

CHINA INCHES TOWARD THE EXASCALE

The powerful Sunway TaihuLight supercomputer makes some telling trade-offs in pursuit of efficiency

➤ **In June, the ranks of the Top500** list were rearranged, and the title of world's most powerful supercomputer was handed off to a new machine—China's Sunway TaihuLight.

The Wuxi-based machine can perform the Linpack Benchmark—a long-standing arbiter of supercomputer prowess—at a rate of 93 petaflops, or 93 quadrillion floating-point operations per second. This performance is more than twice that of the previous record holder, China's Tianhe-2. What's more, TaihuLight achieves this capacity while consuming 2.4 megawatts less power than Tianhe-2.

Such efficiency gains are important if supercomputer designers hope to reach exascale operation, somewhere in the realm of 1,000 Pflops. Computers with that capability could be a boon for advanced manufacturing and national security, among many other applications. China, Europe, Japan, and the United States are all pushing toward the exascale range. Some countries are reportedly setting their sights on doing so by 2020; the United States is targeting the early 2020s. But two questions loom over those efforts: How capable will those computers be? And can we make them energy efficient enough to be economical?

We can get to the exascale now “if you're willing to pay the power bill,” says Peter Kogge, a professor at the University of Notre Dame. Scaling up a supercomputer with today's technology to create one that is 10 times as big would demand at least 10 times as much power, Kogge explains. And the difference between 20 MW and 200 MW, he says, “is the difference [between having] a substation or a nuclear power plant next to you.”


Kogge, who led a 2008 study on reaching the exascale, is updating power projections to cover the three categories of supercomputers built today: those with “heavyweight” high-performance CPUs; those that use “lightweight” microprocessors that are slower but cooler, and so can be packed more densely; and those that take advantage of graphics processing units to accelerate computation.

TaihuLight follows the lightweight approach, and it has made some sacrifices in pursuit of energy efficiency. Based on its hardware specs, TaihuLight can, in the-



AN EFFICIENT ENGINE: Sunway TaihuLight's computations-per-watt improvement is even more impressive than its raw computing power.

ory, crunch numbers at a rate of 125 Pflops. The machine reaches 74 percent of this peak theoretical capacity when running Linpack. But it does not fare as well on a new alternative benchmark, High Performance Conjugate Gradients (HPCG), which is designed to reflect how well a computer can perform more memory- and communications-intensive, real-world applications. When it runs HPCG, TaihuLight utilizes just 0.3 percent of its theoretical peak abilities, which means that only 3 out of every 1,000 possible floating-point operations are actually used



SAVE TIME. DESIGN MORE.

MagNet v7 2D/3D
electromagnetic
field simulation
software. For designing:

- ✓ Transformers
- ✓ Actuators
- ✓ Sensors/NDT
- ✓ Magnetic Levitation
- ✓ Electric Machines
- ✓ More

MotorSolve v5 Electric
Machine Design Software

- ✓ All types of Brushless DC
- ✓ Induction Machine
- ✓ Switched Reluctance
- ✓ Brushed
- ✓ Thermal/Cooling Considerations



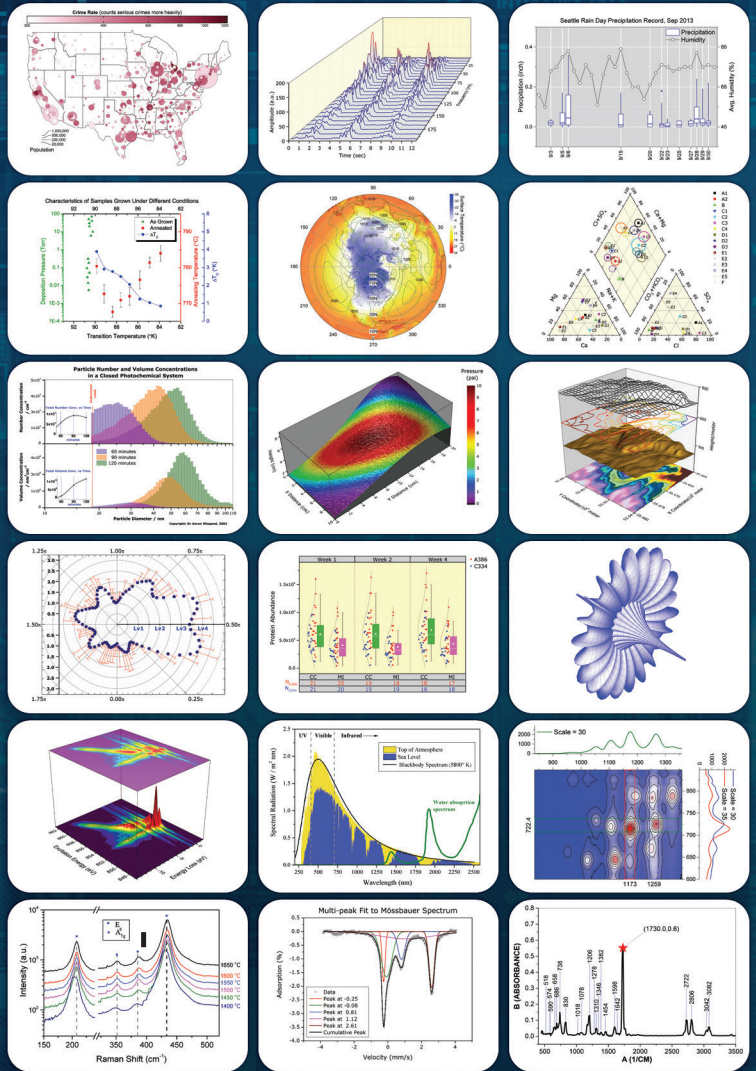
infolytica
corporation

866 416-4400
Infolytica.com

NEW VERSION!

ORIGIN[®] 2016

Graphing & Analysis



by the computer. By comparison, Tianhe-2 and the United States' Titan, the second- and third-fastest supercomputers in the Top500 rankings, respectively, can take advantage of just over 1 percent of their computing capacity. Japan's K computer, currently ranked fifth on the list, achieved 4.9 percent with the HPCG metric.

"Everything is a balancing act," says Jack Dongarra, a professor at the University of Tennessee, Knoxville, and one of the organizers of the Top500. "They produced a processor that can deliver high-arithmetic performance but is very weak in terms of data movement." But he notes that the TaihuLight team has developed applications that take advantage of the architecture; he says that three projects that were finalists for this year's ACM Gordon Bell Prize, a prestigious super-computing award, were designed to run on the machine.

TaihuLight uses DDR3, an older, slower memory, to save on power. Its architecture also uses small amounts of local memory near each core instead of a more traditional memory hierarchy, explains John Goodacre, a professor of computer architectures at the University of Manchester, in England. He says that while today's applications can execute between 1 and 10 floating-point operations for every byte of main memory accessed, that ratio needs to be far higher for applications to run efficiently on TaihuLight. The design cuts down on a big expense in a supercomputer's power budget: the amount of energy consumed shuttling data back and forth.

"I think what they've done is build a machine that changes some of the design rules that people have assumed are part of the requirements" for moving toward the exascale, Goodacre says. Further progress will depend, as the TaihuLight team has shown, on end-to-end design, he says. That includes looking not only at changes to hardware—a number of experts point to 3D stacking of logic and memory—but also to the fundamental programming paradigms we use to take advantage of the machines. —RACHEL COURTLAND

Over 100 new features & improvements in Origin 2016!

FOR A FREE 60-DAY EVALUATION, GO TO ORIGINLAB.COM/DEMO AND ENTER CODE: 8547



Over 500,000 registered users worldwide in:

- 6,000+ Companies including 120+ Fortune Global 500
- 6,500+ Colleges & Universities
- 3,000+ Government Agencies & Research Labs



20+ years serving the scientific & engineering community