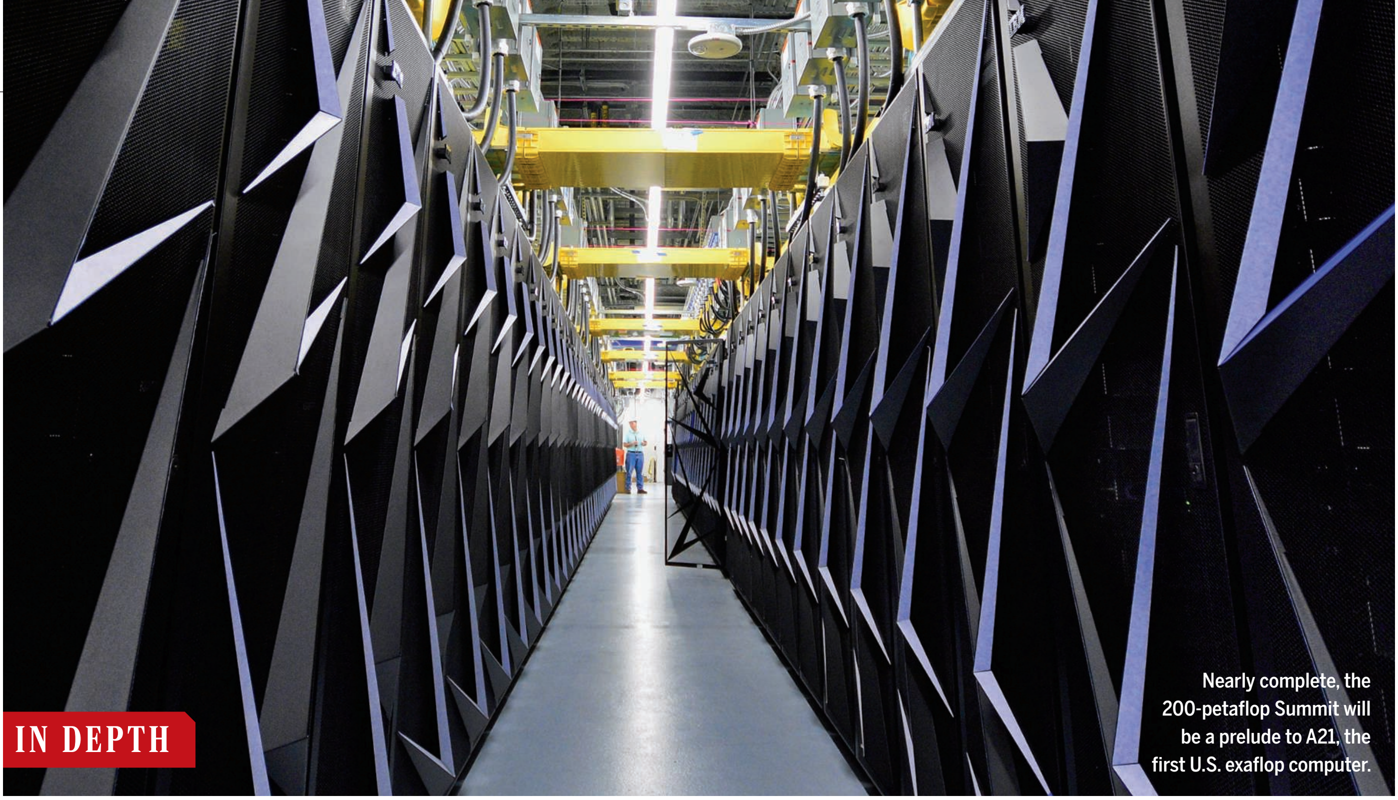

IN DEPTH

Nearly complete, the 200-petaflop Summit will be a prelude to A21, the first U.S. exaflop computer.

SUPERCOMPUTING

Design for U.S. exascale computer takes shape

Competition with China accelerates plans for next great leap in supercomputing power

By **Robert F. Service**

In 1957, the launch of the Sputnik satellite vaulted the Soviet Union to the lead in the space race and galvanized the United States. U.S. supercomputer researchers are today facing their own Sputnik moment—this time with China. After dominating the supercomputing rankings for decades, the United States is now so far behind that the combined power of the top two machines in China easily outpaces that of all 21 supercomputers operated by the U.S. Department of Energy (DOE), the country's top supercomputing funder.

But now, U.S. supercomputing researchers are striking back. Engineers at DOE's Oak Ridge National Laboratory in Tennessee have nearly completed Summit, a computer with twice the power of the top Chinese machine, the Sunway TaihuLight in Wuxi. When fully commissioned this summer, Summit will churn out 200 million billion floating-point operations per second (petaflops). Even more promising, scientists are meeting in Knoxville, Tennessee, this week to get their first detailed look at designs for the next U.S. behemoth, its first 1000-petaflop—1 exaflop or exascale—supercomputer, to be built by 2021 at Argonne National Laboratory in Lemont, Illinois. That's 2 years earlier than planned. "It's a pretty exciting time," says Aiichiro Nakano, a physicist at the University of Southern California in Los Angeles who uses supercomputers to model materials made by layering stacks of atomic sheets like graphene.

Called A21, the Argonne computer will be built by Intel and Cray and is expected to supercharge simulations of everything from the formation of galaxies to the turbulent flows of gas in combustion. "With exascale we can put a lot more physics in there," says Choong-Seock Chang, a physicist at the Princeton Plasma Physics Laboratory in New Jersey who plans to use A21 to model the plasma physics inside a fusion reactor.

> ***"With exascale we can put a lot more physics in there."***
> **Choong-Seock Chang**, Princeton Plasma Physics Laboratory

China and possibly Japan are still likely to reach the exascale promised land first (see sidebar, p. 618). But if it's completed on schedule, A21 could keep the United States from slipping too far behind. The faster pace reflects a change of strategy by DOE officials last fall. Initially, the agency set up a "two lanes" approach to overcoming the challenges of an exascale machine, in particular a potentially ravenous appetite for electricity that could require the output of a small nuclear plant.

The agency had been funding two machines, both stepping stones to the exascale, that would take different approaches to cutting the energy demand. IBM and its partner NVIDIA, the makers of Summit, have focused on marrying central processing units (CPUs) with graphical processing units, which are faster and more efficient for calculations involved in complex visual simulations. Intel and Cray, meanwhile, have long aimed to increase the number of CPU "cores" operating in parallel and creating fast links between them. Their strategy was meant to lead to a 180-petaflop sister for Summit, called Aurora, to be built at Argonne.

In 2015, DOE expected Aurora to be finished this year, with the first U.S. exascale machine appearing in 2023. Then China announced a 5-Year Plan that spelled out the goal of an exascale machine by the end of 2020. The United States wasn't just falling behind, it was about to be lapped.

"There was a lot of stress in the U.S. DOE, National Nuclear Security Agency, and in-

PHOTO: LYNN FREENY/DEPARTMENT OF ENERGY VIA FLICKR dustry," Chang says. DOE changed tacks. It scrapped plans for Aurora, and replaced it with A21, a machine five times bigger. That pushed the launch date back to 2021, but because it was to be the first U.S. exascale machine, it also effectively pushed up the U.S. timeline by 2 years.

Skipping the intermediate step of Aurora is risky, says Kenneth Jansen, an aerospace engineer at the University of Colorado in Boulder. "It means one of the stepping stones is not going to be there." Still, others say it's a risk worth taking. "This is the right way to do it," says Thom Dunning, a computational chemist at the University of Washington in Seattle.

Details of A21's architecture remain closely guarded to protect proprietary technology. But scientists writing software for the new machine will be given detailed briefings on the new architecture after they sign nondisclosure agreements. Some of the first briefings are taking place this week in Knoxville at the second annual Exascale Computing Project meeting.

Researchers already familiar with the plans say the machine is unlike any they've ever seen before. "A21 is a very different architecture," Chang says. In general terms, he says, the design focuses on decreasing the need to move data long distances between processors, an energetically expensive process. He says the new machine will likely require 25 to 30 megawatts of power, only about twice that of Summit. Asked whether he thinks Intel will be able to pull off the new architecture, Chang says, "I am confident they will."

Scaling up

Since 2013, China has operated the most powerful supercomputer in the world. Summit is likely to reclaim the title for the United States this year. China is on track to unveil the first exascale computer in 2020.

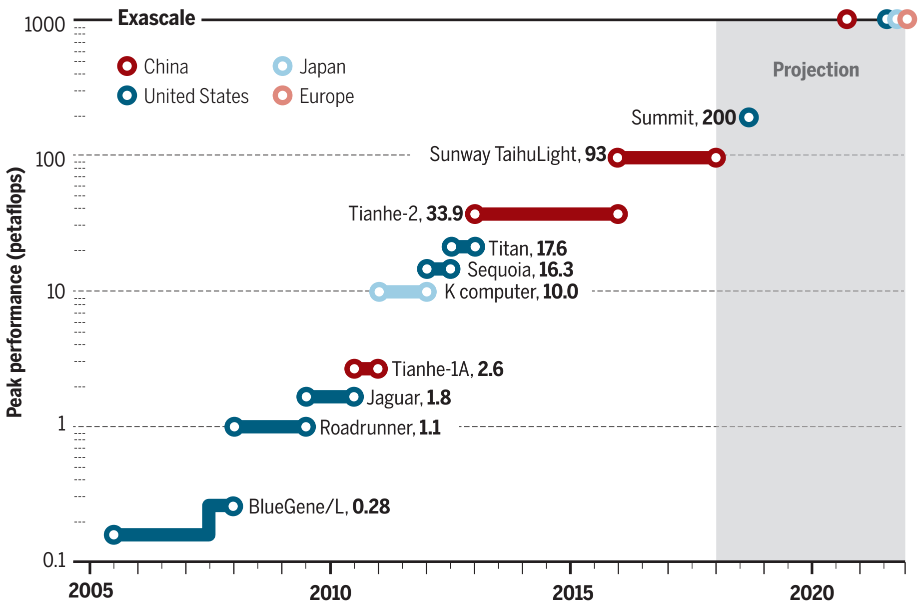

One outside challenge could be money. Congress has yet to pass the fiscal 2018 budget, and instead has funded the government through a series of continuing resolutions that keep funding levels the same as the prior year while forbidding the launch of new projects, such as building the A21 machine. For now, that's not a problem, because DOE is still able to support the underlying scientific developments as part of its existing Exascale Computing Project, says Jack Dongarra, a supercomputing expert at the University of Tennessee in Knoxville. But soon it will be time to start fabricating chips for A21, which is expected to cost between $300 million and $600 million, according to market research firm Hyperion Research. "In 2021 will the budget be there to do this?" asks Horst Simon, a supercomputing expert and deputy director of the Lawrence Berkeley National Laboratory in Berkeley, California. "I don't know." ■

China's planned exascale computer threatens Summit's position at the top

This summer, when engineers flip the switch on Summit, a supercomputer being assembled at Oak Ridge National Laboratory in Tennessee, the machine is expected to be the most powerful in the world. That would return the United States to the top of the supercomputing rankings for the first time since June 2013, when it lost the top spot to Tianhe-2, a machine housed at China's National Supercomputer Center in Guangzhou.

"Of course, we hoped we could have [the top machine] for a longer time," says Depei Qian, a computer scientist at Beihang University in Beijing. "But the government agencies understand that no country can be No. 1 forever."

In the global game of supercomputing leapfrog, China is likely to take back the title from the United States when it builds the first exascale computer: a machine capable of 1 billion billion floating-point operations per second, or 1 exaflop. Under the country's 13th 5-Year Plan, released in 2015, China is committed to launching its first exascale supercomputer by the end of 2020. That could put it a full year ahead of A21, the first U.S. exascale supercomputer, planned for launch in 2021 (see main story, p. 617). Japan is also aiming for an exascale machine with a successor to its K supercomputer at the RIKEN Advanced Institute for Computational Science in Kobe, though the head of the project has said that its delivery could slip to 2021 or 2022. Finally, the European Union could cross the exascale threshold in 2021, according to market research firm Hyperion Research.

But that won't be the end of the race. The four supercomputing powers are convinced they need to push the frontiers in order to compete in a wide range of scientific disciplines, defense technology, industrial technology, and computer products. "Everybody is moving as fast as they can," says Jack Dongarra, a supercomputing expert at the University of Tennessee in Knoxville who keeps close tabs on international supercomputing efforts. And once they cross one threshold, he says, "then it's on to the next one." *—Robert F. Service*

With additional reporting by Dennis Normile.

CREDITS: (GRAPHIC) K. SUTLIFF/*SCIENCE*; (DATA) TOP500

Science

Design for U.S. exascale computer takes shape

Robert F. Service

Science **359** (6376), 617-618.
DOI: 10.1126/science.359.6376.617