WELCOME TO THE WORLD OF TSUBAME2.0!

The next-generation "supercomputer for all"

In the fall of 2010, the successor to the highly successful TSUBAME1.0 began operating at Tokyo Institute of Technology. TSUBAME2.0 is the first supercomputer in Japan to achieve both petaflop-scale processing and low power consumption. In the following pages, we will explore the secrets to its world-class performance.

Tokyo Tech
Supercomputer
UBiquitously
Accessible
Massstorage
Environment

Tokyo Institute of Technology

The World's "Greenest" Supercomputer

The designers of TSUBAME2.0 confronted the problem of increasing power consumption head-on. Many technical improvements, such as full-scale use of modern GPUs that are most effective in reducing power consumption, and using a new cooling system, resulted in a dramatic decrease in power consumption.
In the Green500 list as of November 2010*4, TSUBAME2.0 earned the title of "Greenest Production Supercomputer in the World."

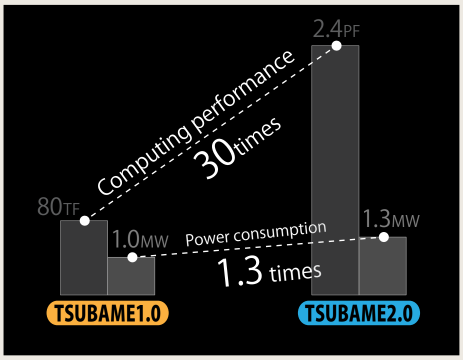

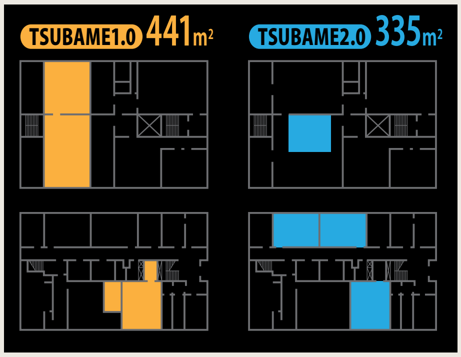

A petaflop-scale supercomputer that is more energy-efficient than a PC

TSUBAME2.0 incorporates a variety of technologies to increase computing capacity while limiting power consumption. These innovations enable TSUBAME2.0 to offer 30 times better performance than TSUBAME1.0 with only a slight increase in power consumption. As a result, TSUBAME2.0 became a petascale machine that is three-times more efficient than ordinary laptop PCs in performance per watt.

Improved performance within a small form factor

Using GPUs, which offer greater per-chip computing performance than a CPU, also enables high-density packaging. As a result, the space needed for installation is 1/3 less than that for TSUBAME1.0. Normally a space as large as a gymnasium is needed to house a large-scale supercomputer, but TSUBAME2.0 can be accommodated in a space equivalent to two classrooms.

Water-cooled supercomputer

TSUBAME2.0 uses a state-of-the-art enclosed water cooling system† that incorporates heat exchangers within the racks. This system not only greatly lowers the power usage effectiveness (PUE)*6, it also seals in the noise produced by the cooling fans in the rack, making the system significantly quieter.The rack doors are also designedtoward this goal.

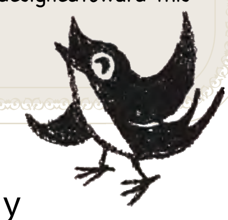

Supercomputer for everyone with massive storage capacity with extremely high-speed data access

TSUBAME2.0's hierarchical storage capacity is over 11PB, more than 10 times that of TSUBAME1.0. Students in the Tokyo Institute of Technology can access the TSUBAME2.0's storage via the network at any time by logging in to their user account.

• Power of hierarchical storage

In the first level of storage, each node has solid state drives (SSDs) with a total of 190 TB for use as a scratch drive, achieving ultra-fast speeds of more than 660 GB per second. The second level is a parallel file domain with more than 6 petabytes that increases the speed of petabyte-class data transfers. The third level is a 4 PB tape library system that allows transparent on-demand data access between the file system and the tape library. Offering shared use to people not only on-campus but also off-campus, TSUBAME2.0 is truly a "supercomputer for all."

• Campus cloud *7storage

The 1.2 PB of cloud-based home storage can be accessed from various locations. Many students in the Tokyo Institute of Technology use the cloud storage as if it was their own PCs via the Wi-Fi network in cafeterias and convenience stores in the campus. Each student can use up to 50GB for storage volume, 25 GB for data storage, and 25 GB for computing.

***4 Green 500 list**
The Green500 list is a world ranking of energy-efficient supercomputers. The ranks are determined by "gas mileage" in which supercomputer computing performance in Top500 is divided by power consumption. In the November 2010 Green 500 ranking, TSUBAME2.0 was the only petaflops supercomputer to be ranked in the top ten. Figuratively speaking, TSUBAME2.0 is like a racing car whose gas mileage is better than that of a compact car.

***5 Enclosed cooling water system**
This system has heat exchanger systems built into the racks. A uniform jet of cooling air is blown into the intake ports of servers, and then 94 - 97% of the heat of the exhaust air is removed by the cooling water. This system maintains the ideal temperature with fully automated temperature control, and also keeps the cooling power consumption to the absolute minimum. Using this system is much more energy-efficient than cooling the entire supercomputer room.

***6 PUE**
This is an abbreviation for "power usage effectiveness" . The PUE is determined by the sum total of the power for computing equipment plus the power for cooling and other purposes, divided by the power for the computing equipment. As the amount of power for cooling decreases, PUE becomes closer to 1.0, which means better efficiency. Tokyo Institute of Technology is pursuing further research and development into power conservation technologies toward the next generation TSUBAME3.0.

***7 Cloud**
Cloud, or Cloud computing is a form of computing in which users are able to use the functions provided by a computer, such as various services, and data storage, via the Internet, without being particularly aware of the software and hardware that make this possible. It is called cloud computing because users perceive the Internet and the computer system as an amorphous cloud overhead. TSUBAME2.0 uses a mechanism for dynamically changing the OS to provide cloud services using excess computing resources, for example, on-campus educational and business systems operate using TSUBAME2.0 cloud services.

New frontiers of research opened up by TSUBAME2.0

TSUBAME1.0 garnered high praise for its contributions to academic research.
With even faster GPU computation and memory access, TSUBAME2.0 makes possible the most advanced research in the world. Applications that can take full advantage of its performance are now being developed, and these efforts will undoubtedly make it even easier to use TSUBAME.
This section will introduce the state-of-the-art research efforts currently underway at Tokyo Institute of Technology.

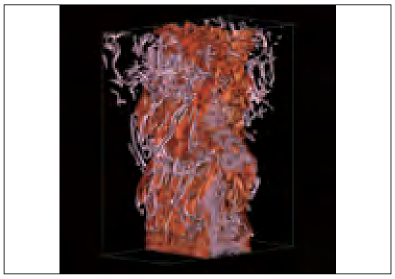

Flame simulations for next-generation energy propulsion systems

Several million years ago, human beings discovered the secret of fire, thereby gaining a source of high-density energy. In the engines of modern rockets, aircraft, automobiles, and so on, as well as thermal power generation systems, a combustion reaction (a sudden and violent chemical reaction) is produced in the midst of extremely complex fluid movements (turbulent flow) to efficiently extract energy. This phenomenon, known as turbulent combustion, is one of the most difficult physical phenomena to simulate. TSUBAME2.0 is being used to conduct the world's most precise turbulent flow simulations, taking into consideration turbulence, complex combustion reactions, and all of the other elements of turbulent combustion phenomena.

Future Expectations

Turbulent combustion is so complex that the details of the phenomenon are almost completely unknown. Explaining this phenomenon will make it possible to develop propulsion engines that offer greater efficiency. More effective use of energy will reduce carbon dioxide emissions and help to resolve global environmental problems.

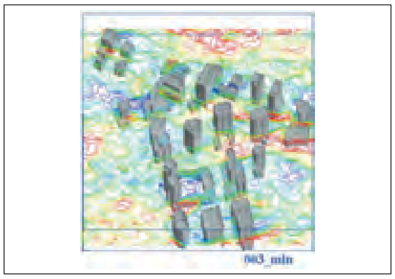

Simulations for environment-friendly urban planning

Urban planning is the process of creating cities through the design of building layouts, constituent materials, wind pathways, green zones, waterfront areas, and so on. Urban planning has a major impact on living comfort and environmental load. Unlike the process of designing an automobile, for example, it is not easy to create a miniature city in a laboratory and subject it to testing. Enter TSUBAME2.0. The supercomputer can perform quasi-experiments through environment simulations that assume a variety of different urban environments, to determine how the temperature and breezes will change as a result of those environments. These efforts will help to create more eco-friendly cities.

Future Expectations

This research will help to promote the design of cities with good air flow that make it difficult for exhaust fumes and other pollutants to stagnate. This will help to alleviate the "heat island" phenomenon that causes rising temperatures in summer, particularly in urban areas.

Real-time tsunami simulation

A tsunami is caused when an earthquake brings about the displacement of the ocean floor. As the speed of the tsunami is much lower than the speed of the seismic wave propagating the earth's crust, quick and accurate warning would make it possible to prevent great damage. The GPU computing is very suitable for tsunami simulations and can complete a computation of 4000 x 8000 mesh with 100-m resolution in less than three minutes. If proper initial conditions of the tsunami are provided when an earthquake happens, it would enable highly accurate tsunami forecasting long before the tsunami arrives at the shore.

Future Expectations

When a big earthquake occurred in Chile about 50 years ago, a tsunami propagated across the Pacific Ocean and caused 142 deaths in Japan. The TSUBAME GPU is available to forecast the arrival time and wave height of tsunamis that come from faraway locations with considerable accuracy.

Air-water violent flow simulation with wave breaking

When water launches into a wet floor, breaking waves occur like the one in the famous Hokusai' s ukiyo-e print "Wave behind Mt. Fuji" . Computation of water flow mixing with air is extremely difficult and one of the grand challenges in computational fluid dynamics. In recent years, researchers have used mesh calculation instead of particle calculation, even in computer graphics for water scenes in Hollywood movies. GPUs are suitable for mesh computation and have recently been applied to sparse matrix solvers. This has made it possible to perform large-scale simulations for air-water violent flows.

Future Expectations

Water and air are essential for life. Future calculations will reveal the essential way in which living organisms — which have been optimized through evolution — relate to water and air and inspire that are ideas useful in our daily lives as well.

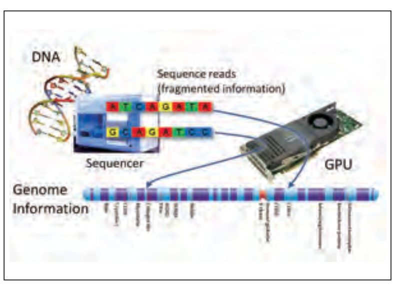

Supercomputer heralds ultrafast genome analysis

The blueprint for life is expressed in the genome, which is made up of DNA molecules. The human genome comprises around 3 billion DNA base pairs. The entire genome can now be decoded in a few days because of dramatic improvements in the performance of DNA sequencers. However, to actually use these data, researchers need to perform extensive data processing including mapping and assembly. Our goal, therefore, is to develop ultrafast algorithms to perform this data processing efficiently on the TSUBAME supercomputer and thereby achieve the desired results from genome analysis in a short period of time.

Future Expectations

This research is expected to lead to "individually tailored medicine," in which the ideal treatment for each patient can be selected (such as administering the most effective anticancer drug based on the genetic information for that person). Another anticipated result is progress in metagenomic analysis, in which the genomes of microorganisms in soil, rivers, or other areas of the environment are analyzed together.

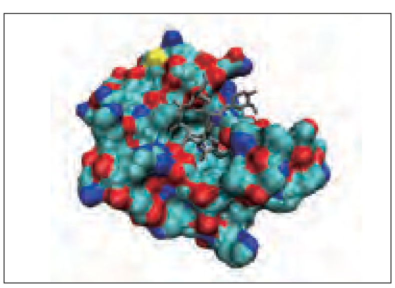

Analysis of the interactions between proteins and ligands

Proteins serve various functions inside the human body by binding with other proteins, chemical compounds, or peptides. Investigations into the interactions between these proteins and the ligands to which the proteins bind can provide information on the protein functions in the body and play an extremely important role in the development of new drugs. In this research, we are building a program for the first ever analysis of binding free energy between proteins and ligands through the use of our supercomputer, with the goal of supporting drug discovery.

Future Expectations

As this research progresses, we expect to be able to design new drugs based on computer simulations, rather than searching for drug compounds in enormous libraries of potential candidates as before. This will reduce research time and cost.

Exploring the mysteries of supercomputers ! !

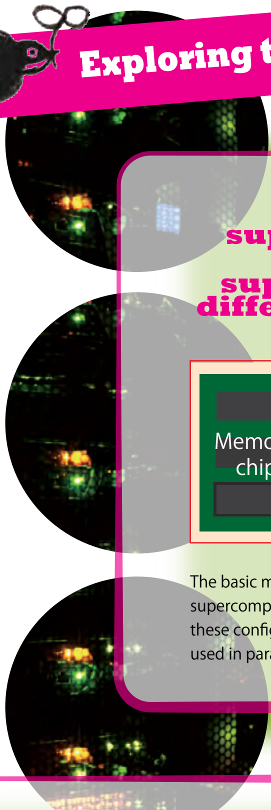

What is a supercomputer? How is a supercomputer different from a PC?

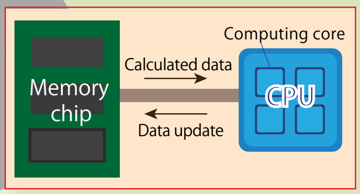

The basic mechanism of PCs and supercomputers. Supercomputers have many of these configurations, connected together and used in parallel. PCs generally have only one.

Fundamentally, supercomputers have the same mechanisms as the PCs that we use in our daily lives. The computing cores in the chip known as the central processing unit (CPU) perform repeated calculations to process and update the data on the memory chip. The results and progress of computation are displayed or transmitted in the form of images or numbers.
Up until a decade ago, supercomputers used extremely expensive special-purpose CPUs and memory, and they were designed to perform extremely high-speed calculations that are not possible on an ordinary PC. However, we are approaching the limit to which the speed of electronic components can be increased, and the chips used in ordinary PCs are equal to those that once powered supercomputers.
Modern supercomputers do not simply use a single sophisticated chip. They use many of them, connected in parallel, to increase computing capacity. It is the same as when all of the people in a group divide up the work into separate tasks; the work can get done more quickly that way. TSUBAME2.0 uses approximately 17,000 computing cores (CPU cores). Each CPU is also connected to extremely large-capacity memory. This requires a mechanism that enables extremely fast communication among CPUs. The important thing is for the design to achieve a good balance of computing, memory, and communication. This is what makes possible ultra high speed, ultra high performance computing.

Do supercomputers need to become even faster?

Simulations are the most important usage of supercomputers. Typhoons, tsunami and other natural phenomena in the world around us can be expressed using mathematical equations. However, analytically solving these equations can be very difficult, and in fact is virtually impossible for mathematical models of real-world complex phenomena such as predicting how big a typhoon will become or what path it will take. For this reason, computer simulations in which an area (for example, Tokyo) is divided into 5 km x 5 km grid squares and cloud generation and wind direction in each grid square are calculated numerically and depicted in visual form will become more and more important. Speed is essential for such simulations. To reduce typhoon damage, we must know what path the typhoon will take in advance; simulation results must be obtained before typhoons hit the area. This is why a supercomputer is needed. Moreover, the faster computations can be done, the greater the quantity of data that can be handled during the same period of time will become. In the example of the typhoon, dividing an area into smaller grid squares and performing calculations yields simulation results that are closer to reality. The figure on the right shows an area that has been divided into grid squares measuring 500 m x 500 m. Compared to the top figure, the bottom figure is able to render the typhoon with greater realism. This is what TSUBAME2.0 makes possible.
In recent years, enormous sizes of data have been obtained from the latest technology. One example is next-generation DNA sequencing technologies, which can rapidly produce vast quantities of data. They revolutionize genomics, and their effects are becoming increasingly widespread. TSUBAME2.0 can analyze enormous size of data in a realistic time. Here as well, further increases in speed are needed to process the rapidly increasing amounts of information.

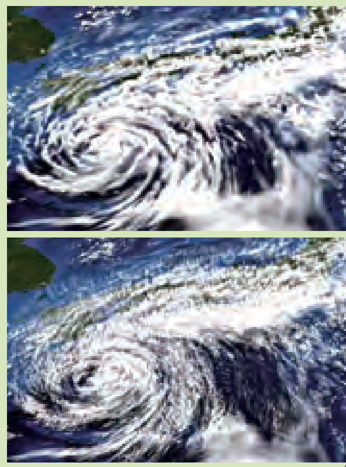

Computation with a fine mesh enables the clouds to be described in greater detail.

TSUBAME features: What is a GPU?

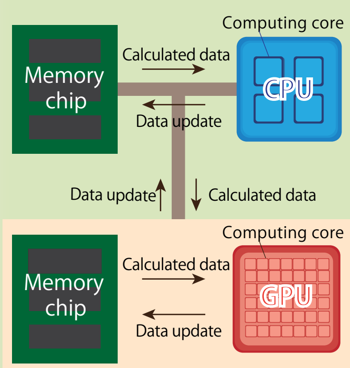

The GPU receives the calculated data from the CPU and processes it at high speed using its many computing cores.

If more and more CPUs are used to increase speed, the result is a supercomputer that steadily grows larger in size. Today's leading-edge supercomputers require a space the size of a gymnasium for installation, and they consume enormous amounts of power.
Therefore, the question is: aren't there any smaller, efficient computing cores that can be used for supercomputers? As it turns out, there are.
Graphics processing units (GPUs) are used by the video cards that control the screens of PCs. GPUs have an extremely large number of computing cores and are designed to process image data in parallel at high speed. They are not as versatile as CPUs, but hundreds of computing cores fit in an extremely small space.
The previous generation of TSUBAME was the first supercomputer in which designers tried using GPUs. The experiment was a complete success. Accordingly, TSUBAME2.0 was designed as a supercomputer with fully-fledged use of GPUs. TSUBAME uses a total of 4,224 GPU cards, and each GPU card has 448 computing cores (CUDA cores), giving a total of 1.89 million computing cores. This is the reason that TSUBAME 2.0 is able to offer computing capabilities that are among the best in the world, yet it can still be housed in a space the size of two classrooms.

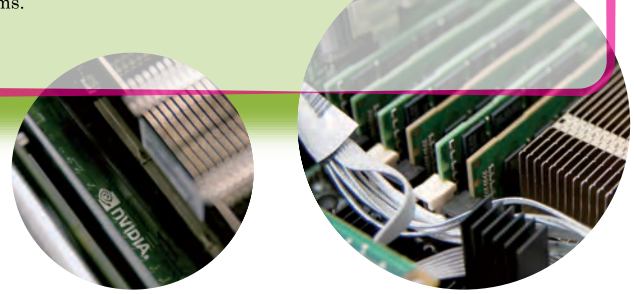

Petaflop-scale supercomputer offers unrivaled computing performance

The 1,442 computing nodes of TSUBAME2.0 offer a total computing performance of 2.4 petaflops, *１ enabling it to perform 2.4 quadrillion floating-point calculations per second.
This is a 30-fold improvement in computing performance over TSUBAME1.0, and it was achieved in a mere four years.
TSUBAME2.0 exceeded the combined performance of all of the other supercomputers operating in Japan in 2010.

Supercomputers offer 100,000 times the computing performance of an ordinary personal computer. Supercomputers are an essential part of national infrastructure in the areas of scientific research and technical development. For this reason, there is intense international competition to develop faster supercomputers. As of 2009, however, the supercomputers of Japan that were formerly number one in the world had fallen to 31st place in the world ranking. With TSUBAME2.0, Japan is now rocketing back up to the top.
TSUBAME2.0, which began operating in November 2010, offers theoretical performance of 2.4 petaflops and is Japan's first petaflop-scale supercomputer. It achieved 4th place on the Top500*2 ranking of supercomputer performance. The secret to its performance lies in the hybrid architecture that realizes extremely efficient vector processing by the GPU. TSUBAME2.0 has more than 1,400 computing nodes, each made up of two CPUs, three GPUs, and high-performance memory. These powerful computing nodes are connected by a large-capacity, high-speed network.

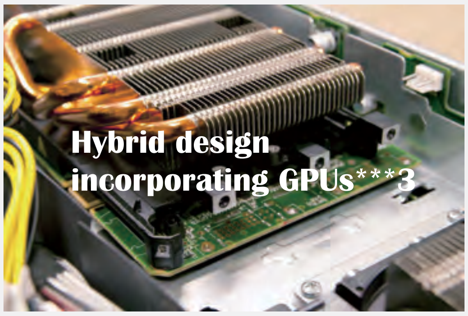

The most notable characteristic of TSUBAME2.0 is its hybrid architecture, which combines scalar processing by industry-standard CPUs with vector processing by GPUs within a single large-scale supercomputer. This conceptual shift resulted in unparalleled computing performance.

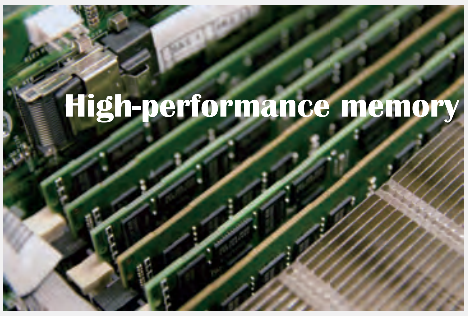

High performance is needed for both the CPUs in the computing nodes and the memory that supplies data to the GPU. The combined memory bandwidth is 720 TB/s, a 40-fold increase over the 17 TB/s of TSUBAME1.0.

The computing nodes are connected in a network that is 80 times faster than that of a high-end PC. The fiber-optic network has a bisection bandwidth of 200Tbps, which is the world top class performance, so no network congestion will occur even if all nodes are transmitting data simultaneously.

****1 Flops (mega-, giga-, tera-, peta- and exa-)***

This is an abbreviation for "floating point operations per second."
This metric is used primarily to express the computing performance of supercomputers that perform scientific and technical calculations, simulations, and so on. Mega represents 10 to the 6th power (million), giga represents 10 to the 9th power (billion), tera represents 10 to the 12th power (trillion), peta represents 10 to the 15th power (quadrillion), and exa represents 10 to the 18th power (quintillion). Accordingly, one petaflop is the ability to conduct one quadrillion floating-point calculations per second.

****2 Top500***

The Top500 Project ranks the world's fastest computer systems. The list is updated twice each year, in June and November. The LINPACK Benchmark, which performs linear algebraic calculations on the computers, is used as the standard for the rankings. Computer systems are evaluated based on their floating-point calculation performance. The Top500 rankings do not necessarily reflect all performance aspects of supercomputers, but rather are used as a metrics for evaluating supercomputers. The Top500 evaluation is conducted under the severest possible conditions — in automobile terms, they are comparable to testing performance by sending the car around and around the track at full speed.

****3 GPU***

This is the abbreviation for graphics processing unit. As the name implies, these units were originally used for image processing. In image processing, numerous calculations must be performed. The GPU contains many computing cores that are simpler and consume less power than a CPU. For this reason, the GPU has high parallel computing capability, and also offers energy-efficient performance. TSUBAME2.0 combines GPUs and CPUs to increase both computing speed and energy efficiency.

Creating the future: the challenge of TSUBAME

Applications in the field of education

TSUBAME2.0 is being used widely not only in university education but also in informatics classes in Japan.

"Supercon" — summer supercomputing contest

Supercon is a supercomputer programming contest that is held every summer, in which teams of two to three high school or technical college students compete using TSUBAME. The ten teams that pass the preliminaries work for several days to create programs to solve assigned problems. The teams compete to be the best in terms of both accuracy and speed. The contest has a relatively long history, having been held since 1995.

Both PCs and supercomputers operate on the same basic principle. However, if the PC is thought of as an ordinary passenger vehicle, the supercomputer is a racecar. The only question is whether the ideas and technology behind the programming are capable of drawing out this "super" performance. The problems in the contest are cutting-edge issues from a variety of scientific and technical fields. Although the problems are designed for high school students and presented in an easy-to-understand manner, they are quite difficult to solve.

Supercomputing Contest http://new-web.gsic.titech.ac.jp/supercon/main/attwiki/

Problems and winning schools in the past five Supercon contests	【First place】	【Second place】	【Third place】
2011(17th) Nakuron: Puzzle Game	PANAI (Kaisei Senior High School)	YAMERO (Koyo Gakuin High School)	H1T0HA (Yachiyo Shoin Gakuen Senior High School)
2010(16th) Bricklaying	zatoriku (Junior & Senior High School at Komaba, University of Tsukuba)	warosu (Junior & Senior High School at Komaba, University of Tsukuba)	nemui (Yachiyo Shoin Gakuen Senior High School)
2009(15th) Search for photographs of fixed stars	zatoriku (Junior & Senior High School at Komaba, University of Tsukuba)	potassio (Junior & Senior High School at Komaba, University of Tsukuba)	H5N1 (Yachiyo Shoin Gakuen Senior High School)
2008(14th) Simple polygon with minimum area	potassio (Junior & Senior High School at Komaba, University of Tsukuba)	fungaa (Tokyo Tech High School of Science and Technology)	maruchan (Hiroshima Gakuin Senior High School)
2007(13th) Layout of starfield monitoring base	snowdrop (Koyo Gakuin Senior High School)	YamaHa (Tokyo Tech High School of Science and Technology)	KuMa (Rakunan Senior High School)

Teaching program that enables even undergraduate students to use a supercomputer

Students at Tokyo Institute of Technology use IC cards as student ID cards. They can use this student ID card to gain access to the computer systems on campus. Naturally, TSUBAME2.0 is one of these on-campus computer systems. Therefore, from their very first year at Tokyo Institute of Technology, students can log on to TSUBAME2.0. The computer systems that are used in undergraduate education are also supported by TSUBAME2.0, forming an environment in which the PCs in the classroom and TSUBAME2.0 can be used seamlessly.

Using this environment, Tokyo Institute of Technology has initiated a program that will provide instruction in the techniques of using supercomputers for various types of research to all schools and departments. Initially, the program will be offered to graduate students, and eventually the university plans to create a curriculum that offers systematic instruction in these techniques starting from the undergraduate level onward.

Extension courses in new computing techniques

The GPUs used in TSUBAME2.0 offer diverse possibilities, but technical expertise is needed to use these capabilities.

The GPU computing seminar offered by the Global Scientific Information and Computing Center (GSIC) at Tokyo Institute of Technology includes a GPU computing (CUDA) workshop that gives students the opportunity to learn the technology from the ground up. The workshop is open to both high school and technical college students.

Global Scientific Information and Computing Center GPU Computing Workshop
http://gpu-computing.gsic.titech.ac.jp/

Toward the next generation of supercomputers

Efforts have already begun in cooperation with universities and research institutions throughout Japan to use TSUBAME2.0 to develop the next generation of supercomputers.

Coordination with next-generation supercomputers

In 2012, the K Computer, a 10 petaflop class supercomputer, will be born in Kobe, Japan. Many researchers are now working together to develop mechanisms to utilize the capabilities of the K Computer to the fullest. Computing centers at nine national universities nationwide that have supercomputers are working to create a supercomputer network with the K Computer at the top. TSUBAME2.0 will also play an active role in this network.

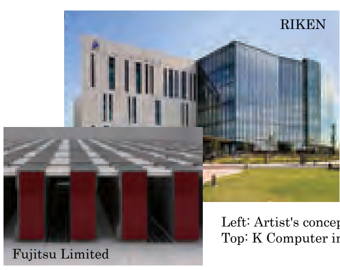

Left: Artist's conception of completed K Computer layout
Top: K Computer installed at Advanced Institute for Computational Science (AICS)

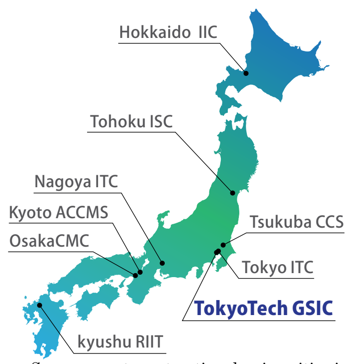

Supercomputers at national universities in Japan

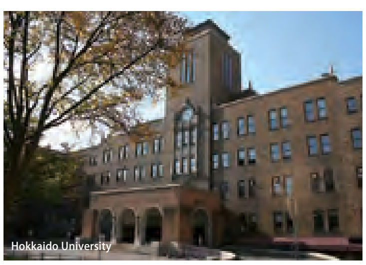

Joint research to develop TSUBAME3.0

The key to supercomputer development from now on will be finding ways to reduce power consumption. TSUBAME3.0, the successor to TSUBAME2.0, aims to reduce the power needed for supercomputer cooling to as close to zero as possible. To this end, GSIC initiated a research and development project for supercomputer cooling technologies starting from fiscal year 2011. GSIC will work with Hokkaido University (cooling technology) and the National Institute of Informatics (NII) (academic network) to pursue this joint research project.

Tokyo Institute of Technology
Global Scientific Information and Computing Center

2-12-1 O-okayama, Meguroku, Tokyo 152-8550 JAPAN
TEL：03-5734-2087　FAX：03-5734-3198　MAIL：office@gsic.titech.ac.jp
URL：http://www.gsic.titech.ac.jp/

TOKYO TECH Pursuing Excellence