K computer

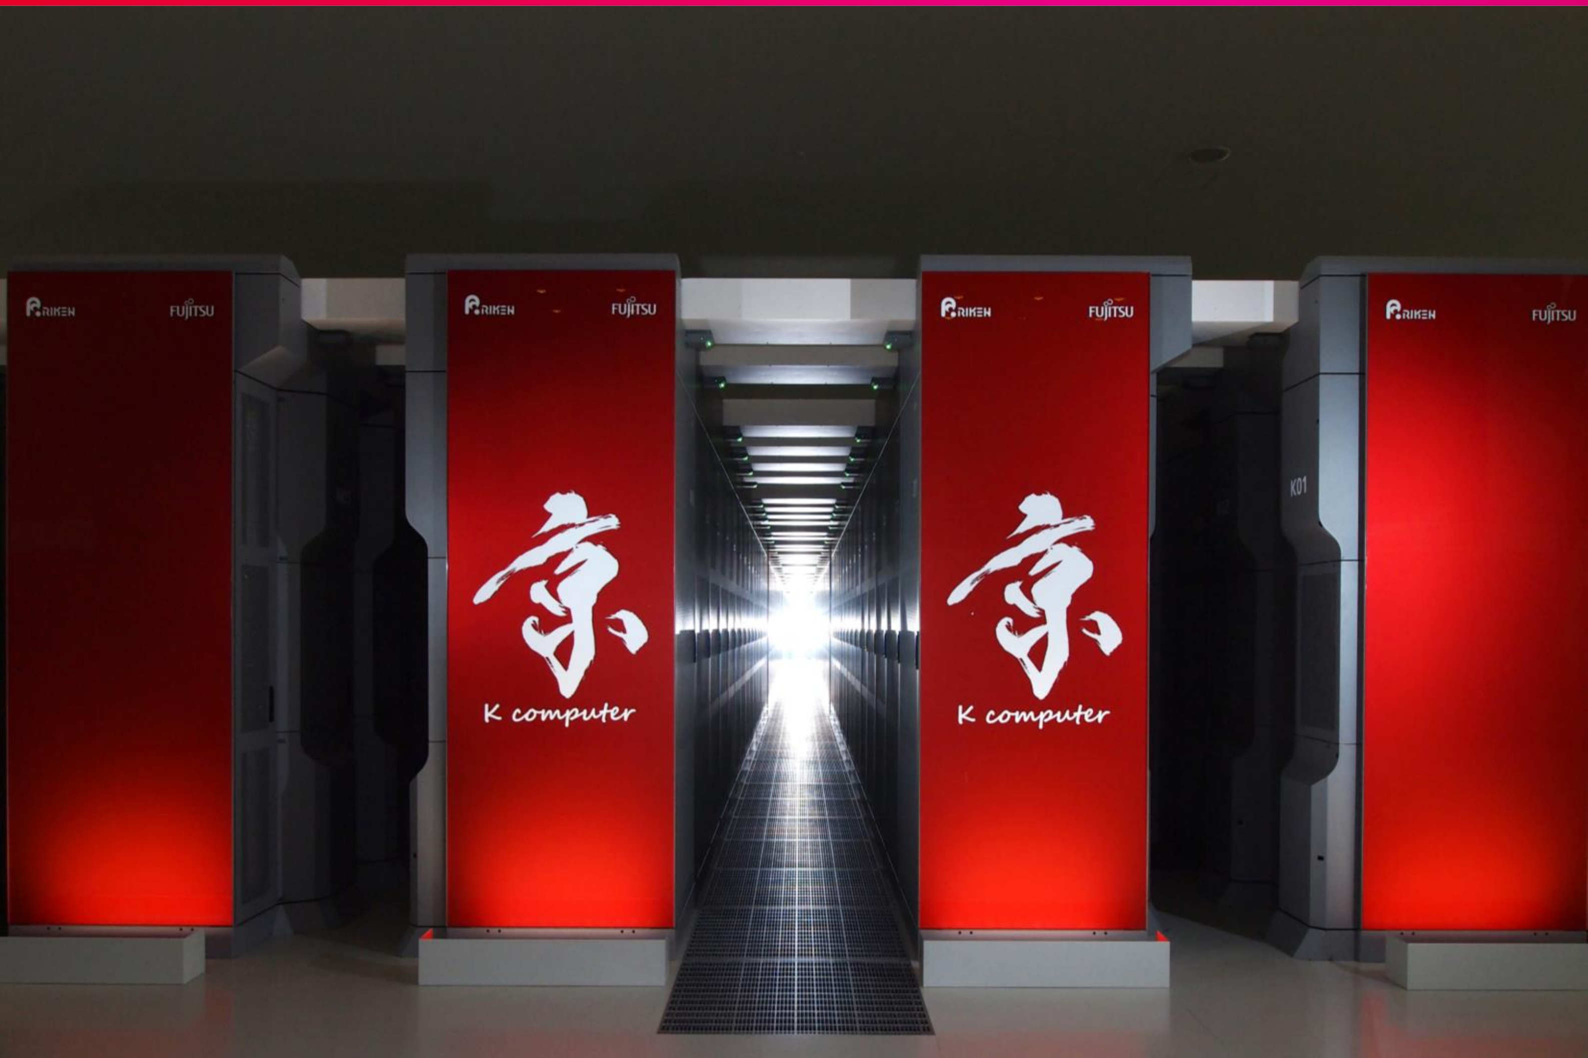

Photo courtesy of RIKEN

Index

This document is a PDF archive of the "K computer" information from the Fujitsu website.

K computer

Operation of the K computer has ended in August, 2019.

Society faces a number of urgent challenges, especially in the areas of climate change and sustainability. To solve these dilemmas, leading research and ICT players need to pool knowledge and accelerate the development of cutting-edge supercomputing power.

Japan's leading research institute RIKEN chose Fujitsu to develop one of the world's most powerful supercomputers, the K computer. As part of the High Performance Computing Infrastructure Initiative led by the Japanese Ministry of Education, Culture, Sports, Science and Technology (MEXT), the K computer will be used to solve the energy, sustainability, healthcare, climate change, industrial and space challenges facing society today.

Technical Computing Related Topics

more Topics »

- June 20, 2017
  K computer takes first place for the second consecutive time on HPCG benchmark

- New insights through computer simulations
- Nurturing next-generation scientists
- Five Strategic Fields

- What is the national development project of the K computer?
- The K computer is incredibly fast
- "SPARC64™ VIIIfx": A Fast, Reliable, Low-Power CPU
- Innovative "6-Dimensional Mesh/Torus" Topology Network Technology

- Why Fujitsu
- Computing the ideal future
- Thoughts from the engineers

Note: About the K computer

The possibilities

FUJITSU Global

New insights through computer simulations

Why do we need computer simulations?

Invaluable in science and industry, computer simulations involve building a virtual model inside a computer and then observing the behavior of that model under various conditions. This makes it possible to reproduce and investigate phenomena which are difficult or impossible to experiment in real life due to the risk, scale or cost involved. Computer simulations are also used to optimize aircraft performance, map the impact of earthquakes and tsunamis, model new drugs and develop advanced materials.

However, the creation of 3D models and digital landscapes requires massive volumes of data. Conventional computers are not capable of handling compute operations on this scale. That is where supercomputers come in. Today's supercomputers deliver the horsepower needed to perform the most sophisticated, compute-intensive calculations and simulations.

A world pioneer in many fields of science and research, Japan has a long-standing tradition in the advancement of supercomputing technologies that deliver next-generation simulation and computational capabilities. In fact, the world's fastest supercomputer was developed by Japanese company Fujitsu.

Driving industrial innovation

A common example of where supercomputers are advancing industrial innovation is automotive design, where computer simulations help engineers to optimize driver and passenger safety in the event of a crash. While crash tests are also beneficial here, they have limitations. It is not financially viable to crash sufficient vehicles to cover all accident variables and there are mechanical limitations to the insights that can be gained from dummies ? as humans obviously cannot be used for crash tests. Highly accurate computer simulations can, however, give designers information on a scale that would otherwise be too hard to obtain, enabling them to build safer vehicles.

Looking beyond safety, computer simulations can also be used to optimize industrial design and environmental performance in general. The latest generation of supercomputers is capable of simulating an entire aircraft or vehicle based on the exact shape, size and properties of each individual component or device. This allows engineers to monitor and capture the precise behavior of those parts and components under a whole range of circumstances and from a variety of angles. Analyzing results otherwise invisible to the human eye allows designers to predict previously unknown phenomena. This reduces development cost and time, while also enabling performance enhancements and improving safety.

Health & Safety: For a better quality of life →

Changing face of industrial innovation →

Unlocking the mysteries of outer space →

- **New insights through computer simulations**
- **Nurturing next-generation scientists**
- **Five Strategic Fields**

Quick guide to supercomputing

Health & Safety: For a better quality of life

Accelerating the development of cancer therapies

Cancer and cardiac disease are major causes of death worldwide and top causes in countries such as Japan. In the search for new cures and therapies to treat cancer, researchers recently discovered that it is possible to capture phenomena not visible to the human eye through computer simulations. If the supercomputers that currently support these computer simulations become much more powerful, researchers expect huge advances in the development of medical treatments for cancer and cardiac disease.

Various medicines are used to treat cancer. But there is still a lot of uncertainty surrounding the basic mechanisms used by drugs to attach to and act on cancer cells. In the development of new medicines, experiments are conducted one by one on vast numbers of candidate chemicals over long periods of time. The efficiency and side effects of these chemicals are then examined. So it takes a very long time and involves massive financial outlay before new medicines can be released to market.

The availability of high-performance supercomputers may dramatically accelerate this discovery process. For example, they allow researchers to observe medicines and cancer cells at atomic and molecular levels. By reproducing this activity moment-by-moment, supercomputers can show, in a form viewable to researchers, how a medicine attaches to cancer cells and how it suppresses the activities of those cells. These simulations have the potential to accelerate the development of effective therapies and reduce the cost involved.

For better protection against natural disasters.

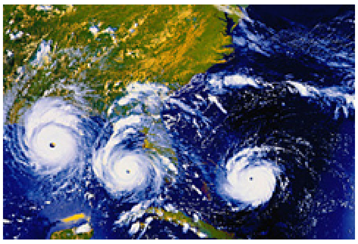

Scientists are hoping that supercomputers will allow them to forecast the weather and predict natural disasters with much greater speed and accuracy, thus enabling more effective disaster-prevention technologies and response plans.
At present, the surface of the Earth is divided into grids ranging in size from a few kilometers to several dozen kilometers. The values for temperature, atmospheric pressure, wind and other variables are calculated for each grid based on parameters measured by monitoring sensors. This data is then used for weather forecasting.
Smaller grids are needed to improve the accuracy of meteorological forecasting. However, the amount of the data processed increases exponentially as grids get smaller. Even with today's supercomputers, it takes a long time to process the vast amount of data gathered.
High-performance supercomputers will allow scientists to divide the surface of the earth into finer grids. This will allow meteorologists, for example, to predict local storms and torrential rains more accurately than ever before.

- New insights through computer simulations
  - Health & Safety: For a better quality of life
  - Changing face of industrial innovation
  - Unlocking the mysteries of outer space

Changing face of industrial innovation

Engineers can now get inside engines!

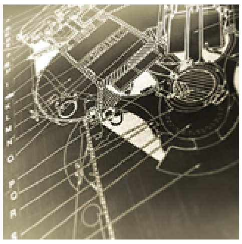

In the past, a skilled engineer typically relied on his or her finely tuned sense of hearing to determine the internal condition of an engine. Now, thanks to advances in supercomputing, they can actually step inside virtual engines and see exactly how complex parts work and interact. Next-generation supercomputers can simulate the precise shapes and properties of components and the behavior of ignited gases inside a combustion chamber. As a result, phenomena invisible to the human eye can now be captured. For example, researchers can monitor pressure and temperature at various locations within an engine and see how they change under different conditions. These insights allow engineers to resolve performance issues and optimize the overall design and energy efficiency of all kinds of aircraft and vehicles.

Simulating an entire aircraft

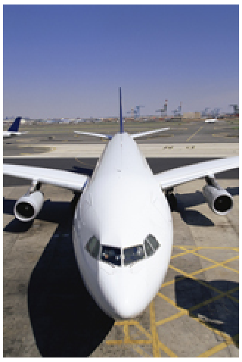

Currently, aircraft designers use various computer simulation methods to analyze the complex air flows which occur during flight. During this process, scale models of parts such as the body and wings of a prototype aircraft are tested in a wind tunnel. Full-scale designs and facilities would be too expensive to build and maintain.
Advances in supercomputing power are paving the way for a fully numeric wind tunnel. In other words, a virtual massive wind tunnel facility powered by high-performance computer simulation capabilities. Data from the full-scale aircraft under development can then be entered in the tunnel to explore performance under various conditions.
The availability of numeric wind tunnels will provide low-cost, rapid design and development capabilities with the added bonus of improved safety. Numeric tunnels would also allow researchers to investigate performance during high-speed flight and during low-speed approach and takeoff. A better understanding of air flows and vortices around a prototype aircraft will translate into more energy-efficient, quieter aircraft designs.

- New insights through computer simulations
  - Health & Safety: For a better quality of life
  - Changing face of industrial innovation
  - Unlocking the mysteries of outer space

FUJITSU Global

Unlocking the mysteries of outer space

Supercomputers as a guide to the world which human cannot reach

Our understanding of outer space is limited. There are a variety of theories on how the distant stars, galaxies, galactic clusters, black holes and supernova were born. However, it is impossible with current technologies for researchers to fly out from the Earth and research the distant stars and galaxies to verify these theories.

Researchers therefore use high-performance equipment such as large telescopes and space probes to observe space. In addition, researchers use computer simulations to validate their theories. However, the scope of these simulations (number of stars, speed) is limited by computational power available. Supercomputers offering higher levels of performance can give researchers more detailed insights into the various phenomena of space.

Today's supercomputers are already exploring various theories on the birth, development and future of galaxies. Ultimately, this will give researchers a better understanding of the Earth's formation and about human evolution into the distant future.

Nurturing next-generation scientists

Innovation hub

The main objective of the K computer development project is to achieve science and technology breakthroughs through the supercomputing capabilities. However, the project also aims to build and foster a world-leading research and development hub.

This project brings together the world's greatest scientific and engineering minds. Researchers are pooling their talent and insights to tackle the grand challenges facing the 21st century ? particularly the need to combat global warming and develop sustainable growth models. The K computer is creating a technical environment that invites tomorrow's talent to share knowledge and wisdom in the search for answers to these challenges.

- **New insights through computer simulations**
- **Nurturing next-generation scientists**
- **Five Strategic Fields**

Quick guide to supercomputing

Five Strategic Fields

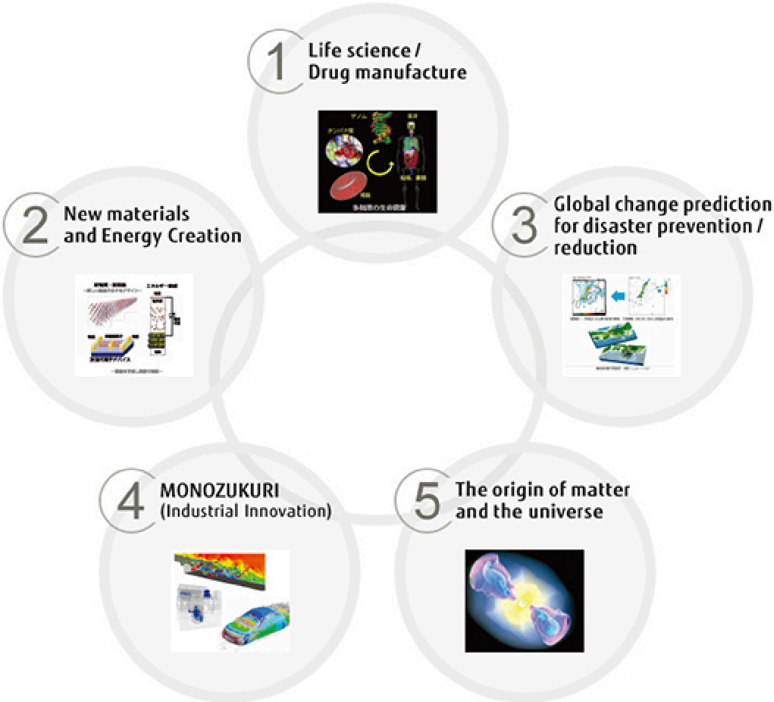

Field 1: Life science / Drug manufacture

Simulations will be performed to understand and predict new life phenomena in life-science computation, to support new drug developments that facilitate greater medical and pharmacological understanding, and to achieve predictable medical treatment. We can contribute to building a foundation for a healthy society by expanding the analysis of large-scale life data based on genome research, analyzing the behavior of biological macromolecules within cells, and conducting dynamic analysis of the layers comprising cells, organs and organisms.
[RIKEN (Collaborating organizations: School of Engineering / e Institute of Medical Science, e University of Tokyo, and others)]

Field 2: New materials and Energy Creation

Industry as we know it today is based on materials such as semiconductors and polymers that emerged from basic research. Thus the following activities will produce a torrent of industrial innovations in the future: ideas by which to search for new quantum phase and new materials, achieve a deep understanding of material functions, apply discoveries to the mainstream of basic science and next-generation electronic devices offering high functionality and performance, and the efficient generation of earth-friendly, renewable energy.
[Institute for Solid State Physics, e University of Tokyo, Institute for Molecular Science, National Institute of Natural Sciences, Institute for Materials Research, Tohoku University]

Field 3: Global change prediction for disaster prevention / reduction

By computing the movement of clouds on a global atmospheric scale, it becomes possible to accurately predict the path of a typhoon or a localized torrential downpour. Consequently, this will help us reduce the damage resulting from a disaster. Moreover, it will be possible to more accurately predict earthquakes and tsunamis, as well as to predict structural damage and the ramifications of complex disaster situations in which these elements are intertwined. This will establish the foundation for the next-generation mapping of earthquake and tsunami activity.
[Japan Agency for Marine-Earth Science and Technology (Collaborating organizations: Atmosphere and Ocean Research Institute / Interfaculty Initiative in Information Studies / Earthquake Research Institute, e University of Tokyo, Meteorological Research Institute, Tohoku University, and others)]

Field 4: MONOZUKURI (Industrial Innovation)

The new technologies may be created from large-scale but highly detailed analysis of complex phenomena. The combination of optimal technologies may be discovered that contributes to the development of new products that deal with critical future topics, such as resource recycling products. It may become possible to evaluate multifaceted characteristics--such as performance,

efficiency, comfort, and reliability--by diagnosing the product as a whole, and thereby contribute to a new, more efficient manufacturing sector.
[Institute of Industrial Science, e University of Tokyo, Japan Atomic Energy Agency, Japan Aerospace Exploration Agency]

Field 5: The origin of matter and the universe

The mission of understanding is, for mankind, inextricably linked to the universe and our existence within it. Therefore, to unlock the origins of various materials or the complex structure of the universe, events from the big bang up to the current day will be recreated through numerical simulations, and the future will be predicted. Using the fundamental equations of physics, elementary particles atomic nuclei, atoms and molecules, as well as stars, interstellar matter and dark matter, are described. This will reveal the true character and mystery of birth, both in the outer sphere and our own life.
[Center for Computational Sciences, University of Tsukuba, High Energy Accelerator Research Organization, National Astronomical Observatory of Japan]

- New insights through computer simulations
- Nurturing next-generation scientists
- Five Strategic Fields

Quick guide to supercomputing

Global

About K computer

Quick guide to supercomputing

FUJITSU

Global

What is the national development project of the K computer?

RIKEN and Fujitsu completed the development of the K computer in June 2012 and it started full-service from September 2012. This is in accordance with the High Performance Computing Infrastructure Initiative promoted by Japan's Ministry of Education, Culture, Sports, Science & Technology (MEXT).

System Development Schedule

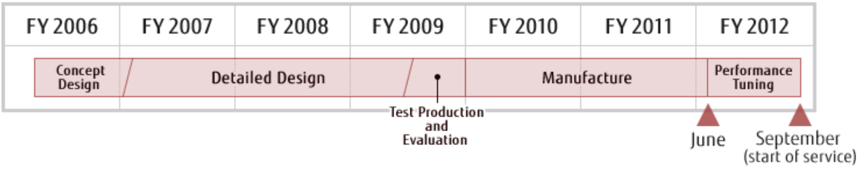

Note on name

The K computer is the nickname RIKEN has been using for the supercomputer of this project since July 2010. "K" comes from the Japanese word "Kei" which means ten peta or 10 to the 16th power. The logo for the K computer based on the Japanese character for Kei, was selected in October 2010.

The K computer facility

The site is located on Port Island, Kobe

The K computer is installed at the RIKEN Advanced Institute for Computational Science (AICS), which is located on Port Island, Kobe, in Hyogo Prefecture. Kobe is a port town in Japan. Port Island is an artificial island built in the Kobe harbor. Constructed in 1981 it is connected to the center of Kobe by the Kobe Bridge and the Minatojima Tunnel.

AICS provides the central base

The AICS facility contains the Computer Building, the Research Building, and the Heat Source Management Building.

The Computer Building
The building where the K computer is installed.

The Research Building
The many research rooms for use by Japanese and overseas researchers.

The Heat Source Management Building
Equipped with the cooling facility that removes the heat generated by the K computer.

(Provided by RIKEN)

Image of the K computer installed in the computer room

Each computer rack is equipped with about 100 CPUs. In the Computer Building, 800 or more computer racks are installed for the K computer.

Facility for environment and the K computer capability optimization

To operate the supercomputer under optimal conditions, it is necessary to have a facility where careful attention is paid to the research environment. In this facility, various features are implemented to efficiently remove the heat generated during computations; and to enable optimum setup of the system within the facility. The facilities structure always assures the performance of the K computer and that it is setup as necessary for stable operation. There is also a research and education environment that promotes research interactions and the fusion of various types of knowledge. It is expected this will greatly contribute to enhancing the research and development foundation and sustain and advance technical capabilities in Japan.


Quake-absorbing device

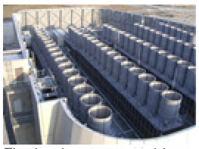
The heat source machine building


Cold-water pipe


The research building

(Provided by RIKEN)

Project Overview

Simulations are becoming increasingly important as the third major research and development methodology along with experimentation and theory. Supercomputers that can conduct more advanced and accurate simulations are required to tackle the unsolved challenges that face the world's populations and researchers.

Therefore, as a foundation of competitive power in science and technology in Japanese industry, supercomputing technology is designated as a "key technology of national importance" and the project is being promoted by MEXT. Fujitsu developed the K computer jointly with RIKEN, the main authority for this development.

The development of the K computer is designated as an important technology in long term national strategy (key technology of national importance), and we are aiming to commence full-service in 2012.

- **What is the national development project of the K computer?**
- **The K computer is incredibly fast**
- **"SPARC64™ VIIIfx": A Fast, Reliable, Low-Power CPU**
- **Innovative "6-Dimensional Mesh/Torus" Topology Network Technology**

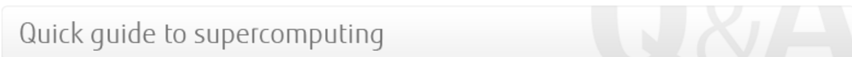

The K computer is incredibly fast

Designed for speed

The K computer is the world's first supercomputer that broke the 10 petaflops barrier. So how fast is 10 petaflops?

The number ten "peta," or 10 quadrillion corresponds to 1 followed by 16 zeros. In Japanese, this is expressed as one "Kei." That is why this supercomputer is called the K computer.

10 quadrillion worth of computations is equivalent to the world's 7 billion people each performing one computation per second, 24 hours a day for about 17 days. The K computer is able to do all of those computations in just one second.

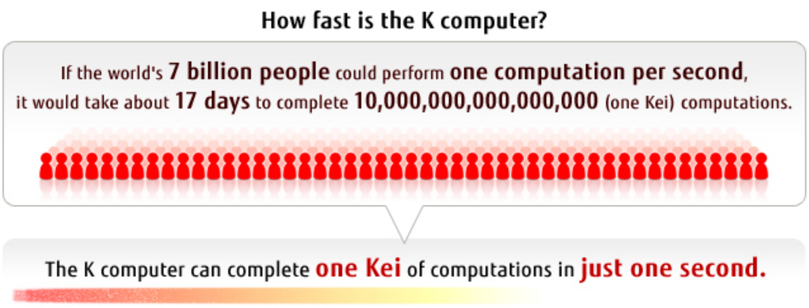

The enabling technologies

What made the K computer with its tremendous computational power of 10 petaflops possible?

There is a limit to the computational power and throughput of a single CPU. To scale computational power, a massive number of CPUs have to be interconnected. In the case of the K computer, a CPU with world-class performance and the technology to connect more than 80,000 CPUs were specially developed.

The CPU is the heart of any computer. However, the K computer cannot conduct large-scale calculations at high speed just on the strength of CPU power ? just as we humans cannot act just using our hearts and nothing else. In large-scale computations, various repeating calculation processes are performed, using previous calculation results. Such processes are allocated to multiple CPUs, making data communication between CPUs very important. Since there are 80,000 or more CPUs, a network is needed with the ability to efficiently manage the vast streams of data traffic.

Memory capabilities are also essential in supercomputer design. The vast amounts of data required for the large-scale computations conducted by the K computer at the ultra-high speed of 10 petaflops need to be read in and out of memory. In addition, a storage system is required. This must be capable of holding the vast volumes of computational results that are continuously output.

Fujitsu designed and developed an innovative new technology by combining its know-how of the various building block required to build an ultra-high speed computation machine. This was incorporated in the K computer. The K computer can achieve 10 petaflops because building blocks in the K computer operate at very high efficiency levels in the world.

How to build a system with high availability

The K computer consists of a vast number of components including 80,000 or more CPUs and 200,000 or more cables. This makes system resilience a critical success factor. It is essential that

- Individual components are highly fault tolerant
- The system as a whole continues to operate even in the event of partial failure
- Failed components can be replaced while the K computer is in use

The heat generated by the CPUs during operation is managed by a water-cooling system to lower the CPU failure rate.

Another important element - a high-performance, high-reliability unique network to exchange data between the CPUs featuring a 6-dimensional mesh/torus topology was developed. The 6-

dimensional mesh/torus topology supports flexible data exchange thanks to a multidimensional CPU interconnection design. Therefore, even if one component fails during operation, the system simply bypasses the failed component and continues to operate.

Quick guide to supercomputing

"SPARC64™ VIIIfx": A Fast, Reliable, Low-Power CPU

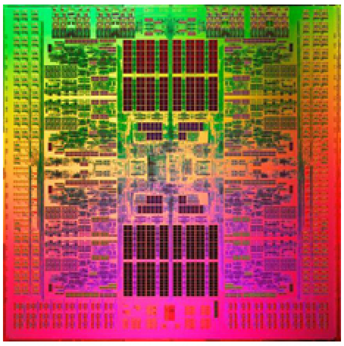
SPARC64™ VIIIfx

The K computer, uses a "SPARC64™ VIIIfx" CPU designed and developed by Fujitsu. The features of this CPU are described below.

45 nanometer process technology is employed in the "SPARC64™ VIIIfx CPU to provide a number of advantages. If you drew complex electric circuit using 1-mm wide wires, it would turn out to be the size of an athletic field. But if you make the wire as thin as possible the same circuit can be reduced to the size of a postage stamp. As the level of integration increases the wire length in the circuit also becomes shorter and the number of cores, the base units that perform calculations, increase. This increases the overall performance. At the same time, since this miniaturization also decreases the number of components, the failure rate also decreases. Furthermore, the smaller the physical circuit size, the greater the reduction in power consumption.

Let's look at the approaches to achieving "high performance," "high reliability," and "low power consumption" in the SPARC64™ VIIIfx.

[Column]

How big is "45nm"?

The size of each transistor (gate length) within the "SPARC64™ VIIIfx" is 45 nanometer (45 nm). So, how big is a nanometer? The term "nano" means "dwarf" in Latin and as a unit of length represents "one billionth". One nanometer is one billionth of one meter, i.e. one millionth of one millimeter. The average diameter of a human hair is about 80 millionths of a meter, usually written as 80μm. This makes 45nm about 1700th the thickness of the average human hair.

How is the CPU created?

How on earth is a CPU with such high computation power created?
Let's check the procedures for creating a CPU using the silicon ingot process.

See How to Create a CPU »

10 petaflops achieved using 80,000 or more high performance CPUs

The heart of the K computer consists of 80,000 or more "SPARC64™ VIIIfx" CPUs. Each CPU is equipped with eight cores, which are "the minimum units that perform calculations", and performs 128-gigaflops. Moreover the 80,000 or more CPUs (640,000 or more cores) combined achieve the tremendously high speed of 10-petaflops computational performance per second (one Kei in Japanese).

Schemes to improve reliability

With the K computer equipped with 80,000 or more CPUs, one of the most important challenges is ensuring optimal performance from such a large-scale system. This is done by making all of the CPUs perform in a stable state.
Circuits in each SPARC64™ VIIIfx CPU are equipped with an "error recovery function". This is a mechanism for automatic recovery. If there is an error during computation, that instruction is automatically executed again so that there is no impact on system operation.
In addition, the CPUs in the K computer are efficiently cooled with water. Ensuring the CPUs always operate at low temperatures provides a low failure rate, which also improves component life.
These schemes assist with the great improvement in reliability of the whole system.

World's top performance per unit of power consumption

Another challenge with the K computer is how much power consumption can be reduced. In a normal large-scale system with many interconnected high performance CPUs, increases in the number of CPUs are proportional to increases in the amount of power consumed. To reduce the power consumption of the whole system, it is necessary to reduce the power consumption of each CPU.

It is difficult to achieve both high performance with many interconnected CPUs and reduce power consumption at the same time. However, we achieved both in the development of the K computer.

Generally speaking, the power consumption in an electric circuit is reduced more when its temperature is lower. The K computer reduces its power consumption by decreasing CPU temperatures and heating values through its efficient water cooling system.
In addition, a design scheme is employed whereby any unused circuits during CPU computations do not consume power unnecessarily.
Based on these innovations, the "SPARC64™ VIIIfx" CPU achieves world top speeds of 2.2 billion computations per second per watt (2.2-gigaflops per watt) in a single chip.

It can be said that K computer is an environmentally friendly system with very high energy efficiency.

- **What is the national development project of the K computer?**
- **The K computer is incredibly fast**
- **"SPARC64™ VIIIfx": A Fast, Reliable, Low-Power CPU**
- **Innovative "6-Dimensional Mesh/Torus" Topology Network Technology**

Quick guide to supercomputing

How to Create a CPU

How is a CPU, the heart of a supercomputer, created?

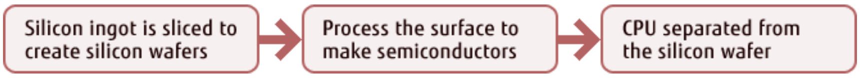

Silicon ingot is sliced to create silicon wafers

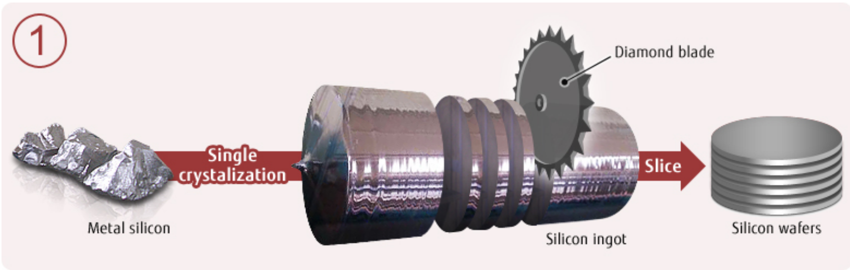

A large silicon monocrystal or ingot (99.999999999% pure) is sliced by a diamond blade to create thin silicon wafers.

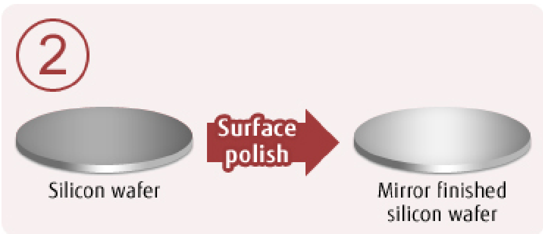

The surface of the silicon wafer is polished to create a mirror finish.

Process the surface to make semiconductors

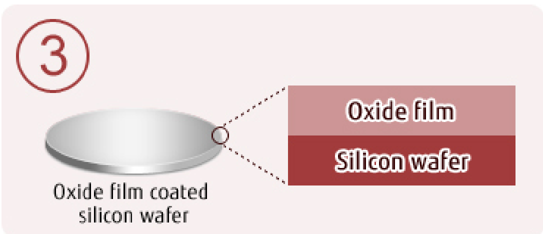

An oxide film is grown onto the wafer.

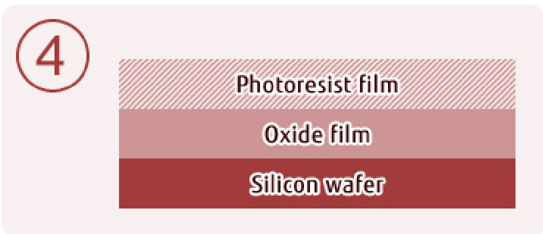

A photoresist film is coated on the wafer surface.

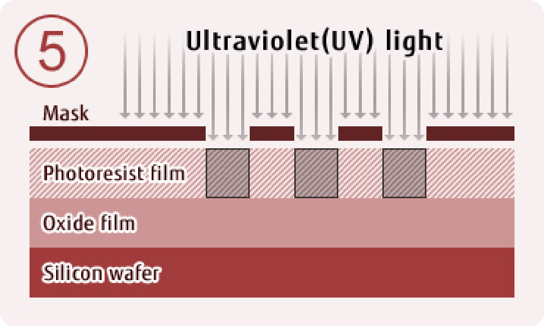

The photoresist film is exposed to ultraviolet (UV) light through the pattern on the mask.

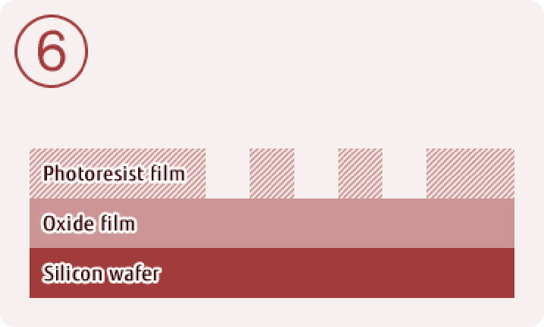

The photoresist region that was exposed to the UV light is removed using a developer. (The exposed area changes to a substance that is dissolved by the developer.)

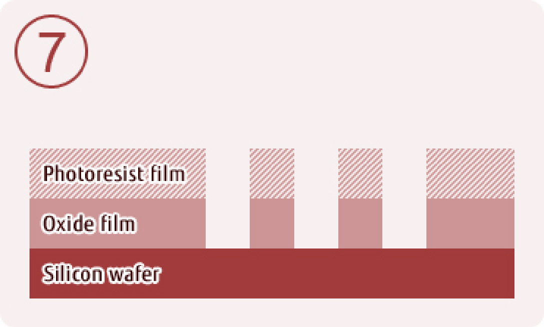

The oxide film is then removed using a caustic agent. This exposes the silicon surface.

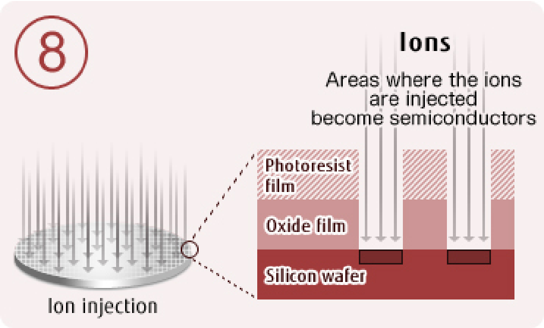

The required ions are then injected into the silicon surface. The character of the silicon then changes into a semiconductor, which is a state where elements with electrical characteristics can be created.

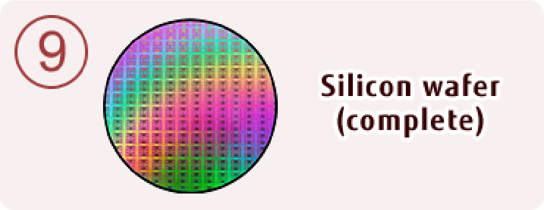

Elements with electrical characteristics are then created by connecting wires and creating circuits. Many chips can be created on the same silicon wafer. Wafer probers are used to

electrically test the chips and define good or bad chips.
Note: This image has been simplified for illustration purposes. On an actual wafer, solder bumps are formed for connection to the package board.

CPU separated from the silicon wafer

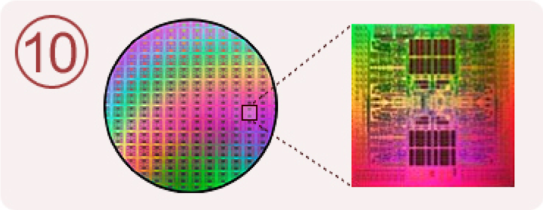

Good chips are separated from the silicon wafer.
Note: This image has been simplified for illustration purposes. On an actual wafer, solder bumps are formed for connection to the package board.

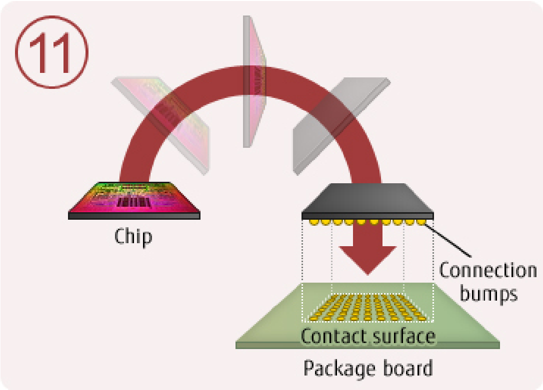

The solder bumps on each separated chip are matched to the ceramic or resin package board.
Note: This image has been simplified for illustration purposes. On an actual wafer, solder bumps are formed for connection to the package board.

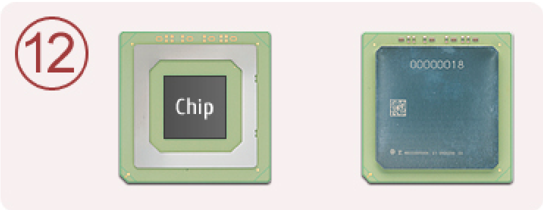

A heat sink to efficiently remove heat is attached onto the package board where the chip is mounted. The CPU is now complete. After various tests, only the good products are selected for use.

- What is the national development project of the K computer?
- The K computer is incredibly fast
- "SPARC64™ VIIIfx": A Fast, Reliable, Low-Power CPU
- Innovative "6-Dimensional Mesh/Torus" Topology Network Technology

Innovative "6-Dimensional Mesh/Torus" Topology Network Technology

Tofu interconnect ("6-Dimensional Mesh/Torus" Topology Network Technology)

In the K computer, a tremendously large system containing more than 80,000 CPUs, the network that exchanges data such as computational results between CPUs plays a very important role. The K computer's network, called Tofu, uses an innovative structure called "6-dimensional mesh/torus" topology. This enables the mutual interconnection of more than 80,000 CPUs.

How fast can CPUs exchange data?

Each CPU in the K computer can perform numerous computations in a short period of time. There are cases where data exchanges take place between CPUs during the computations. If these data exchanges are slow, the high computational power of all those CPUs cannot be fully utilized. The design of the 6-dimensional mesh/torus topology in the K computer provides many communication routes between neighboring CPUs. Execution of data communications between CPUs via the shortest route and over the shortest period of time is enabled to ensure the network can fully draw out the world top-class CPU computational power.

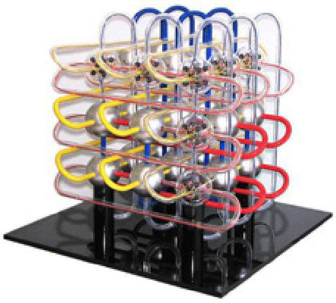

"6-dimensional mesh/torus" topology (model)

Immediate detection of CPU failure and organized data traffic

For the K computer to always maintain its highest performance, it is important that failures do not occur. Further even if a partial failure does occur, its overall impact must be minimized. The K computer is configured with alternate routes in the network between CPUs, and with a mechanism that bypasses failed CPUs so data exchange can continue. This means computational processing cannot be stopped.

Management functions to ensure efficient and maxim CPU performance

The K computer is expected to process computations (jobs) from many users simultaneously. The K computer therefore assigns the required CPUs to each respective job, based on their process contents from its 80,000 or more CPUs. But when assigning CPUs to many jobs, it is more efficient to ensure there is no unnecessary data communication between the CPUs. In the K computer, the job management software assigns respective jobs to the CPUs and controls the order of processes. This resembles the situation where you pack as many boxes as possible in a warehouse efficiently, leaving as little wasted space as possible, even when the boxes are of different shapes and sizes. Because the 6-dimensional mesh/torus topology network provides many communication routes between neighboring CPUs, the shapes of the jobs assigned to each CPU group can be flexibly changed. In other words multiple jobs can be assigned very flexibly and efficiently within the overall CPU population.

The K computer, using the features of the 6-dimensional mesh/torus topology, can maximize the computational power of its 80,000 or more CPUs without waste.

- What is the national development project of the K computer?
- The K computer is incredibly fast
- "SPARC64™ VIIIfx": A Fast, Reliable, Low-Power CPU
- **Innovative "6-Dimensional Mesh/Torus" Topology Network Technology**

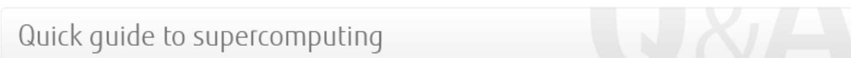

Global

Global challenge and Fujitsu

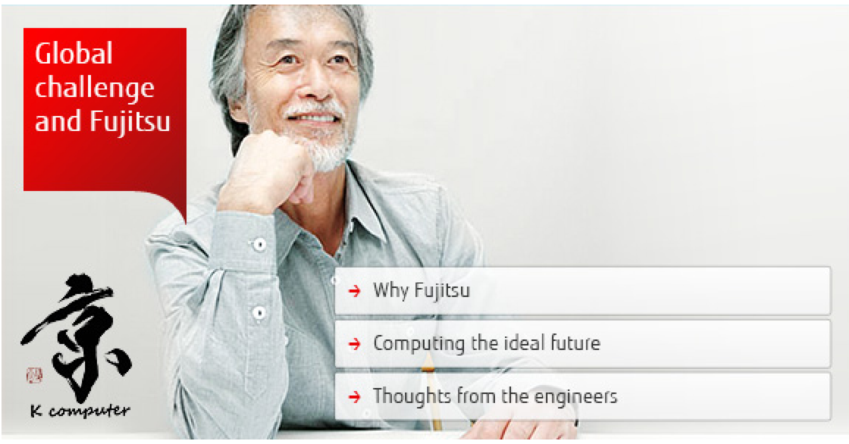

Quick guide to supercomputing

Why Fujitsu

In 2006, the Japanese government defined supercomputing as a key technology of national importance. To advance research in this strategic area and maintain Japan's leading edge, the government also earmarked investment funds for large-scale promising projects. As part of the High-Performance Computing Infrastructure (HPCI) initiative led by the Japanese Ministry of Education, Culture, Sports, Science and Technology (MEXT), RIKEN – Japan's leading research institute – was entrusted with the task of developing a next-generation supercomputer.

To support it on this national strategic project, RIKEN needed a trustworthy, experienced partner. It turned to Fujitsu. Not only does Fujitsu have over 40 years of experience in the mainframe business, it has also gained invaluable insights pioneering the development of supercomputers over the past 30 years. Even more importantly, RIKEN valued Fujitsu's engineering spirit and unique approach to innovation. It was confident that Fujitsu would have the courage and vision to take a risk and try something totally new in order to break the existing speed and performance barriers. In addition, very few companies can match Fujitsu's technical depth and breadth, spanning core competencies in processor development, interconnect controller chips and system administration/management software for large-scale systems.

Both partners had been working on the K computer since 2006 with full operation and shared use scheduled for November 2012.

When we launched the K computer project, we knew we would need an exceptional partner at our side – a partner with the vision, determination and technical depth to accompany us into uncharted terrain and come out the other side with the world's fastest computer. Fujitsu was the obvious choice for us, bringing an invaluable combination of technical expertise and innovation to the table.

Quick guide to supercomputing

FUJITSU Global

Computing the ideal future

The world is facing difficult challenges hard to solve immediately, such as the environment, energy sustainability, and food issues. However, we need to promptly and properly solve these issues in order to bring about an ideal future for the Earth and its inhabitants.

To enable continuous economic development, while solving these challenges, it will be essential to create innovations based on the wisdom in the world and leading-edge technologies.

Supercomputers can compute tremendous amounts of information at ultra-high speeds to support the analyses and choices that will create such innovations. This will more quickly and objectively enable us to discover ways to deliver that future world, that as yet we have been unable to perceive.

Fujitsu continues to work on realizing such an ideal world through supercomputer development and its applications with you.

Quick guide to supercomputing

Thoughts from the engineers

Strive: Means aiming at being the world's best.

Toshio Yoshida

Rising to the challenge of difficult hurdles
The energy consumption of computers has become a huge issue globally. The performance demanded from the CPUs was three times or more than previous models, but power consumption had to be half or less.
This project challenge let us rise up through the term "world's fastest computer." By working for that distant goal, we were able to achieve it.
[CPU development team member]

Naoyuki Shida

World-class supercomputer fully created in Japan
The tough challenge of reducing memory consumption, while maintaining high performance, is a necessity in the software for large-scale supercomputers.
In the development team there are many young engineers, who are very eager. With energetic members like them, I would like to proceed with further development toward the finalized form of the software for this supercomputer wholly made in Japan. [Software development team member]

Enthusiastic engineers come together.

Kouichi Hirai

Starting from zero
I started my project participation in 2008. There were many young engineers among the team members. Experienced engineers like me joined this group and we engaged in vigorous discussions to bring forward the design. Under the shared goal of becoming the world's best, we were able to unite as one team. [Middleware development team member]

Masatoshi Haraguchi

Endless performance improvement
I participated as a compiler engineer in a new CPU design, and helped complete the SPARC64™ VIIIfx. It all depends on the compiler whether you use or kill a great design idea. The endless effort towards performance improvement has given me a sense of satisfaction. Although we had a nonstop schedule unlike anything I had ever experienced before, I worked to accomplish the product and its performance to a level that I could be proud of. [Compiler software development team member]

Yoshiro Ikeda

Cooperation with the verification team
The interconnect was specifically designed to efficiently connect tens of thousands of CPUs. Since the scale was at an unprecedented level, there were not many places where we could employ our accumulated development know-how within the company. However, through close cooperation with the verification team, we were able to achieve the requested high performance, high reliability and low power consumption design. [Interconnect development team member]

Across the organizational divide

Gou Sugizaki

Cheers from our competitors
When I introduced our technologies at the supercomputer exhibition in U.S.A., I even received words of encouragement from our competitors. That encouragement came as a fresh reminder of the importance of the project and of Fujitsu's technical capabilities. I am profoundly grateful that I was able to strive to a higher goal through the project. We continue to work toward establishing the system while cherishing the cheers from our competitors. [CPU development team member]

Yuuichirou Ajima

Valuable experience
When I explained the interconnect specification to the in-house and external team members involved in the project, I was shocked to receive comments such as, "I don't understand" and "It's difficult to understand." I then received valuable advice from many people including the leading authority in this field. While I was improving the explanatory documents, I learned the importance and the difficulties of creating understandable descriptions. [Interconnect development team member]

A desire to realize the dream

Kenichirou Sakai

A system no one had ever experienced before
Every team member tackling the development was beating their brains out and working hard. Fujitsu is one of the few companies in the world that can develop such a supercomputer even from the software level. In playing a part in the development team and elevating Fujitsu's presence in the world, I would also like to think we are contributing to the development of Japanese science and technology and the realization of a more prosperous society. [File system software development team member]

Takeharu Kato

Supporting future social infrastructure
The large-scale structure is unprecedented in the world, so is the technology to detect failures anywhere and recover from faults. Those are technologies that can be applied to other highly reliable systems supporting databases and social infrastructure in the future. Daily I continue to devote my energies to this development while dreaming that I can contribute to other fields through this project. [OS development team member]

Yoshihiro Kusano

Joining Fujitsu made my dream come true
I joined Fujitsu with the desire to develop world top class computing machines. Now, I am deeply thrilled with my involvement in the K computer development project. I would like researchers to use the K computer to their hearts' content and to create innovations that will make many people happy. [System development team member]

- Why Fujitsu
- Computing the ideal future
- Thoughts from the engineers

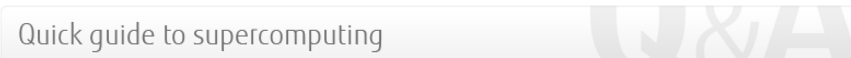

Quick guide to supercomputing

Throughout the ages, humans have always sought to understand the world around them, often using tools to adapt to their environment and solve problems. As the challenges facing humans have become more sophisticated, so too have the tools they use. Early stone hammers have ultimately given way to ultra-high-speed supercomputers. These supercomputers are being used to address many of the pressing challenges facing society today...

Read more

Questions

What is a supercomputer?

A supercomputer is able to perform advanced scientific and technical calculations quickly and on a scale that would be difficult to achieve using regular computers. Supercomputers are mainly used by research institutions and industrial companies.

Questions

What is a supercomputer used for?

The world faces a variety of problems that are complex and difficult to solve quickly. These include the need to conserve resources, protect the environment and develop new medical therapies. Supercomputers perform calculations at super-high speeds, processing massive amounts of information. This allows them to recreate or model a wide variety of phenomena inside the computer, such as the unobservable internal structure of the earth, large-scale disasters that cannot be reproduced and the effects of pharmaceuticals on the human body that are not clearly understood. These models give scientists unique insights into these events and phenomena. This process is called computer simulation, and has become the third pillar of science, flanking theory and experiments. Supercomputers have become an essential tool in all fields of research and development, supporting everything from fundamental research right up to manufacturing.

Questions

Are there any other applications for supercomputers?

Leading-edge simulations using supercomputing power are widely used in a variety of fields beyond research. Common applications include forecasting share prices in financial transactions, creating animations in computer graphics, developing swimwear with lower resistance for competitive swimming, developing washing machines that remove dirt using less water and designing equipment for manufacturing potato chips.

Questions »

Who developed the world's first supercomputer?

Computers were originally developed to perform scientific and technical calculations. In other words, the first generation of computers were what we now call supercomputers. The very first computer was ENIAC, which was created in the USA in 1946. ENIAC was developed for military purposes, and was mainly used for calculating the ballistics of artillery. The first computer in the world to be referred to as a supercomputer was the "CDC6600" developed in 1964 by Seymour Cray, who is known as the "father of supercomputing".

Questions »

How many supercomputers are there in the world?

Supercomputers are currently used for a wide range of applications by research institutions and industrial corporations. They range in size from super-large scale to small-scale devices. There are so many supercomputers in use today that it is extremely difficult to know the exact number.

Questions »

What makes a supercomputer fast?

Most modern supercomputers are equipped with a large number of CPUs.

There are two key speed factors:

1) each CPU must perform calculations quickly and

2) data must be exchanged quickly between the CPUs.

The speed of each CPU can be accelerated by increasing the operating frequency, increasing the number of cores, increasing the number of internal computation units and supplying the required data at high speed. Data exchange between the CPUs is accelerated by developing fast inter-CPU networks and developing interconnect technology.

Both the software and hardware must be optimized for high speed.

Questions »

What data is communicated between the CPUs?

The information (or data) needed to perform a calculation is communicated between the CPUs.

For example, consider a weather forecasting computer simulation in which the region being forecast is divided into a fine grid of cells, with the calculations for each cell performed by a single CPU responsible for that cell. In this situation, the next state is calculated by comparing meteorological conditions such as the temperature in the neighboring cells. In a supercomputer, a single job is tackled by dividing it up among a large number of CPUs. Data is communicated between the CPUs so each CPU can find out what the other CPUs are doing and tell the other CPUs what it is doing.

Questions »

Why do CPUs get hot?

Internally, a CPU contains many components that are made of semiconductor materials and metal. When electrical current flows through the CPU in order to perform a calculation, heat is generated by the electrical resistance (resistance to the flow of electrical current) of these components and wiring. That is what makes the CPU hot.

Questions »

Where is the K computer?

The K computer is located in Japan on an artificial island called Port Island in Kobe city in the Hyogo prefecture. The decision to pick Kobe from the 15 candidate locations across Japan was made in March 2007 as the result of an investigation that took more than 6 months. The investigative team looked at a variety of factors, including the effects of natural disasters, stability of the electrical power supply, cost of outfitting and running the facilities, ease of access and quality of infrastructure to support a research and education center.

Questions

What are petaflops (PFLOPS)?

FLOPS means floating-point operations per second. It is used to express the processing power of a computer. FLOPS can be prefixed by various letters to indicate the following orders of magnitude:

- 1 G(giga = 10^{9})FLOPS: 1,000,000,000 calculations per second
- 1 T(tera = 10^{12})FLOPS: 1,000,000,000,000 calculations per second
- 1 P(peta = 10^{15})FLOPS: 1,000,000,000,000,000 calculations per second

The K computer is capable of handling 10,000,000,000,000,000 (10 quadrillion) calculations per second (10 petaflops = 10^{16}), which corresponds to one "kei" in Japanese. Which is why it was called the K computer.

Questions

How many people were involved in the development of the K computer?

The K computer was developed in partnership between RIKEN and Fujitsu with over 1,000 people involved in the project.

Questions

How much power does the Supercomputer "K computer" consume?

A lot of electrical power is needed to run a supercomputer.

The electrical power used by the Supercomputer "K computer" when measuring the benchmark for TOP500 list in November 2011 was 12.65989 MW/h. Considering that average power used per hour by a typical household is 400 W/h, the power used by the K computer is equivalent to approximately 30,000 homes.

However, the performance per unit of power is in the top class among the supercomputers in the world, and the K computer can be said to be a computer that has high environmental performance.

Questions

Look into the K computer

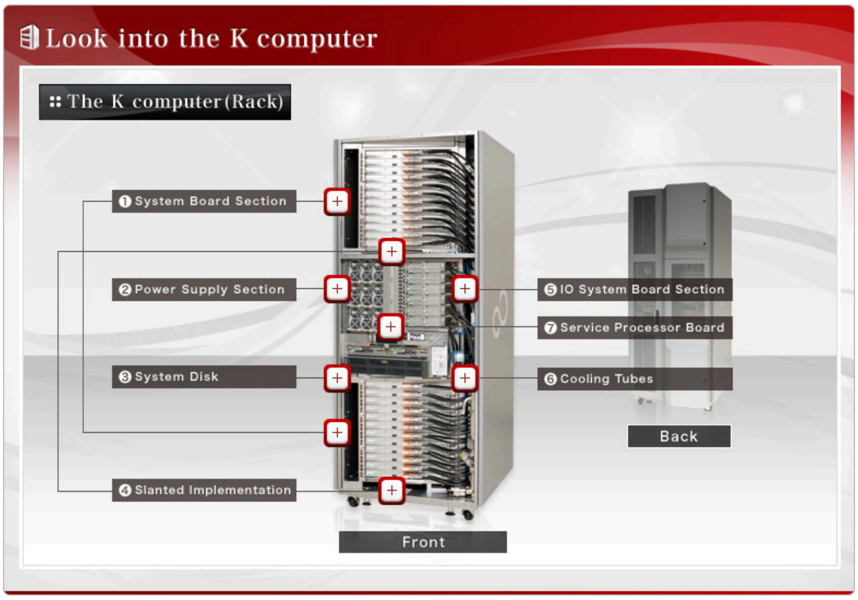

1. System Board Section

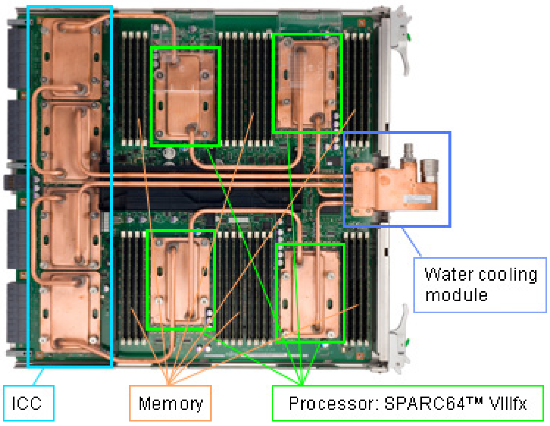

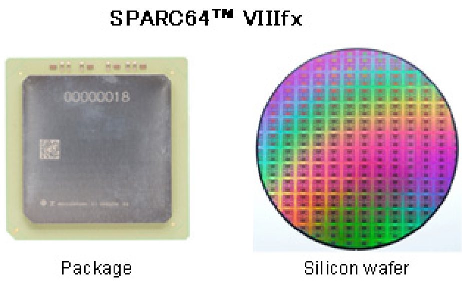

Each system board is equipped with four CPUs (SPARC64™ VIIIfx). In the SPARC64™ VIIIfx, eight cores are integrated per CPU. In terms of power consumption, this board delivers the world's highest level of performance to power use ratio at 2.2 gigaflops per watt.
In addition, the heat generated from CPU and ICC during operation is removed by water cooling, which leads to a reduction in CPU and ICC power consumption and adds to longer component life.

Note: ICC (interconnect controller) is a chip that controls the network (interconnect).

2. Power Supply Section

There are nine power supply units in each rack which supply power to the system. They have a fully redundant design for the highest reliability.
This means operations can continue even if one power supply unit fails.

3. System Disk

The disk stores the Operating System (OS) that controls the system.

Note: Aside from the internal hard disks, an external storage system for computational data and results is connected to the rack.

4. "Slanted Implementation" supports cooling and high density

For stable operation of densely racked equipment, efficient removal of heat generated during operation is necessary.
Therefore, the K computer uses air cooling as well as water cooling.
Optimization of the air flow within the equipment is an important element in achieving high performance.
To create appropriate air flow paths within the rack, the system boards are mounted in a slanted manner.

5. IO System Board Section

To exchange data required for computations and computational results with the external storage system, six IO system boards are installed.

6. Cooling Tubes

These tubes supply water to cool the CPU and ICC. They have sensors that constantly monitor water pressure, temperature and condensation, etc.

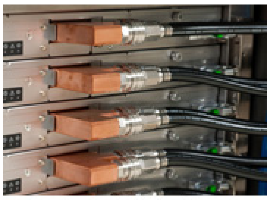

System board with Joint unit

Water cooling pipe safety valve

7. Service Processor Boa

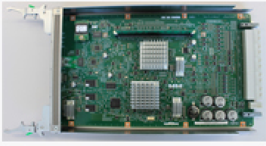

The service processor board controls the rack, it performs the initialization of system boards and monitors errors, failures and anomalies. The service processor boards are configured redundantly for higher reliability, and two boards are installed per rack.

Related Information

- **"SPARC64™ VIIIfx": A Fast, Reliable, Low-Power CPU**
- **Innovative "6-Dimensional Mesh/Torus" Topology Network Technology**

Supercomputing explained in three minutes

New tools for new challenges

Throughout the ages, humans have always sought to understand the world around them, often using tools to adapt to their environment and solve problems. As the challenges facing humans have become more sophisticated, so too have the tools they use. Early stone hammers have ultimately given way to ultra-high-speed supercomputers. These supercomputers are being used to address many of the pressing challenges facing society today. These include the need to:

- protect our environment more effectively and conserve natural resources
- develop new, renewable sources of energy
- mitigate the effects of climate change
- predict natural disasters such as earthquakes, tsunamis and storms and develop appropriate response plans for these increasingly frequent events, and
- develop new drugs to treat debilitating diseases.

The scale and complexity of these problems make progress difficult – traditional problem-solving techniques are simply too costly or time-consuming. Supercomputers have the ability to speed up this process.

About a future that delivers supercomputers and simulations

It's all about speed

A supercomputer is simply a computer that can perform incredibly fast calculations. One of the fastest supercomputers in the world – the K computer – was developed by Fujitsu in partnership with RIKEN under the High Performance Computing Infrastructure Initiative promoted by the Japanese Ministry of Education, Culture, Sports, Science and Technology.

The K computer operates at a speed of 10 quadrillion calculations per second. To give you an idea of just how fast that is, imagine the following: All seven billion people in the world have a calculator and are asked to perform one calculation per second 24 hours a day, non stop. It would take the world's population approximately 17 days to do what the K computer can do in just one second. If we think of a regular desktop as a snail, then the K computer is a jet airplane.

This incredible speed makes supercomputers ideal for compute-intensive calculations and simulations. To simulate events such as car crashes, earthquakes and tsunamis, supercomputers need to handle massive volumes of data and that calls for exceptional computational horsepower.

About the tremendous calculation speed of supercomputers

A race

Around the globe, many countries are competing with each other to develop the world's fastest supercomputers. This race is inspired by the desire to advance research in a variety of fields. But it's not just research and development that benefit from supercomputers. They also drive innovation across a variety of industries by enhancing safety, design and environmental performance. In addition, a world-leading high-performance computing development environment attracts top-class scientists and fosters a culture of innovation. This – in turn – raises the bar for research and can even make a country more competitive if research findings are channeled into industry innovations. It's a bit like Formula One racing – the company with the fastest car attracts the best drivers and this technology ultimately finds its way into commercial and passenger vehicles.

Development of the next generation of supercomputers for nurturing the next generation of scientists

Computing the answers to today's challenges

At the heart of any computer is the CPU – the central processing unit. Regular computers have a single CPU. By comparison, the K computer is equipped with over 80,000 CPUs. These CPUs were designed for maximum performance, high reliability and low power consumption. CPUs are only part of the picture, however. Equally important is the technology used to interconnect these processors and transmit data between them at ultra-high speeds. In total, over 200,000 cables extending over 1200 km were used for the innovative K computer interconnect network, which is managed by the LSI interconnect controller. The high-performance CPUs and innovative network technology played a key role in making the K computer the fastest machine in the world today.

The Japanese government plans to make the K computer available to a wide range of customers and research organizations. It will be used to support and accelerate the search for viable solutions to today's climate and environmental challenges, advance cutting-edge research and industry innovation and help make the world a safer place.

Memoirs from engineers incorporated into the Supercomputer "K computer"

Note: About the K computer

