CCS Report: PART I

Overview of Center for Computational Sciences

Summary of Activities

2008 – 2013

February 2014

Center for Computational Sciences
University of Tsukuba

1 Introduction

This report summarizes the main events and activities of the Center for Computational Sciences (CCS), University of Tsukuba, from April 2008 to January 2014. The report has been prepared for the external review of CCS, which is scheduled from February 18 to 20 in 2014.

The purpose of the review is to receive an outside examination of the research activities and their outcomes at the CCS in view of the founding objectives, and to incorporate the recommendations from the review into future developments of the CCS.

Since the details of the research results are prepared as a separate report, this report focuses on an overview, which includes a brief overview and the mission of the CCS (Section 2), a chronicle (Section 3), organizational aspects (Section 4), and a brief summary of the research activities and collaborations of the Center (Section 5). More details of the research activities and results in each division and group are described in the CCS Report: PART II. The future plan and vision are described in CCS Report: PART III.

2 Mission and Overview of the Center for Computational Sciences

Computational science, which is defined as the exploration of science by means of computers, is an indispensable research methodology in the basic and applied sciences, and contributes significantly to the progress of a wide variety of scientific research fields. The mission of the Center for Computational Sciences (CCS) is the promotion of "Multidisciplinary Computational Science" through enhanced cooperation between, and the fusion of, computational and computer sciences. To that end, the CCS works toward the development of high-performance computing systems and networks, conducts sophisticated simulations in a variety of scientific research fields, and endeavors to expand the frontiers of Big Data analysis and innovative information technology. The scientific research areas of our Center encompass particle physics, astrophysics, nuclear physics, nano-science, life science, environmental science, and information science. To realize the high-speed, large-scale simulations required in these research areas, the CCS works continuously to develop new state-of-the-art computing systems and networks, while also striving to advance cutting-edge knowledge in computational intelligence, computational media, and database technology. In FY2013, the CCS was authorized as one of the two prime research centers in the University of Tsukuba by "Organization for the Support and Development Strategic Initiatives". Also, the CCS was reorganized towards more strategic goals with setting up Research Project Offices, and 8 new faculty members have been recruited.

The CCS has a history of 22 years, which is shown in the chronology of CCS (Section 3). Originally formed as the Center for Computational Physics (CCP) in 1992, the CCS was reorganized and relaunched under its current name in April 2004. Currently, the CCS is known as one of the world's leading research institutes engaged in the pursuit of the abovementioned fields, and is also a notable joint-use facility for outside researchers. In all of these areas, the CCS plays a significant role in the development of the Multidisciplinary Computational Science. Since 2010, the CCS has been approved as a "national core-center" under the Advanced Interdisciplinary Computational Science Collaboration Initiative (AISCI) launched by the Ministry of Education, Culture, Sports, Science and Technology (MEXT) of Japan, and has provided the use of its computational facilities to researchers nationwide as part of the Multidisciplinary Joint-use Program. Furthermore, as part of efforts aimed at supporting collaboration in multidisciplinary computational science, the CCS welcomes applications for scientific meetings, hosts talented researchers from overseas, and strives to locate and retain short-term auxiliary supporters. All these efforts combine to create a strong bridge between domestic and international collaborations, while promoting the interchange of researchers and students. Also in 2010, the CCS was recognized as an affiliated institute of the High Performance Computing Infrastructure (HPCI) Strategic Program Field 5 "Origin of Matter and the Universe", which aims to further advance the computational sciences using "K" supercomputer technology. Furthermore, the CCS is active in the HPCI "Study on Exascale Heterogeneous Systems with Accelerators", which is being conducted to realize exascale computers. In this project, being conducted through collaboration between computational and computer scientists, the hardware and software best suited for advanced scientific research are being explored.

So far, the CCS has developed computers best suited for scientific exploration based on cooperation between computational and computer scientists, not solely providing computational facilities to such scientists. We define this approach as "Multidisciplinary Computational Science" and use it to promote research in a wide variety of fields related to computational sciences, while supporting researchers nationwide by providing computational resources and technical advice from computer scientists in joint-use programs. In 1996, we developed the massively parallel supercomputer named the Computational Physics by Parallel Array Computer System (CP-PACS), which was ranked as the No. 1 system on the Top 500 List of November 1996. Large-scale computations carried out on CP-PACS have resulted in significant progress in particle physics and astrophysics. Since then, a large-scale cluster system named the Parallel Array Computer System for Computational Sciences [PACS-CS (2006)], and a special-purpose parallel system for astrophysics, the Cosmo-simulator FIRST (2007), have been developed through interdisciplinary collaboration. The T2K-Tsukuba system (2008) is a large-scale general-purpose cluster system designed by the University of Tokyo and Kyoto University researchers through the T2K open supercomputer alliance. The "T2K" acronym itself refers to the alliance between the University of Tsukuba, the University of Tokyo, and Kyoto University. Recently, we have started a project aimed at post-petascale computing using accelerators, which is known as Highly Accelerated Parallel Advanced system for Computational Sciences [HA-PACS (2011)]. Now, we are developing the 9th generation in the PACS series, called COMA (PACS-IX), which is based on the many-core architecture. This approach has capitalized on "Study on Exascale Heterogeneous Systems with Accelerators" to realize exascale computing beyond current petaflop computing levels.

Collaborations and alliances in computational science research are important factors for promoting and accelerating interdisciplinary computational science. To that end, we have strengthened international collaborations through alliances with The University of Edinburgh in the UK and Lawrence Berkeley National Laboratory in the US. Additionally, CCS computational material scientists collaborate with Vanderbilt University and Washington University in the US, under a bilateral program sponsored by the Japan Society for the Promotion of Science (JSPS). We are also promoting the International Lattice Data Grid (ILDG), which is an international project aimed at the development of data grids for sharing lattice quantum chromodynamic (QCD) configurations worldwide. An XML-based markup language, QCDml, which describes metadata for QCD configurations and ensembles (sets of configurations with common physics parameters), has been developed. The construction of regional ILDG grids has been finalized in the US, UK, Germany, Australia, and Japan [where it is known as the Japan Lattice Data Grid (JLDG)]. Regional grid interoperability has been achieved for download operations, and valuable configurations have already been archived in the grid. The JLDG is supported by the "Cyber Science Infrastructure Project" carried out by the National Institute of Informatics as a part of grid infrastructure developments for academic research between the universities and research institutions.

As for nationwide alliances, the CCS has entered the T2K alliance, which is an alliance aimed at the acquisition of a T2K supercomputer system. In 2013, we joined with the University of Tokyo in establishing the Joint Center for Advanced HPC (JCAHPC) as the first interuniversity joint-use center in Japan, which is targeting at the construction of an unprecedentedly large-scale computational facility. Furthermore, since our university is located in Tsukuba Science City, which is home to many government research organizations, we already have firm connections to a number of major research institutions located there. These include the High Energy Accelerator Research Organization (KEK), the Advanced Institute for Science and Technology (AIST), the National Institute for Materials Science (NIMS), the Meteorological Research Institute (MRI), and the National Institute for Environmental Studies (NIES).

Picture of CCS

3 Organization and Administration

3.1 Organization & Operation

The research activities of the CCS are organized under the leadership of the Director with the support provided by the Administrative Office. The organization and the roles of meetings in the CCS are as follows.

Organization of CCS

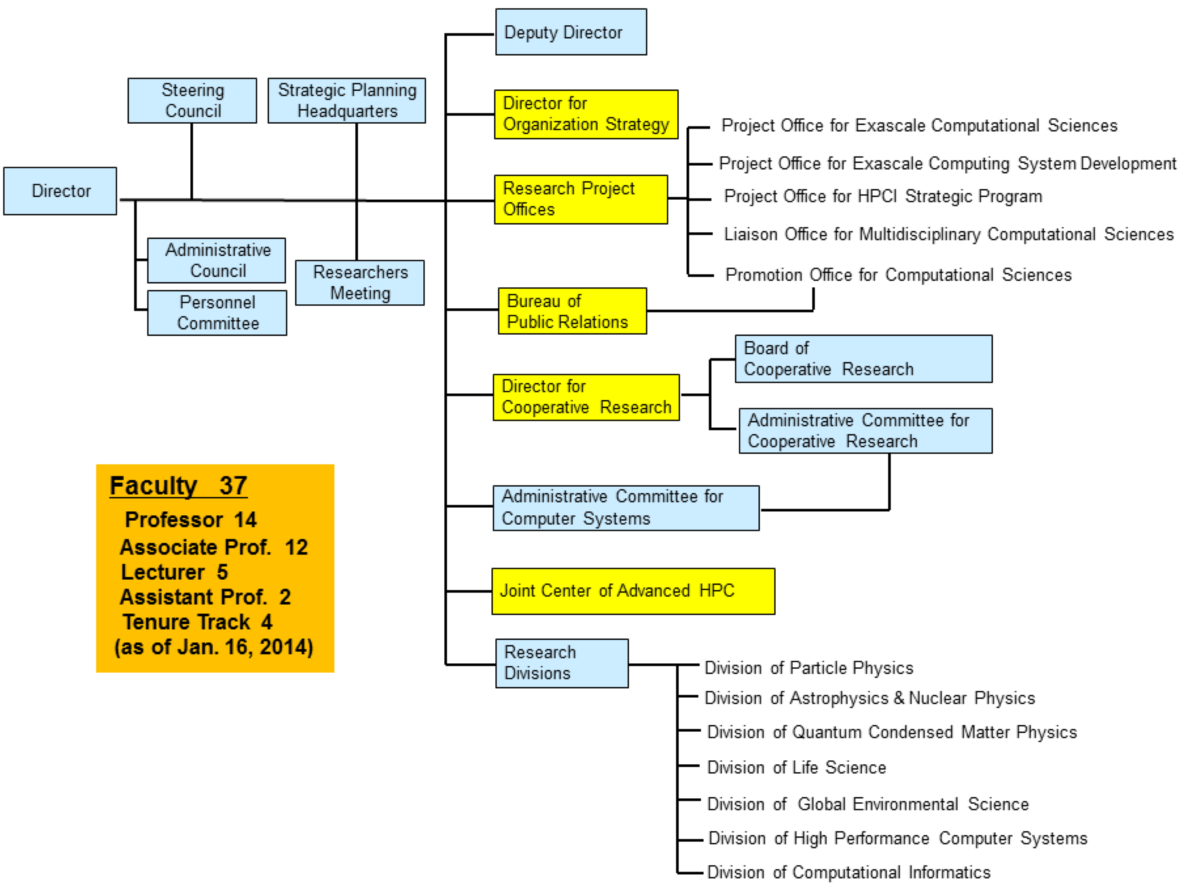

✧ Director

The Director of the Center has to hold the rank of professor of University of Tsukuba. The Director is appointed by the President of University of Tsukuba; in practice, a candidate is selected by a ballot of the Center faculty. The appointment is for two years and can be repeated up to six years. The past directors are Prof. Yoichi Iwasaki (particle physicist) from April 1992 to March 1998, Prof. Akira Ukawa (particle physicist) from April 1998 to March 2007, and Prof. Mitsuhisa Sato (computer scientist) from April 2007 to March 2013. The current director since April 2013 is Prof. Masayuki Umemura (astrophysicist).

✧ Steering Council

The council is held twice a year, and the chair is selected from outside of U. Tsukuba. The council reviews the research activities in each field and the collaborations based on the multidisciplinary joint-use program, and discusses the direction of research of the Center and matters related to joint-use of the Center facility.

✧ Administrative Council

The Director of the Center chairs the council, which holds a meeting every month. The committee discusses significant issues for operating the Center, which include matters related to the organization of the Center, selection of faculty members, budget planning, and confirmation of expenditures.

✧ Strategic Planning Headquarters

The headquarters is planning the requests for budgetary appropriations and faculty members for the enhancement of research activities as well as international collaborations.

✧ Researchers Meeting

This meeting consists of the entire Center faculty and the Associated Research Fellows. This meeting is chaired by the Director of the Center and is held every month. At the meeting, all aspects of research are discussed, such as the status of ongoing projects, procurement of equipment and operation of the Center computer system.

✧ Research Project Offices

These offices propel the missions regarding the challenge to novel computer technology, the multidisciplinary collaborations, and the nationwide/worldwide promotion of computational science. The following offices are set up:

- Project Office for Exascale Computational Sciences
- Project Office for Exascale Computing System Development
- Project Office for HPCI Strategic Program
- Liaison Office for Multidisciplinary Computational Sciences

✦ Promotion Office for Computational Sciences

(see Section 5.9 for the details)

✧ Administrative Committee for Cooperative Research

The CCS is calling for applications twice a year for the joint-use program of computer facilities. Since 2007, these activities have been reinforced by enhanced links with computer sciences through the promotion of the Multidisciplinary Cooperative Research Program (MCRP). Since 2010, the Center has been recognized as a national core-center under the Advanced Interdisciplinary Computational Science Collaboration Initiative (AISCI) by the MEXT. The administrative committee manages this joint-use program.

(see Section 5.8 for the details)

✧ Bureau of Public Relations

In 2010, the bureau was initiated aimed at improving relations with society in order to advance the mission of the CCS and to expand the public outreach. The bureau organizes the opening to the public every year, and also guides visitors at various kinds of events.

(see Section 5.9 for the details)

✧ Research Divisions

The specific researches in each field are promoted in the following divisions:

- ✦ Particle Physics
- ✦ Astrophysics & Nuclear Physics
- ✦ Quantum Condensed Matter Physics
- ✦ Life Science
- ✦ Global Environmental Science
- ✦ High Performance Computer Systems
- ✦ Computational Informatics

(see Part II for the details)

3.2 Faculty of the Center

At University of Tsukuba, all faculty members belong to one of seven graduate schools. The faculty members of the Center "belong" to a graduate school and "work" at the Center. The affiliation of a faculty member is written as "Graduate School (Center for Computational Sciences)".

The faculty members as of January 2014 are listed in Appendix A.

3.3 Faculty Selection

The Center appoints new faculty members when openings become available among the Center faculty positions. The selection procedure is as follows.

- **Choice of field**

When an opening of a position becomes available, the field of the position is chosen by the Personnel Committee based on consideration of the strategic purposes of the Center and the balance among research fields.

- **Selection of candidate**

Once the field is chosen, the actual selection process is handled by the Faculty Selection Committee under the Personnel Committee. The opening is publicly announced to invite applications. Candidates are initially screened by the Faculty Selection Committee. This screening includes reading the candidates' main papers and interviews. When the screening at the committee converges toward a candidate, the conclusion is reported to the Personnel Committee. If the Personnel Committee approves the candidate, the result is reported both to the Administrative Council and the University Personnel Committee for official approval.

4 Chronicle of CCS

The chronicle of the Center for Computational Sciences goes back to 1992, when the Center for Computational Physics (CCP) was founded at the University of Tsukuba. The purpose of the CCP was (i) to carry out research and development of parallel computers suitable for large-scale computational physics calculations, and (ii) to carry out research with the developed system in the area of particle physics, astrophysics, and condensed matter physics. In particular, the CCP was to serve as the base for the project entitled "Research of Field Physics with Dedicated Parallel Computers" (FY1992-FY1996) funded by the Ministry of Education of the Japanese Government. This project developed a massively parallel system CP-PACS (Computational Physics with Parallel Array Computer System). This system achieved 368.2 Gflops for the Linpack Benchmark in September 1996, and was ranked as the No. 1 system in the November 1996 Top 500 List. CCP had a faculty of 10 researchers in four divisions: Computational Particle Physics, Computational Astrophysics, Computational Condensed Matter Physics and Parallel Computer Engineering. The simultaneous existence of physics and computer science faculties within the CCP, and actual close collaborations between them for the development of parallel computer systems and their usage, marked the unique feature of the Center. In fact, this interdisciplinary approach goes back to the early 1980s, when research and development of parallel computer systems started at University of Tsukuba under Professor Tsutomu Hoshino. The CP-PACS is the 6th generation of PACS series. In the late 1990s, large-scale calculations using CP-PACS produced a number of well-known results. The major scientific achievements are as follows:

1. Calculation of the quenched light hadron spectrum in lattice quantum chromodynamics (QCD). (S. Aoki et al., Physical Review Letters 84 (2000) 238)
2. First systematic lattice QCD calculation including light dynamical up and downquarks, showing that up and down quark masses are much lighter than previously thought. (S. Aoki et al., Physical Review Letters 85 (2000) 4674)
3. First astrophysical radiation hydrodynamical calculation of the reionization of the Universe including interaction of radiation and matter. (T. Nakamoto et al., Monthly Notice of the Royal Astronomical Society 321 (2001) 593)
4. First many-body quantum mechanical calculation of the phases of solid hydrogen. (H. Kitamura et al., Nature 404 (2000) 259)

In April 2004 the Japanese national university system, of which University of Tsukuba is a member, underwent a major transition. While previously the national universities were part of the Japanese government and were strictly controlled by the Ministry of Education, Science, Culture and Sports, their legal status was changed to that of independent institutes. While the budget still came from the government, each university was granted much wider freedom of action on its own. The reason behind this change was pressure toward reformation of the Japanese government, including not only a reduction of employees and government spending, but also reformation of the university system to be more effective in the era of globalization and worldwide competition. Preparations for the transition started in FY2002 at all national universities across Japan. The Center for Computational Physics considered the transition as a welcome opportunity to reexamine its future plan. While the Center for Computational Physics, with its small faculty of a dozen members, had functioned well in pursuing research projects within a limited area of fundamental physics, it was strongly felt that the interdisciplinary approach toward computational physics pursued at the Center could be extended to wider areas of science, and that the Center had a potential to significantly expand its activity by doing so. Discussions with the Executive Office of the University on a possible reorganization and expansion started in the early summer of 2002. The Executive Office basically welcomed the Center's proposal, and provided strong support during negotiations with the Ministry of Education, Science, Culture and Sports to secure governmental approval from the winter of 2002 to spring of 2003.

The Center for Computational Physics was formally reorganized as the Center forComputational Sciences on 1 April 2004. The faculty expanded threefold from 11 to 31 academic members, and the following reorganization and expansion of divisions were accomplished:

- Division of Particle Physics and Division of Astrophysics were combined into a single division with an increased faculty of 7 members.
- Division of Condensed Matter Physics and Division of Biophysics were combined and significantly expanded to 10 faculty members in the Division of Materials and Life Sciences.
- A new Division of Global Environment and Biological Sciences was introduced with 3 faculty members.
- Division of Parallel Computer Science was expanded from 3 to 5 faculty members, and was renamed Division of High Performance Computing Systems.
- A new Division of Computational Informatics with 6 faculty members was set up to carry out research on computational intelligence and media.

Of the 20 members added to the original 11 for the Center for Computational Sciences, 14 were faculty members who were actively working in various areas of computational science at graduate schools of the university, while the remaining 6 were approved by the university and filled by new recruitments in 2004-2005.

Gordon Bell Prize 2011 for Peak Performance (University of Tsukuba, University of Tokyo, RIKEN) in an atomic state simulation at actual semiconductor device scale. Gordon Bell Prize 2012 (University of Tsukuba, RIKEN, Tokyo Institute of Technology) for scalability and sustained performance in a world's largest scale of a dark matter simulation. In 2013, we established the Joint Center for Advanced HPC (JCAHPC) with the University of Tokyo, which is aiming at the construction of an unprecedentedly large-scale computational facility. This year, the CCS was authorized as one of the two prime research centers in the University of Tsukuba by "Organization for the Support and Development Strategic Initiatives". Also, the CCS was reorganized towards more strategic goals with setting up Research Project Offices.

Chronology of CCS

1992	April	CP-PACS Project begins (5-year project). Center for Computational Physics (CCP) founded (10-year term).
1993	August	Computer building completed
1995	March	First research building completed
1996	September	CP-PACS (2048 PU) completed and installed
	November	Ranked as No. 1 on the Top 500 World Supercomputer List
1997	April	Research for the Future Program "Development of Next-Generation Massively Parallel Computers" begins (5-year term)
2002	April	Second 10-year term of the Center for Computational Physics begins
2004	April	CCP is reorganized, expanded and relaunched as the Center for Computational Sciences (CCS)
2005	April	Development of Massively Parallel Cluster PACS-CS in the project begins (3-year term)
2007	March	Second research building completed
	April	Cosmo-Simulator FIRST completed and installed
2008	June	Operation of T2K-Tsukuba begins.
2010	April	Approved under the Advanced Interdisciplinary Computational Science Collaboration Initiative (AISCI) CCS is reorganized from five to seven divisions
2011	November	Gordon Bell Prize 2011 for Peak Performance (University of Tsukuba, University of Tokyo, RIKEN)
2012	February	HA-PACS starts operation
	November	Gordon Bell Prize 2012 (University of Tsukuba, RIKEN, Tokyo Institute of Technology) for Scalability and Sustained Performance
2013	March	Joint Center for Advanced HPC established in alliance with the University of Tokyo
	April	Reorganized towards strategic goals
	August	Authorized as one of the two prime research centers in U. Tsukuba

5 Overview of Research Activities from 2008 to 2013

5.1 PACS-CS Project

5.1.1 Motivation and Objective

After the success of CP-PACS Project at CCP (Center for Computational Physics, former organization before CCS) from FY1992 to FY1996, commodity technology had been rapidly introduced as a major part of HPC solutions. Especially, the large scale PC clusters were widely used instead of proprietary MPP system by their high performance/cost ratio. A number of highly parallelized applications were ported to them as cache-aware scalar codes. In Japan, the latter half research period of RWCP under MITI (1992-2002) mainly focused on R&D on the operating system and middleware named SCore for massively PC clusters.

In CCS, we also considered to introduce PC cluster technology for the follow-up system to contribute the next generation applications, however, one of the most serious issues on PC clusters was the lack of bandwidth both on memory and interconnection network provided by commodity PC server technology. As a solution, we designed a new type of PC cluster based on the concept of "bandwidth-aware PC cluster by the combination of commodity hardware with specially designed HPC middleware." Most of PC clusters in those days focused on to increase the peak performance with multi-core/multi-socket solution regardless the bandwidth balance. As a more powerful solution for wide variety of computational scientific problems, we decided to build a PC cluster with small number of cores and network trunk to enhance the memory and network bandwidth balanced with the peak FLOPS. The machine is named PACS-CS (Parallel Array Computer System for Computational Sciences) and funded by MEXT as three year project from FY2005 to FY2007.

While PACS-CS aims to promote wide variety of computational sciences, we specially focused on the development of first-principles quantum simulations in materials and life sciences capable to treat O(10,000) atoms, thereby allowing an exploration of the connection between the spatial structure and function characteristic of nano materials such as carbon nano-tubes and large bio molecules such as proteins. Another emphasis is advancement of full QCD simulations in particle physics and astrophysics research which forms the core of our fundamental understanding of the creation and history of our Universe.

5.1.2 Project Formation

PACS-CS is a massively parallel cluster to be developed to carry out computational breakthroughs described above. It targets the peak performance of 14.3 Tflops with 2560 nodes connected by a 3-dimensional Hyper-Crossbar Network. The operation started on July 2006.

The development of the PACS-CS computer was being carried out through a close collaboration of the computer scientists and scientists of the target areas (materials and life sciences, particle and astrophysics, geoenvironmental and biological sciences). A working group containing members from the two areas were set up in each of the target areas, and regular meetings were held in which technical points for the hardware and software development of the PACS-CS were discussed and decided. Optimization of applications codes, for each node and its parallelization, was an important issue. The PACS-CS design details were being fixed step by step through these procedures. Fig. 5.1.2 shows the schedule on the design and preparation on hardware system and application codes under the collaboration of system and application scientists at CCS.

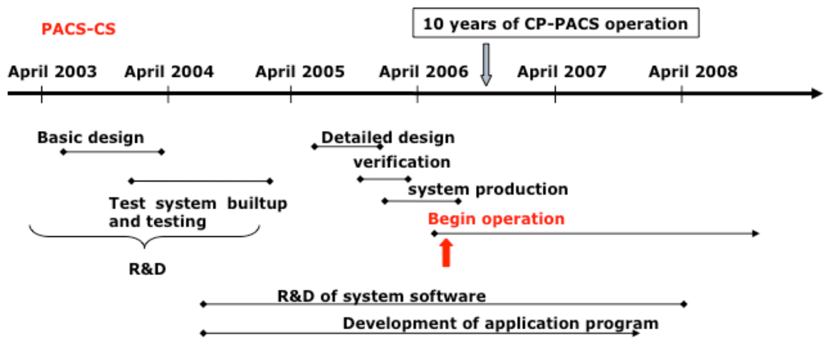

Fig. 5.1.2 Schedule of design, preparation and operation of PACS-CS

For manufacturing of the PACS-CS computer, Hitachi Ltd. was selected as the primary system construction including special mother board to realize the concept of bandwidth-aware cluster. For developing the Hyper-Crossbar Network driver software, Fujitsu Ltd. was selected. The Project members are working in close collaboration with these companies to develop the computer.

5.1.3 System Overview

Our choice to resolve the issues on bandwidth-awareness, we equipped each node with a single processor and connect it with the fastest bus available to the main memory with a matched I/O speed. Specifically, we selected FSB800 for the bus and two banks of PC3200 DDR2/400 memory, which means the memory to processor bandwidth of 6.4GByte/s. For the processor, we used Intel Low Voltage Xeon for a better error check and correction as well as wider variety of chipset to carry enough I/O bandwidth than the Pentium4 series. The frequency of CPU clock is 2.8GHz (5.6Gflops peak) which is not the highest one in those day’s technology since higher rates would not provide higher effective performance, and the low-voltage version to suppress power consumption.

The interconnection of 2560 nodes forms a 3-D HyperCrossbar Network topology to combine both characteristics of mesh-torus network and crossbar network. It is conceptually same as that used in CP-PACS, but on PACS-CS, all the network links are made by the trunk of commodity GbE and all the swithes are also commodity many-port GbE switches. Each node is equipped with 6 of physical links of GbE and the mother board is specially designed to support single-socket CPU, two banks of DDR2 memory and 8 ports of on-board GbE (six for data communication, one for system control and one for IPMI), and two separate nodes are implemented in a space of 1-U rack-mountable chassis. Fig. 5.1.3 shows the system construction and the chassis containing two computation nodes.

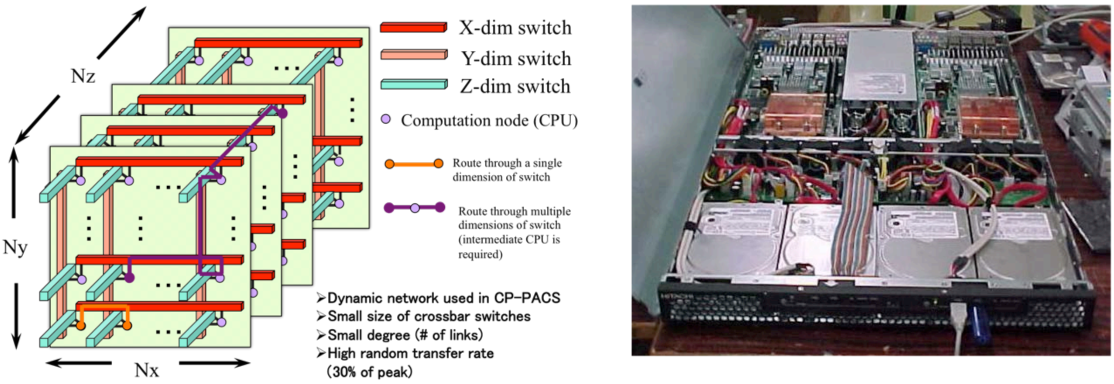

Fig. 5.1.3 3-D HyperCrossbar Network construction (left) and 1-U sized chassis with two computation nodes of PACS-CS (right)

With a specially designed network driver for multi-trunk GbE to manage 3-D HyperCrossbar Network named PM-Ethernet/HXB, any 3-D orthogonal problem mapping can be fit to the physical construct of PACS-CS to reduce the communication overhead as well as 3-D simultaneous data send/receive supported by 6 of GbE links which provides higher bandwidth than SDR InfiniBand. Fig. 5.1.4 shows the entire system overview of PACS-CS.

Fig. 5.1.4 PACS-CS overview (left) and GbE trunk network with 3-D connection (left) where blue, red and green colored cables correspond to X-, Y- and Z-dimension of network, respectively.

Table 5.1.1 shows the overall specification of PACS-CS.

Table 5.1.1 Specification of PACS-CS

# of nodes	2560 (16 x 16 x 10)
peak performance	14.3 Tflops
node configuration	single CPU / node
CPU	Intel LV Xeon EM64T, 2.8GHz, 1MB L2 cache
memory	2GB/node (5.12 TB/system), DDR2 interleaved
network for parallel processing	3-dimensional Hyper-Crossbar Network
link bandwidth	one-sided: 250MB/s/dim. one-sided: 750MB/s (3-D simultaneous trans.)
local HDD	160 GB/node（RAID-1）
total system size	59 rack
power consumption	550 kW

The operating system of PACS-CS is Red Hat Enterprise Linux equipped with SCore cluster management middleware developed by RWCP. For distributed memory programming we employed MPICH2 based on PM-Ethernet/HXB. Intel Fortran/C/C++ compilers are also introduced.

PACS-CS was ranked at #20 in the TOP500 list issued on June 2006 with 10.35 TFLOPS of HPL score.

5.1.4 Operation

From July 2006, PACS-CS starts its full size operation and especially dedicated for the most important application development and running as planned in PACS-CS Project. From April 2007 to September 2011, PACS-CS has been fully dedicated to the Collaborative Interdisciplinary Program of CCS for nationwide open use for advanced computational sciences.

The operation of PACS-CS was very stable and the average utilization ratio is more than 80% through the entire operation years. On March 2011, however, by the earthquake attack in north-east Japan, it was down due to lack of electricity. It took about one month for checking the entire system including UPS backup and recovered on the end of April. For the last three months (July to September 2011), the system could not be available on daytime for electricity consumption shift, and we shut-down and booted-up the system every day. Although such a frequent power on/off operation provides heavy stress to the system especially on the file system, PACS-CS continued its stable operation without any special trouble.

5.2 T2K-Tsukuba Project

5.2.1 Motivation and Objective

In University of Tsukuba, the operation of supercomputers has been distributed to two organizations, Center for Computational Sciences and Academic Computing & Communications Center (ACCC) before FY2007. CCS had been mainly operating PACS series machines (CP-PAPCS and PACS-CS) for nationwide utilization while ACCC had been operating Fujitsu's vector machine for internal use in the university. After the last system in ACCC, Fujitsu VPP-5000, the university decided to gather all the supercomputer resources and operation to CCS for efficient and effective procurement and operation of university's supercomputer resources.

For the procurement of next supercomputer after VPP-5000, CCS decided to collaborate with University of Tokyo and Kyoto University for the strategic plan of these three university's supercomputer replacement. This collaboration was based on an alliance named "T2K Alliance" where T2K stands for Tsukuba-Tokyo-Kyoto. Under the alliance, three universities introduced the same architecture of large scale PC clusters for general utilization by nationwide researchers. Actually, after these systems started the operation, MEXT decided to start a new nationwide supercomputer utilization program named HPCI and these machines took roles to dedicate the same system architecture of machines where users can easily port their codes for these machines. These machines were introduced to T2K Universities at the same time, June 2008, with individual system names of T2K-Tsukuba, T2K-Todai and T2K-Kyoto. ("Todai" is the nickname of University of Tokyo in Japanese.)

The procurement itself was performed individually at T2K Universities. Finally, three T2K machines with the same basic system configuration but different node counts were delivered. The system vendors of them were different, Appro International for T2K-Tsukuba, Hitachi for T2K-Todai and Fujitsu for T2K-Kyoto. The rental period of each of them also differs as 4, 5 and 6 years for Kyoto, Tsukuba and Tokyo, respectively. However, the period for T2K-Tsukuba was extended one year to fit to our next future plan. Finally, the operation period of T2K-Tsukuba was almost 6 years and it will be finished on February 2014.

5.2.2 System Overview

At the timeframe of three T2K machines' procurement, the commodity CPU technology entered to the era of quad-core per chip, and several high-end CPUs were equipped with the feature to connect more than two sockets of them in a node. To accept a number of traditional applications which require a large capacity of memory as well as CPU cores per node, we decided to configure a computation node with at least 16 CPU cores and 32GB of memory.

The most appropriate CPU for our purpose was AMD Opteron with quad-core configuration known as “Barcelona” series. For the cost effectiveness, we decided to introduce 2.3 GHz quad core, quad socket CPUs per node.

The most cost effective interconnection solution to configure a system with hundreds of nodes was Ssystem Area Network (SAN) such as InfiniBand or Myrinet. On T2K-Tsukuba, we employed quad rail trunk of InfiniBand DDR to provide 8 GByte/sec of theoretical peak bandwidth. For the freedom of node management and scheduling, we employed full-bisection bandwidth of Fat-Tree network supporting quad-rail connection (four InfiniBand HCA cards per node). To save the total cost on network switches and cables, the Fat-Tree network consists of a number of medium size switches with 24 ports per switch, connecting them in 3-stage network without introducing big spine switch.

Fig. 5.2.1 shows the system overview of T2K-Tsukuba. Fig. 5.2.2 shows the computation node and its configuration. Fig. 5.2.3 shows the network configuration of 3-stage Fat-Tree.

Fig. 5.2.1 Overview of T2K-Tsukuba

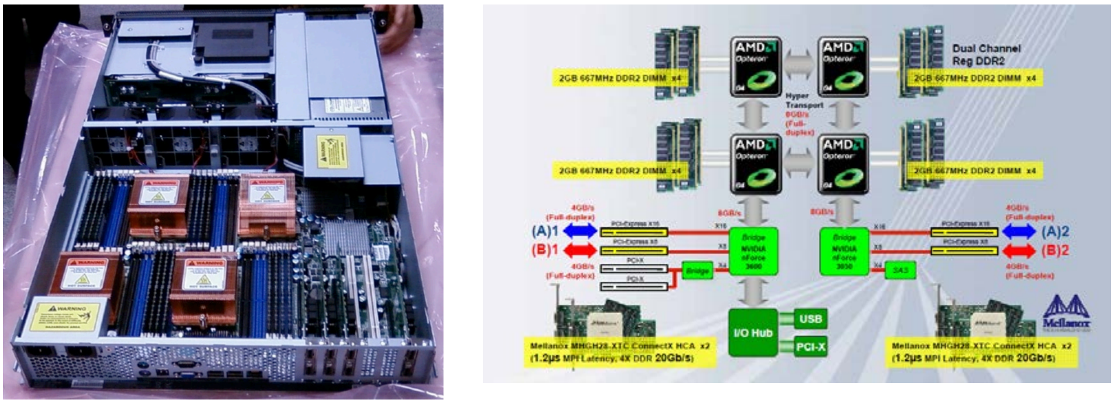

Fig. 5.2.2 Computation node (left) and Fat-Tree network construction (left)

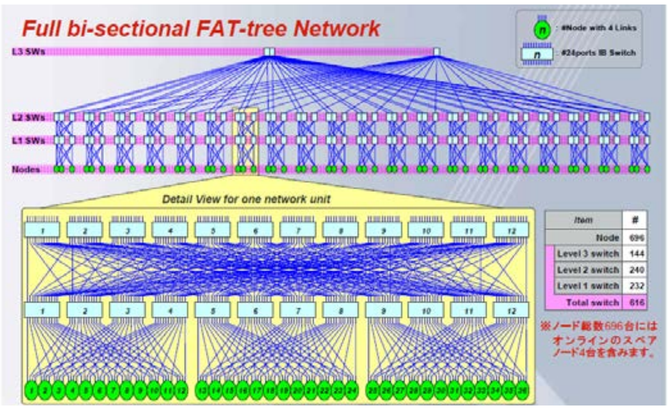

Fig. 5.2.3 3-Stage Fat-Tree network configuration

The system is equipped with a redundant file system shared by all the computation nodes. It is based on RAID6 storage system with Lustre cluster file system, which is provided by Data Direct Network.

Table 5.2.1 shows the basic specification of T2K-Tsukuba to provide 92 TFLOPS of theoretical peak performance.

Table 5.2.1 Overall specification of T2K-Tsukuba

Computation node	CPU	AMD quad-core Opteron “Barcelona” 2.3GHz x4 socket
	# of cores	16 (4 cores / CPU)
	Main memory	32GB (DDR2 800MHz x 8)
	Network HCA	InfiniBand DDR x 4 rail (Mellanox)
	Peak performance	147.2 GFLOPS
	Peak network b/w	8 GB/s
Number of computation nodes		648
Interconnection configuration		3-Stage Fat-Tree with full bisection b/w
Peak performance		95.4 TFLOPS
Network bisection b/w		5.18 TB/s
Shared file system		RAID6 with Lustre (DDN)
File system capacity		800TB (user), 1PB (physical)

The operating system of T2K-Tsukuba is Red Hat Enterprise Linux. Most of users of the system write their codes in a hybrid manner with MPI+OpenMP. We employed Intel OpenMP C/C++/Fortran compiler and MVAPICH2 library developed in Ohio State University since it is equipped with the feature of network trunk with multiple rails of InfiniBand links. The entire system is managed by ACE (Appro Cluster Engine) and SGE batch/scheduler.

T2K-Tsukuba was ranked at #20 in the TOP500 list issued on June 2008 with 76.46 TFLOPS of HPL score.

5.2.3 Operation

After four months of shakedown and test use from June to September 2008, T2K-Tsukuba has been dedicated to the Collaborative Interdisciplinary Program by CCS from October 2008 to February 2014, and also partially used under fee-charged program by several users groups. From October 2012, it is partially dedicated as an HPCI resource which is a nationwide supercomputer utilization program under MEXT. Through these three programs, the total number of users reached more than 400.

For 5 years and 9 months of continuous operation, T2K-Tsukuba has been quite stable without any serious fault even at the time of earthquake attack on March 2011 except its shut down due to lack of electricity. Thanks to the feature of flat configuration with 648 computation nodes and full-bisection bandwidth of Fat-Tree network, the freedom of node allocation and scheduling is very high and the high utilization ratio with more than 80% in average has been kept.

Originally, the operation period of T2K-Tsukuba was scheduled as finishing on May 2013. However, the post-project of T2K in University of Tsukuba and University of Tokyo has been launched on April 2013 (Kyoto University does not join at this time), and we agreed to introduce our next system in FY2015 (see the section on JCAHPC Project). In order to adjust the machine operation period, we decided to extend the operation of T2K-Tsukuba 10 months then introduce a "bridge machine" for the preparation of JCAHPC joint system installation. As a result, T2K-Tsukuba will be shut down in the end of February 2014.

5.3 FIRST Project

This project is the collaboration between astrophysicists and computer scientists, aiming at the elucidation of the evolution of first generation objects through large-scale radiation hydrodynamic (RHD) simulations with a high-performance hybrid computer system. The project is initiated by *Grant-in-Aid for Specially Promoted Research* in 2004 (approved by The Ministry of Education, Culture, Sports, Science and Technology (MEXT) in Japan), and continued by *Grant-in-Aid for Scientific Research (S)* in 2008. The duration of the project is 10 years (2004~2013) with the total budget of JP405.4 million (US$4.2 million). Under this project, we have developed a new type of hybrid computer dedicated for astrophysical RHD, called *FIRST* simulator, in the collaboration with the Division of High Performance Computing Systems.

◆ *FIRST* Simulator

For the realization of 3D RHD, we have built up a new generation of hybrid computer system, called *FIRST* (Fusional Integrator for Radiation-hydrodynamic Systems in Tsukuba University) simulator. The *FIRST* simulator is a hybrid PC cluster, where a newly-developed board for gravity calculations, called Blade-GRAPE, is embedded in each node. The Blade-GRAPE is composed of four GRAPE-6 chips that are dedicated chips based on the pipeline architecture, and is

Blade-GRAPE X64

FIRST simulator (256 node, 36.1TFLOPS)

designed for PCI-X bus in a PC cluster. The theoretical peak performance of Blade-GRAPE is 136.8GFLOPS. Each board has 16MB memory and can calculate the self-gravity of 260,000 particles simultaneously at the maximum. The Blade-GRAPE is directly connected via PCI-X bus, and occupies the space of two PCI-X bus slots. The electric power supply is from the PCI-X bus (3.3V) as well as from the cluster server board, +12V (54W). Each server PC is equipped with multi-port Gigabit Ethernet NIC to be connected to a special interconnection network with commodity Ethernet switches.

The first version of Blade-GRAPE works with 32bit and 33MHz (PCI). Then, it has been improved to 64bit and 100MHz version (PCI-X), Blade-GRAPE X64. Using Blade-GRAPEs, we have constructed a 256 node hybrid PC cluster system, that is, *FIRST* simulator. The system possesses 16 Blade-GRAPE boards and 224 Blade-GRAPE X64 boards. The host PC cluster node is a 2U-size of 19-inch rack mountable server PC (HP ProLiant DL380 G4) that has dual Xeon processors in SMP configuration. The peak performance of *FIRST* simulator is 36.1TFLOPS, where the host PC cluster is 3.1 TFLOPS and the Blade-GRAPEs are 33 TFLOPS. All nodes are connected uniformly with each other via multi-port Gbit ether interconnect switch. The total memory of *FIRST* simulator is 1.6TB. Also, the *Gfarm* Grid file system, which is the commodity-based distributed file system that federates local disk of each node

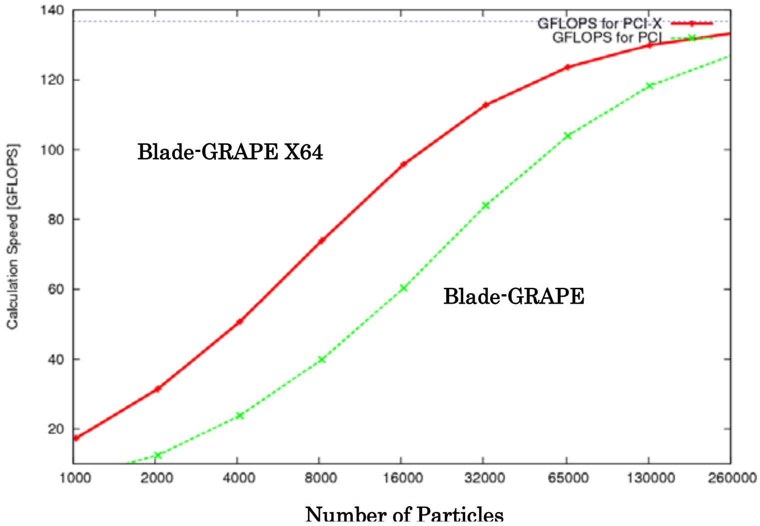

The dependence of performance on the number of particles in a board is shown for Blade-GRAPE and Blade-GRAPE X64.

(http://datafarm.apgrid.org/index.en.html), is installed. With *Gfarm*, the storage of 22TB is available as a seamless file server.

With *FIRST* simulator, we have investigated the acceleration of gravity calculations by Blade-GRAPEs. First, we have tested the two types of Blade-GRAPEs. The main deference between Blade-GRAPE and Blade-GRAPE X64 is the bus clock, where a Blade-GRAPE works at 33MHz and a Blade-GRAPE X64 does at 100MHz. The dependence of performance on the number of particles in a board is shown in the above figure. If 260,000 particles are calculated with a board, the performance of Blade-GRAPE is 93% of the peak speed, and that of Blade-GRAPE X64 is 98%. This weak dependence shows that for this number of particles the communication with the host computer is negligible compared to the force calculations in GRAPE chips. But, the performance is reduced with a decreasing number of particles. For 16,000 particles, the performance of Blade-GRAPE is 44%, and that of Blade-GRAPE X64 is 74%. For this number of particles, the communication with the host computer is noticeable, and therefore Blade-GRAPE X64 has a roughly double performance of Blade-GRAPE.

Next, in a simulation using 256^3 particles, we have measured the run time with and without Blade-GRAPEs. The results are shown in the figure below. In this simulation, the dominant part of calculations is the self-gravity calculations. The acceleration turns out to be 5 times with Blade-GRAPE, and 10 times with Blade-GRAPE X64. In this simulation, roughly ten thousand particles are sent to each Blade-GRAPE board on average. Therefore, the performance with Blade-GRAPE X64 is roughly twice of that with Blade-GRAPE due to the communication bandwidth.

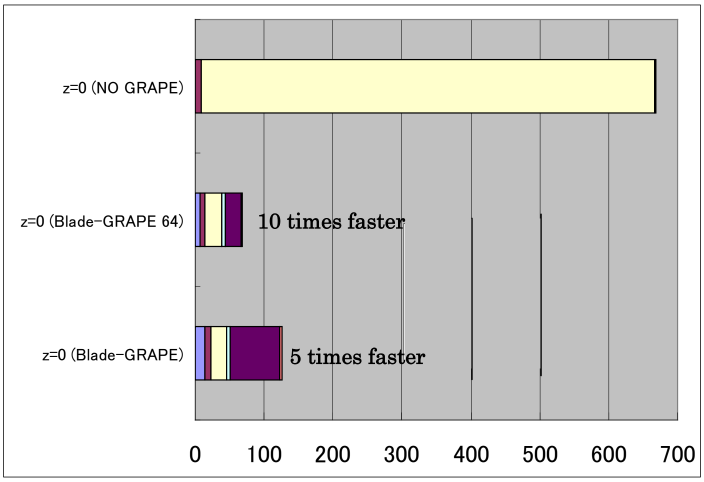

The acceleration of calculations for 256^3 particle simulations.

5.4 HA-PACS Project

5.4.1 Motivation and Objective

After the end of PACS-CS Project, the accelerated computing represented by GPGPU technology has been focused as one of the solutions to high performance/power ratio in the future HPC systems. GPU became the most inexpensive device to provide very high peak performance and memory bandwidth with relatively low power consumption especially after NVIDIA started to provide CUDA environment for easy and portable programming method. However, GPU solution implies a number of problems such as low bandwidth on communication with CPU and external communication devices due to the bottleneck by PCIe (PCI Express) channel, limited memory capacity, long latency of communication, complicated programming, etc.

As a solution to communication performance issues and programmability, we launched a new project named HA-PACS (Highly Accelerated Parallel Advanced system for Computational Sciences) as MEXT supported project from FY2011 to FY2013. We focused on the commodity communication channel of PCIe to connect all the peripherals including GPU, Ethernet, HCA or HDD. PCIe is originally a communication channel under master control by host CPU (RootComplex) where all the peripheral devices work as slaves (Endpoints). The communication among GPUs on different computation nodes must be performed through CPU memory and its support. For example, MPI function is called for external communication among nodes by host CPU with before/after communication between CPU and GPU within the node. If we have an active and intelligent PCIe switch instead of passive switch, it is theoretically possible to enable the active communication between GPUs without overhead of data copy, and low latency and high bandwidth on inter-GPU communication are realized. For the next generation of large scale GPU clusters, such a sophisticated interconnection network is required, especially for strong scaling parallelization.

This concept is named TCA (Tightly Coupled Accelerators) as the basic idea of HA-PACS system. Fig. 5.4.1 shows the difference between ordinary inter-node communication among GPUs and communication based on TCA.

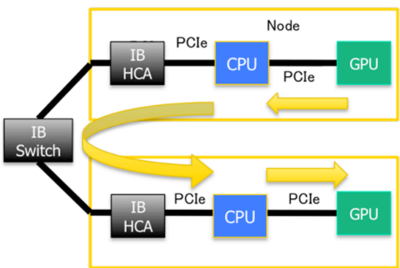

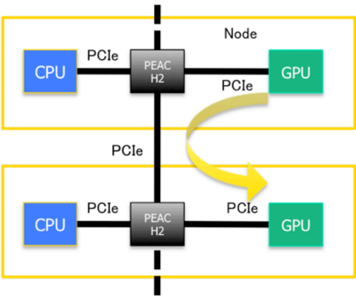

Fig. 5.4.1 Traditional inter-node communication between GPUs on GPU clusters (left) and the method in TCA (right)

The objective of HA-PACS project is (1) establishing the elementary technology to enable accelerator device to device (ex. GPU) direct communication on intelligent PCIe switch, and (2) developing the basic algorithm and code for several specially focused application fields under TCA concept. Especially, the second objective is important as the basic algorithmic enhancement or renewal of traditional simulation codes to be ready for next generation supercomputers with advanced accelerated computing. As the Feasibility Study on next generation supercomputer, CCS and collaborative team are proposing very wide SIMD accelerator architecture for the computing element (see the section of Feasibility Study in this report), and we need a drastic model/algorithm change for next generation. We consider that the algorithms modified based on TCA concept take important roles there.

Not only for development of hardware/software on TCA concept, we also introduce a large scale GPU cluster system for basic development, evaluation and produc-running of various codes developed in CCS. Since TCA architecture research is undergoing, we introduce two systems in HA-PACS Project. The first one is named HA-APCS Base Cluster and the latter is named HA-PACS/TCA.

5.4.2 System Overview

Since the development of the communication devices based on TCA concept takes approximately one year, we first installed HA-APCS Base Cluster (“Base Cluster” hereafter) as a commodity GPU cluster with the latest technologies. It was delivered and started the operation on February 2012. This delivery timing allowed us to introduce the advanced CPU of Intel Xeon E5-2670 (SandyBridge) as the host CPU and NVIDIA M2090 (the last version of Fermi) as GPUs on each node. We employed two sockets of CPUs and four GPUs on each node because SandyBridge architecture provides PCIe gen.3 x40 lanes at maximum per CPU and two sockets of them theoretically provide x80 lanes in total. That is enough to connect four GPUs each of them requires PCIe gen.2 x16 lanes. We also employed dual rail InfiniBand QDR HCA supported by PCIe gen.3 x8 lanes to provide 56Gbps of maximum performance for interconnection network.

Fig. 5.4.2 shows the overview of Base Cluster and TCA part, and Fig. 5.4.3 shows the picture of computation node and its block diagram. Table 5.4.1 shows the basic specification of Base Cluster.

Fig. 5.4.2 Overview of HA-PACS Base Cluster and TCA part (five racks of front side are for TCA part and others are for Base Cluster)

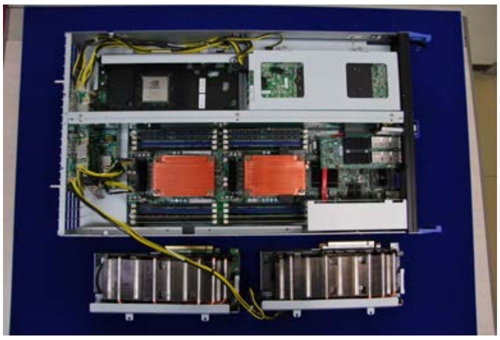

Fig. 5.4.3 Computation node and its block diagram of HA-PACS Base Cluster

HA-PACS/TCA was introduced later on October 2013 with a dedicated PCIe extension card for TCA architecture and connected to the central InfiniBand switch with Base Cluster. The basic node configuration of TCA part is same as that of Base Cluster, except CPU is Intel Xeon E5-2670v2 (Ivy Bridge) and GPU is NVIDIA K20X (Kepler) for greater performance enhancement. Main memory bus speed is also upgraded to 1866MHz. In total, the peak performance of computation node in TCA part is 5.69 TFLOPS (348 GFLOPS by CPU and 5320 GFLOPS by GPU), and the total peak performance of TCA part with 64 nodes reaches 364 TFLOPS. We can run Base Cluster and TCA part as a single system although these CPUs and GPUs are different, and the entire system performance is 1.166 PFLOPS as peak.

Each computation node of TCA part consists of a PCIe card named “TCA Board” to contain an FPGA chip and its peripherals including three ports of PCIe extended connection link to connect with other nodes. We employed FPGA technology to realize TCA architecture because we need to enhance its function and performance.

Table 5.4.1 Basic specification of HA-PACS Base Cluster

Computation node (Base Cluster)	CPU	Intel Xeon E5-2670, 2.6 GHz x2 socket
	# of cores	16 (8 cores / CPU)
	GPU	NVIDIA M2090 (Fermi) x4
	Main memory	128GB (DDR3 1600MHz x 16)
	GPU memory	24GB (6GB / GPU)
	Peak performance	332.8 GFLOPS (CPU) + 2660 GFLOPS (GPU)
Computation node (TCA)	CPU	Intel Xeon E5-2670v2, 2.8 GHz x2 socket
	# of cores	20 (10 cores / CPU)
	GPU	NVIDIA K20X (Kepler) x4
	Main memory	128GB (DDR3 1866MHz x 16)
	GPU memory	24GB (6GB / GPU)
	Peak performance	448 GFLOPS (CPU) + 5240 GFLOPS (GPU)
Computaion node (common)	Network HCA	InfiniBand QDR x 2 rail (Mellanox)
	Peak network b/w	7 GB/s
Number of nodes		268
Interconnection configuration		Fat-Tree with full bisection b/w
Peak performance		802 TFLOPS
Network bisection b/w		1.88 TB/s
Shared file system		RAID6 with Lustre (DDN)
File system capacity		500TB (user)

This FPGA chip for TCA architecture is named “PEACH2 (PCI Express Adaptive Communication Hub ver.2)”. The name of PEACH is inherited from our previous research in JST/CREST to utilize PCIe link as low power and high performance reliable link for embedded systems (see HPCS Division’s section in this report). Fig. 5.4.4 shows the picture of PEACH2 card.

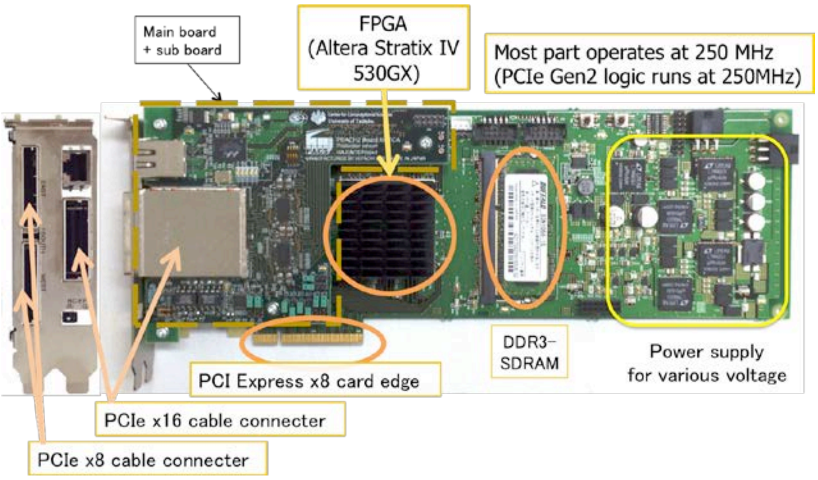

Fig. 5.4.4 PEACH2 Board for TCA architecture

The operating system of HA-PACS is Red Hat Enterprise Linux and we employ MVAPICH2, NVIDIA CUDA, CUBLAS, PGI compiler (C/C++/Fortran for CUDA), Intel compiler for GPU, OpenMP and MPI programming environment.

For the programming on TCA communication, we provide a special API to allocate and map the transfer target memory, set-up the communication such as RDMA transaction, and invoke the communication. To perform the memory to memory copy between GPUs on different nodes, we need the information on CUDA bios and memory map on system side. For this purpose, we made an NDA with NVIDIA to share the information required for R&D of TCA and PEACH2.

We measured HPL performance on Base Cluster and TCA part individually since the installation time differs and it is difficult to make a good balance on load between two types of GPUs where the new one provides approximately doubled performance of the older although entire system including two parts can be operated as a big cluster. We recorded HPL score individually to TOP500 site, and HA-PACS Base Cluster was ranked at #41 on TOP500 list on June 2012 with 421.6 TFLOPS of HPL score. TCA part was ranked at #134 on TOP500 list on November 2013 with 277.1 TFLOPS. The score by Basic Cluster is reasonable as a commodity GPU cluster but it is shown that PCIe to support four GPUs per node is not the bottleneck in the performance thanks to rich configuration of Intel Xeon Sandy Bridge or Ivy Bridge. The HPL efficiency is especially high on TCA part corresponding to 76% of the peak performance. Thanks to its low power consumption, HA-PACS/TCA was ranked at #3 in Green500 list on November 2013. Since the PEACH2 board is under tuning, we did not employ it for parallel processing on HPL measurement, and we used MVAPICH2 with dual rail of InfiniBand QDR.

5.4.3 Operation

Base Cluster started its operation on February 2012 and we took 7 months for system shakedown and rapid development of its important target applications including QCD, Astrophysics and Bioscience at CCS. Then HA-PACS Base Cluster has been opened for public use from October 2012 under the Collaborative Interdisciplinary Program by CCS. 100% of system resources are dedicated for this program.

For the past 2 years of continuous operation, HA-PACS Base Cluster has been almost stable, but the failure rate of GPUs in the node is not negligible. In average, we count approximately 3 or 4 GPUs failure out of 1072 of them in the system. Thanks to the feature of flat configuration with 268 computation nodes and full-bisection bandwidth of Fat-Tree network, the freedom of node allocation and scheduling is very high and the high utilization ratio with more than 70% in average has been kept.

HA-PACS/TCA started its operation from November 2013 and we have been still on the stage of development and tuning of TCA (PEACH2). Some groups of application development for their domain science started the preparation to port their communication part to TCA. For the performance evaluation of PEACH2 on

HA-PACS/TCA, please refer "5.9.2 Project Office for Exascale Computing System Development" part in this report.

5.5 ILDG/JLDG Project

5.5.1 ILDG: International Lattice Data Grid

ILDG is a grid of 5 regional data grids (see Fig.1) for sharing lattice QCD (computational elementary particle physics) configurations (fundamental data) worldwide. The Center for Computational Sciences (CCS) worked as a representative of Japanese grid and contributed to the ILDG in various ways. Construction of interoperable ILDG system was completed in 2007. Since then, the CCS maintains a gateway of the ILDG and Japanese grid under close collaboration of lattice QCD physicists and computer scientists.

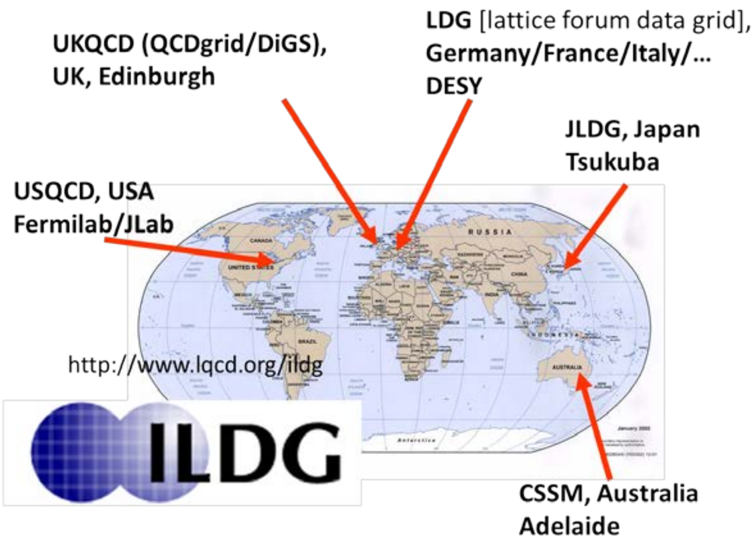

Figure 1: ILDG regional grids

CCS contribution to the ILDG for 2008-2013 is summarized as follows. 1) The ILDG gateway has been continuously maintained. Down time of the system was less than 2 months for 6 years. 2) Japanese grid (Lattice QCD Archive) was restructured to a new grid system JLDG (see below) in 2009. A part of the JLDG can be accessed from the ILDG with ILDG interface. 3) Faceted navigation system was developed and released for public use in 2009. The web application enables users to search ensemble (a set of configurations with common physics parameters) specifying facets such as physics parameters and other metadata (regional grid, collaboration name etc.). 4) The CCS hosts two ILDG components (Plone official web page and Inca monitoring system) since 2013. 5) PACS-CS Collaboration released their configuration (10 ensembles, 3800 configurations, 8 TB) for public use in 2009 and 2013. Data access from outside Japan counts over 2500 times. (Access from inside Japan is not recorded). 6) A member of CCS took a role of the

ILDG board chair and hosted two ILDG video workshops in 2009. 7) A member of CCS joins the ILDG metadata working group, in which revision of QCDml, (an XML-based markup language developed by the working group) was discussed. 8) A member of CCS joins the ILDG middleware working group and takes a part of ILDG system maintenance and upgrades. For more details, refer to http://www.lqcd.org/ildg

5.5.2 JLDG: Japan Lattice Data Grid

JLDG is a data-grid for the lattice QCD (LQCD) community in Japan. Several large LQCD collaborations in Japan have been working on lattice QCD simulations using supercomputers distributed over distant sites. The JLDG provides such collaborations with an efficient method of data management and sharing. File servers installed on JLDG sites (7 sites as of Jan. 2013, see Fig.2) are connected to the NII SINET VPN called HEPnet-J/sc and are bound into a single file system with the GFarm, a grid-base file system software. Because the file system looks the same from any sites, users can do analyses (measurement of physical quantities) on a supercomputer on a site, using data generated and stored in the JLDG at a different site. In this way, the JLDG improves performance and reduces human costs of data sharing over supercomputers. System development started in 2005 and first installation and test operation were done in 2007.

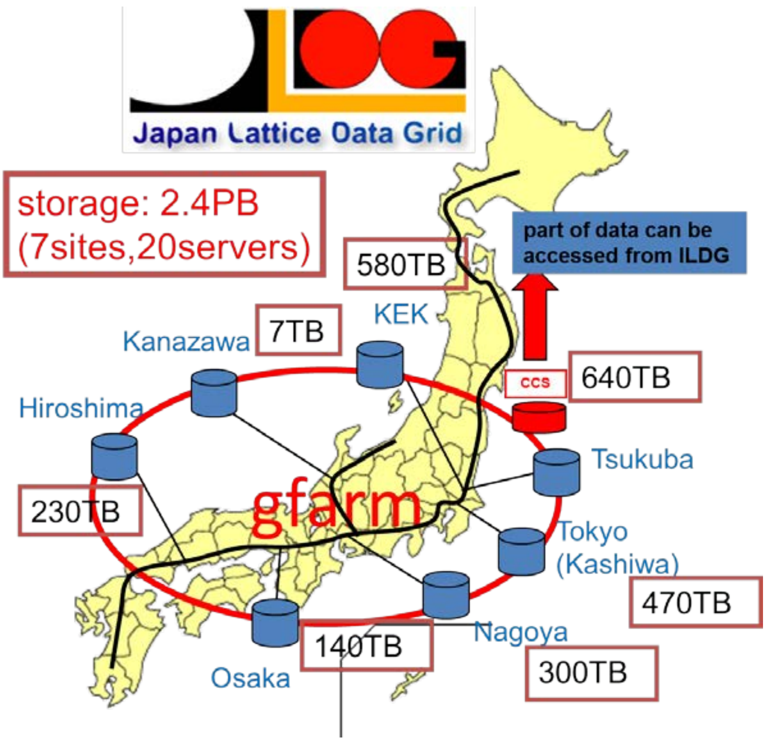

Figure 2: Conceptual view of JLDG

Since the official start of operation in 2008, the JLDG has been improved in various ways. Among others, major improvements are: 1) The JLDG was connected to the ILDG in 2008. A part of the JLDG file system can be accessed from the ILDG via grid-ftp protocol. (2) Research groups started to store daily research data in 2009, because file access control based on user and group was implemented. 3) FUSE mount (user level mount as a unix file system) was supported in 2011. (Users can access the JLDG files as if they were stored on a local disk.) 4) New 2 sites joined the JLDG in 2012. 5) Fast data copy system between the JLDG and the HPCI Shared Storage was developed in 2013. (HPCI: national project started in 2011 for constructing High Performance Computing Infrastructure. For details, see https://www.hpci-office.jp/). 6) More 2 sites will join the JLDG within FY2013 (March 2014). These results are achieved by an extensive collaboration of all JLDG sites, which has been continuously led by the CCS.

As of Jan. 2013, 11 research groups (67 users in Japan) store their daily research data of 1.6PB (57 Million files) in total. Number of publications for works used the JLDG is 66 (as of Feb. 2012). The large number of publications and recent rapid increase of disk usage (see Figure below) convince us that the JLDG has grown up into a useful infrastructure for LQCD community in Japan. For more details, refer to http://www.jldg.org.

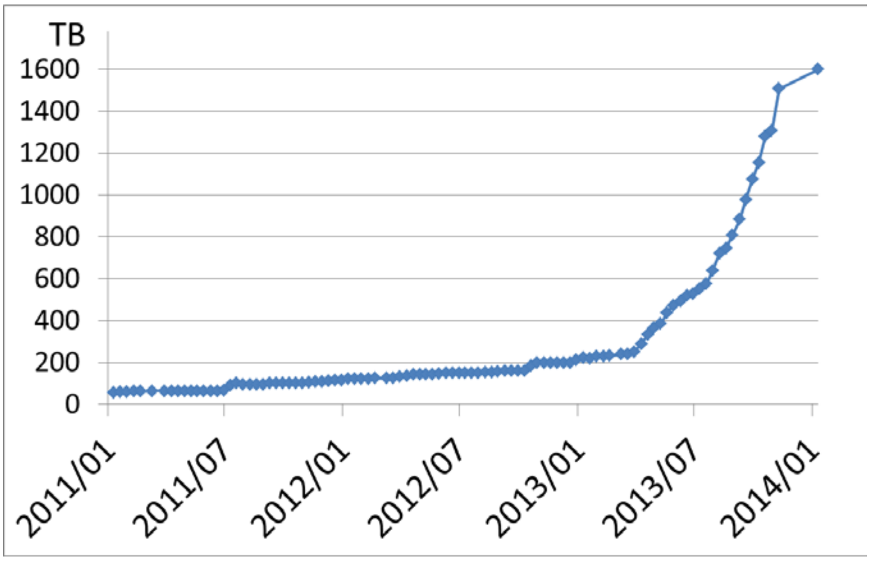

JLDG disk usage

5.6 Multidisciplinary Cooperative Research

The mission of the CCS focuses on promoting research efforts in computational sciences where scientists of various fields do not simply employ computers, but also work in close collaboration with computer scientists to develop computers that are most suitable for scientific investigations. Since 2002, the Center has been making computer resources available to collaborative research efforts in order to facilitate nationwide progress in computational sciences. Since 2007, these activities have been reinforced by enhanced links with computer sciences through the promotion of the Multidisciplinary Cooperative Research Program (MCRP). Since 2010, the Center has been recognized as a national core-center under the Advanced Interdisciplinary Computational Science Collaboration Initiative (AISCI), through which it has been active in the MCRP, while also providing cooperative assistance and support in the areas of research meetings, visitor information, cooperative research travel and short-term employment.

The MCRP consists of the following programs:

(1) Multidisciplinary Pioneering Program (MPP)
(2) Priority Subjects Promoting Program (PSPP)
(3) Cooperative Research Promoting Program (CRPP)

The MDP promotes combining different fields, while the PSPP is expected to promote computational science subjects via large-scale computations. The CRPP aims at promoting advancements in computational approaches to sciences and technologies through close collaboration between researchers inside and outside the Center.

The Center has been soliciting proposals for the MCRP in the fields of computational sciences, including particle and nuclear physics, astrophysics, material sciences, life sciences, global environmental sciences, and biology, as well as in the fields of computer sciences including high performance computer systems, computational informatics, and numerical analysis. Computational facilities that can be used freely by MCRP associates include the Parallel Array Computer System for Computational Sciences [PACS-CS (2007-2011)], FIRST (2009-2011), T2K-Tsukuba (2008-2013), and the Highly Accelerated Parallel Advanced system for Computational Sciences [HA-PACS (2013-)].

The number of MCRP-approved projects conducted over the last seven years is shown in the table. Approximately 30 to 60 projects are approved each year. A breakdown of the 2013 projects by research fields is shown in the graph, and demonstrates how the projects extend across a variety of natural science and engineering fields.

A breakdown of the 2013 projects by research fields

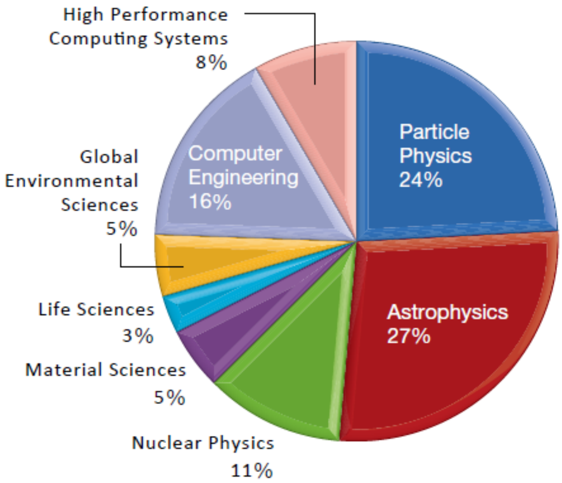

The number of MCRP-approved projects

Year	# of projects	Computer
FY2007	13	PACS-CS
FY2008	36	PACS-CS、T2K-Tsukuba
FY2009	53	PACS-CS、T2K-Tsukuba、FIRST
FY2010	24	PACS-CS、T2K-Tsukuba、FIRST
FY2011	31	PACS-CS、T2K-Tsukuba、FIRST
FY2012	48	T2K-Tsukuba、HA-PACS
FY2013	59	T2K-Tsukuba、HA-PACS

Past Symposiums Organized by the CCS
(including MCRP Report Meetings)

- 5th symposium on “Discovery, Fusion, Creation of New Knowledge by Multidisciplinary Computational Sciences”
  Nov. 5-6, 2013, University of Tsukuba

- 4th symposium on “Discovery, Fusion, Creation of New Knowledge by Multidisciplinary Computational Sciences” – Report Meeting on Multidisciplinary Cooperative Research Program (MCRP)
  Oct. 25, 2012, University of Tsukuba

- 3rd symposium on “Twenty years anniversary of Center for Computational Sciences, University of Tsukuba” – Development of Multidisciplinary Computational Sciences towards Exa-scale
  Sept. 7, 2012, International Congress Center Epochal Tsukuba

- 2nd symposium on “Discovery, Fusion, Creation of New Knowledge by Multidisciplinary Computational Sciences” – Progresses of Computational Sciences by PACS-CS and Developments towards Next-generation Computing
  Sept. 12-13, 2011, University of Tsukuba

- 1st symposium on “Discovery, Fusion, Creation of New Knowledge by Multidisciplinary Computational Sciences” – Development of Multidisciplinary Computational Sciences towards Post Peta-scale Computing
  May 5-6, 2010, University of Tsukuba

5.7 Research Project Offices

5.7.1 Project Office for Exascale Computational Sciences

Supported by the dramatic growth of computer power, the Computational Sciences have seen remarkable developments in recent years, and are now recognized – along with theory, experiment, and observation – as indispensable partners in the promotion of the cutting-edge scientific frontiers of the 21st century.

Nowadays, computers characterized by a new architecture are appearing in rapid succession and include large-scale systems equipped with Graphics Processing Units (GPUs) as accelerators, massively-parallel computers equipped with many-core processors, which may play a central role in future large-scale computing, to say nothing of ordinary massively-parallel computers combining a multitude of central processing units (CPUs). Despite such advances, it is by no means an easy task to make computations for researchers, even with the full machine power usage of such diverse computers.

By means of close collaborations among researchers of physics and those of computer sciences, the Center for Computational Science (CCS) has successfully achieved a number of breakthroughs by developing optimum computers for computational sciences in the fields of particle and astrophysics. Furthermore, the Center now has a Project Office for Exascale Computational Sciences dedicated to close collaborations among computational and computer scientists in a wide variety of fields, and to establishing cooperation among researchers belonging to divergent scientific fields.

The services of this office are not restricted to the Center alone, but extend outside to promote nationwide collaborations. To date, it has conducted analyses and examinations of hot spots for major applications to ensure that they show high efficiency in the Highly Accelerated Parallel Advanced system for Computational Sciences (HA-PACS), which is a massively-parallel supercomputer equipped with GPU accelerators and is operated under close cooperation with the Project Office for Exascale Computing System Development, which is the organization promoting the development of new computers.

The Center's future plans include tackling innovative technologies such as the Tightly Coupled Accelerator (TCA) technology that will be incorporated in the extended HA-PACS system and a massively-parallel computer equipped with processors of Many Integrated Core architecture, which will be introduced in Joint Center for Advanced High Performance Computing (JCAHPC).

Activities on HA-PACS Undertaken by the Project Office for Exascale Computational Sciences

Field	Scientific Target	Numerical Calculations employing GPU
Particle Physics	•Multiscale physics •Physics at finite temperature and/or finite density	•Solution of large-scale sparse linear equations •Matrix-matrix product calculation of dense matrices
Astrophysics	•Realization of computational astrophysics in six dimensions •Simulation of collisional self-gravitating systems	•Ray-tracing calculations of light-intensity and chemical reactions •Calculations of gravitational acceleration and its derivatives
Nuclear Physics	•Research on nuclear response and reaction dynamics •First-principles calculations for light-matter interaction	•Real-time and real-space finite difference calculations •3D Fast Fourier Transformations
Material Science	•Investigations on structure and dynamical processes of matter, and control of quantum processes	•Real-time and real-space finite difference calculation of time-dependent Schrödinger Equations
Life Science	•Enzymatic reaction mechanisms •Dynamical structures of biomolecules	•Quantum mechanical calculations •Molecular dynamics simulations
Global Environmental Science	•Next-generation atmospheric circulation model (NICAM) •Calculation of urban climate	•3D normal-modes expansion •LES
Computational Informatics	•Knowledge discovery from large-scale database	•Association rule mining •Time-series data retrieval •Clustering

5.7.2 Project Office for Exascale Computing System Development

The theoretical peak performance of massively parallel supercomputers is represented as the processor performance of a single node and the number of those nodes. The performance improvement of supercomputers, to date, has been mainly achieved by the "scaling system", which refers to increasing the number of computing elements in a system. However, faced with limitations related to electric power consumption and/or the thermal or fault ratio of each part, "weak-scaling" strategies are facing their achievable limits.

To overcome these limitations, it will be necessary for the next generation of Exascale computing to improve the single node performance by several tens TFlops, as well as to establish fault tolerant technology that will enable several million nodes to be incorporated into a system. For single node performance improvements, one of the most promised methods is using accelerators to reduce the time-to-solution, rather than enlarging the computation problem size. This is known as the strong-scaling problem.

In the Project Office for Exascale Computing System Development, we have been researching the next generation of accelerated computing systems under the new concept of accelerated computing named Tightly Coupled Accelerators (TCA). The Highly Accelerated Parallel Advanced System for Computational Sciences (HA-PACS) is an experimental testbed for TCA architecture, as well as the application development platform with the latest graphics processing unit (GPU) computing technology. The system consists of two parts, a base cluster part with commodity GPU technology, and a TCA part that incorporates our original GPU-direct communication technology.

The HA-PACS base cluster part has been operating since February 2012 using 268 nodes, each of which consists of two Intel E5 SandyBridge CPU sockets and four NVIDIA M2090 (Fermi) GPUs. From November 2013, the additional TCA part, equipped with 64 nodes each consisting of two Intel E5 IvyBridge CPUs sockets and four NVIDIA K20X (Kepler) GPUs, connected by our original named PCI-Express Adaptive Communication Hub ver.2 (PEACH2) technology will be added to facilitate the development of TCA system software and its applications.

■ Base Cluster part

GPGPU (General Purpose GPU) is a computing method that applies GPUs to general purpose scientific computing as the computation-accelerating device. The latest GPUs provide extremely high theoretical performance in floating point operations, which can then be applied to a wide variety of high performance computing (HPC) problems.

■ HA-PACS/TCA part

Current GPU clusters often encounter performance bottlenecks caused by poor communication performance during inter-node GPU to GPU communication on parallel processing systems. TCA architecture provides a fundamental solution to this problem by using hardware that enables true GPU-direct communication over nodes by achieving strong scaling on various HPC problems that are often subject to communication latency bottlenecks.

■ PEACH2

As the experimental TCA concept testbed for developing a communication system to enable GPU to GPU direct communication over nodes, this system consists of a field programmable gate array (FPGA) PEACH2 chip and its peripheral component interconnect (PCI) attachable board. Since it is constructed FPGA technology, it is easy to enhance the performance and function of PEACH2 as the TCA concept evolves. PEACH2 realizes the following functions:

- Autonomous communication by GPUs without depending on host CPUs → high performance parallel processing of general purpose code and accelerated code
- Direct data copy between GPUs on different nodes → processing by GPUs only
- Direct I/O between GPU and other peripheral I/O devices → fault tolerant I/O system in the event of CPU failure

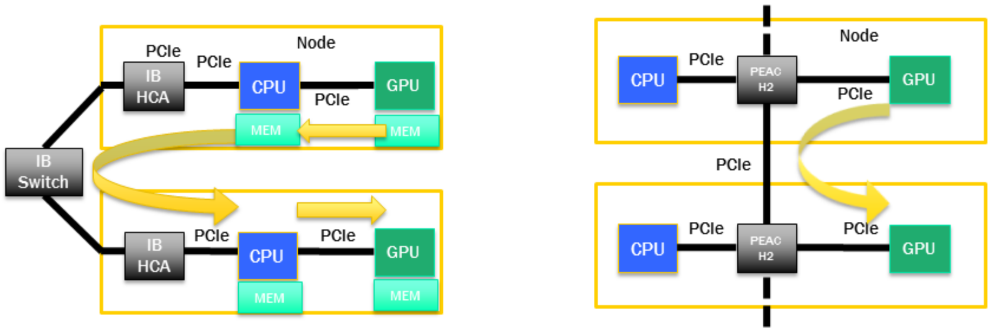

Inter-GPU communication on general GPU clusters (left) and true GPU-direct communication by TCA (right)

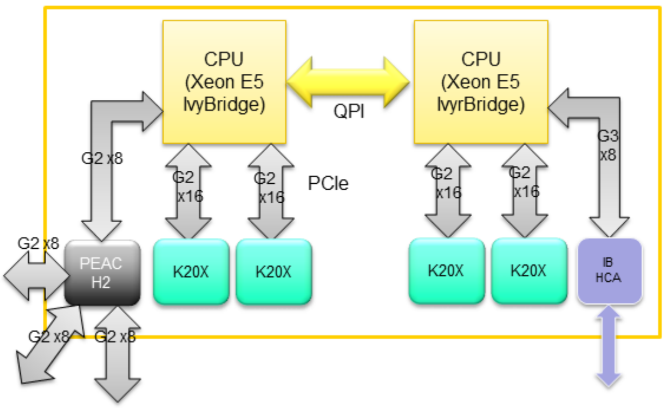

Block diagram of a computation node of HA-PACS/TCA

5.7.3 Liaison Office for Multidisciplinary Computational Sciences

The Center for Computational Science (CCP) conducts research activities grounded in the central concept of "Multidisciplinary Computational Science" through collaboration and cooperation among various computer-related and scientific fields. The Liaison Office for Multidisciplinary Computational Sciences is a key part of the Center's efforts aimed at maintaining its position as an international hub for computational science information.

■ Cooperation among Different Fields

In order to promote multidisciplinary computational science, cooperation among different fields is essential. In 2011, using the research results of first-principles electronic state calculation for 100,000 atom silicon nanowires using "K" supercomputer, a research group from RIKEN, the University of Tsukuba, the University of Tokyo, and Fujitsu Limited won the Gordon Bell Prize for Peak-Performance, which is recognized as an outstanding achievement in high-performance computing. This is just one of the results of cooperation among different fields that have taken place at the Center in recent years.

■ International Cooperation

As for international cooperation, partnership agreements have been concluded between University of Tsukuba and the University of Edinburgh in the UK, as well as with Lawrence Berkeley National Laboratory in the US. Furthermore, as part of the efforts to deepen the computational science research exchanges, the CCP and the University of Edinburgh Parallel Processing Center (EPPC) are currently hosting joint symposiums, and the Center is cosponsoring workshops with Lawrence Berkeley National Laboratory. Additionally, bilateral exchange joint research projects into computational materials science are being conducted with Vanderbilt University and the University of Washington in the US under the sponsorship of the Japan Society for the Promotion of Science (JSPS).

■ Multidisciplinary Pioneering Program

The Multidisciplinary Pioneering Program encourages the promotion of research efforts that require collaboration and cooperation among divergent scientific fields. For example, there is a current project underway that requires cooperation between the computational physics (science) and numerical analysis (engineering) fields in order to implement calculation code with high execution efficiency, and to perform the necessary large-scale simulations using the code. There are also projects that aim at new research deployments and developments by facilitating organic cooperation among different fields utilizing the same techniques in computational science.

This Multidisciplinary Pioneering Program has been carried out in conjunction with the collaborative multidisciplinary program of the center.

5.7.4 Promotion Office for Computational Sciences

The Promotion Office for Computational Sciences promotes the maintenance of the personnel development program associated with computational sciences in alignment with the graduate school research departments, conducts activities to raise awareness related to summer school and other matters, and contributes to society by disseminating information concerning the advanced research conducted at this Center.

■ Personnel Development

・Graduate School Dual Degree Program

The graduate school consists of a pre-doctoral course (master's) and a doctoral course (doctorate). Doctorate degrees are awarded upon completion of a specialized field through both courses. However, since sophisticated computer skills and experience are necessary to advance research in scientific fields such as physics, the global environment, and biology, a broad level of expertise in both science and computers is required for computational science candidates.

The graduate school promotes a dual degree program to meet this need. This program makes it possible to obtain a doctorate in a scientific field along with a master's degree of computer science simultaneously, by engaging in, concurrently and in parallel to, scientific field doctoral courses and pre-doctoral computer research courses.

・Computational Science English Program

In response to the increasing numbers of international students from overseas, the Computational Science English Program was established for the pre-doctorate computer science course that must be completed in English. As shown in the table, this course provides English lectures in 13 subjects, along with English language research guidance aimed at assisting master degree candidates in successfully obtaining their degrees.

・Campus-wide Courses for Graduate Students

Since computers are the foundation of all research at this graduate school, classes in computational science literacy and high-performance parallel computing technology for computational science have been established as common subjects for all students. These are summer school subjects, and are also available to outside researchers and students from other universities.

■ Social Contribution

・Cooperation with other institutions

Domestically, research cooperation in the field of high performance computational technology is advanced under the T2K Alliance that connects Tsukuba University with the University of Tokyo and Kyoto University, along with the close research cooperation ties with other research organizations within the

City of Tsukuba. We have also built an international research cooperation framework through exchange with the International Lattice Data Grid (ILDG), which is an international data-sharing project of particle physics, the University of Edinburgh (UK), Lawrence Berkeley National Laboratory (US), and among others. We intend to strengthen and develop the domestic and international research cooperation achieved to date, and to utilize those relations as a bridge between researchers and exchange students to pursue collaborative research.

・Mt. Tsukuba Project

This project supports research and educational activities, while contributing to society, by continuously observing, recording, and disseminating weather observations (temperature and humidity) from the 877 m summit of Mt. Tsukuba.

・Dissemination of Information

Our office produces three types of brochures. One aimed at high school students and above, another aimed at graduate students and above, and the third written in English to meet the needs of our diverse applicants and students. The office also coordinates visits from high school students within Japan and overseas university students, as well as “open house” events that are open to the public. Through these efforts, we work to deepen exchanges with people around Tsukuba City and in the Kanto region. In addition, as teaching individuals, the faculty actively offers and promotes off-campus classes and lectures.

5.7.5 Bureau of Public Relations

In FY2010, the bureau initiated activities aimed at improving relations with society in order to advance the mission of the CCS and to promote the computational sciences themselves. To that end, three public relations concepts underlying the purposes behind the establishment of the Center have been set forth for use in guiding daily operations.

■ Public Relations Concepts

- We will conduct bilateral public relations to foster enhanced relations between the Center and society.
- We will aim at increasing the visibility of the Center and the computational sciences.
- We will enhance the level of understanding of the fundamental policies of the Center, and work to improve public confidence in its actions.

■ Activity Policy for the Bureau of Public Relations

- We will support public relations and public hearings for researchers associated with the center.
- We will create content, including websites, in order to provide information to stakeholders.
- We will actively gather information and create feedback mechanisms for researchers within the center.
- We will create a risk management organization.
- We will examine and implement methods to measure visibility, levels of understanding, and degrees of confidence.

■ Activity Results

- Website: http://www.ccs.tsukuba.ac.jp/CCS/

Provides up to date information on researchers, projects, computer basics, database disclosures, symposiums/seminars, disaster preparedness and response, etc.

- Brochures

We prepare three types of brochures based on the intended readership: “Computational Sciences and Supercomputers” for the general public (this is aimed at high school readers and older); a Japanese text for professionals (aimed at graduate students and older); and an English text for professionals (for use by international graduate students and others).

- Movies

Productions of 12 movies (Japanese and English), including an introduction to the Center, the computing facility, and 10 research fields, have been produced to date.

- Tour and Open House

The number of visitors was 351 in FY2010, 1,056 in FY2011, and 844 in FY2012. We created a FAQ section on the Center's website.

- Media Promotions

Seven press releases were disseminated to the mass media in fiscal year 2011, and five were released in fiscal year 2012. The Center is also active on the social media Twitter (@CCS_PR), where it has gathered 3,000 submissions and 750 followers (as of September 2013).

Appendix A. Faculty Members of CCS

Director: Masayuki Umemura

Deputy Director: Taisuke Boku

Masayuki Umemura, PhD, Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Theoretical Astrophysics. In particular, the study with radiation hydrodynamics on the formation of first generation objects, supermassive black holes, and galaxies

Taisuke Boku, PhD, Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Large scale parallel processing, high performance interconnection, cluster computing, hybrid parallel processing system

(1) Division of Particle Physics (Chief: Yoshinobu Kuramashi)

Faculty members

Yoshinobu Kuramashi, Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Theoretical and computational studies on strong interactions using lattice field theories

Akira Ukawa, Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Computational particle physics, in particular lattice QCD, and development of computers for such studies

Tomoteru Yoshie, Associate Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Elementary Particle Physics

Naruhito Ishizuka, Associate Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Computational studies on strong interactions using lattice field theories

Ishii Noriyoshi, Associate Professor(HPCI Strategic Program)

Graduate school:

Research topics:
Nuclear force by lattice QCD

Nemura Hidekatsu, Associate Professor(HPCI Strategic Program)

Graduate school:

Research topics:
Lattice QCD study of hyperonic nuclear forces

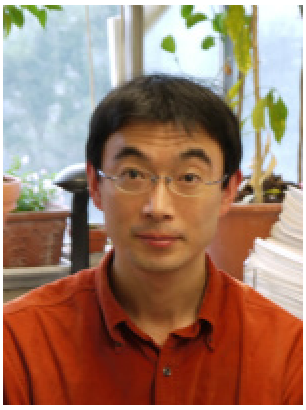

Taniguchi Yusuke, Lecturer

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Lattice gauge theory

Sinya Aoki, Visiting Professor

Graduate school:
Yukawa Institute for Theoretical Physics, Kyoto University

Research topics:
Lattice Field Theories, Lattice QCD and its applications to Nuclear Physics

Collaborative fellow

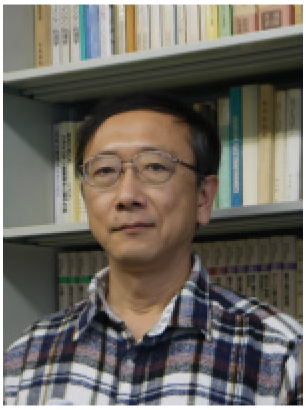

Kazuyuki Kanaya, Professor, Executive Officer, Provost of Faculty of Pure and Applied Sciences

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Theoretical particle physics, lattice field theory, parallel computers

Research staff

- Tomoya Nagai
- Naoya Ukita
- Yusuke Namekawa
- Kenji Sasaki

(2) Division of Astrophysics and Nuclear physics (Chief: Kazuhiro Yabana)

Astrophysics Group (Group Leader: Masayuki Umemura)

Faculty members

Masayuki Umemura, Professor, Director

Graduate school:
Graduate school of Pure and Applied Sciences
Research topics:
Theoretical Astrophysics. In particular, the study with radiation hydrodynamics on the formation of first generation objects and galaxies

Masao Mori, Associate Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Theoretical astrophysics, galaxy formation, black-hole accretion flow, astrophysical numerical simulation

Kohji Yoshikawa, Lecturer

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Astrophysics. Theoretical studies on observational cosmology, the formation of galaxies and galaxy clusters, and intergalactic medium

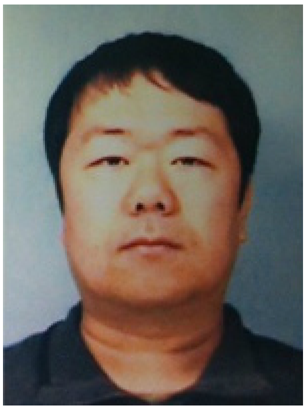

Naohito Nakasato, Visiting Associate Professor

Graduate school:
School of Computer Science and Engineering, Univ. of Aizu

Research topics:
Astrophysical Simulations and High Performance Computing

Research staff

- Tomoaki Ishiyama
- Daisuke Namekata
- Kenji Hasegawa
- Alexander Wagner
- Takatoshi Sibuya

Theoretical Nuclear Physics Group (Group Leader: Kazuhiro Yabana)

Faculty members

Kazuhiro Yabana, Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Nuclear physics, computational sciences on atomic, molecular,and optical sciences

Yukio Hashimoto, Lecturer

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Microscopic theory of nuclear collective motions

Jun Terasaki, Associate Professor(HPCI Strategic Program)

Graduate school:

Research topics:
Research of nuclear physics

Research staff

- Yasutaka Taniguchi

(3) Division of Quantum Condensed Matter Physics (Chief: Kazuhiro Yabana)

Faculty members

Hiroyasu Koizumi, Associate Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Elucidation of mechanism of high temperature copper oxide superconductor and Realization of quantum computer using it.

Tong Xiao-Min, Associate Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Atoms, molecules interaction with intense laser field and time-dependent density functional theory

Nobuya Maeshima, Lecturer

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Photoinduced phenomena in strongly correlated electron systems

Atsushi Oshiyama, Visiting Professor

Graduate school:
Graduate school of Engineering, Univ. of Tokyo

Research topics:
Theoretical Condensed Matter Physics, Computational Materials Science

Collaborative fellow

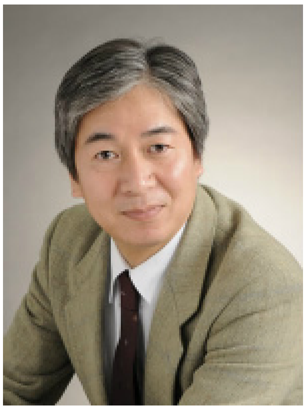

Hino Ken'ichi, Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Photophysics of Condensed Matters

Susumu Okada, Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Condensed matter physics
Computational material sciences

(4) Division of Life Sciences (Chief: Masayuki Umemura)

Biological Function and Information Group (Group Leader: Masayuki Umemura)

Faculty member

Mitsuo Shoji, Assistant Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Theoretical investigation on the reaction mechanisms of enzymes

Collaborative fellow

Jun Yanagisawa, Professor

Graduate school:
Graduate school of Pure and Applied Sciences

Research topics:
Analysis of molecular mechanism for epigenetic regulation in mammaliancells

Biological Science Group (Group Leader: Yuji Inagaki)

Faculty member

Yuji Inagaki, Associate Professor

Graduate school:
Graduate school of Life and Enviromental Sciences

Research topics:
Molecular phylogeny of eukaryotes Investigation of lateral genetransfers at the whole-gene and sub-gene levels Estimation of protein functions by combining structural and evolutionary parameters Artifacts in molecular dataanalyses

Collaborative fellow

Tetsuo Hashimoto, Professor

Graduate school:
Graduate school of Life and Enviromental Sciences

Research topics:
Molecular evolutionary studies on the origin and early evolution of eukaryotes

Hideko Urushihara, Professor

Graduate school:
Graduate school of Life and Enviromental Sciences

Research topics:
Studies on developmental programs described in the genome

Research staff

- Takuro Nakayama

(5) Division of Global Environmental Science (Chief: Hiroshi L. Tanaka)

Faculty members

Hiroshi L. Tanaka, Professor

Graduate school:
Graduate school of Life and Enviromental Sciences

Research topics:
General circulation, dynamics and energetics of the atmosphere

Hiroyuki Kusaka, Associate Professor

Graduate school:
Graduate school of Life and Enviromental Sciences

Research topics:
Urban climate, Applied Meteorology, Numerical simulationof mesoscale weather using the WRF model

Collaborative fellow

Hiroaki Ueda, Professor

Graduate school:
Graduate school of Life and Enviromental Sciences

Research topics:

Yasutaka Wakazuki, Assistant Professor

Graduate school:
Graduate school of Life and Enviromental Sciences

Research topics:
Precipitation and Disaster Meteorology, Applied Meteorology and Climatology
Akio Kito, Senior Researcher

Graduate school:
Graduate school of Life and Enviromental Sciences

Research topics:

Research staff

- Yuko Akimoto
- Ryosaku Ikeda

(6) Division of High Performance Computing Systems (Chief: Taisuke Boku)

Faculty members

Taisuke Boku, Professor, Deputy Director

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Large scale parallel processing, high performance interconnection, cluster computing, hybrid parallel processing system

Mitsuhisa Sato, Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
High performance parallel computing, compilers and performance evaluation, grid computing

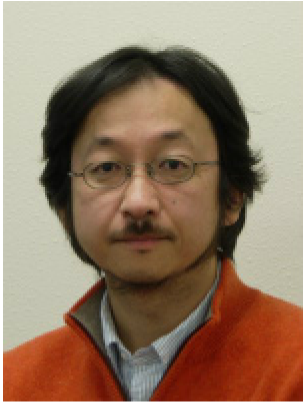

Yuestu Kodama, Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Parallel Computer Architecture, Reconfigurable System, Fat Long Distance Network, Low Power Architecture

Daisuke Takahashi, Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
High-performance computing: High-performance numerical algorithms on parallel computers and performance evaluation

Osamu Tatebe, Associate Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
High Performance Computing, Grid Computing, Distributed File System

Hideyuki Kawashima, Lecturer

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Database Systems and Data Engineering: DBMS Architecture, Sensor Networks

Hiroto Tadano, Assistant Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Numerical analysis: Numerical algorithms for large scale linear systems. Parallel computing for eigenvalue problems.

Yutaka Ishikawa, Visiting Professor

Graduate school:
Information Technology Center, the University of Tokyo

Research topics:
System Software for Parallel and Distributed Systems

Collaborative fellow

Moritoshi Yasunaga, Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
VLSI Engineering, Evolvable Hardware, Dependable Systems

Kouichi Wada, Professor , Director of Academic Computing and Communications Center

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Parallel and Distributed Computing, Network Architecture for Clusters, Multimedia Processor Architecture

Tetsuya Sakurai, Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Numerical algorithms and simulation, Mathematical software for GRID computing

Yoshiki Yamaguchi, Assistant Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Reconfigurable System, Energy-efficient computer system and architecture, Dependable computer system

Akira Imakura, Assistant Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Numerical linear algebra, Algorithms for solving linear systems (Krylov subspace methods and preconditioning techniques)

Research staff

- Masahiro Tanaka
- Hiroaki Umeda
- Hideo Nuga
- Mohamed Amin Jabri
- Kazuya Matsumoto

(7) Division of Computational Informatics (Chief: Kitagawa Hiroyuki)

Database Group (Group Leader: Hiroyuki Kitagawa)

Faculty members

Hiroyuki Kitagawa, Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Database systems, Data engineering Information integration, WWW and databases, Knowledge discovery, XML databases, Multimedia information retrieval, and DBMS architecture

Toshiyuki Amagasa, Associate Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Database System, Data Engineering, XML Databases, Database Applications in e-Science

Computational Media Group (Group Leader: Yuichi Ohta)

Faculty members

Yuichi Ohta, Executive Director/Vice President

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:

Yoshinori Kameda, Associate Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Massive Sensing, Intelligent Video Making based on Scene Understanding,Model Based Vision,Coorperative Distributed Vision, Lecture Archiving and Distance Learning, Visual Surveillance, Human Interface to Virtual Reality

Itaru Kitahara, Associate Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Image media with ubiquitous sensor networks

Collaborative fellow

Tomonori Shirakawa, Professor

Graduate school:
Graduate school of Systems and Information Engineering

Research topics:
Distributed Systems

Research staff

- Senya Polikovsky

Appendix B. Facilities of CCS

B.1 Computer Systems Overview

The computational facilities of the center mainly consist of the Highly Accelerated Parallel Advanced system for Computational Sciences (HA-PACS) and T2K-Tsukuba. In addition, there are Cray XK6m (Blazar), mini HA-PACS cluster, and mini TCA cluster. There are additional medium to small size clusters for data analysis and general networking services.

HA-PACS is the 8th generation of PACS/PAX series supercomputer. For the development and product-run on cutting edge scientific computations toward next generation accelerated computing, it is equipped with the latest GPUs and CPUs connected by new generation of PCI-express to provide rich I/O bandwidth. The Base Cluster was delivered on February 2012 with 802 TFLOPS of peak performance. HA-PACS/TCA, which is the extended part of HA-PACS was introduced on Oct. 2013, and entire HA-PACS system became 1.166 PFLOPS of GPU cluster system for peak performance. HA-PACS/TCA was ranked 3rd in the Green 500 list on Nov. 2013.

T2K-Tsukuba is configured as a single PC cluster with 648 of computational nodes connected by Fat-Tree network, and the total number of CPU cores is 10,368 to provide 95.4 TFLOPS of peak performance. The high-bandwidth of flat Fat-Tree network also realizes the effective parallel processing and flexible job scheduling. T2K-Tsukuba was ranked at number 21th in the TOP500 list on June 2008. The system was installed under the T2K Open Supercomputer Alliance with the University of Tokyo and Kyoto University, and the grid operation is available with other T2K systems in these universities sharing the same basic system architecture with T2K-Tsukuba.

Cray XK6m (Blazar) is used for the development of our original OpenACC compiler based on Omni Compiler. We are evaluating the performance and functionality of it in comparison with Cray's OpenACC compiler for NVIDIA GPU.

The large file server for Lattice Data Grid (JLDG) is located in CCS. JLDG is also connected to International Lattice Data Grid (ILDG). JLDG is a nation-wide data base for Lattice QCD configuration file shared by Japanese researchers in that field. It is implemented by Gfarm grid file system which is developed in Division of High Performance Computing System at CCS. ILDG is a world-wide data base for Lattice QCD configuration file. All the meta-data of ILDG is encoded in XML format shared by that community to provide easy searching and accessing feature for researchers in all of the world.

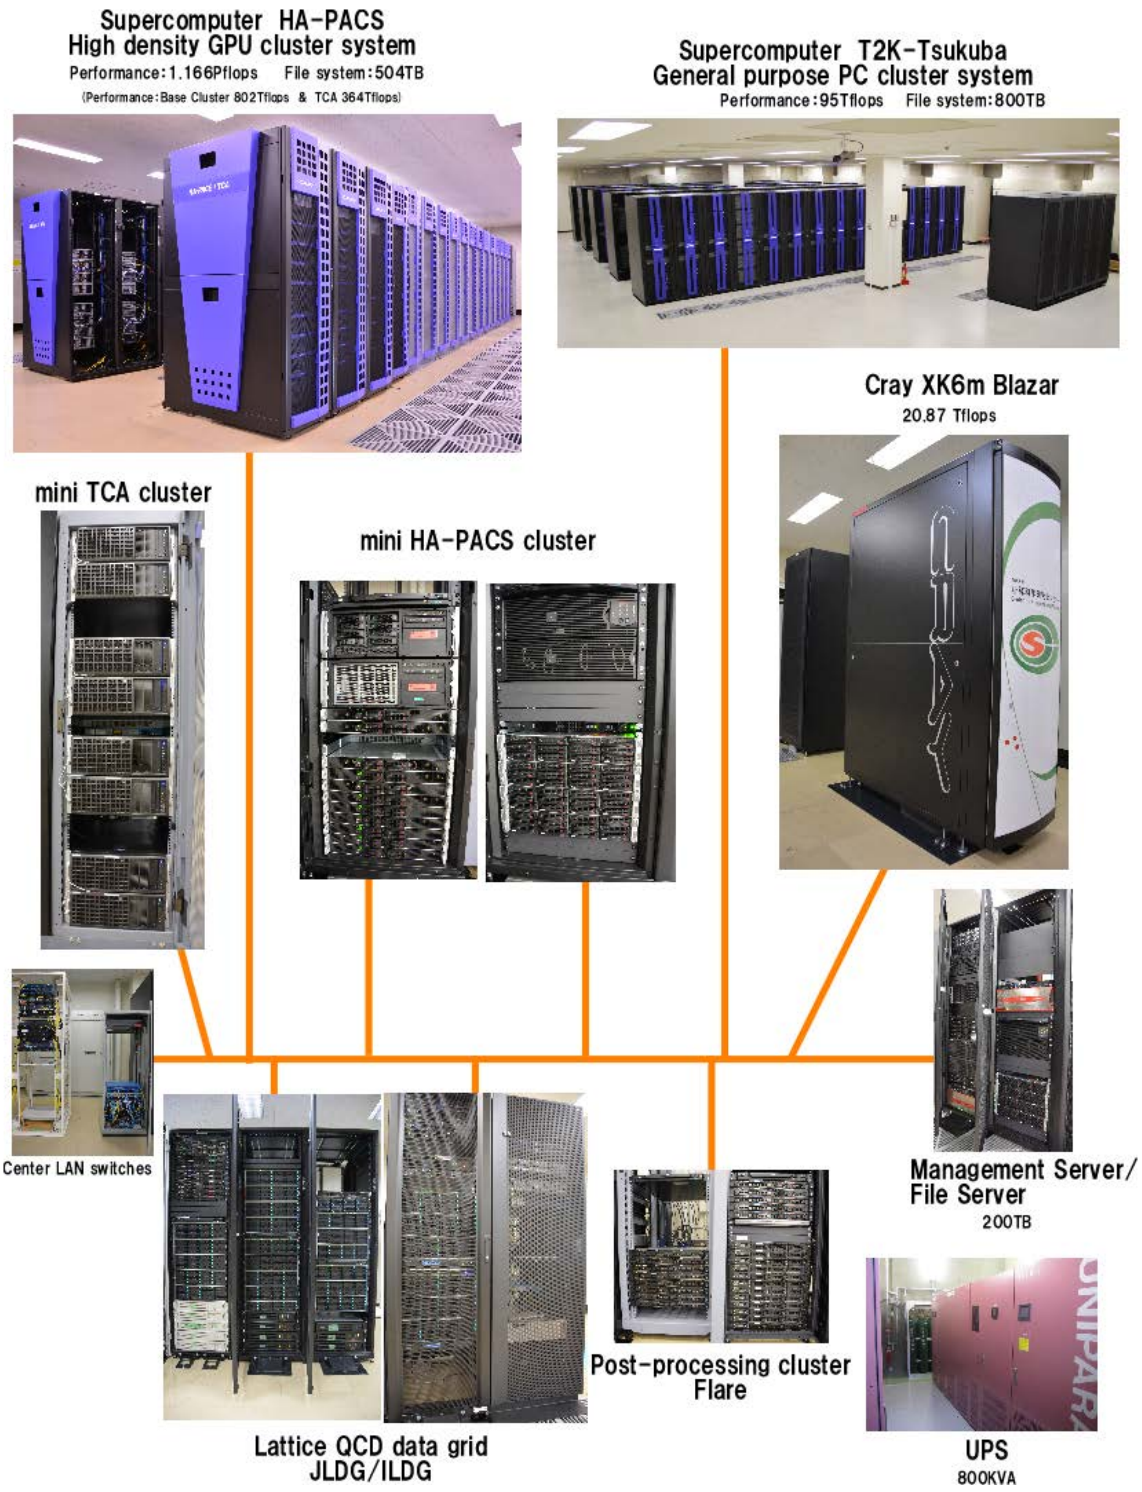
Supercomputer HA-PACS
High density GPU cluster system
Performance:1.166Pflops File system:504TB
(Performance:Base Cluster 802Tflops & TCA 364Tflops)
Supercomputer T2K-Tsukuba
General purpose PC cluster system
Performance:95Tflops File system:800TB
Cray XK6m Blazar
20.87 Tflops
mini TCA cluster
mini HA-PACS cluster
Center LAN switches
Management Server/
File Server
200TB
Post-processing cluster
Flare
Lattice QCD data grid
JLDG/ILDG
UPS
800KVA

B.2 Network Environment

University of Tsukuba has been connected 10 Gbps of nation-wide network for research and education named SINET3 under MEXT since April 2007. In addition, the university is also connected to major research institutes in Tsukuba City by a high-speed WAN with 20 Gbps bandwidth named TSUKUBA-WAN. Our center is connected to SINET3 with 2 Gbps of general purpose network and three of 1 Gbps special purpose research network.

Moreover, we are also connected to the nation-wide research network JGN2 with 10Gbps link to stimulate various research activities in the center. In such network environment, we share the valuable data computed by HA-PACS and other center facilities as well as high-performance Grid computing research toward new generation of computational sciences.