception to operational.

LEONIDAS leverages Laboratory-directed Research & Development (LDRD) investments in radiation testing, computing hardware design and algorithm development and provides an in-situ testbed for ongoing LDRDs exploring space resiliency in the STARCS Strategic Initiative IAT. This work was funded by the National Nuclear Security Administration’s Space Flight Systems Program. [▲](#)

NEXT-GENERATION CONCENTRATING SOLAR POWER PARTICLE RECEIVERS ENABLE INEXPENSIVE ELECTRICITY GENERATION WITH ENERGY STORAGE

How can HPCs enable solar energy storage?

*Team: Brantley Mills,
Reid Shaeffer, Clifford K. Ho,
Kevin Albrecht, Jeremy Sment,
Nathan Schroeder, and
Henk Laubscher.
Contributing Writer:
Laura Sowko*

To help propel new forms of renewable energy, the Department of Energy's Solar Energy Technologies Office has been encouraging development of next generation concentrating solar power (CSP) systems. CSP systems offer a distinct advantage over photovoltaics by enabling energy storage through means other than batteries.

CSP uses a field of mirrors (heliostats) to concentrate sunlight onto a target (receiver) to irradiate a heat transfer medium. That heated medium is then used to generate electricity or drive other thermochemical reactions. Researchers at Sandia are investigating using particles (specifically, a ceramic bauxite particle) inside the receiver within a gravity-driven, tower system. The receiver concept being developed to facilitate this system is called a falling particle receiver (FPR). In an FPR, particles are dropped within a cavity as a curtain past a beam of concentrated sunlight that heats them. The heated particles can be stored for future use when the sun is not shining or be used immediately in conjunction with a highly efficient supercritical-CO₂ Brayton cycle to generate electricity. The particles are then lifted to the top of the tower to be dropped in the FPR and heated again. The biggest advantage of particles is that they enable the FPR to reach very high temperatures (>800°C) where traditional CSP heat transfer mediums do not. This allows use of the efficient supercritical CO₂ cycle, making the electricity much cheaper.

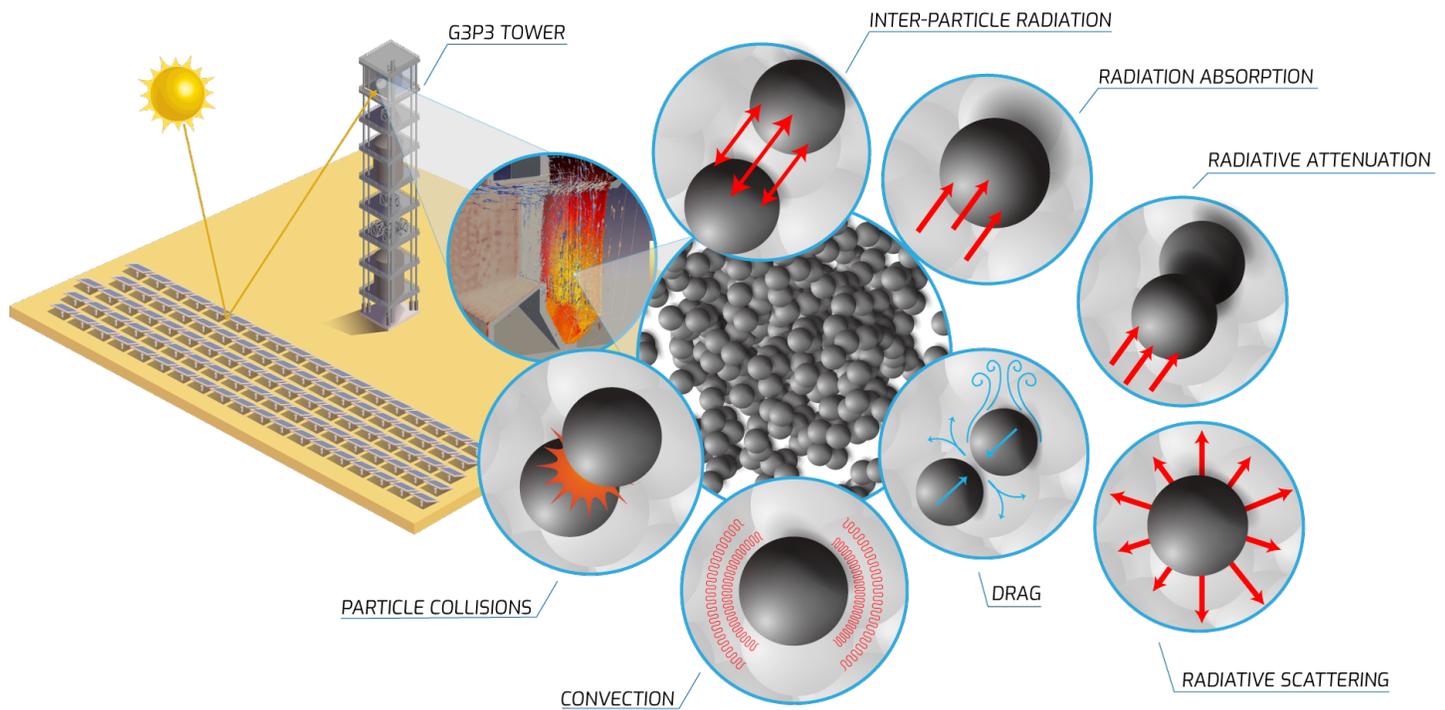


Figure 1: The G3P3 falling particle receiver's simulated thermal response under nominal conditions

Making this concept economical requires high efficiency particle receiver designs. Modeling and simulation can help drive this development (minimizing the number of complex experiments needed), but it can be very computationally expensive to model all the necessary physics to predict a particle receiver's efficiency. Sandia is unique in the amount of HPC support available for these types of problems. Thousands of CPU hours were used to design and understand receiver operation under parameters such as weather variability and receiver response to transients. This information was compiled into a compelling case for yearly receiver performance and key data points for evaluating feasibility. "The simulations help us maximize the amount of sunlight the particles absorb where small changes in the geometry, operating conditions, or environment can affect that absorption," says researcher Brantley Mills. "The less energy the particles absorb, the lower the efficiency and the more expensive it is to generate electricity. These models are helping us identify and minimize all the thermal losses that may contribute to lower efficiency."

Falling particle receivers have been simulated before, but this is the first time they have been simulated at this level of geometric and physical fidelity that provides more confidence in the design's expected performance (i.e., efficiency). Typically, various levels of physics have been excluded (e.g., particle fines, wind, particle bouncing) to bring down the computational cost of these simulations, but Sandia's HPC resources have allowed researchers to account for these additional details. HPC systems allowed coupling high-fidelity computational fluid dynamics models with discrete particle tracking, highly directional radiation transport from the heliostats, and energy transport through the insulative receiver walls. The solutions to these physics allowed precise tracking of how energy was flowing through the system, which has been invaluable for minimizing losses for high efficiencies.

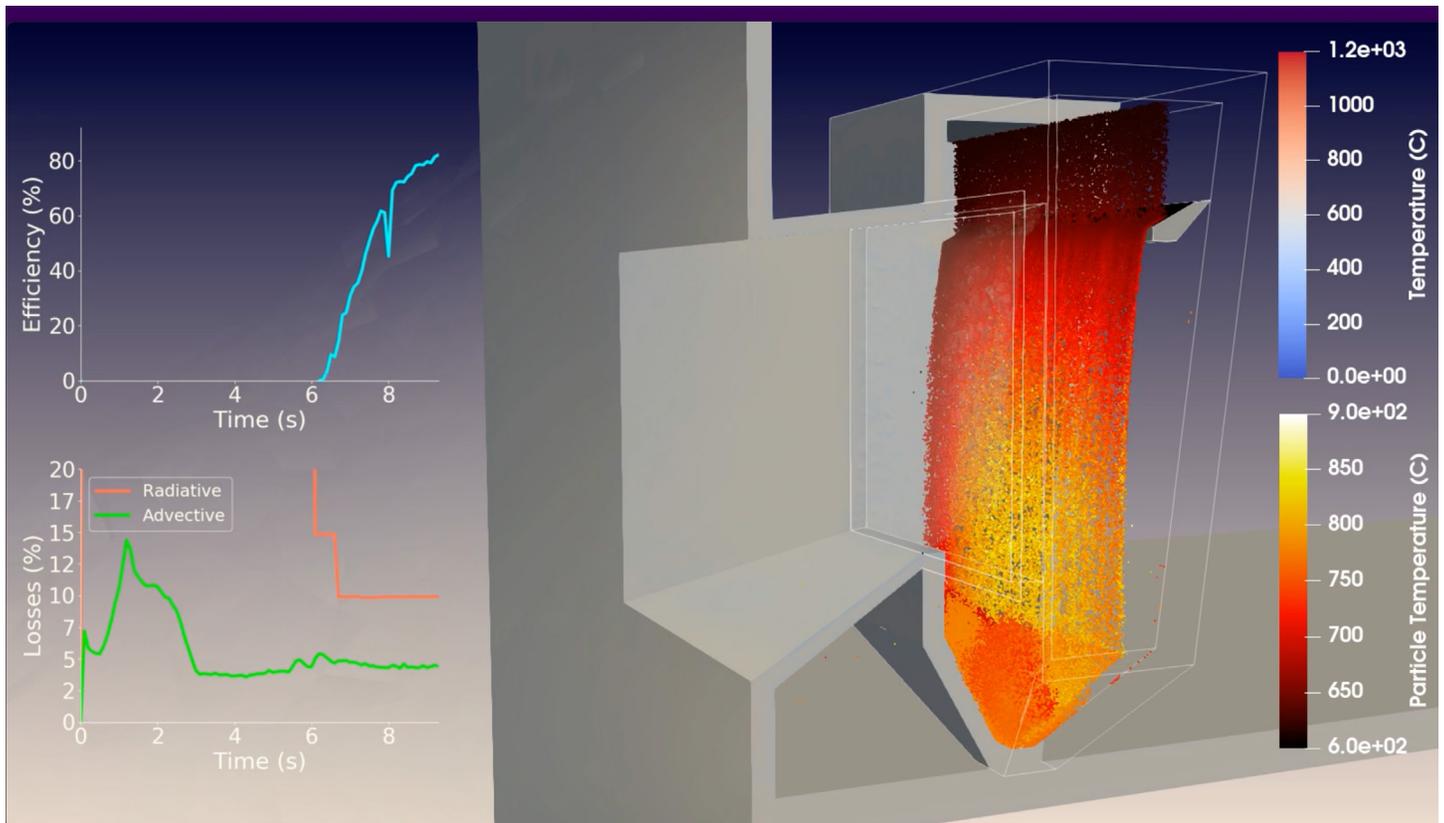


Figure 2: Critical particle physics that must be modeled to accurately predict a falling particle receiver's efficiency



View video with SimMagic app

The researchers are using a wide array of physical models that have exercised the code coupling capabilities within Sandia's Sierra simulation software at scale. Specifically, they are using the Sierra modules: Aria (to model energy transport through the walls), Fuego (to model the air continuum inside and outside the receiver and the Lagrangian particles), and Nalu (to model the radiation transport in a discrete ordinates framework). They have also worked with developers to incorporate needed features into the code suite. For example, the ability to model non-grey radiation transport with walls, reduced order models for the treatment of the particle troughs, and the ability to describe complex radiation beams like those used from heliostat fields. These models are executed using meshes with millions of elements to provide sufficient geometric fidelity.

In addition to the complex coupling, many of the physics have different timescales associated with them (radiation — effectively instantaneous; fluids and particles — seconds; solids — minutes). "We have had to come up with creative solutions to extract the necessary observations from the simulations. For example, we execute the simulation for a period of time without radiation to minimize numerical instabilities as the flow field forms. We have also used non-physical thermal masses in the wall materials at the onset of the simulation to decrease the time it takes for the temperature field to arrive at a steady-state condition," says Mills.

Sandia's HPC models have also helped the researchers better understand the flow of small particle fines within the receiver cavity. The particles are normally ~350 microns, but very small particle fines are generated in operation (as small as 1 micron) and Sandia has used the models to predict the flow of those particles in the receiver cavity. Models have also been able to design and explore different geometric features designed to capture these particle fines, minimizing the hazards they pose to the environment and personnel.

The modeling and simulation also help inform a receiver design that will be constructed and integrated into the Generation 3 Particle Pilot Plant (G3P3) tower (which is planned to be completed in 2024). The completed 1 MWth system constructed at the National Solar Thermal Test Facility at Sandia will be a fully integrated particle pilot plant that can help resolve questions on the integration and scaling of this technology. If successful, this technology will be scaled up to plant sizes as large as 100 MWe with the ability to store energy for times when sunlight is not readily available. HPC models have been used to evaluate the final G3P3 receiver design at very high geometric fidelity.

HPC models have also offered the opportunity to test out new concepts, for which patents are pending, such as a set of intermediate troughs that "catch and release" the particles. These troughs slow the particles down as they fall through the cavity, increasing the residence time of the particles within the concentrated sunlight while not exposing any surfaces directly to the beam. This allows for more sunlight absorption and higher particle temperatures and has been shown in tests to increase the curtain's uniformity. Further, the troughs help protect the back wall of the cavity, which can be damaged in spots if directly exposed to the beam. This technology will be an essential component of future 100 MWe systems to control the particle curtain in falling particle receivers with cavities as tall as 20 m.

HPC models have played a critical role in designing FPRs for next-generation CSP systems using particles. The researchers intend to extend this modeling capability to investigate larger receiver systems and other receiver designs that could be built in the future. The models will also be leveraged to help predict FPR performance variability throughout a typical day at the G3P3 tower. [▲](#)

Falling particle receivers have been simulated before, but this is the first time they have been simulated at this level of geometric and physical fidelity that provides more confidence in the design's expected performance (i.e., efficiency).

VALHALLA: DESIGNING CUSTOM SATELLITES FOR NATIONAL SECURITY MISSIONS WITH HIGH-PERFORMANCE COMPUTING IN WEEKS, NOT YEARS

*Team: James Meub, Matt Russell, Barry Holland, Forest Danford, Leonardo Le
Contributing Writer: Kristen Meub*

*Is it possible to design custom
satellites in weeks, not years?*

*Figure 1: Simulation screen
capture for arctic domain
awareness analysis
produced by Valhalla.*



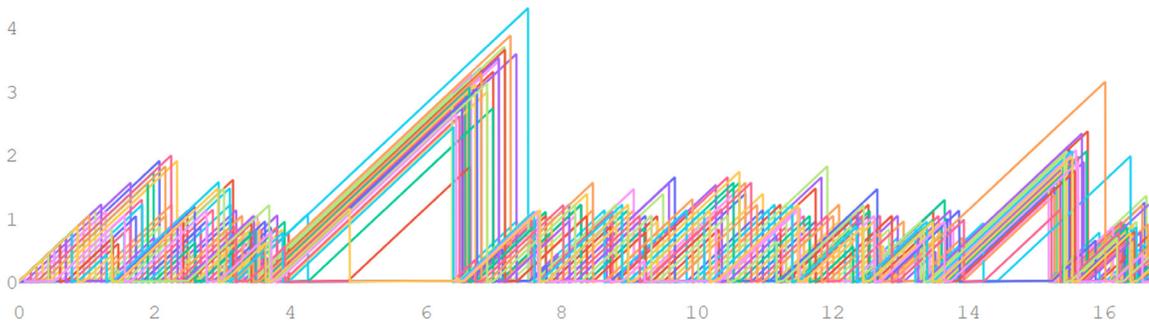


Figure 2: Plot of time between revisits for 300 arctic targets (x-axis: minutes, y-axis: minutes) for a concept arctic domain awareness constellation. Revisits refer to how often a space system observes its targets.

Remote sensing can be used to address many national security concerns, but until now it could take a team years to explore existing technology, develop reference designs and study how to best address the challenge.

Valhalla, a Python-based performance modeling framework developed at Sandia, uses HPC to run thousands of mission simulations to produce a preliminary design for a satellite or space system in just weeks.

Increasingly, a constellation of small satellites working as a system instead of one large satellite is being considered for remote sensing missions. Valhalla designs a satellite and then populates multiple copies into a constellation and simulates the system executing the specified mission. Valhalla can also model how groups of non-identical satellites with different features work together to execute a given mission.

Designing a space system with Valhalla is similar to customizing a new car online. A user visits Sandia's internal Valhalla website and selects the features and parts needed to execute the mission. The first step is to provide Valhalla with mission requirements that can be parsed into quantitative goals, including:

- Determining the level of fidelity and detail needed by specifying if the satellite system will need to detect, identify, categorize or associate objects or events.
- Selecting the class of object or event to search for, including vehicles, explosions, aircraft, boats, space debris and other satellites.
- Selecting the frequency of data collection — how often the system should revisit the target objects or events.
- Selecting the target's motion model — will the system need to observe the trajectory of trucks traveling from one location to another, or will it need to observe a city, region, or neighborhood of interest or monitor objects or events across the globe?
- Specifying any constraints — for example, the system may need to be compatible with a specific external network.
- Selecting a mission operation center to feed the images and data captured by the satellite.

With the mission requirements clearly defined, the next step is to select parts such as on-board computers, radios, batteries, reaction wheels, sensors, propulsion tanks and thrusters and GPS receivers from a component catalog with about 2,500 parts to choose from.

There are thousands of variables and relationships to consider when designing a space system for a specific national security mission. Valhalla visualizes the relationship between variables so the user can see how the overall space system is affected when parameters that are relevant to each other change.

Valhalla takes all the selected parts and creates a baseline satellite design by using encoded interface control documents, engineering-best practices and rules-based logic. It sizes the payload, structure, solar arrays and batteries and optimizes the mass of the satellite and its compatibility with various launch vehicles.

Valhalla then populates the baseline design into a constellation and simulates the space system executing the specified mission by running a day-in-the-life power model simulation. The simulation compares the mission performance metrics against space systems goals and repeats that process thousands of times using a Monte Carlo algorithm.

Valhalla also models the space system’s orbital dynamics and the gravity harmonics of the earth, moon and sun to simulate the space environment each satellite will experience over the course of the day.

There are thousands of variables and relationships to consider when designing a space system for a specific national security mission. Valhalla visualizes the relationship between variables so the user can see how the overall space system is affected when parameters that are relevant to each other change. For example, the desired altitude of the spacecraft will affect the size of the telescope needed and vice versa. As the altitude of the satellite increases, the telescope also needs to become more powerful, but the system will need fewer satellites to provide enough sensing coverage to execute its mission.

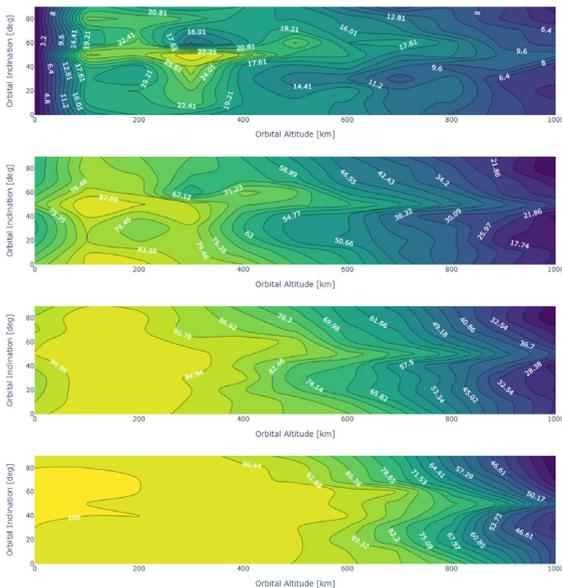


Figure 3: Daily data throughput (Gigabytes per day) analysis for four space-to-space transceiver concepts across all orbital altitudes and inclinations using TeleSat-LEO relay constellation.

A critical component of this process is autonomous operations. Typically, satellite systems require a team of operators to build a task list for the constellation. Valhalla uses a novel set of autonomy algorithms to simulate data fusion, sensor-target optimization and signature collection performance. To run faster than real-time, which is required to run these simulations thousands of times, the satellites need to be able to task themselves within the simulation. Valhalla does this through the Shortstack code base developed at Sandia to optimize sensor placement. Shortstack is an optimization routine that produces sensor target pairing to maximize information gain.

Tracktable, a code base developed at Sandia, is another routine integrated into Valhalla. It takes billions of data points in historical data and uses neural networks to fit hypothesized target trajectories based on data collected by satellites in the simulation.

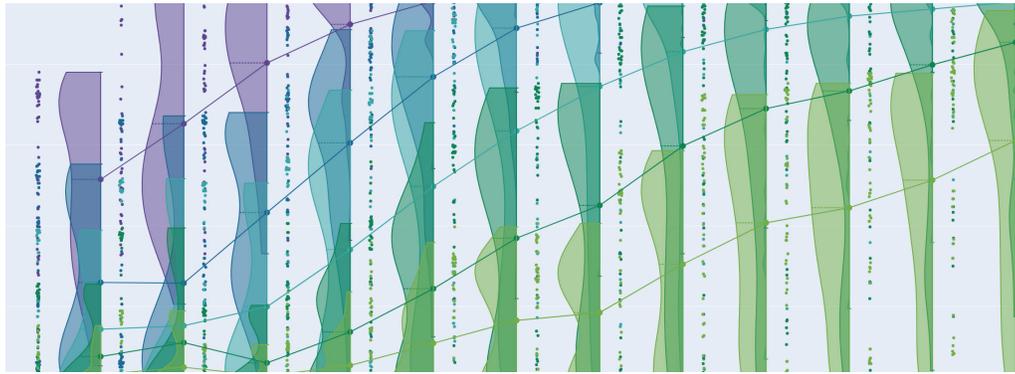


Figure 4: Daily data throughput (Gigabytes per day) analysis for four space-to-space transceiver concepts across all orbital altitudes and inclinations using TeleSat-LEO relay constellation.

Valhalla includes a cost analysis tool, developed by the Aerospace Corporation, that estimates the program cost of a satellite system based on the selected components, orbits and sub-system mass ratios.

The simulation data is post-processed in a GPU cluster. The GPU takes all the data from the simulation and combs through it to correlate performance metrics across variable sets to produce an N-dimensional hypersurface that provides the user with hundreds of plots, including a 3D visualization of the constellation executing its mission, revisit and collection statistics, launch vehicle compatibility and simulated satellite telemetry feeds.

All the data is presented in a visual, digestible format that allows users to see why certain simulation runs produce better performance metrics than others. The user then selects the solution that best fits the mission based on the data and plots provided. Valhalla gives the user the information they need to rapidly understand the national security challenge in a mission context and design a space system that meets the unique mission requirements. This enables teams to focus on delivering national security solutions in months rather than years.

Valhalla has been used to:

- Explore space systems for enhanced arctic security.
- Study the effects of neuromorphic computing on space system design.
- Simulate tracking and monitoring the rescue of refugee vessels in the Mediterranean.
- Explore satellite fault tolerance and identification of single-point and dual-point failure paths.
- Analyze satellite rendezvous and proximity operations.
- Perform satellite threat resilience analysis.
- Study novel intra-satellite communication networks.

A future goal of Valhalla is to use cost as a performance metric and build a mapping of satellite requirements to program cost, allowing national security programs to identify the cost-critical features of a specific mission or to see how to improve performance while maintaining the cost profile constraints from customers. ▲

ACKNOWLEDGEMENTS

EDITORS

*Sophia Corwell
Ariana Stern*

DESIGN

Johanna Hartenberger

VIDEO

Judith Preston

SIMMAGIC

*Danny Gomez
Aaron Moreno
Chadwick Johnson*

PROOFREADER

Whitney Lacy

MANAGEMENT

*Tom Klitsner
Scott Collis
Manoj Bhardwaj
Sophia Corwell
David Littlewood
Kendall Pierson*



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2022-10819 M



Sandia National Laboratories



U.S. DEPARTMENT OF ENERGY

NNSA
National Nuclear Security Administration



sandia national laboratories
high-performance computing

