

Nascent exascale supercomputers offer promise, present challenges

Adam Mann, *Science Writer*


Sometime next year, managers at the US Department of Energy's (DOE) Argonne National Laboratory in Lemont, IL, will power up a calculating machine the size of 10 tennis courts and vault the country into a new age of computing. The \$500-million mainframe, called Aurora, could become the world's first "exascale" supercomputer, running an astounding 10^{18} , or 1 quintillion, operations per second.

Aurora is expected to have more than twice the peak performance of the current supercomputer record holder, a machine named Fugaku at the RIKEN Center for Computational Science in Kobe, Japan. Fugaku and its calculation kin serve a vital function in modern scientific advancement, performing simulations crucial for discoveries in a wide range of fields. But the transition to exascale will not be easy. "As these machines grow, they become harder and harder to exploit efficiently," says Danny Perez, a physicist at Los Alamos National

Laboratory in NM. "We have to change our computing paradigms, how we write our programs, and how we arrange computation and data management."

That's because supercomputers are complex beasts, consisting of cabinets containing hundreds of thousands of processors. For these processors to operate as a single entity, a supercomputer needs to pass data back and forth between its various parts, running huge numbers of computations at the same time, all while minimizing power consumption. Writing programs for such parallel computing is not easy, and theorists will need to leverage new tools such as machine learning and artificial intelligence to make scientific breakthroughs. Given these challenges, researchers have been planning for exascale computing for more than a decade (1).

Multiple countries are competing to get to exascale first. China has said it would have an exascale machine



Rows of cabinets hold incredible processing power for one of the world's best supercomputers, Summit, at Oak Ridge National Laboratory in TN. Exascale computing will surpass these existing computers by leaps and bounds. Image credit: Flickr/Oak Ridge National Laboratory, licensed under [CC BY 2.0](https://creativecommons.org/licenses/by/2.0/).

Published under the [PNAS](https://www.pnas.org/) license.

First published September 9, 2020.

www.pnas.org/cgi/doi/10.1073/pnas.2015968117

PNAS | September 15, 2020 | vol. 117 | no. 37 | 22623–22625

www.manaraa.com

by the end of 2020, although experts outside the country have expressed doubts about this timeframe even before the delays caused by the global severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) pandemic. The United States aims to have Aurora operational sometime in 2021.

Engineers in Japan and the European Union are not far behind. "Everyone's racing to exascale," says

"Everyone's racing to exascale. Who gets there first, I don't know."

—Jack Dongarra

Jack Dongarra, a computer scientist at the University of Tennessee in Knoxville. "Who gets there first, I don't know." Along with bragging rights, the nations that achieve this milestone early will have a leg up in the scientific revolutions of the future.

Computational Boost

The increase in the power of computers has long followed Moore's Law, named after Intel cofounder Gordon Moore, who observed in 1965 that the processing power of computer chips was doubling roughly every two years (2). Supercomputers shifted from being able to do thousands of operations per second to millions, then billions, then trillions per second, at a cadence of roughly a thousand-fold increase in ability per decade.

Such powerful computing required enormous amounts of electricity. Unfortunately, much of this power it was getting lost as wasted heat—a considerable concern in the mid-2000s, as researchers grappled with petascale computing capable of 10^{15} calculations per second. By 2006, IBM partly solved this problem by designing chips known as graphics processing units (GPUs), meant for the newly released Sony PlayStation 3. GPUs are specialized for rapidly rendering high-resolution images. They divide complex calculations into smaller tasks that run simultaneously, a process known as parallelization, making them quicker and more energy-efficient than generalist central processing units (CPUs). GPUs were a boon for supercomputers.

In 2008, when Los Alamos Laboratory unveiled Roadrunner, the world's first petascale supercomputer, it contained 12,960 GPU-inspired chips along with 6,480 CPUs and performed twice as well as the next best system at the time. Besides GPUs, Roadrunner included other innovations to save electricity, such as turning on components only when necessary. Such energy efficiency was important because predictions for achieving exascale back then suggested that engineers would need "something like half of a nuclear power plant to power the computer," says Perez.

For such highly interconnected supercomputers, performance might be pinched by bottlenecks, such as the ability to access memory or store and retrieve data quickly. Newer machines in fact try to avoid shuffling around information as much as possible,

sometimes even recomputing a quantity rather than restoring it from slow memory. Issues with memory and data retrieval are only expected to get worse in exascale. Should any link in the chain of computation have bottlenecks, it can cascade into larger problems. This means that a machine's peak performance, the theoretical highest processing power it can reach, will be different from its real-world, sustainable performance. "In the best case, we can get to around 60 or 70 percent efficiency," says Depei Qian, an emeritus computer scientist at Beihang University in Beijing, China, who helps lead China's exascale efforts.

Hardware is not the only challenge—the software comes with its own set of problems. Before the transition to petascale, Moore's law brought performance improvements without having to completely rethink how a program was written. "You could just use the old programs," says Perez. "That era is over. The low hanging fruits—we've definitely plucked them."

That's partly because of those GPUs. But even before they came along, programs were parallelized for speed: They were divided into parts that ran at the same time on different CPUs, and the outputs were recombined into cohesive results. The process became even more difficult when some parts of a program had to be executed on a CPU and some on a GPU. Exascale machines will contain on the order of 135,000 GPUs and 50,000 CPUs, and each of those chips will have many individual processing units requiring engineers to write programs that execute almost a billion instructions simultaneously.

So running existing scientific simulations on the new exascale computers is not going to be trivial. "It's not just picking out a [simulation] and putting it on a big computer," says L. Ruby Leung, an atmospheric scientist at the Pacific Northwest National Laboratory in Richland, WA. Researchers are being forced to reexamine millions of lines of code and optimize them to make use of the unique architectures of exascale computers, so that the programs can reach as close to the theoretical maximum processing power as possible.

Teams around the world are wrestling with the different tradeoffs of achieving exascale machines. Some groups have focused on figuring out how to add more CPUs for calculations, making these mainframes easier to program but harder to power. The alternative approach has been to sacrifice programmability for energy efficiency, striving to find the best balance of CPUs and GPUs without making it too cumbersome for users to run their applications. Architectures that minimize the transfer of data inside the machine, or use specialized chips to speed up specific algorithms, are also being explored.

Critical Calculations

Despite all these challenges, researchers are intent on harnessing the power of exascale machines. Science has technically already entered the exascale era with the distributed computing project Folding@home. Users can download the program and allow it to commandeer tiny bits of available processing power on their home PCs to solve biomedical conundrums.

Folding@home announced in March that when all of its 700,000 participants were online, the project had the combined capacity to perform more than 1.5 quintillion operations per second. These simulation abilities have been put to use during the pandemic to search for drugs effective against COVID-19.

Indeed, the field of biochemistry makes heavy use of supercomputers, which can act like “computational microscopes” to let researchers peer closely at the otherwise invisible ways that molecules interact. In the 1990s, researchers were only able to study a single organic chemical *in silico* for a few virtual trillionths of a second. But today’s best machines can routinely model the movement of complex entities, such as viruses, over timescales of milliseconds.

Exascale supercomputers will enable simulations that are more complex and of higher resolution, allowing researchers to explore the molecular interactions of viruses and their hosts with unprecedented fidelity. In principle, the boost in computing power could help researchers better understand how lifesaving molecules bind to various proteins, guide biomedical experiments in HIV and cancer trials, or even aid in the design of a universal influenza vaccine. And whereas current computers can only model one percent of the human brain’s 100 billion neurons, exascale machines are expected to be able to simulate 10 times more of the brain’s capabilities, in principle helping to elucidate memory and other neurological processes.

On a massively grander scale, the next generation of computers promise to offer insight into the potentially disastrous effects of climate change. Weather phenomena are prototypical examples of chaotic behavior in action, with countless minor feedback loops that have planetary-scale consequences. A coordinated effort is ongoing into building the Energy Exascale Earth System Model (E3SM), which will simulate biogeochemical and atmospheric processes over land, ocean, and ice with up to two orders of magnitude better resolution than current models (3). This should more accurately reproduce real-world observations and satellite data, helping determine where adverse effects such as sea-level rise or storm inundation might do the most damage to lives and livelihoods. Exascale power will allow climate forecasters to swiftly run thousands of simulations, introducing tiny variations in the initial conditions to better gauge the likelihood of events a hundred years hence.

Chemistry, cosmology, high-energy physics, materials science, oil exploration, and transportation will likely all benefit from exascale computing. Paired with machine learning, exascale computers should enhance researchers’ capacity for teasing out important patterns in complex datasets. For instance, experimental nuclear fusion reactors, where superheated plasma is contained within powerful magnetic fields, have artificial intelligence (AI) programs on supercomputers that indicate when the plasma might be on the verge of becoming unstable. Computers can then adjust the magnetic fields to shepherd the plasma and keep it from breaching its constraints and hitting the walls of a reactor. Exascale machines should allow for faster reaction times and greater precision in such systems.

“Artificial intelligence is helping to identify relationships that are impossible to find using traditional computing,” says Paresh Kharya, who is responsible for data center product management at the AI computing platform company NVIDIA in Santa Clara, CA.

Computing power figures to increase considerably in the coming years. Following Aurora, the DOE plans to bring online a \$600-million machine named Frontier at Oak Ridge National Laboratory in TN in late 2021 and a third supercomputer, El Capitan, at Lawrence Livermore National Laboratory in CA, two years later, each of which will be more powerful than their predecessor. The European Union has a range of exascale programs in the works under its European High-Performance Computing Joint Undertaking, whereas Japan is aiming for the exascale version of Fugaku to be available to users within a couple years. China—which had no supercomputers as recently as 2001 but now boasts the fourth and fifth most powerful machines on Earth—is pursuing three exascale projects (4). China has said that it expected the first, Tianhe-3, to be complete this year, but project managers say that the coronavirus pandemic has pushed back timelines.

Ironically, it is just these sorts of urgent, seemingly intractable problems—swiftly developing vaccines and therapeutics to address COVID-19, for example—for which exascale computers are meant. If groups can solve the technical challenges, there should be an impressive array of applications. Says Qian, “Supercomputing is supposed to benefit ordinary people in their daily lives.”

- 1 Kogge, P. et al., Exascale computing study: Technology challenges in achieving exascale systems. *Defense Advanced Research Projects Agency Information Processing Techniques Office (DARPA IPTO), Technical Representative* **15** (2008).
- 2 G. Moore, Cramming more components onto integrated circuits. *Electronics (Basel)* **38**, 114–117 (1965).
- 3 J. Golaz et al., The DOE E3SM coupled model version 1: Overview and evaluation at standard resolution. *J. Adv. Model. Earth Syst.* **11**, 2089–2129 (2019).
- 4 Y. Lu, Paving the way for China exascale computing. *CCF Transactions on High Performance Computing* **1**, 63–72 (2019).