



News & Highlights

First Supercomputer Breaks Exascale Barrier, with More Expected Soon

Mitch Leslie

Senior Technology Writer



The world's fastest publicly acknowledged supercomputer, known as Frontier, sprawls across 372 m² of floor space in a building at the Oak Ridge National Laboratory (ORNL) in Tennessee, USA (Fig. 1) [1]. In May of 2022, the 269 000 kg behemoth became the first computer to cross the so-called exascale barrier and reach a top speed of more than one exaflop, or over one quintillion floating-point operations per second, in the semiannual TOP500 rankings of supercomputer performance [2]. Maxing out at 1.1 exaflops, Frontier was more than twice as fast as its nearest competitor, and it repeated the feat in November of 2022 in the next TOP500 standings [3,4].

Frontier's speed record heralds a surge in exascale computing. The United States will soon debut two more exascale machines, and a coalition of European countries is building its own [5,6]. China may already have one or two exascale supercomputers, although the country has not confirmed their existence and did not enter any such machines in the recent TOP500 evaluations [3,4]. However, according to some predictions, China could be operating as many as 10 exascale supercomputers by 2025 [7].

Breaking the exascale barrier “is a next-level achievement,” said Simon McIntosh-Smith, professor of high-performance computing at Bristol University in the United Kingdom. He and other experts expect this explosion of number-crunching power to revolutionize a variety of scientific fields, allowing researchers to develop more detailed, realistic, comprehensive, and informative models and simulations. Among the areas that could benefit from the new machines are climate prediction, materials science, astrophysics, energy research, and vaccine development and testing. “Science today is driven by simulation,” said Jack Dongarra, professor of computer science at the University of Tennessee, Knoxville, TN, USA, and one of the experts behind the TOP500 list. “We need exascale computing to help push science forward.”

Dongarra, who has been tracking supercomputer performance since the 1970s, wrote the series of linear equations, called the Linpack benchmark, which serves as the standard for the machines' capabilities [8]. In 1993, when he and colleagues launched the TOP500 list [9], the fastest supercomputer was the CM-5 from the now-defunct Thinking Machines Corporation; that machine clocked in at just under 60 gigaflops [10]. Dongarra noted that supercomputer speed has increased by about nine orders of magnitude since then.

Still, developing a working exascale machine “has been an enormous engineering challenge on every level,” said McIntosh-Smith. Several design features allow Frontier to be so fast. The machine contains 9408 central processing units (CPUs), each of which boasts 64 cores, or individual subprocessors [11]. This arrangement is what computer scientists call massively parallel—each core can work on part of a problem, speeding the solution [12]. Frontier's 37 632 graphics processing units (GPUs) also contribute to its record performance [11]. Originally designed to deliver high-end visuals for applications such as video games, GPUs turned out to be particularly good at scientific processing, providing “five to ten times more number-crunching power than the fastest CPUs,” said McIntosh-Smith. Frontier contains about 10 000 more GPUs than its ORNL predecessor, Summit, a 200-petaflop machine that held the speed record for about 18 months [13]. Not all supercomputers include GPUs. The second-place competitor on the TOP500 list, Fugaku from the Riken Center for Computational Science in Kobe, Japan, only contains CPUs—almost 159 000 of them [13–15]. But Frontier's GPUs are one reason it is more than 2.5 times faster than Fugaku [5].



Fig. 1. A view from outside of Frontier shows several of the 74 Cray cabinets, each weighing more than 3600 kg, that make up the supercomputer. The machine features 2.0 GHz central processing units (CPUs) and 1.7 GHz graphics processing units (GPUs), all made by Advanced Micro Devices. Credit: Oak Ridge National Laboratory (CC BY 2.0).

Even the fastest cores will underperform if they cannot access the data they need. In supercomputers, “the computational bottleneck tends to be data movement. It is orders of magnitude slower than doing the arithmetic,” said Dimitrios Nikolopoulos, professor of engineering at Virginia Tech University in Blacksburg, VA, USA. Several of Frontier’s hardware features accelerate the process. To improve data access, the GPUs carry 128 gigabytes of their own high-bandwidth memory [11]. In addition, high-speed interconnects link CPUs to GPUs and shuttle data among each of the supercomputer’s blades, processing units that hold two CPUs and eight GPUs [11,16].

Like other supercomputer designers, Frontier’s team had to deal with two challenges that stem from cramming in so many processors. The first is heat production [17]. To keep Frontier from getting too hot, its cooling system can circulate more than 151 000 L of water per minute through the machine and then direct the water to cooling towers that allow the heat to dissipate (Fig. 2) [5,18].

Supercomputers also require vast amounts of power—Frontier’s annual electricity consumption is 40 MW, about the same as 30 000 houses [19]. But its designers took several steps to cut its power use. For instance, its cooling system requires that water only be “chilled” to 29 °C before entering the machine—by contrast, some supercomputers use water cooled to 15 °C [5,20]. Frontier’s reliance on GPUs also leads to power savings, said McIntosh-Smith, because they can be more energy efficient than CPUs [21]. “That is why it was possible to build an exascale machine that did not need its own power station next door,” he said. The designers’ focus on cutting power use also made Frontier the most efficient supercomputer in the TOP500 list, achieving 52 gigaflops per watt [12].

More exascale supercomputers should begin crunching numbers soon. The second US machine, known as Aurora and located at the Argonne National Laboratory in Lemont, IL, USA, is set to come online in 2023 [22], with a peak speed of two exaflops [23]. The Lawrence Livermore National Laboratory in California, USA will host a third exascale supercomputer, El Capitan, which will begin operating later in the decade [21]. JUPITER, Europe’s entry in the exascale supercomputer race, is being built in Germany and should switch on in 2023 [6]. And much faster machines are in the works. The United States is planning a successor to Frontier that will be capable of 5–10 exaflops [24], and Japan is looking to build a machine by the end of the decade that could reach 20 exaflops, said McIntosh-Smith.

It took researchers 14 years to boost supercomputer speed 1000 times, from one petaflop to one exaflop, Dongarra said. He predicts that surpassing the next big barrier, zettaflops, or 1000 exaflops,



Fig. 2. Part of Frontier’s cooling system, which can circulate more than 151 000 L of water per minute through the machine. The system is 30%–40% more efficient than the cooling system of its predecessor supercomputer, Summit. Credit: Oak Ridge National Laboratory (CC BY 2.0).

will take longer because the rate of improvement in computer chips is slowing [25]. Nikolopoulos agreed. “We are seeing hard limits on what the traditional semiconductor computer can do,” he said. The astronomical cost of the machines may also slow their rate of improvement. Building three exascale supercomputers will cost the United States around 1.8 billion USD, Dongarra said.

Whatever the future of supercomputers holds, researchers are eager to discover what they can do with the machines that are coming online. Two of the US supercomputers will be publicly accessible through competitive research grants. As with many existing supercomputers, scientists will be able to submit proposals to conduct research on the machines, said Dongarra. The supercomputer at Lawrence Livermore National Laboratory will be used to conduct classified research.

The enhanced simulation abilities of exascale supercomputers may allow scientists to probe potential new sources of energy, such as nuclear fusion, and design more efficient solar panels and wind turbines. In medicine, exascale supercomputers could allow researchers to virtually test reformulated vaccines against new variants of severe acute respiratory syndrome coronavirus 2, the virus that causes coronavirus disease 2019, dramatically reducing development time. And by creating better models of the Earth’s climate and weather, researchers might gain a better understanding of the effects of climate change.

With several countries pursuing these supercomputers, the fact that the United States was the first to break the exascale barrier—at least according to the TOP500 rankings—is not that important, said Nikolopoulos. “The real question is not who delivers the first machine, but who makes the best use of these machines to benefit society and humanity.”

References

- [1] Brodtkin J. 1.1 quintillion operations per second: US has world’s fastest supercomputer [Internet]. New York City: Ars Technica; 2022 May 31 [cited 2022 Dec 25]. Available from: <https://arstechnica.com/information-technology/2022/05/1-1-quintillion-operations-per-second-us-has-worlds-fastest-supercomputer/>.
- [2] Conover E. The world’s fastest supercomputer just broke the exascale barrier [Internet]. Washington, DC: Science News; 2022 Jun 1 [cited 2022 Dec 25]. Available from: <https://www.sciencenews.org/article/supercomputer-exascale-frontier-speed-record-computing-calculations>.
- [3] Clark D. U.S. retakes top spot in supercomputer race [Internet]. New York City: New York Times; 2022 May 31 [cited 2022 Dec 25]. Available from: <https://www.nytimes.com/2022/05/30/business/us-supercomputer-frontier.html>.
- [4] Choi CQ, Anderson M. TOP500: Frontier still No. 1. Where’s China? [Internet]. New York City: IEEE Spectrum; 2022 Nov 15 [cited 2022 Dec 25]. Available from: <https://spectrum.ieee.org/frontier-exascale-top500-export-controls>.
- [5] Rimmel A. What exascale computing could mean for chemistry [Internet]. Washington, DC: Chemical & Engineering News; 2022 Sep 2 [cited 2022 Dec 25]. Available from: <https://cen.acs.org/articles/100/i31/exascale-computing-mean-chemistry.html>.
- [6] Dumiak M. Europe gets an exascale supercomputer [Internet]. New York City: IEEE Spectrum; 2022 Dec 18 [cited 2022 Dec 25]. Available from: <https://spectrum.ieee.org/europe-s-exascale-supercomputer>.
- [7] Waters R. US rushes to catch up to China in supercomputer race [Internet]. London: Financial Times; 2022 May 17 [cited 2022 Dec 25]. Available from: <https://www.ft.com/content/9ec4c04c-d71d-4d54-87fe-eef4ff92ddc9>.
- [8] Metz C. Turing Award won by programmer who paved the way for supercomputers [Internet]. New York City: New York Times; 2022 Mar 30 [cited 2022 Dec 25]. Available from: <https://www.nytimes.com/2022/03/30/technology/turing-award-jack-dongarra.html>.
- [9] Introduction and objectives [Internet]. Sinsheim: TOP500; [cited 2022 Dec 25]. Available from: <https://www.top500.org/project/introduction/>.
- [10] Celebrating 25 years of TOP500 [Internet]. Sinsheim: TOP500; [cited 2022 Dec 25]. Available from: <https://www.top500.org/25years/>.
- [11] Choi CG. The beating heart of the world’s fastest exascale computer [Internet]. New York City: IEEE Spectrum; 2022 Jun 24 [cited 2022 Dec 25]. Available from: <https://spectrum.ieee.org/frontier-exascale-supercomputer>.
- [12] Hu C. The new Frontier supercomputer will be the fastest in the world [Internet]. New York City: Popular Science; 2022 Jun 1 [cited 2022 Dec 25]. Available from: <https://www.popsoci.com/technology/frontier-supercomputer-ranks-fastest-world/>.
- [13] Palmer J. More super supercomputers. *Engineering* 2019;5(3):357–8.
- [14] About Fugaku [Internet]. Kobe: Riken Center for Computational Science; [cited 2022 Dec 25]. Available from: <https://www.r-ccs.riken.jp/en/fugaku/about/>.

- [15] Skibba R. Japan's Fugaku supercomputer crushes competition, but likely not for long. *Engineering* 2021;7(1):6–7.
- [16] Frontier [Internet]. Oak Ridge: Oak Ridge National Laboratory; [cited 2022 Dec 25]. Available from: <https://www.olcf.ornl.gov/frontier/>.
- [17] Cooling the world's fastest supercomputer Fugaku a feat for operators in Japan [Internet]. Tokyo: The Mainichi; 2021 Feb 20 [cited 2022 Dec 25]. Available from: <https://mainichi.jp/english/articles/20210220/p2a/00m/0na/037000c>.
- [18] Currin G. Director of the world's fastest supercomputer explains how they did it [Internet]. San Francisco: Interesting Engineering; 2022 Jun 14 [cited 2022 Dec 25]. Available from: <https://interestingengineering.com/innovation/director-worlds-fastest-supercomputer>.
- [19] Frontier supercomputer debuts as world's fastest, breaking exascale barrier [Internet]. Oak Ridge: Oak Ridge National Laboratory; 2022 May 30 [cited 2022 Dec 25]. Available from: <https://www.ornl.gov/news/frontier-supercomputer-debuts-worlds-fastest-breaking-exascale-barrier>.
- [20] Greenemeier L. Warm water flowed through supercomputers to cool down their heat [Internet]. New York City: Scientific American; 2010 Apr 15 [cited 2022 Dec 25]. Available from: <https://www.scientificamerican.com/article/microchannel-warm-liquid-cooling/>.
- [21] Caulfield B. What's the difference between a CPU and a GPU? [Internet]. Santa Clara: Nvidia; 2009 Dec 18 [cited 2022 Dec 25]. Available from: <https://blogs.nvidia.com/blog/2009/12/16/whats-the-difference-between-a-cpu-and-a-gpu/>.
- [22] Condon S. Intel's highest density' new GPU will help power a supercomputer [Internet]. San Francisco: ZDNet; 2022 Nov 9 [cited 2022 Dec 25]. Available from: <https://www.zdnet.com/article/intel-unveils-new-chips-that-will-power-the-upcoming-aurora-supercomputer/>.
- [23] At the Frontier: DOE supercomputing launches the exascale era [Internet]. Washington, DC: United States Department of Engineering; 2022 Jun 7 [cited 2022 Dec 25]. Available from: <https://www.energy.gov/science/articles/frontier-doe-supercomputing-launches-exascale-era>.
- [24] Kan M. US eyes supercomputers 10X faster than current exascale machines [Internet]. New York City: PC Magazine; 2022 Jun 29 [cited 2023 Jan 5]. Available from: <https://www.pcmag.com/news/us-eyes-supercomputers-10x-faster-than-current-exascale-machines>.
- [25] Edwards C. Where now for the transistor? [Internet]. Stevenage: Engineering and Technology; 2022 Dec 16 [cited 2022 Dec 25]. Available from: <https://eandt.theiet.org/content/articles/2022/12/at-75-where-now-for-the-transistor/>.