

This three-dimensional simulation of an idealized inertial confinement fusion implosion shows turbulent mixing in a spherical geometry. Livermore's Sierra supercomputer makes high-fidelity calculations like this routine, yielding crucial insights into physical phenomena.

LAURENCE Livermore's high-performance computing (HPC) facilities house some of the fastest supercomputers in the world, including the flagship Sierra machine. Online for more than a year, Sierra primarily runs simulations for the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing (ASC) Program. Sierra substantially increases the Laboratory's ability to support ASC's stockpile stewardship work by providing more accurate, predictive simulations.

Sierra was formally dedicated in October 2018 and opened to NNSA users for classified work in the spring of 2019. Leading up to those milestones was a yearslong effort by CORAL—a collaboration of Oak Ridge, Argonne, and Lawrence Livermore national laboratories—and Livermore's Sierra Center of Excellence to prepare applications for the first major heterogeneous system at the Laboratory (see *S&TR*, March 2015, pp. 11–15; March 2017, pp. 4–11). Through CORAL, the Department of Energy (DOE) requested proposals for extreme-scale computing systems. After

hardware vendors were selected, the Center of Excellence provided resources to modify Livermore's large code base for execution on the new machine's architecture.

Likewise, Sierra had to be optimized for DOE workloads. "We have been continuously engaged with HPC vendors to help steer development of future computing systems," explains Chris Clouse, Livermore's acting program director for Weapons Simulation and Computing. "Collaborating with other DOE laboratories helps capture vendors' attention so they can appreciate the larger context of the country's scientific computing needs."

This lengthy preparation—designing advanced computing hardware, shifting the software programming paradigm, and pushing the industry standard—has culminated in unprecedented simulation capabilities at Livermore. As NNSA's most powerful supercomputer, Sierra ushers in a new era of computing architectures, software development, and scientific applications.



Processing Power

When the Sequoia supercomputer came online in 2012, it put Lawrence Livermore squarely in the petaflop (10^{15} floating-point operations per second) era of computing (see *S&TR*, July/August 2013, pp. 4–13). It had taken a decade of computing innovations to achieve the factor of 1,000 gains in performance from the first teraflop (10^{12} floating-point operations per second) systems. All that extra computing power, and resulting computing gains over the last decade, have come at the expense of increased power demands. Looking ahead to the next generation of computers, Laboratory researchers knew a more energy-efficient approach would be needed. “Any extreme-scale machine must balance operating costs with computing performance gains,” says Clouse. “More powerful regimes cannot be achieved by simply scaling out current, heavyweight core technology. We need a new approach that offers an affordable electricity bill.”

Sierra’s advanced heterogeneous, or hybrid, architecture uses more than one type of processor or core. It combines 17,280 NVIDIA Tesla V100 (volta) graphics processing units (GPUs), which increase parallel processing power, and 8,640 IBM Power9 central processing units (CPUs). Clouse says, “Some parts of our large multiphysics applications simply will not run well on a solely GPU-based system.” Sierra’s GPU-to-CPU ratio helps balance the machine’s application workload. Sierra’s 100-gigabit-per-second network swiftly transfers data between these processors, and its memory capacity for data-intensive calculations reaches 1.38 petabytes. In comparison to other GPU-based machines at the Laboratory, Sierra also has significantly more high-bandwidth memory per processor. “Performance is heavily tied to the high-bandwidth memory associated with the GPUs. The more data we

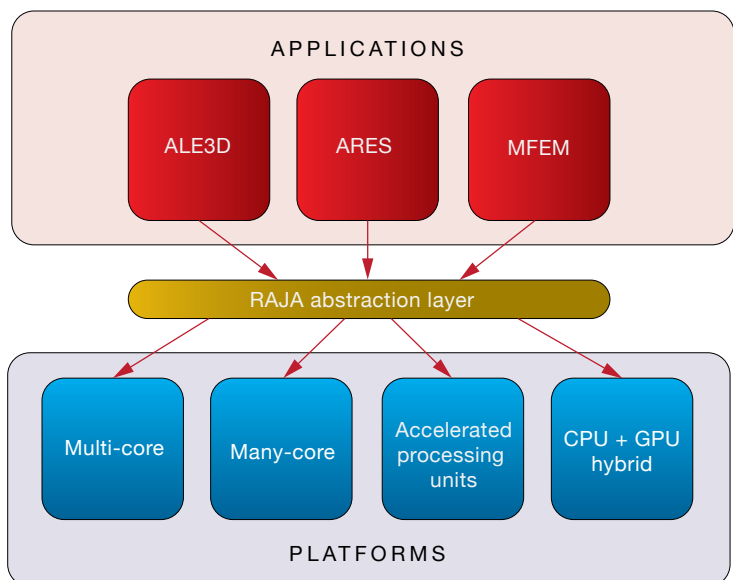
Advanced computer architectures such as Sierra (above) combine graphics processing units with central processing units and require a new programming paradigm to ensure that simulation codes run efficiently and effectively. (Photo by Randy Wong.)

can fit into that type of memory, the better our codes will perform,” states Clouse.

Sierra’s sophisticated architecture enables the machine to register a peak performance of 125 petaflops using only 11 megawatts of electricity. In other words, Sierra is six times faster than Sequoia but uses only one-third more wattage. Terri Quinn, associate program director for Livermore Computing, explains, “GPUs sip energy compared to traditional server processors. With leading HPC systems incorporating tens of thousands of processors and consuming multiple megawatts of power each year, GPUs keep power consumption in check, require a smaller footprint, and cost less than a CPU-only system of comparable performance—if you can take advantage of them.”

Abstraction and Memory Solutions

According to Quinn, pursuing Sierra’s hybrid architecture was a difficult decision. She says, “The ASC teams have put in an extraordinary amount of effort, often with IBM and NVIDIA experts alongside, to ready our codes (which did not run on GPUs) for Sierra.” Clouse adds that NNSA has led the advancement of portability solutions for new architectures. He says, “Rather than completely rewrite these complex codes, we are working to make codes run well on a wide range of platforms with relatively small, isolated modifications.”



The Livermore-developed RAJA software abstraction layer systematically isolates machine-specific optimizations without disrupting the source code, enabling the latter to run on different platforms. For example, ALE3D and ARES are two of Livermore's proprietary hydrodynamics codes, and MFEM is a high-order finite element library.

Global Distinction

Heterogeneous computing architectures combine graphics processing units with central processing units to achieve higher performance with more computing flexibility and less energy consumption. Supercomputing centers worldwide recognize the value of these architectures, and the Department of Energy (DOE) aims to stay at the forefront of developing leading-edge machines.

The biannual TOP500 list evaluates high-performance computing (HPC) systems using a linear algebra processing benchmark designed for distributed-memory architectures. In June 2020, Livermore's Sierra ranks third on this prestigious list. In fact, DOE laboratories house four of the top 20 supercomputers, including Lassen—Sierra's smaller, unclassified counterpart located at Lawrence Livermore. Sierra also ranks 12th on the Green500 list of energy-efficient supercomputers.

Terri Quinn, associate program director for Livermore Computing, emphasizes the importance of top-tier HPC capabilities to the Laboratory's national security mission. "We provide scientists and engineers with unique and powerful HPC resources to give Livermore, the National Nuclear Security Administration, and the United States a competitive advantage. World-class systems induce code advances, offer new simulation possibilities, and attract top computational and computer science talent to the Laboratory. Who wouldn't want to work on the most powerful computers in the world?"

Livermore computer scientists and software developers explored numerous portability and optimization solutions during the years leading up to Sierra's installation (see *S&TR*, September 2016, pp. 4–11). For example, new algorithms exploit data parallelism and memory access to help ensure that codes capitalize on GPUs. Novel tools combine simulation and visualization routines to maximize data processing while in memory, and innovative memory-management models automate data movement between memory locations with minimal disruption to the source code. Clouse adds, "Small parts of each code were ported over and optimized to the new system before the entire code was considered ready for Sierra."

One key GPU-portability innovation is Livermore's RAJA abstraction layer. Large multiphysics codes typically contain millions of lines of code and thousands of calculation loops. RAJA provides abstractions at the loop level, separating platform-independent and platform-specific code for streamlined execution. Many of Livermore's production codes have adopted RAJA, and codes under development will include abstraction layers from inception.

Memory-allocation tools, such as the Livermore-developed Umpire and CHAI (Copy-Hiding Application Interface), are crucial partners of abstraction layers. Memory movement between Sierra's GPUs and CPUs is coherent, which means that both processor types share and communicate any changed values in memory regardless of which processor recorded the changes. Clouse elaborates, "In this setup, GPUs and CPUs remain in sync when they access data from memory."

Tools such as CHAI and Umpire allow more control over Sierra's memory allocations to improve its performance. "They also provide portability on machines that, unlike Sierra, do not guarantee memory coherence," says Clouse. Together, abstraction and memory management simplify how codes run on Sierra. Programmers do not need to explicitly allocate data on GPUs because these software tools do it for them.

Simulation Speed-Ups

Before Sierra, researchers primarily relied on two-dimensional (2D) approximations of three-dimensional (3D) simulations, which were computationally expensive and therefore performed sparingly. The machine can run 3D simulations more efficiently and effectively. Furthermore, Sierra can run complex calculations with fewer nodes, which means dozens of simulations can run concurrently.

Clouse explains, "For many of our 3D applications, Sierra's architecture brings speed-ups on the order of 5 to 20 times what our (older) commodity clusters can do. With its extraordinary resolution, we can begin replacing our daily 2D workload with 3D simulations." In one example, Sierra produced a 3D



inertial confinement fusion simulation in 60 hours compared with a 30-day estimate on Livermore's multi-core CPU system, Sequoia. The resulting data set provides further understanding of turbulence models.

With Sierra, Livermore scientists are seeing significant impacts on programmatic work. For the W80-4 life-extension program, simulations run on Sierra help assess the warhead's new and refurbished components in 3D (see *S&TR*, October/November 2018, pp. 4–11). In another effort, a research team runs machine-learning algorithms on Sierra to analyze data from simulations and experiments (see *S&TR*, March 2019, pp. 4–11). Researchers at NNSA's other national laboratories, Sandia and Los Alamos, also use Sierra for stockpile stewardship applications. Quinn notes, "Users tell us they can run calculations they would never have dreamed of running before Sierra."

From Petaflops to Exaflops

Today's fastest computing technologies will be considered slow tomorrow. In August 2020, DOE announced a partnership with Cray Inc. (now Hewlett Packard Enterprises), to build NNSA's first exascale computer at Livermore. El Capitan is expected to come online in 2023 with a new peak performance standard of at least 1.5 exaflops, or 1.5 quintillion (1.5×10^{18}) floating-point operations per second, ushering in the next factor of 1,000 gains in computing power. Clouse states, "El Capitan will also have a GPU-based architecture. Our portability and optimization work for Sierra will benefit us greatly when the time comes, but no doubt El Capitan will present unique challenges."

Quinn points out that DOE's first three exascale-class supercomputers—El Capitan, Argonne National Laboratory's

Bruce Hendrickson (right), Livermore's associate director for Computing, talks with former U.S. Secretary of Energy Rick Perry during a tour of the Sierra machine room. (Photo by Julie Russell.)

Aurora, and Oak Ridge National Laboratory's Frontier—will use GPUs. She states, "GPUs are becoming more popular for scientific and engineering workloads, and I expect this trend to continue for the remainder of this decade."

Meanwhile, Sierra churns through the Laboratory's physics codes, improving simulation fidelity and prediction while laying the groundwork for exascale machines. The intervening years will provide abundant opportunities to leverage Sierra's capabilities for NNSA's weapons program and nuclear counterproliferation and counterterrorism efforts. "The ability to process crucial simulations efficiently and more realistically means 3D resolution is becoming routine," states Clouse. "Sierra is a game-changer for computational scientists."

—Holly Auten

Key Words: Advanced Simulation and Computing (ASC) Program, central processing unit (CPU), Copy-Hiding Application Interface (CHAI), Department of Energy (DOE), exascale, graphics processing unit (GPU), high-performance computing (HPC), memory, National Nuclear Security Administration (NNSA), portability, RAJA, Sierra, simulation, supercomputer, Umpire.

For further information contact Chris Clouse (925) 422-4576 (clouse1@llnl.gov) or Rob Neely (925) 423-4243 (neely4@llnl.gov).