Exascale Breakthroughs

By John Benson

December 2023

1. Introduction

Early in my career I worked in a position where I modeled hydrodynamic and thermodynamic systems. The technique we used in that work was generally called finite-element modeling, where we broke each system down into a large number of cells, and used equations for each cell that described the liquid-flow (hydrodynamic) and/or heat-transfer (thermodynamic) in and out of each cell.

Many of the techniques we used back then (late 1970s) were very primitive, but the basic finite-element technique remains in place, and in use for modeling the largest system we deal with currently – our planet's biosphere.

My across-town neighbors (and others) recently won a major award for computer model-building, and described some breakthroughs they made in winning this prize below.

A Lawrence Livermore National Laboratory (LLNL)-led effort that performed an unprecedented global climate model simulation on the world's first exascale supercomputer has won the first-ever Association for Computing Machinery (ACM) Gordon Bell Prize for Climate Modelling, ACM officials announced Thursday (Nov 16).[1]

2. Climate Modeling

The Simple Cloud Resolving E3SM[2] *Atmosphere Model (SCREAM) team, led by LLNL staff scientist Peter Caldwell and including researchers from Sandia National Laboratories (Sandia) and six other Department of Energy (DOE) national laboratories, took home the prestigious prize at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC23) in Denver. The prize — the highest honor in high performance computing for the climate modeling category — was established to recognize "innovative parallel computing contributions toward solving the global climate crisis," according to ACM. The award will be given out annually for the next 10 years to recognize the contributions of climate scientists and software engineers in addressing climate change.*

"The Gordon Bell Prize has been an inspiration to DOE's climate scientists for many years (and doubly so for this new climate award)," Caldwell said. "In that sense, Gordon Bell has accomplished his goal of pushing climate science forward."

The SCREAM team's winning Gordon Bell submission, led by Energy Exascale Earth System Model (E3SM) chief computational scientist Mark Taylor of Sandia, presents the efficient and performance-portable implementation of SCREAM. SCREAM is the first global cloud-resolving model to run on an exascale supercomputer and the first model to break the one-simulated-year-per-day barrier for realistic cloud-resolving simulations at less than five kilometers of horizontal resolution: achieving 1.26 simulated years per day using 32,768 graphics-processing units (GPUs) on 8,192 nodes of Oak Ridge National Laboratory's Frontier supercomputer.

[1] Lawrence Livermore National Laboratory (LLNL), "LLNL-led SCREAM team clinches inaugural Gordon Bell Prize for climate modeling," Nov 16, 2023, https://www.llnl.gov/article/50631/llnl-led-scream-team-clinches-inaugural-gordon-bell-prize-climate-modeling

[2] Energy Exascale Earth System Model

This model is also is the first to run on both AMD and NVIDIA GPUs, as well as conventional CPUs. This benchmark time is about 50% faster than the initial prize submission, thanks mainly to improved compile-time settings, researchers said. The paper also discusses the benefits of running the model on an exascale system and how it can be used to improve scientific understanding of climate change and weather patterns.

Members of the Simple Cloud Resolving E3SM Atmosphere Model (SCREAM) team, led by Lawrence Livermore National Laboratory staff scientist Peter Caldwell and including researchers from Sandia National Laboratories and six other Department of Energy national laboratories, won the first-ever Gordon Bell Prize for Climate Modelling at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC23) in Denver on Nov. 16.

“This award is a testament to teamwork,” Caldwell said. “It took us five years and the equivalent of six full-time employees spread across about 15 staff members to get where we are. We come from an interesting mix of backgrounds: half of the people are atmospheric physicists and the other half are computer scientists or applied mathematicians. One of the major accomplishments of this effort was just learning how to work together across disciplines...”

Caldwell singled out Sandia’s Andy Salinger for coming up with SCREAM’s performance-portability strategy, and LLNL scientist Dave Bader, leader of the E3SM project that SCREAM is part of, for his “unwavering support,” as well as DOE’s Office of Biological and Environmental Research for sponsoring the work. Other LLNL staff named in the Gordon Bell entry include scientists Aaron Donahue, Chris Terai and Renata McCoy.

As a full-featured atmospheric general-circulation model developed for very fine-resolution simulations on exascale machines, SCREAM incorporates state-of-the-art parameterizations for fluid dynamics, microphysics, moist turbulence and radiation. SCREAM can explicitly resolve large convective circulations and other important atmospheric phenomena, thereby avoiding critical sources of uncertainty in traditional climate models. By running SCREAM at exascale, E3SM can provide more trustworthy and higher-fidelity predictions of future climate for longer durations than were previously possible, researchers said.

Sandia's Luca Bertagna, who presented the paper at the SC23 conference, noted that by rewriting SCREAM in Kokkos, a C++ library for on-node parallelism, the team was able to run efficiently across the spectrum of computing architectures, including heterogenous supercomputers using either AMD or NVIDIA GPUs (with support for Intel GPUs on the way).

Bertagna emphasized that "what we did for Gordon Bell was not just for Gordon Bell, we actually need this stuff… All of these simulations we run we use the same preparation that we use for the Gordon Bell submission, so we're using that configuration to answer actual climate-modeling questions." In particular, SCREAM has already completed a set of 40-day simulations for all four seasons and has almost finished a pair of 13-month simulations aimed at estimating climate sensitivity.

The team is currently gearing up to launch a pair of multi-decadal simulations used to capture the statistics of rare high-impact events.

SCREAM researchers also said the work is a significant milestone, establishing E3SM as a reference point for all future climate models. LLNL's Caldwell said the team has "pioneered a performance-portable path for running climate models on exascale machines. Portability has huge advantages because it allows the team to take advantage of compute time on any available machine and provides protection against future disruptive changes in computing architecture." He added that the impacts of the paper will be felt for years to come, as computing moves into the exascale era.

"I predict that all climate models are going to need to move to higher resolution in order to provide state-of-the-art predictions," Caldwell said. "We've been trying to parameterize these sub-grid scale effects for more than 50 years and I'm not sure we will ever get the accuracy we need without instead explicitly resolving the things that matter; computing is giving us the ability now to do that. Our work has received a ton of excitement from a lot of different centers looking to move to these global cloud-resolving models. I think this class of model is just going to become more of the norm and the expectation of people looking for accurate climate predictions."

3. Gordon Bell

Chester Gordon Bell (born August 19, 1934) is an American electrical engineer and manager. An early employee of Digital Equipment Corporation (DEC) 1960–1966, Bell designed several of their PDP machines and later became Vice President of Engineering 1972–1983, overseeing the development of the VAX computer systems. Bell's later career includes entrepreneur, investor, founding Assistant Director of NSF's Computing and Information Science and Engineering Directorate 1986–1987, and researcher emeritus at Microsoft Research, 1995–2015.[3]

[3] https://en.wikipedia.org/wiki/Gordon_Bell

During the 1980s Bell became involved with public policy, becoming the first and founding Assistant Director of the CISE Directorate of the NSF, [4] and led the cross-agency group that specified the National Research and Education Network.

Bell also established the ACM Gordon Bell Prize (administered by the ACM and IEEE) in 1987 to encourage development in parallel processing. The first Gordon Bell Prize was won by researchers at the Parallel Processing Division of Sandia National Laboratory for work done on the 1000-processor nCUBE-10 hypercube.

3.1. Bell's Law of Computer Classes

Bell's law of computer classes was first described in 1972 with the emergence of a new, lower priced microcomputer class based on the microprocessor. Established market class computers are introduced at a constant price with increasing functionality and performance. Technology advances in semiconductors, storage, interfaces and networks enable a new computer class (platform) to form about every decade to serve a new need. Each new usually lower-priced class is maintained as a quasi-independent industry / market. Classes include: mainframes (1960s), minicomputers (1970s), networked workstations and personal computers (1980s), browser-web-server structure (1990s), palmtop computing (1995), web services (2000s), convergence of cell phones and computers (2003), and Wireless Sensor Networks aka motes (2004) …

4. Other Exascale Breakthroughs

To really understand how a material behaves, researchers need to simulate its whirling electrons, which govern most of its chemical and electronic properties. But they have traditionally faced a trade-off. They could simulate up to a couple of hundred electrons with near-perfect accuracy. Or they could simulate a much larger number—while accuracy fell off a cliff.[5]

The world's most powerful supercomputers, operating at the far frontier of speed known as the exascale, have now begun to eliminate that trade-off.

At SC23, a supercomputing conference here, researchers last week reported simulating the behavior of up to 600,000 electrons within a microscopic chunk of a magnesium alloy with nearly the accuracy of a quantum Monte Carlo simulation, the gold standard for much smaller numbers of electrons. "We broke through the accuracy-length scale barrier," says Sambit Das, a mechanical engineer at the University of Michigan and a member of the team presenting the work.

The simulations showed how defects form in the alloys, which could open the way to designing novel lightweight alloys for fuel efficient cars and airplanes. Applying the techniques to materials called quasicrystals, ordered solids that lack repeating atomic arrangements, also revealed why they take the unusual shapes they do, an advance that could lead to novel magnetic materials and superconductors.

---

[4] Computer and Information Science and Engineering Directorate of the U.S. National Science Foundation

[5] Robert F. Service, Science, "Exascale computers show off emerging science," Nov 24, 2023, https://www.science.org/content/article/world-s-fastest-supercomputers-are-helping-sharpen-climate-forecasts-and-design-new

In other feats of computation, researchers at SC23 reported efforts to predict air flow and noise from a fuel-efficient jet engine design, and how heat would pulse through the core of a small modular nuclear reactor, an advance that could inform safer designs. All are among the early results emerging from one of the world's first exascale supercomputers, Frontier, at the Department of Energy's (DOE's) Oak Ridge National Laboratory. Capable of 1.1 exaflops (1018 flops), or 1.1 billion billion operations per second, Frontier is more than twice as fast as the fastest machine from just 2 years ago.

These results and others on the way from exascale machines coming online in the next few years promise to open a new window into materials, climate science, biology, and medicine. "There is a new science era that is unfolding," says Ceren Susut, associate director for DOE's Advanced Scientific Computing Research program.

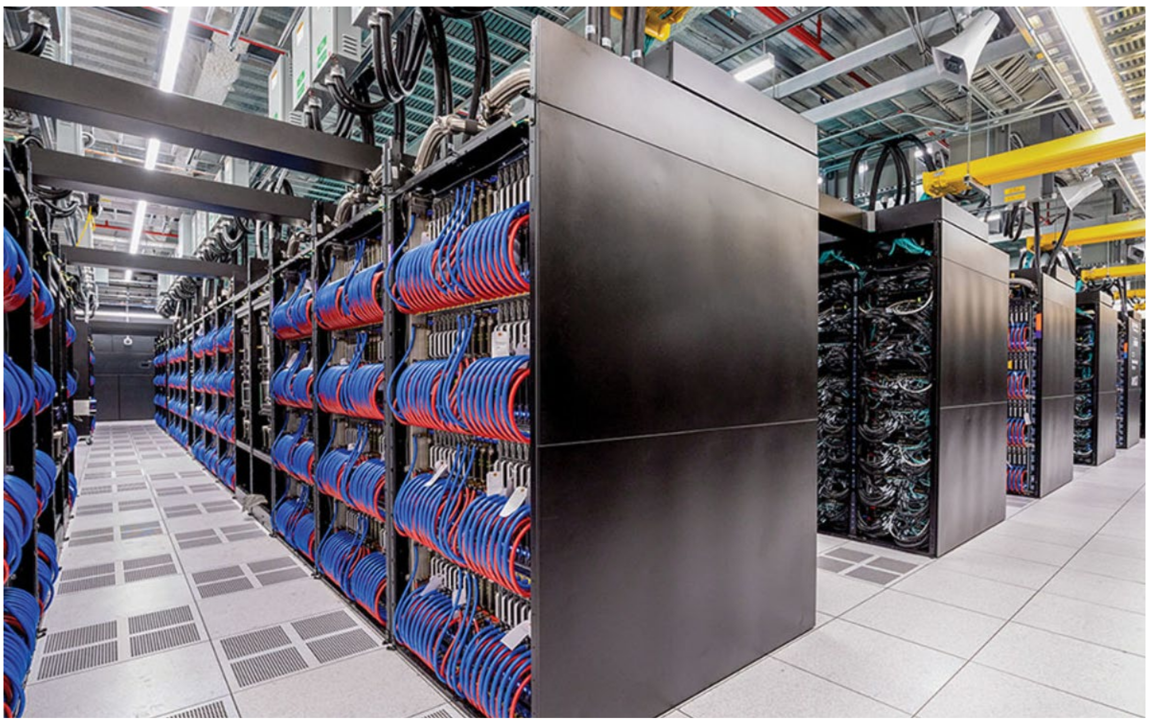

Frontier, at Oak Ridge National Laboratory, can make more than 1.1 billion billion computations per second. PHOTO: OAK RIDGE NATIONAL LABORATORY

5. El Capitan: The Next Step

Continuing its tradition as a dominant Top500 high-performance computing (HPC) site, Lawrence Livermore National Laboratory (LLNL) began installing components in May 2023 for the National Nuclear Security Administration's (NNSA's) first exascale supercomputer, El Capitan. An exascale supercomputer can calculate at least one quintillion (1,000,000,000,000,000,000+) Double Precision (64-bit) operations per second (1 exaflop). When it is deployed in 2024, El Capitan is projected to be the world's most powerful supercomputer capable of performing more than 2 exaflops per second.[6]

[6] Lawrence Livermore National Labs, Exascale, "El Capitan", https://asc.llnl.gov/exascale/el-capitan

Funded by NNSA's Advanced Simulation and Computing (ASC) program, El Capitan was a collaboration among the three NNSA labs—Livermore, Los Alamos, and Sandia. El Capitan will ensure the safety, security, and reliability of the nation's nuclear stockpile in the absence of underground testing. It will be essential for the design and stewardship of a modernized stockpile and other critical national security missions. Research performed on El Capitan will also support unclassified mission areas of interest to national security, including material discovery, high energy-density physics, nuclear data, material equations of state, and conventional weapon design.

To ensure the system achieves its full computing potential, LLNL is investing in cognitive simulation capabilities such as artificial intelligence (AI) and machine learning (ML) techniques that will benefit both unclassified and classified missions.

El Capitan's details and distinguishing features:

- *Siting complete in 2024*
- *Expected peak performance ≥ 2.0 exaflops*
- *Peak power < 40 MW (anticipating ~30 MW)*
- *AMD MI300 accelerated processing unit (APU)-3D chiplet design, which includes a tightly coupled central processing unit (CPU)—graphics processing unit (GPU) in one processing unit*
- *Slingshot interconnect*
- *Innovative Rabbit near-node local storage*

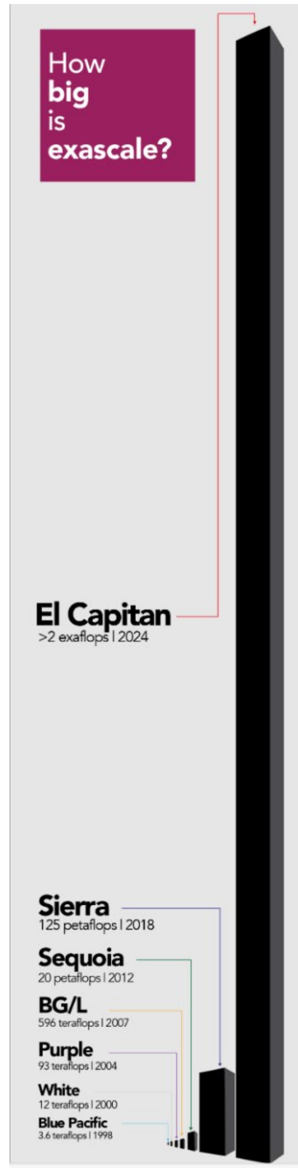

Final author's comment: Although El Capitan will mainly be used for national defense functions, many of these involve advanced system simulation, and this will advance future Exascale computers' techniques to model our climate and other systems. One of the reasons that these national labs support both nuclear weapons development and environmental-science applications is that both are important to our national interests, and both require the most advanced computing capabilities.