

China's supercomputer surprises U.S. experts

John Markoff

Reproduced from
THE HINDU, October 31, 2011



Fast forward: A journalist shoots video footage of the data storage system of the Sunway Bluelight supercomputer at the newly established National Supercomputing Center in Jinan on Saturday. — PHOTO: Xinhua

China has made its first supercomputer based on Chinese microprocessor chips, an advance that surprised high-performance U.S. computing specialists.

The announcement was made this week at a technical meeting held in Jinan, China, organised by industry and government organisations. The new machine, the Sunway BlueLight MPP, was installed in September at the National Supercomputer Center in Jinan, the capital of Shandong province in eastern China.

The Sunway system, which can perform about 1,000 trillion calculations per second - a petaflop - will probably rank among the 20 fastest computers in the world. More significantly, it is composed of 8,700 ShenWei SW1600 microprocessors, designed at a Chinese computer institute and manufactured in Shanghai.

Currently, the Chinese are about three generations behind the state-of-art chip making technologies used by world leaders such as the United States, South Korea, Japan and Taiwan.

“This is a bit of a surprise,” said Jack Dongarra, a computer scientist at the University of Tennessee and a leader of the Top500 project, a list of the world's fastest computers.

Last fall, another Chinese based supercomputer, the Tianhe-1A, created an international sensation when it was briefly ranked as the world's fastest, before it was displaced in the spring by a rival Japanese machine, the K Computer, designed by Fujitsu.

Internal system

But the Tianhe was built from processor chips made by U.S. companies, Intel and Nvidia, though its internal switching system was designed by Chinese computer engineers. Similarly, the K computer was based on Sparc chips, originally designed at Sun Microsystems in Silicon Valley.

Mr. Dongarra said the Sunway's theoretical peak performance was about 74 per cent as fast as the fastest U.S. computer the Jaguar supercomputer at the Department of Energy facility at Oak Ridge National Laboratory, made by Cray Inc. That machine is currently the third fastest on the list.

The Energy Department is planning three supercomputers that would run at 10 to 20 petaflops. And the United States is embarking on an effort to reach an exaflop, or 1 million trillion mathematical operations in a second, sometime before the end of the decade, although most computer scientists say the necessary technologies do not yet exist.

Design principles

To build such a computer from existing components would require immense amounts of electricity roughly the amount produced by a medium-sized nuclear power plant.

In contrast, Mr. Dongarra said it was intriguing that the power requirements of the new Chinese supercomputer were relatively modest about 1 megawatt, according to reports from the technical conference. The Tianhe supercomputer consumes about 4 megawatts and the Jaguar about 7.

The ShenWei microprocessor appears to be based on some of the same design principles that are favoured by Intel's most advanced microprocessors, according to several supercomputer experts in the United States.

But there is disagreement over whether the machine's cooling technology is appropriate for designs that will be required by the exaflop-class supercomputers of the future.

Photos of the new Sunway supercomputer reveal an elaborate water-cooling system that may be a significant advance in the design of the very fastest machines.

“Getting this cooling technology correct is very, very difficult,” said Steven Wallach, chief scientist at Convey Computer, a supercomputer firm based in Richardson, Texas.

“This tells me that this is a serious design. This cooling technology could scale to exaflop. They are in the hunt to win.”

— **New York Times News Service (28th Oct)**

Some technical terms

- Supercomputer
- Petaflop, exaflop
- Top500
- Microprocessor
- Intel, Nvidia, Sparc processors
- Switching system
- Power supply, cooling system

Supercomputer [Wikipedia]

- A computer at the frontline of current processing capacity, particularly speed of calculation.
- Used for highly calculation-intensive tasks such as
 - quantum physics
 - weather forecasting
 - climate research
 - molecular modeling
 - simulation of large/complex systems

Blue Gene



Cray



Param and Eka



Petaflop, exaflop

- FLOP : Floating point operations per second
 - a performance measure for executing numerical computation dominated applications
- Mega : 10^6
- Giga : 10^9
- Tera : 10^{12}
- Peta : 10^{15}
- Exa : 10^{18}

Top500

<http://www.top500.org>

- The main objective : to provide a ranked list of general purpose systems that are in common use for high end applications
- Computers are listed ranked by their performance on the LINPACK Benchmark
- Systems designed specifically to solve the LINPACK benchmark problem or aiming primarily at a high Top500 ranking are disqualified



Microprocessor

- The term originated with Intel's 4004 chip produced in early 1970's containing a full 4 bit CPU
- The technology progressed to 8 bit, 16 bit, 32 bit and 64 bit CPUs on a chip
- Currently chips with multiple CPUs (multicore) are available


Intel, Nvidia and Sparc processors

- Intel: Most successful processor chip manufacturer for general purpose computing
- Nvidia: Leading producer of graphics processors (GPUs). Originally meant for accelerating graphics performance, now being used for high performance computing
- Sparc (Scalable processor architecture): Prominent RISC class architecture, available now as open source

Switching system



Crossbar



Network

Power supply and cooling system

- 100 W per processor x 10,000 processors
= 1 MW
- Higher performance means higher power consumption
- Requires a massive power supply
- Electrical energy gets dissipated as heat
- Without forced cooling, temperatures could reach a level high enough for combustion

Questions to explore

- Explore Top500 website to study supercomputer performance growth over the last couple of years
- Find out the factors which have lead to this
- What is LINPACK benchmark
- What developments have taken place in India
- Find details of ShenWei SW1600 Microprocessor