

Learn the basics of supercomputers, complex and powerful machines

What is a supercomputer, how does it structure and what do people use it for?

In recent times you must have heard a lot about the term supercomputer, or supercomputer. This is a very different kind of computer than the desktops, laptops that you use every day. It is bigger in size, much more powerful than personal computers. That is why supercomputers can never be used to compose text, play pikachu games, but people apply it to scientific research, processing and complex calculations. So what is a supercomputer, how to structure it and what do people use it for?




A small part of IBM Blue Gene / P supercomputer.

History of supercomputers

The term supercomputer first appeared in a New York World article in 1929 referring to the giant spreadsheets that IBM set up at Columbia University. In the 1960s, the supercomputer architecture was designed by an American engineer named Seymour Cray working for Control Data Corporation (CDC) and the CDC 6600 launched in 1964 was considered the world's first supercomputer. Seymour Cray is honored as the father of supercomputers.

After a while, Cray wanted to leave CDC with some of his colleagues but was turned down by the company's CEO William Norris because Cray was working on a project for the US Navy. It was not until 1972 that Cray realized this intention and he founded his own company called Cray Research. Cray Research later changed its name to Cray Inc. and this is one of the leading supercomputer vendors in the world at the present time.

At its own company, in 1976, Cray announced the Cray-1 supercomputer with 80MHz CPU and it is one of the most known supercomputers. By 1985, Cray-2 debuted with eight processors. The system is cooled by liquid and flourinert made by 3M. This machine had a computing speed of 1.9 gigaflops and it was the fastest supercomputer in the world until the 1990s.



Cray-1 supercomputer with cylindrical design, ensure the distance between processor and other components in the system is equal.

If the supercomputers of 1980 only used a few CPUs, by 1990, super computers were equipped with thousands of processors and they began to appear in the US and Japan. From here on, the computing speed of this type of computer began to skyrocket quickly. For example, Fujitsu's wind tunnel supercomputer used 166 vector processors * at 1.7 gigaflops per chip in 1994. Two years later, it was Hitachi SR2201 with 600 gigaflops. Thanks to 2048 processors, it has surprised the whole world. Intel also has its Paragon supercomputer with about 1,000 to 4,000 i860 chips, and it was the fastest title in the world in 1993.

* **Vector microprocessor** is a CPU integrated with special instructions to be able to operate in one-dimensional data arrays. CPU vectors can significantly increase performance in certain tasks, especially in the field of numerical simulations or performing the same tasks. CPU vector was popular in the 1970s and 1980s but today it has almost disappeared completely.

As of today, supercomputers are gradually appearing more and more and they are also available in more countries. The operating system, architecture and hardware used in supercomputers have also changed a lot compared to a few decades ago.



Some cabinets belong to Titan supercomputers.

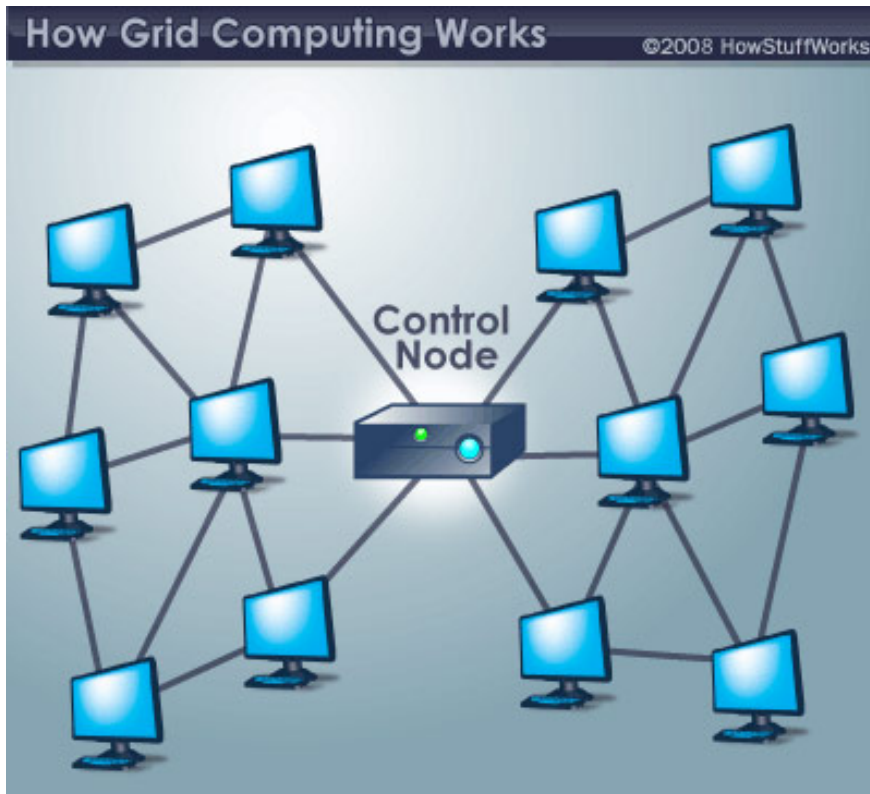
Supercomputer architecture and hardware

A. Architecture


The way people design supercomputers has changed a lot recently. Seymour Cray's early supercomputers worked on parallel computing and compact design to achieve high computing performance.

Talking more about parallel computing (parallel computing), this is the term used to refer to the use of a large number of CPUs to perform a certain set of calculations. All calculations will be performed in parallel. There are two ways that people often apply:

Using a network of multiple computers scattered in many places to process data (grid computing), and dispersion here means that their geographical distance is relatively far apart. Usually when a computer on the network is powered up, it will immediately become part of a parallel computing system, and the more computers that participate, the faster the processing speed will be. There is a main machine (Control Node) in the middle of the task of controlling and allocating tasks to the workstations. There are two small ways of this type, we will learn more below.




Using a large number of CPUs placed close together, and this is called a computer cluster and this is a centralized type of computing. These CPUs are usually located in many similar, neighboring computers (called nodes, card nodes or computer nodes) and they are connected to create a larger, more complete system. People see the system as a single supercomputer. With this measure, designers need to ensure that the speed and flexibility of interconnection between computers is sufficient to meet the work requirements. According to data from TOP500, the number of cluster supercomputers now accounts for 82.2% of the global market share of supercomputers. IBM Blue Gene / Q supercomputer uses cluster form.



From the microprocessor, supercomputer manufacturers / suppliers will integrate it onto a "computer card" with RAM and may have more components to create a near-complete computer.

Many computer cards are attached to a tray, then the trays continue to be put in the rack. Many racks will form a complete supercomputer system.


There is also a Massively Parallel Processors (MPP) supercomputer, which is a huge computer but has thousands of CPUs and RAM modules in it. They are connected to each other by a special high-speed network standard rather than using common things like clusters. In addition, each CPU will have its own memory and a separate copy of operating system / application. MPP currently accounts for 17.8% of the supercomputer market share, according to TOP500. The IBM Blue Gene / L supercomputer (ranked 5th in the world in 2009) is designed in MPP format.



B. Hardware

As mentioned above, in the 1970s, supercomputers used only a small number of processors. But by the 1990s, the number of CPUs had reached thousands and at the present time, the number of tens of thousands of CPUs is very normal for a supercomputer.

In cluster supercomputers, people often combine many small "nodes" (nodes) together to create a large system like the image you saw right above. Each of these buttons can be considered as a nearly complete computer with one or more CPUs, GPUs, RAM modules, cooling fans and a number of other components. The nodes will be connected to each other in many ways, whether using normal copper cables, or switching to fiber optic cables to ensure better bandwidth. The power of a supercomputer will be the combined power of all the nodes.



Connect the wires in the Titan supercomputer.

In supercomputing systems, people also need storage drives, and these HDDs and SSDs are not in device attached storage (DAS) as on PC. Instead, they are usually arranged in a private closet (storage area network (SAN), have a private network connection and the capacity is also "terrible".

Now, in addition to simple CPUs, people also use more GPGPU (general purpose graphic processor unit) to power supercomputers. So far we know the GPU is used to build images, process things related to graphics, but in addition they can process data and do some of the same work as the CPU. Nowadays, the price of GPGPUs has decreased, the efficiency has increased so more and more supercomputers "rely" on this component to increase power but still ensure the cost is not too much. The world's most powerful supercomputer (as of May 23, 2013), named Titan in the US also uses 18,688 NVIDIA Tesla K20 GPUs in addition to 18,688 AMD Opteron 16-core CPUs, so the total CPU core of the machine is 299,008, one a huge number compared to computers with only 2 or 4 cores.




A node of the Titan supercomputer, including 4 AMD Opteron 16 cores and 4 GPUs.

NVIDIA Tesla K20 (construction architecture is Kepler). Each node has 32GB of RAM, plus 6GB of GDDR5 memory for the graphics card. A total of 710 Terabytes of memory in Titan.


In fact, the use of GPGPU to enhance the CPU is still a problem that is being discussed by scientists. Many people believe that the addition of GPGPU does help supercomputers become stronger and score higher


benchmarks, but when put into practice, it takes a lot of effort to adjust the software to have The power of this combined architecture can be squeezed out.



The structure of Cray used in Titan is a combination of both Cray XK6 and XE6, in which XK6 has two NVIDIA cards that come with RAM besides AMD Opteron 64 CPU.

Europe is also developing a supercomputer project using the Tegra 3 v2 GPU GeForce. Specifically, this machine uses 256 NVIDIA Tegra 3 quad-core processors, 256 GeForce 520MX graphics processors and 1TB of DDR3 RAM to operate. With a configuration like this, the machine can calculate 38 teraflops per second but only consume 1W of power for every 7.5 gigaflops. NVIDIA calls this the first supercomputer to use ARM's CPU in combination with its GPU. Mont-Blanc's research and cooperation project to build supercomputers in Europe promises that this product will operate four to ten times higher than other supercomputers.





Heat dissipation, energy consumption

Over the decades, one of the main problems with centralized supercomputers was power consumption and heat dissipation. Each supercomputer eats a huge amount of energy, and most of it is converted into heat. When the system overheats can dramatically reduce performance, component life is significantly shortened as well. Scientists have studied many options to solve this problem, such as pumping Fluorinert coolant, liquid cooling, air cooling, by liquid-to-air hybrid system. Especially, this hybrid system is commonly used in supercomputers made up of many cabinets. In addition, some supercomputers use low-power processors to cope with high temperatures such as IBM's Blue Gene. The IBM Aquasar is even more special when it uses hot water to dissipate heat, and that water is also used to warm the whole building during the cold season.




Titanium water cooling system.

In addition to the amount of electricity used by supercomputer, we can take for example the Chinese Tianhe-1A with a capacity of 4.04 megawatts, that is, the governing agency has to spend 400 USD / hour on electricity. and for the whole year is 3.5 million USD. Therefore, the issue of cost when shipping supercomputers is also a complicated problem.

To measure the energy efficiency of a supercomputer, people use the number of FLOPS per Watt and the higher the index the better, that is, for every Watt of electricity, the supercomputer calculates more calculations. In 2008, IBM's Roadrunner achieved 376 MFLOPS / W. In 2010, the Blue Gene / Q reached 1684 MFLOPS / W. In 2011, the Blue Gene in New York reached 2097 MFLOPS / W. And what is FLOPS, I will talk about below.

Operating system

Since the beginning of the 20th century, the operating system used in supercomputers has also recorded many changes, similar to the way supercomputer architecture changed. The first operating systems were customized for each supercomputer to increase speed, but it was very time consuming, labor and money. Therefore, in the current trend, people will use common operating systems such as Linux for supercomputers, not use UNIX as in the 90s. You can look at the chart right below to see this trend right away.



Of the 500 fastest computers in the world, Linux is the largest operating system and the adoption of this open source OS is increasingly popular. At the time of writing, Linux is accounting for 93.8% of the supercomputer market share. UNIX is also the operating system used on many supercomputers but is not as popular as Linux and it holds 4% market share. Windows and BSD are also present, but they are insignificant because they are not as reliable as Linux and UNIX, and are also affected by licensing costs. Apple's OS X was previously used for supercomputers but mainly in distributed supercomputers.

Software tools

Because of the parallel computing architecture of a supercomputer, people often have to apply special programming techniques to get the most out of it. The tools used for this are API functions such as MPI, PVM, VTL, in addition to open source software solutions such as ERowulf. In most cases, the PVM and MPI environments will be used for clustered systems, while OpenMP is for systems with shared memory. The algorithms also need to be greatly optimized because supercomputers not only run on one but on a lot of CPUs, GPUs, not to mention each supercomputer cabinet is separate from each other. Scientists also have to consider minimizing the amount of CPU free time they can wait for data from another node to transfer. With GPGPU supercomputers, more CUDA models of NVIDIA are available to enhance machine performance.

Distributed supercomputer

We've talked a lot about centralized supercomputer for a while now, let's turn to distributed supercomputer a bit. In distributed form, there are two ways to take advantage of many small computers, the opportunistic method and the quasi-opportunistic method.

A. Method of opportunity

This is a form of grid computing. In this method, a large number of individual computers will voluntarily work together to form a large network to perform computing tasks equally equally. Grid computing has solved many parallel computing problems, but it also has the drawback of not counting some classic tasks like flow simulation.

The fastest supercomputer network as of March 2012 is Folding @ Home (developed by Stanford University). It has power of 8.1 petaflops and uses x86 processors. Among them, 5.8 petaflops are "devoted" from computers with different types of GPUs, 1.7 petaflops are from PlayStation 3 game consoles, the rest are contributed by multiple CPUs. In 2011, the world had a BOINC grid with a power of 5.5 petaflops contributed from 480,000 computers. Last year, Simon Cox built a supercomputer with 64 Raspberry Pi's. The system is called Iridis-Pi and costs only £ 2,500 and has 1TB of memory (16GB SD card for each Raspberry Pi).

B. Opportunity approach

Quasi-opportunistic computing is similar to opportunistic computing, but it has a higher quality of service by increasing control over the tasks that every single machine will do. It also tightly controls the use of distributed resources. Also included is an intelligent system that helps ensure the presence and stability of the member machines. In order for the opportunistic approach to work, computers will need to have a "resource allocation contract" in combination with complex forms of communication and error prevention.

Performance measurement

A. Capacity vs Capacity

Supercomputers are designed to target complex calculations, ie capability computing. This means that supercomputers are used to maximize computing power, thereby solving a problem in the shortest possible time. These calculations are very special and it is so difficult that normal computers cannot handle, such as nuclear explosion simulation, weather forecast, quantum research .

If you want to process a large amount of data, people do not use supercomputer but use a computer called mainframe. The mainframe can handle huge input but the calculations it runs are not as complex as supercomputers. Mainframes can be used to solve many small problems at the same time. We will have a separate mainframe article, see you then to talk more about this type of computer.

B. Measuring supercomputer performance

If on a PC, laptop, tablet, smartphone people will conduct a benchmark to know the power of the machine, the supercomputer is the same. However, the computational power of a supercomputer is measured by FLOPS (FLoating Point Operations Per Second - floating point calculation is done per second), while normal computers are measured in MIPS (instructions per second - number instructions are executed every second). FLOPS can be added with some prefixes in SI measurement systems such as tera- (TFLOPS, ie 10^{12} FLOPS, pronounced teraflops), peta (10^{15} FLOPS).

Currently, the world's leading supercomputers have reached the threshold of Petaflops, for example, the 2008 IBM Roadrunner was 1,105 Petaflops, the 2011 Fujitsu K reached 10.51 Petaflops, while the most powerful Cray Titan is 17.59. Petaflops. It is predicted that only after about 10 years, supercomputers will soon step into Exaflops (10^{18} FLOPS) because the technology of CPU and GPGPU is growing rapidly, the price is cheaper and consumption efficiency is high. Electricity is increasingly being raised.

So where do the FLOPS numbers come from? It was measured by a software called Linpack. However, it should be added that no single number alone can reflect the overall performance of computers in general and supercomputers in particular. There are two numbers expressed when it comes to supercomputers: the processor's theoretical floating point computing performance (Rpeak symbol) and input processing performance (Rmax). Rpeak is almost impossible to achieve in real life, while Rmax is completely attainable when supercomputers run. All the FLOPS numbers you see above are Rmax.

C. TOP500 list


Since 1993, the fastest supercomputers in the world have been listed on a list called TOP500 based on its benchmark scores. TOP500 has their own website at <http://top500.org> and it provides us a lot of useful information. Besides looking at the list of the world's leading supercomputer, you can see statistics about the distribution of operating systems in the world of supercomputers, the number of supercomputers per country, supercomputer architecture (MPP or cluster) . In essence, this list has not been rated as absolutely accurate and unbiased, but it is one of the most common sources that people often take when comparing supercomputer power. at some point.

1. The 10 fastest supercomputers in the world 2019

Supercomputer application

Some applications of supercomputers in today's times:

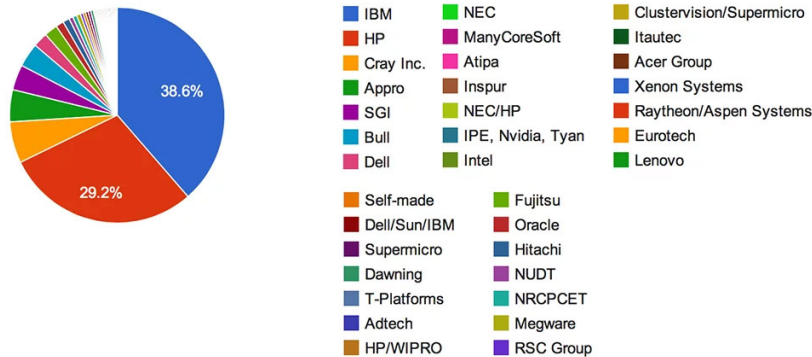
1. Weather forecast, aerodynamic research, climate change study, earthquake simulation
2. Probability analysis, radioactive modeling
3. Simulation of nuclear explosion in 3D
4. Quantum, molecular, cell biology, studying the folding of proteins
5. Simulate the human brain
6. Research and model of physical phenomena
7. Research and simulate artificial intelligence
8. Reproducing the Bigbang explosion (performed by a supercomputer at the Texas Advanced Computing Center), studying dark matter
9. Study astronomy
10. Model the spread of a disease
11. Play chess! (IBM's Deep Blue supercomputer defeated grandmaster Garry Kasparov in 1997)



The simulation image was created by the US Oceanic and Meteorological Agency (NOAA) supercomputer using their supercomputer.

Some manufacturers produce and supply supercomputers

You can see that IBM, HP and Cray are the three leading companies in the field of supercomputers today. Dell, Intel, NEC, Lenovo, Acer, Fujitsu, Oracle also participate in this field.



1. Which country owns the most powerful supercomputers today?

According To Delicate.

You finished reading the article "**Learn the basics of supercomputers, complex and powerful machines**" edited by the [TipsMake](#) team. We hope this article has provided you with many useful tech tips and tricks. You can search for similar articles on tips and guides. Thank you for reading and for following us regularly.