From *SIAM News*, Volume 43, Number 10, December 2010

Chinese Supercomputer Emerges at Head of Latest Top500 List

From time to time, announcement of the top 500 supercomputers in the world, a highlight of the annual Supercomputing conference, contains an element of drama. Release of the latest list at SC 2010 (New Orleans, November 13–19) is such an instance.

Of the top five computers on the list, two are from China, two are from the U.S., and one is from Japan. Number 1, as predicted with near certainty two weeks in advance of the official announcement in an article in *The New York Times*, is the Tianhe-1A, located at the National Supercomputing Center in Tianjin, China. With sustained performance of 2.57 petaflop/s (46% faster than the Jaguar at Oak Ridge National Laboratory, the previous number 1 system), the Tianhe-1A has captured the attention of those who use, analyze, create, and fund supercomputers. Making the prediction was Jack Dongarra, director of the Innovative Computing Laboratory at the University of Tennessee. The Tianhe "blows away the existing No. 1 machine," he told the *Times* reporter. "It is unlikely that we will see a system that is faster soon."

At the beginning of November, with the official release of the list still two weeks away, Dongarra, who also has an appointment at Oak Ridge National Lab and the University of Manchester and has been the driving force behind the twice yearly Top500 list since its 1993 founding, spoke to *SIAM News* by phone about the Chinese supercomputer and its implications. As this issue of *SIAM News* goes to press, the predictions mentioned here have been verified.

■■■

SIAM News: The emergence of a Chinese computer at the top of the list reminds us of the appearance of the Japanese Earth Simulator some years back. Are these developments in fact similar?

Dongarra: The Earth Simulator, which startled the world by topping the list in 2002, was developed in very different circumstances. An astounding 3.4 times faster than the previous top machine, the Earth Simulator was the result of a huge investment. The Japanese "pumped up the level of computing, and deployed a massive system before its time." But it was a single machine, not intended to become part of a defined program in high-performance computing.

As some readers will know, the Earth Simulator was isolated in another way as well: Although available to scientists anywhere in the world, it was not connected to the Internet; users had to travel to the Earth Simulator Center in Yokohama.

The Tianhe-1A, faster than its predecessor but less startlingly so, is to be part of a series of computers capable of petaflop/s performance. It will be put on the Internet; the idea is that anyone will be able to use it (although no mechanism for that seems to be in place at the moment).

One thing the 2002 and 2010 Top500 announcements have in common: the "leapfrog effect," in which one country will de-ploy a large system and become the leader in high-performance computing.

SIAM News: What can you tell us about the architecture of the Tianhe-1A?

Dongarra: The new computer is based on a hybrid architecture: It contains both commodity CPUs, from Intel, and GPUs, from NVIDIA, which is a key to higher performance and lower power consumption today. But what truly sets this machine apart is the interconnect.

The Tianhe-1A is a mix of foreign and home-grown technology. Along with the Intel and NVIDIA processors, which are from the U.S., a small number of multicore processors, developed in China, are used as the front end for the Tianhe-1A system. And perhaps more interesting is the interconnect; the switch and router were developed by Chinese researchers from the National University of Defense Technology. The switch is based on a fat-tree topology, and is twice as fast as the commodity-based QDR InfiniBand used in many supercomputers. (The Nebulae computer system, in Shenzhen, China—which is number 3 on the new Top500 list—also uses Intel and NVIDIA components but has a commodity (InfiniBand) interconnect.)

SIAM News: How is the Tianhe-1A being used in scientific computing?

Dongarra: At the moment the system is just coming up. The list of intended applications for the computer include petroleum reservoir simulation, aerospace and biomedical research, animation design, development of new materials, and weather/climate prediction. In addition, the Chinese are interested in the prediction of sandstorms, radar cross sectioning, cosmology, and CAD/CAE for the automotive industry.

SIAM News: What about software?

Dongarra: In using high-performance machines like these, it's not just a question of the hardware: There's a whole ecosystem of components that make it work.

The Chinese combined open-source and commer-

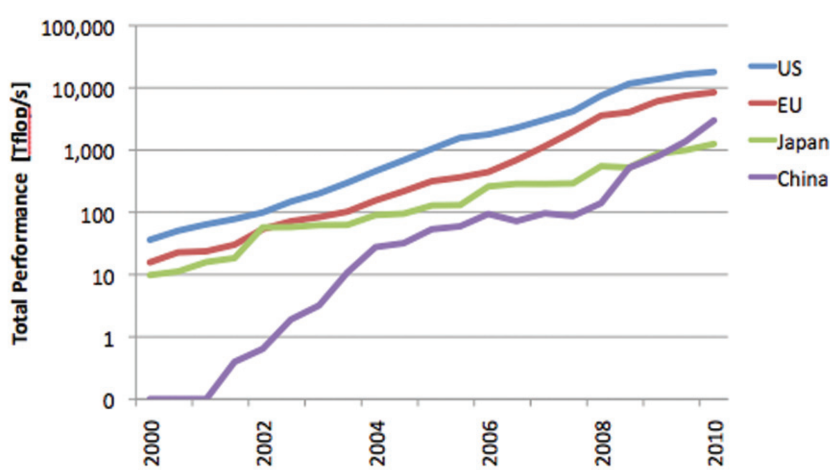

Top500 performance by region over time.

cially available software with components they have developed. Their software stack contains familiar open-source items, including LAPACK, ScaLAPACK, PETSc, AZTEC, TAO, FFTW, and PARPACK. They are also acquiring commercial software, such as TotalView, LSF, Gaussian, Ansys LS-Dyna, and Fluent. They are developing their own software for adaptive mesh refinement, fast multipole methods, and symmetric eigenvalue solvers.

The operating system of the Tianhe is a variant of Linux, called Kylin. Kylin was produced by the University of Science and Technology for National Defense. It is based on standard Linux, but was designed in China and as such is under their control and tamper-proof, with no "back doors." The OS, according to its website, has already achieved one of the highest national data-security standards, and is therefore to be used in critical military and government servers.

I should add that the Chinese provided all requested information about the Tianhe architecture and software. "My interactions with the Chinese have always been very open."

SIAM News: Returning to the expected leapfrogging, how do you see U.S. policy changing?

Dongarra: The long-term implication for the U.S. is that China is seriously interested in high-performance computing, and is developing and deploying computing resources. The accompanying graph, which is based on Top500 data over time, provides some perspective. The rapid rise of China is clearly visible; the Chinese recently overtook Japan, and there's some feeling that they may catch up with the EU countries within six months. If the U.S. is to remain competitive in high-performance computing, this should not go unnoticed.

This isn't a race, however—it's more a sign that China is serious in pursuing high-performance computing to aid the growth of science, engineering, and economic competitiveness. For a growing number of problems in which experiments are impossible, dangerous, or inordinately costly, extreme-scale computing will enable the solution of vastly more accurate predictive models and the analysis of massive quantities of data, leading to quantum advances in areas of science and technology. Economic competitiveness is also significantly enhanced as companies utilize resources to accelerate the development of superior new products and spur creativity arising from modeling and simulation at unprecedented speed and fidelity.

The U.S. Department of Energy is now in the midst of an effort to develop and deploy exaflop/s (10^{18} floating-point operations per second) systems; the response will certainly include a continued push toward that level. There are signs that we may be able to reach exascale computing between 2018 and 2020. But it's important to remember that it is not just about the hardware. The whole ecosystem, including the system software, mathematical theory, numerical algorithms, and scientific software, has to be supported to achieve that goal.

In a blog posted after the announcement, Steve Koonin (DOE undersecretary for science) reiterated the importance of advanced computing to the U.S. Interested readers can find his comments at http://blog.energy.gov/blog/2010/10/29/simulation-and-high-performance-computing.