Updates of CCS
University of Tsukuba

Mitsuhisa Sato, Director
Center for Computational Sciences,
University of Tsukuba

Agenda

- Update of CCS
  - Status after 3.11
  - HA-PACS project
  - Education: Computational Science DD program
  - Status and CCS projects for the K computer
- XcalableMP
  - GPU-extension
  - Experiment of one-sided communication on HECToR
- Future plan
  - Post-T2K: what is the next of T2K system?

CCS, University of Tsukuba

- Center for Computational Sciences http://www.ccs.tsukuba.ac.jp/
  - Founded on 1992 as CCP (Center for Computational Physics), expanded and reorganized to CCS in 2004
- Mission of CCS is
  - to enable **scientific discovery by computational science** through the application of advanced computing technologies,
  - and support researches of computational science in Japanese universities by running **leading-edge advanced computing systems**
- **Collaborative researches** with Computational Scientists (application) and Computer Scientists (system)
  - Needs from applications and Seeds from systems
  - More than 30 faculties and PDs, students: Particle Physics and Astrophysics. Material and Life Sciences, Global Environment and Biological Sciences, High Performance Computing System, Data base & Data Mining
- Major Computing Facilities
  - **T2K Tsukuba Cluster**, 95.4 TF, 648 nodes, four quad-core Opetron processor (Appro XtremeServer-X3) and 4 DDR inifiniband network for each node
  - **FIRST**, 33TF, 256 node hybrid PC cluster with GRAPE accelerators for Astrophysics.
  - **HA-PACS**, 802TF, 268 nodes, 4 GPUs and 2 CPUs for one node connected with 2 QDR InfiniBand network

Development of Massively Parallel Computer Systems in University of Tsukuba

- 1977 research begins (by Hoshino, Kawai)
- 1978 1st machine
- 1996 CP-PACS (top of Top500)
- 2006 7th machine PACS-CS

CP-PACS
- **First large-scale general-purpose MPP system in Japan**
  - **Development supported by ``Research of Field Physics with Dedicated Parallel Computers'' funded by the Ministry of Education of the Japanese Government.**
  - **ranked as No. 1 system in the November 1996 Top 500 List.**
- **Collaboration by physicists and computer scientists**
- **Collaboration with industry, and released as Hitachi SR2201**

1978
1st PACS-9

1980
2nd PAXS-32

1989
5th QCDPAX

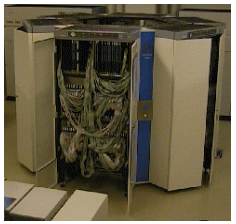

1996
6th CP-PACS (top of Top500 in 1996)

PACS-CS
Massive Parallel Clust

(2006)

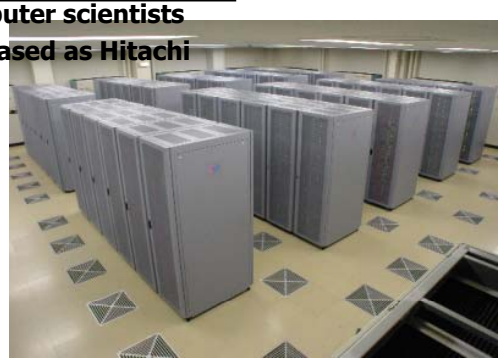

year	system	Performance
1978	PACS-9	7 KFlops
1980	PAX-32	500 KFlops
1983	PAX-128	4 MFlops
1984	PAX-32J	3 MFlops
1989	QCDPAX	1.4 GFlops
1996	CP-PACS	614 GFlops
2006	PACS-CS	14340 GFlops

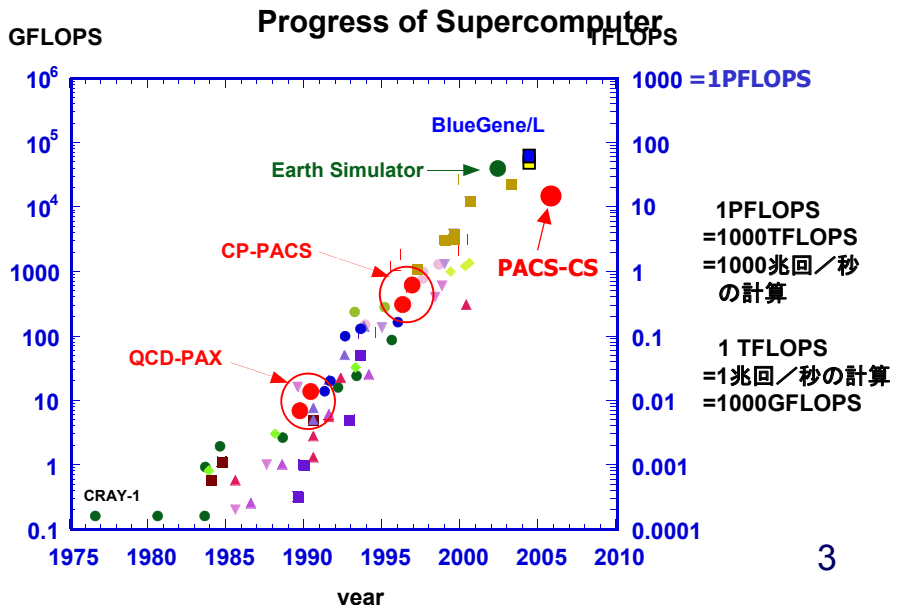

CCS Research divisions and Expertise

- More than 30 faculties (31) and PostDocs, students
- Research organization of computational sciences, not supercomputer "service" center
- 7 Research Divisions and 11 Groups

- **Computational Science**
  - **Division of Particle Physics**
    - **Particle physics – Lattice QCD**
  - **Division of Astrophysics and Nuclear Physics**
  - **Division of Quantum Condensed Matter Physics**
    - **Material science – nano-science, DFT**
    - **Quantum Many-Body Systems**
  - **Division of Life Sciences**
    - **Molecular dynamics, CPMD**
    - **Biological science – molecular phylogenetic analyses**
  - **Division of Global Environment Science**
    - **Global environment – Meteological science, global climate simulation**

- **Computer Science**
  - **Division of High Performance Computing Systems**
    - **System architecture**
    - **Grid computing**
  - **Division of Computational Informatics**
    - **Computational Intelligence - Data Mining & Knowledge Discovery, Large scale database**
    - **Computational Media - Visualization, Computer graphics**

Organization of Education and Research for Computational Science in University of Tsukuba

- **CCS integrates the development of supercomputing facilities and several researches of computational science.**
- **CCS carries out researches in major computational science fields.**
- **The faculty members of CCS "belong" to a Graduate School for education, and "work" for the Center for search.**

Center for Computational Sciences

Founded in 1996, re-organized in 2004

Faculty 32
Professor 11
Associate Prof. 12
Assistant Prof. 9

- **Division of Particle Physics**
- **Division of Astrophysics and Nuclear Physics**
- **Division of Quantum Condensed Matter Physics**
- **Division of Life Sciences**
- **Division of Global Environmental Science**
- **Division of High Performance Computing Systems**
- **Division of Computational Informatics**

Graduate School of Pure and Applied Sciences
- Dept. of Physics
- Dept. of Frontier Science
- Dept. of Materials Science

Graduate School of Life and Environmental Sciences
- Dept. of Geoenvironmental Sciences
- Dept. of Structural Biosciences

Graduate School of Systems and Information Engineering
- Dept. of Computer Science
- Dept. of Intelligent Interaction Technologies

Supercomputing and Power restriction after 3.11

- Just after the big earthquake/tsunami happens, a planned power outage has been scheduled everyday due to the limited power supply (35 GW at disaster).
- In the last summer, Tokyo Electric Power Company estimated power supply capacity 55GW against **peek demand** (more than 60 GW in hot summer!)
- Our government requested more than 15% cut of power consumption for all companies and people in Japan during the summer from July to September.

- Some supercomputers (esp. in east Japan) was stopped during summer time to meet the power cut-off request.
  - In CCS, U Tsukuba, one of supercomputer system (PACS-CS) was stopped from 9AM to 9PM to reduce the peak of power consumption.
  - This situation will continue for several years
  - The power restriction of supercomputers is now a major issue for designing future HPC systems.

- EPCC (UK HPC community) offered HECToR computing resources for Japanese researchers
  - Life science application (Dr. Shouji)
  - Performance study on One-sided communication (Dr. Nakao and me)

Computing resources in CCS

PACS-CS (2006～)

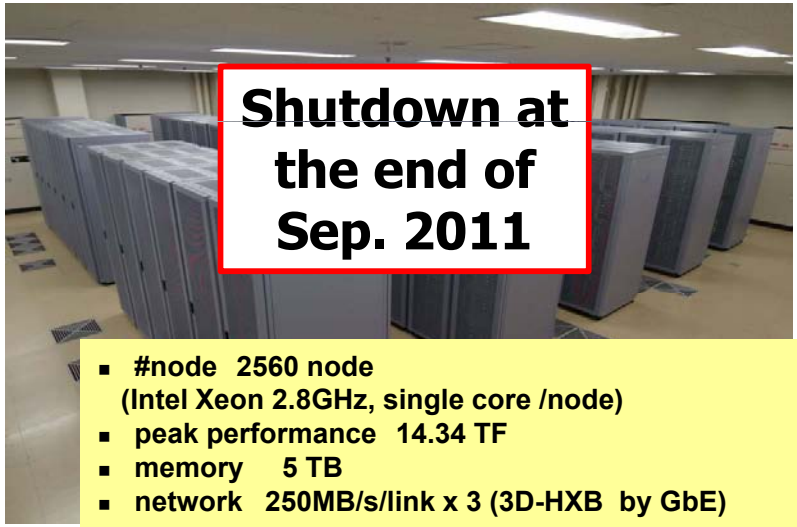

Shutdown at the end of Sep. 2011

- #node 2560 node
  (Intel Xeon 2.8GHz, single core /node)
- peak performance 14.34 TF
- memory 5 TB
- network 250MB/s/link x 3 (3D-HXB by GbE)

FIRST (2007～)

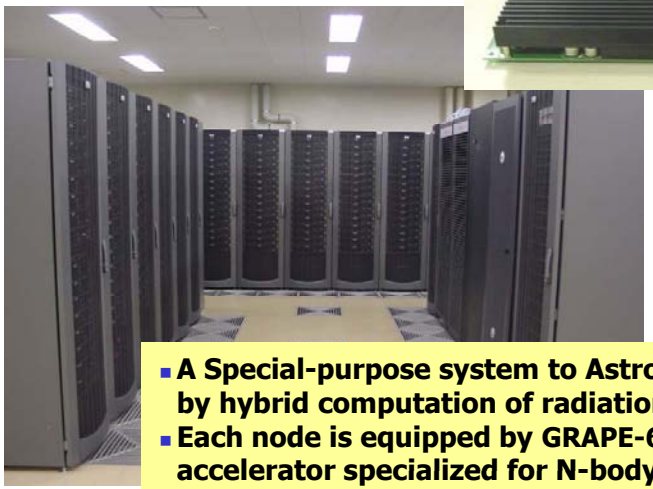

GRAPE-6

- A Special-purpose system to Astrophysics simulation by hybrid computation of radiation and N-body.
- Each node is equipped by GRAPE-6, which is an accelerator specialized for N-body Gravity calculation.
- 256 nodes
- performance: cluster 3.5TFLOPS＋Grape-6 35TFLOPS

T2K-tsukuba (2008～)

Designed by T2K Open Supercomputer Alliance (U. Tokyo and Kyoto U)

Spec;
- 648 nodes
  (quad Opteron, 4sockets/node)
- 10000 cores
- Peak performance 95.4TF
- total memory 20TB
- total disk capacity 800TB

(20th in top 500, June, 2008)

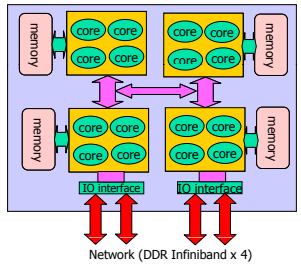

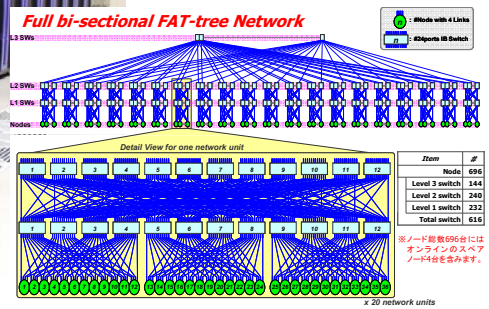

HP-PACS system (base-cluster)

- **System spec.**
  - **268 nodes**
  - **CPU 89TFLOPS + GPU 713TFLOPS = total 802TFLOPS**
  - **Memory 34TByte、memory bandwidth 26TByte/sec**
  - **Bi-section bandwidth 2.1TByte/秒**
  - **Storage 504TByte**
  - **Power 408kW**
  - **26 ranks, Installed on Jan, 2012**
  - **Operation started from Feb, 2012**

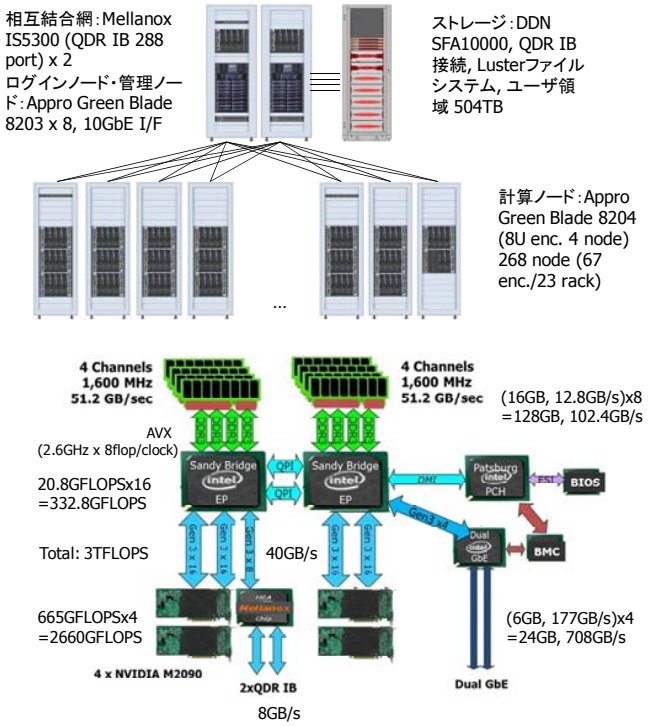

- **Spec of compute nodes**
  - **CPU x2 + GPU x4/node**
    （4 GPU: 2660 GFLOPS + 2 CPU: 332 GFLOPS~ 3 TFLOPS / node）
  - **Advanced CPU：Intel SandyBridge： high-peak performance enhanded by 256bit AVX instruction, and high memory bandwidth by 1600MHz DDR3 (2.6GHz SandyBridge-EP (8 core) = 166.4 GFLOPS, 51.2GB/s memory bandwitdh, 128GB)**
  - **x40 lane for CPU direct I/O of PCIe Gen3**
  - **Advanced GPU：NVIDIA M2090: M2070 512core enhance version：peak performance 665GFLOPS**
  - **Interconnect network**
    - **Infiniband QDR x 2 rail (trunk)**
    - **Connected by PCIe Gen3 x8 lane**

HA-PACS project

- HA-PACS (Highly Accelerated Parallel Advanced system for Computational Sciences)
- Funded by MEXT, The objective is exploiting technologies and applications for exascale computing
- Research topics
  - 1. Code development of the next-generation computational science applications for exascale (3 important area)
  - 2. Design of system architecture for exascale: Direct interconnect between GPU
  - 3. Programming environment for exascale (XcalableMP device extension)
- Project organization
  - Project Office for Exascale Computing System Development
  - Project Office for Exascale Computational Sciences

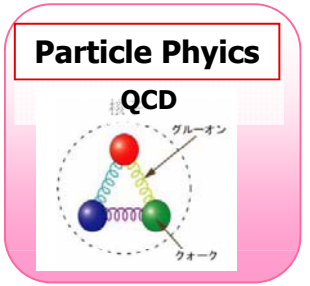

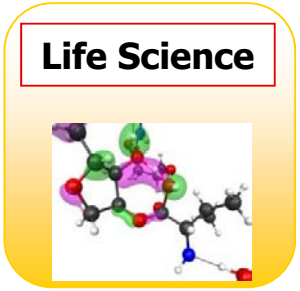

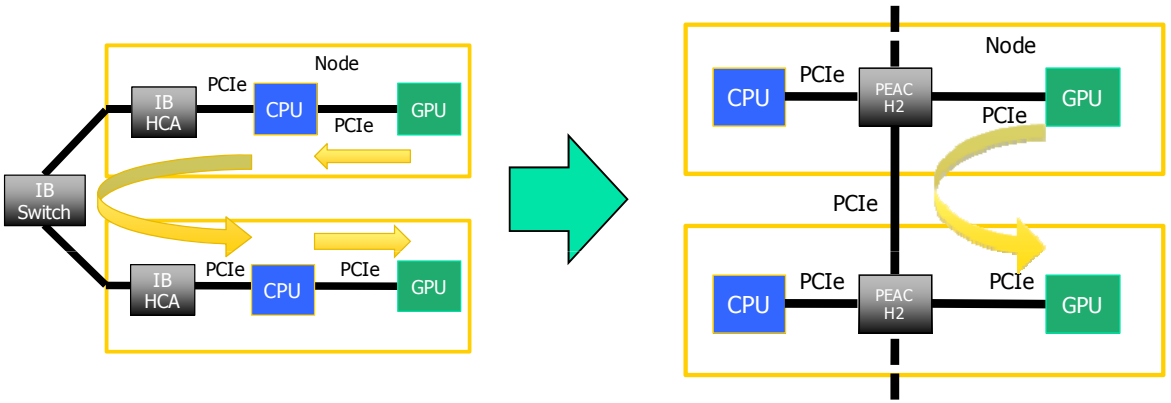

True GPU-direct
With cooperation of NVIDA

Educational Activities in CCS

- HPC Seminar
  - This seminar presents knowledge, methods and techniques for programming modern high performance computer systems, including recent microprocessors, and its performance turning, parallel programming.
  - Participants: researchers and users of computational science (including researchers in companies)
  - Periods: 2 or 3 days in summer season
    - Held since July 2007
  - Also broadcasted via internet

- Campus-wide courses on “computational sciences” for graduate students
  - Faculty members of CCS give lectures
  - Accredited as “unit” in graduate courses.
  - Courses
    - Computational Science Literacy
    - High Performance Parallel Computing Technology for Computational Sciences (overlapped with HPC Seminar)
  - Started from 2008

- Computational Science Dual Degree (double major) Program
  - Enables a graduate student in a doctoral program to simultaneously belong to a masters program of a different Graduate School, and receive both a doctoral degree in science and a masters degree in computer science, or *vice versa,* upon graduation.
  - Design of curriculum and courses for advanced computational science
  - Educate researchers who can push forward new Interdisciplinary computational science from global viewpoints
  - Started in 2009 (Physics in Doctor course and Computer sciences in Master course)
  - To be expanded to “Environmental Science and Biological Science”

- Computational Science Courses in English in “Global 30 Program”
  - To accept International Students.

Computational Science Dual Degree Program

- Computational Science Dual Degree Program fosters qualified researchers in interdisciplinary computational science with research ability in both computational and computer sciences.
  - We had 8 students so far in this course.

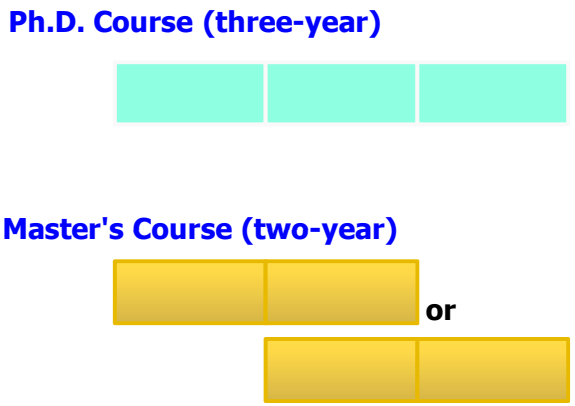

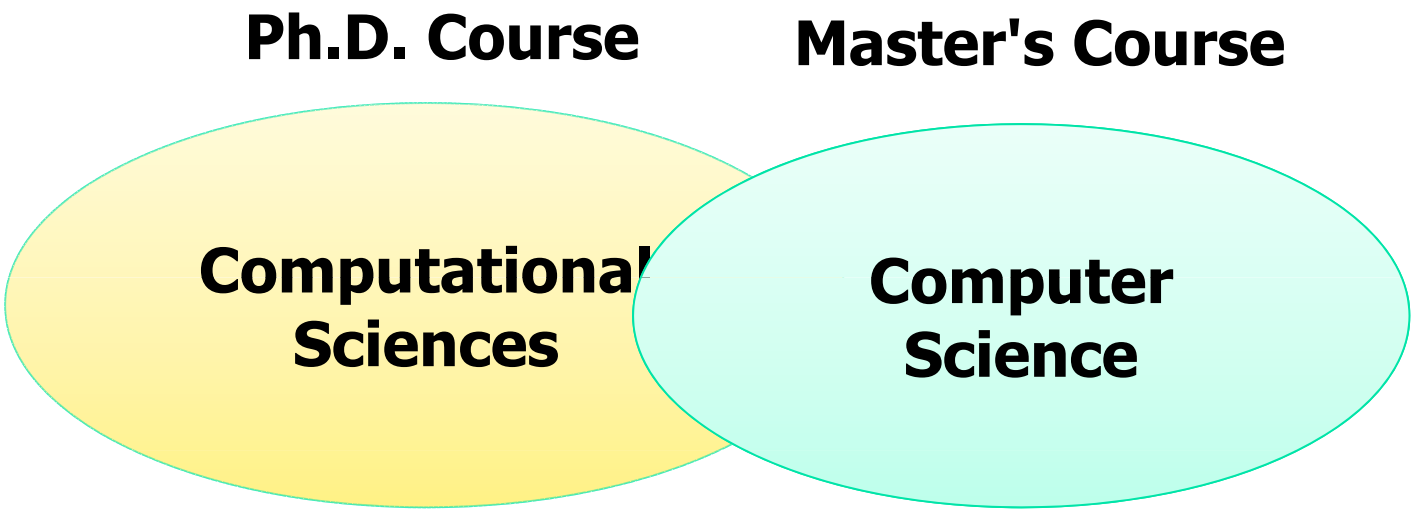

MEXT Educational Project for Establishing Core Universities for Internationalization (Global 30)

- The Ministry of Education, Culture, Sports, Science and Technology (MEXT) has launched the Global 30 Project for Establishing Core Universities for Internationalization.
- In 2009, thirteen universities were selected to lead Japan's internationalization.

Selected Universities

National	Private
Tohoku University	Keio University
University of Tsukuba	Sophia University
The University of Tokyo	Meiji University
Nagoya University	Waseda University
Kyoto University	Doshisha University
Osaka University	Ritsumeikan University
Kyushu University	

Global 30 Education Programs in University of Tsukuba

- Offers the opportunity for undergraduate/graduate international students to take coursework in English.
  20 Course Programs
- Improve campus environment for international students:
  About 4,000 dormitory rooms available, international students are given priority.
- Financial aid for Japanese language and culture learning opportunities, and career support systems will be expanded.
- Furthering internationalization.
  - Exchanging students and scholars.
  - Encourage Japanese students to study abroad.

List of Classes in Global 30 Computational Science DD Program

- Computational Science Literacy
- High Performance Parallel Computing Technology for Computational Sciences
- Advanced Course in Computational Algorithms
- Special Lecture on Numerical Simulation
- Programming Environment
- Data Engineering
- Advanced Course in High Performance Computing
- Data Analysis

SC DD program and G30

- From 2009 Computational Science Dual Degree Program for Ph.D. course
  - Ph.D. Course: Graduate School of Pure and Applied Sciences, Physics & Frontier Science
    Graduate School of Life and Environmental Sciences (Geoenvironmental Sciences)
  - Master's Course:Graduate School of Systems and Information Engineering, Computer Science
- From 2011 Global 30 English Program for international students
- Exchange students: dual degree, internship, fieldwork, credits, etc.
- Promote international collaborations

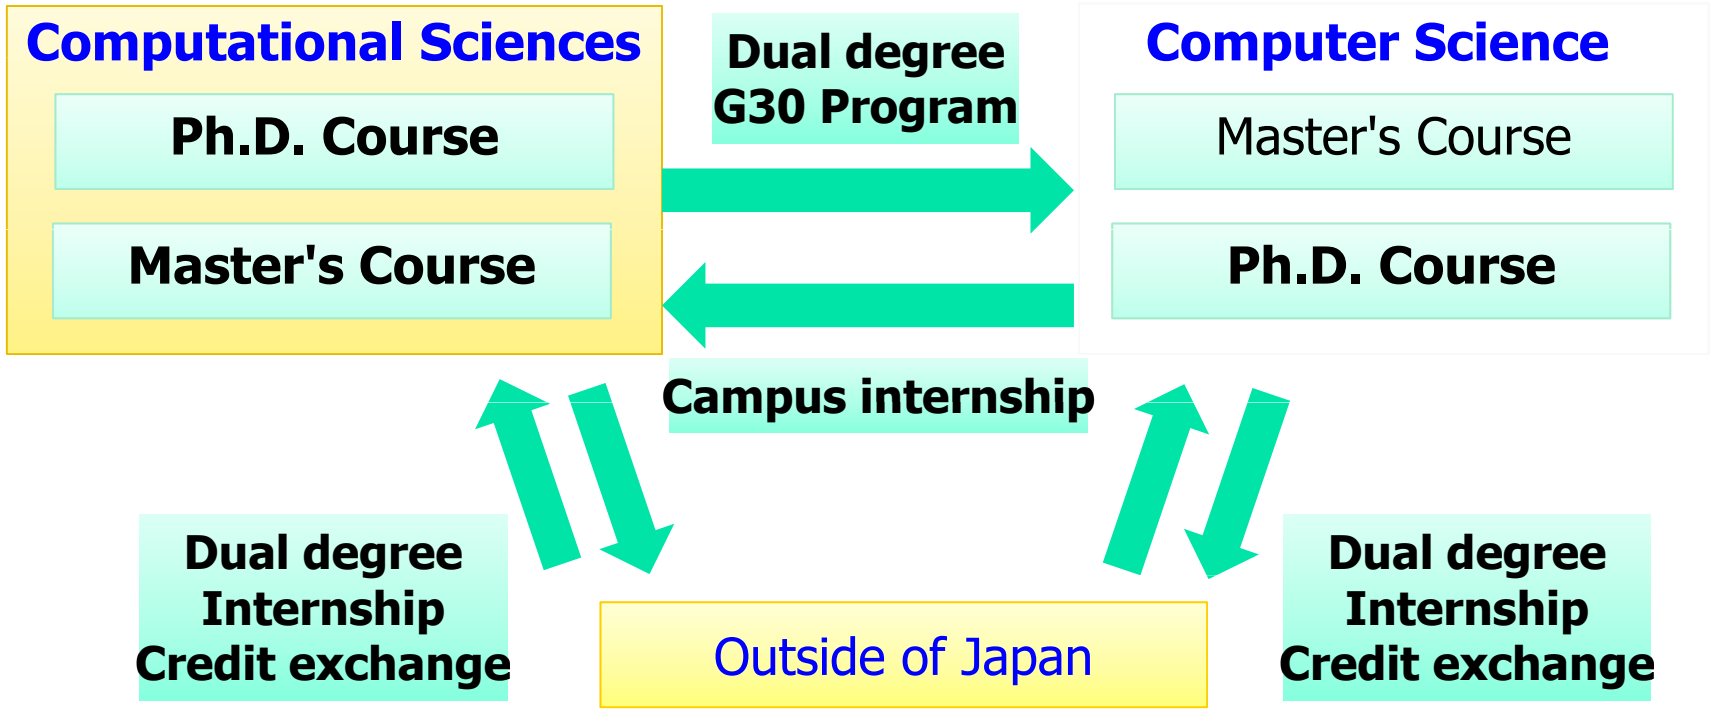

Now, the K computer is fully installed

- The first eight racks of the K computer were delivered to Kobe from Fujitsu on September 28, 2010. More than 800 racks are required for a 10 Peta Flops Performance.
- The system wons 1st of Top 500 List at June and November 2011
- The system has been fully installed and under system adjustment. It will be open to users from 4Q

Photo of First d

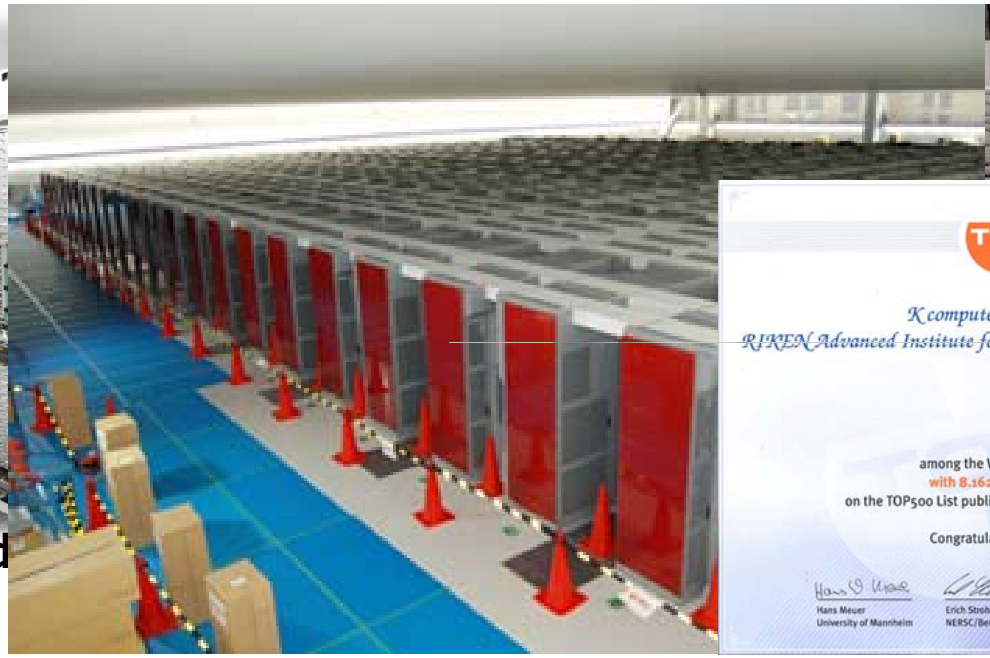

Schedule of development

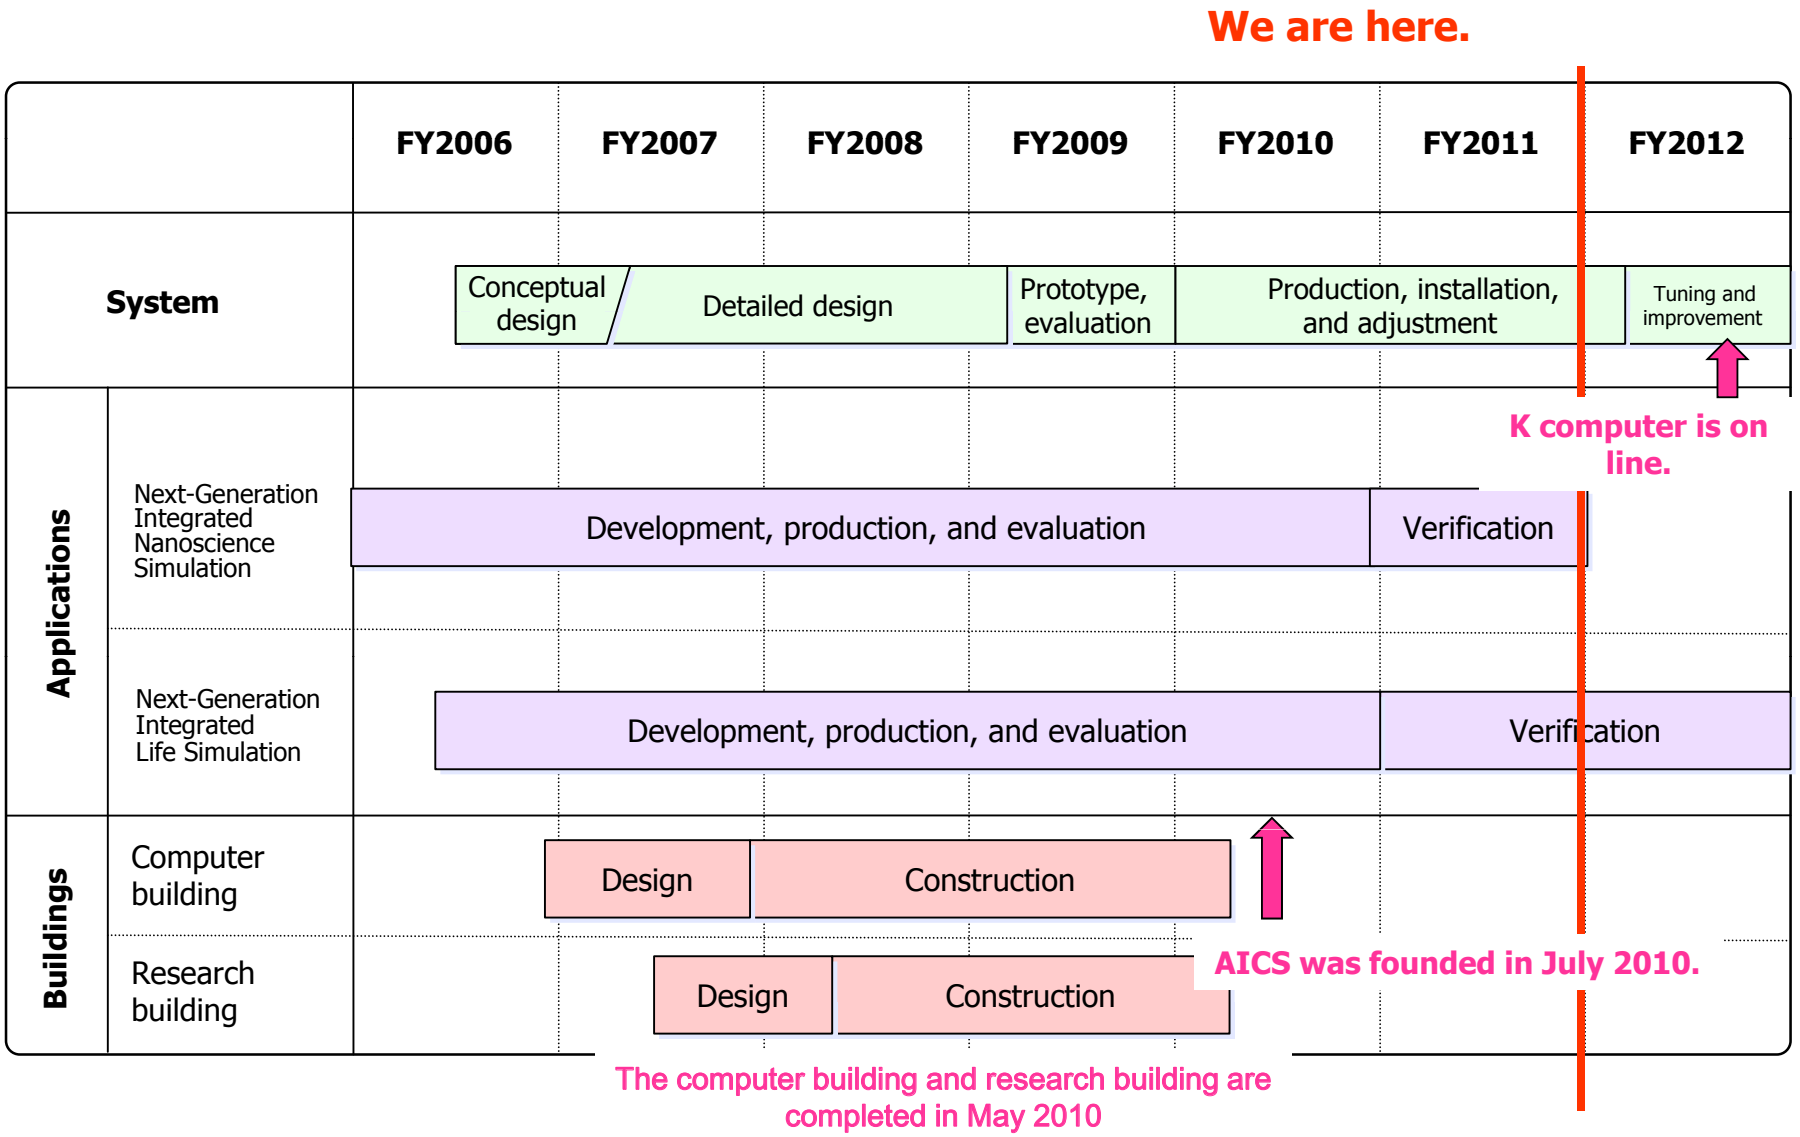

Contributions for NGS (the K computer project)

As a leading research organization, it is very important issue for us to take a leadership concerning to building/operation of the petascale system in national-wide computing infrastructure.

- Collaboration with RIKEN on the petaflops system development
  - Formal agreement between U. Tsukuba and RIKEN signed in September 2006
  - Participation of several CCS faculty in the system design as a concurrent researchers
  - Agreement of actual collaboration in tuning of representative petascale applications (QCD, RS-DFT and FFT)
  - **RS-DFT on the K computer won the SC11 Gordon-Bell Prize !!!**
- CCS is a core organization of NGS strategic-use for "Field 5: The origin of matters and the universe"
- One of site of Japanese High-performance Computing Infrastructure (HPCI)

How to organize users of NGS

- Strategic Use: MEXT selected 5 strategic fields from national viewpoint.
  - Field 1: Life science/Drug manufacture
  - Field 2: New material/energy creation
  - Field 3: Global change prediction for disaster prevention/mitigation
  - Field 4: *Mono-zukuri* (Manufacturing technology)
  - Field 5: The origin of matters and the universe
- MEXT Funds 5 core organizations that lead research activities in these 5 strategic fields

More than 50% resources will go to strategic use

- General Use:

The use for the needs of the researchers in many science and technology fields including industrial use and educational use

CCS was selected as a core organization of NGS for "Field 5: The origin of matters and the universe"

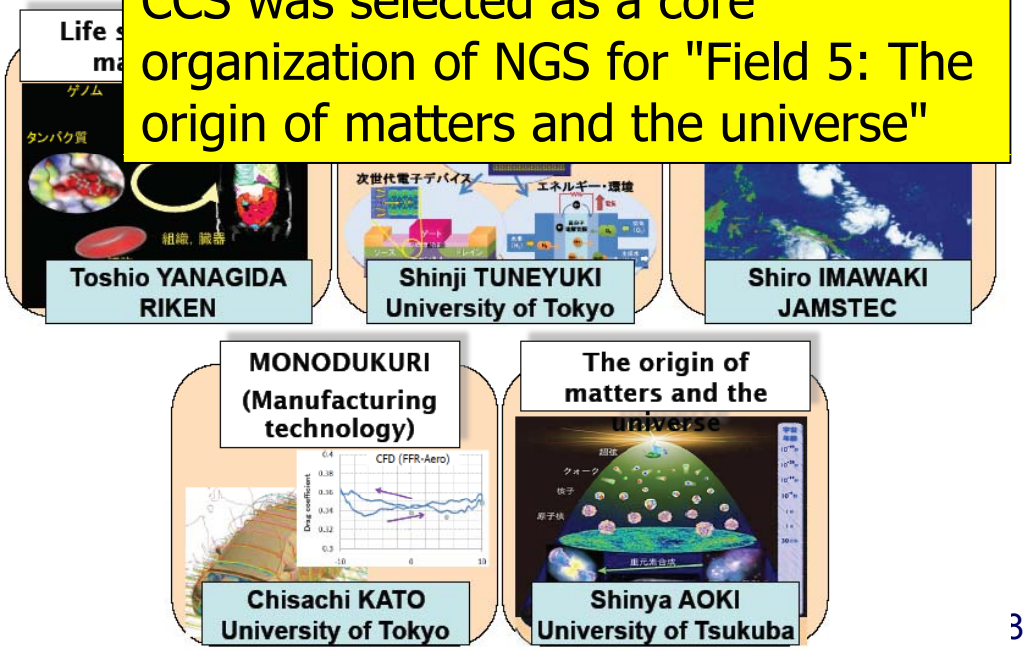

Consortium and High-performance Computing Infrastructure (HPCI)

- Background:
  - The goal of the NGS has been reconsidered by the new government for accountability for "taxpayers" : “Creation of the Innovative High-Performance Computing Infrastructure (HPCI)”.
- HPCI: High-Performance Computing Infrastructure
  - Integrated operation of NGS with other institutional supercomputers
  - Seamless access from supercomputers and user's machines to NGS.
  - Large-scale storage systems shared by NGS and others.
- HPCI (or HPC) Consortium
  - To play a role as a main body to run HPCI (and design HPCI).
  - To organize computational science communities from several application fields and institutional/university supercomputer centers.
    - Including Kobe Center

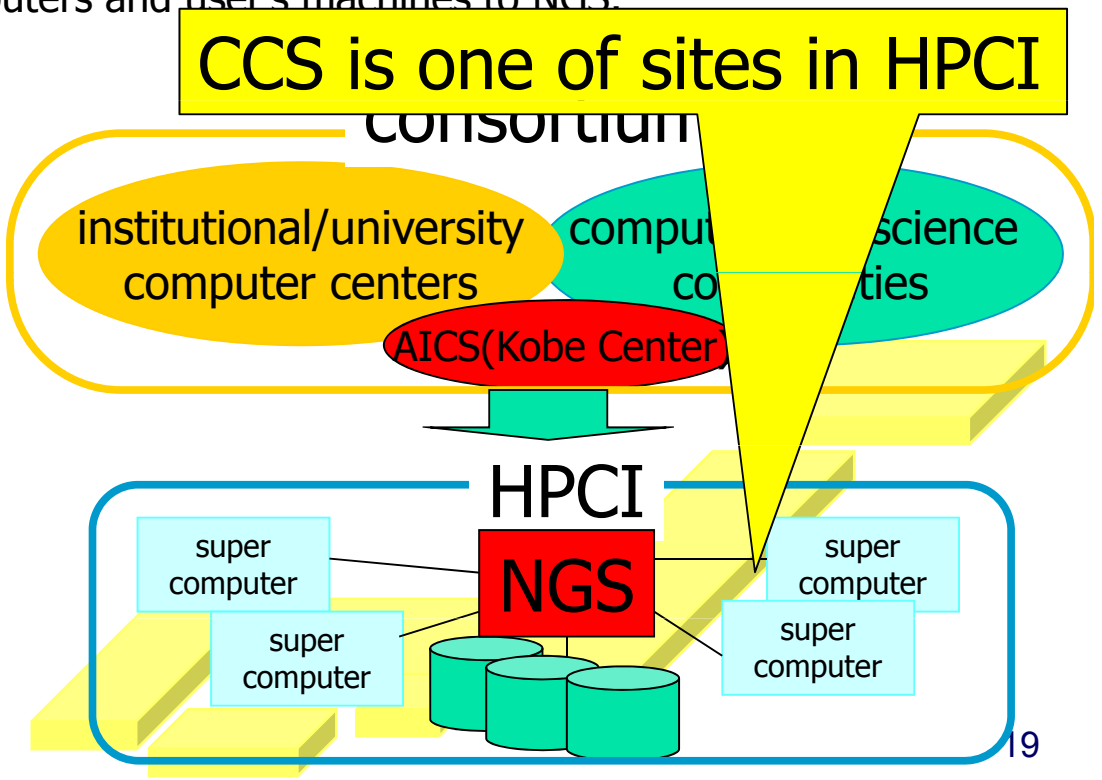

What's the next of the K computer?

Players and Projects in Japan **NOW**

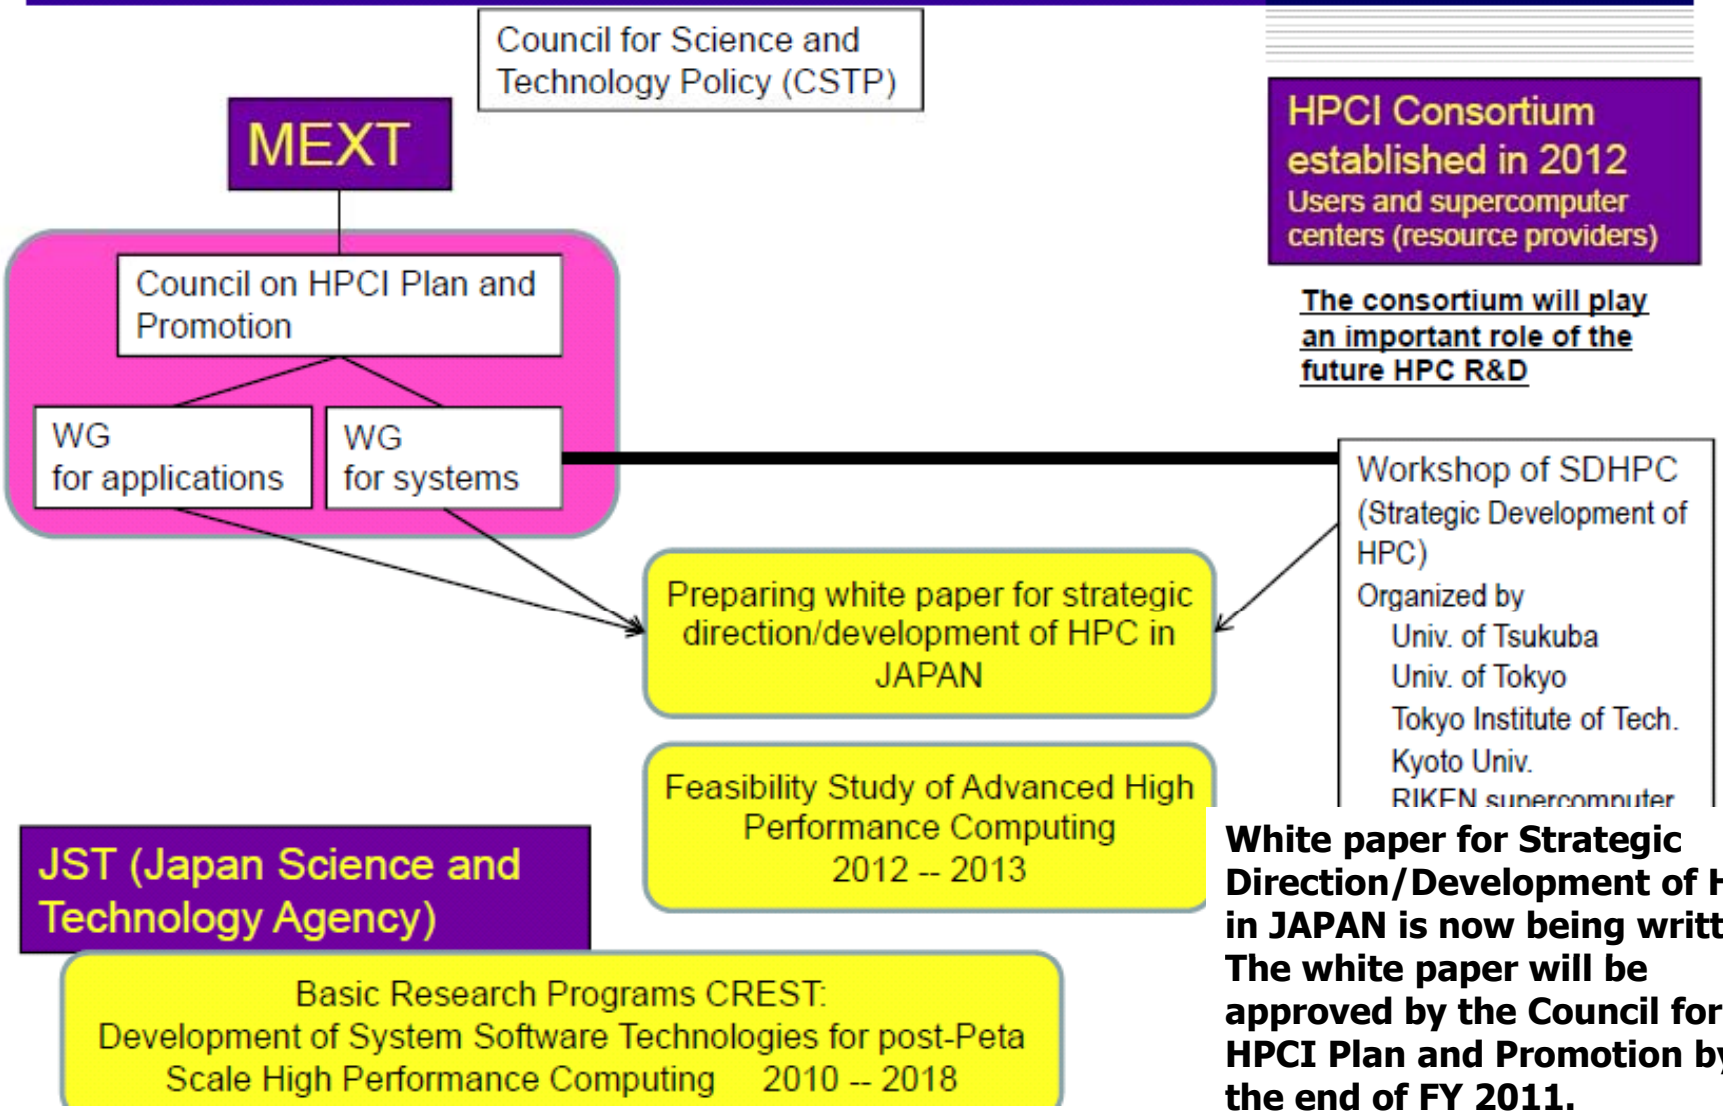

White paper for Strategic Direction/Development of HPC in JAPAN is now being written. The white paper will be approved by the Council for HPCI Plan and Promotion by the end of FY 2011.

Feasibility Study of Advanced High Performance Computing

- MEXT (Ministry of Education, Culture, Sports, Science, and Technology) has just proposed two-year project for feasibility study of advanced HPC which will start in FY2012 (April of 2012)

- Objectives
  - The high performance computing technology is an important national infrastructure
  - Keeping development of top-level HPC technologies is one of Japanese international competitiveness and contributes national security and safety
  - This two-year project is to study feasibilities of such development that Japanese community should focus on

- Budget requested 11.45M USD (859M JPY) / Year

XcalableMP

http://www.xcalablemp.org

XcalableMP : directive-based language eXtension for Scalable and performance-aware Parallel Programming

- A PGAS language. Directive-based language extensions for Fortran and C for the XMP PGAS model
  - To reduce the cost of code-rewriting and education
- Global view programming with global-view distributed data structures for data parallelism
  - A set of threads are started as a logical task. Work mapping constructs are used to map works and iteration with affinity to data explicitly.
  - Rich communication and sync directives such as "gmove" and "shadow".
  - Many concepts are inherited from HPF
- Co-array feature of CAF is adopted as a part of the language spec for local view programming (also defined in C).

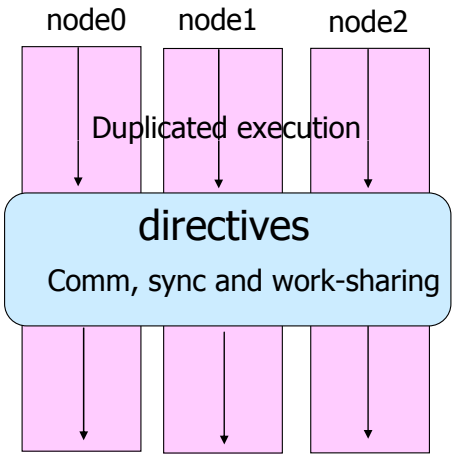


```
int array[N];
#pragma xmp nodes p(4)
#pragma xmp template t(N)
#pragma xmp distribute t(block) on p
#pragma xmp align array[i][ with t(i)

#pragma xmp loop on t(i)  reduction(+:res)
  for(i = 0; i < 10; i++)
     array[i] = func(i,);
     res +=    ;
   } }
```

Extension of XcalableMP

- Extension for mulitcore nodes：clause to specify parallelism in nodes and mixed description with OpenMP
- Parallel Library interface to call high-performance parallel libraries written in MPI
- XMP-IO: IO for distributed array
- XMP-dev: Extension for acceleration devices such as GPU
  - Hide complicated communication by reflect operation on distributed array allocated on GPUs
  - 21.6 times speedup on Laplace 4GPU(Tesla C2050) in Laplace solver.

Break down of execution time of Laplace solver (4K*4K) in XMP-dev

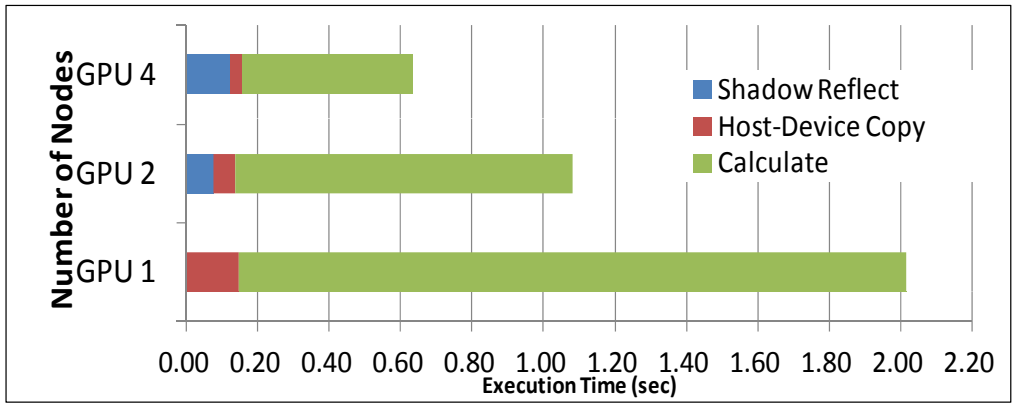


```
#pragma xmp loop on t(i) threads firstprivate (x) num_threads (4)
for (int i = 0; i < N; i++) {  out[i] = x * in[i];   }
```

```
#pragma xmp device replicate(u, uu)
{
#pragma xmp device replicate_sync in (u)
  for (k = 0; k < ITER; k++) {
#pragma xmp device reflect (u)
#pragma xmp device loop (x, y) on t(x, y) threads
    for (y = 1; y < N-1; y++)
      for (x = 1; x < N-1; x++)
        uu[y][x] = (u[y-1][x] + u[y+1][x] +
                    u[y][x-1] + u[y][x+1]) / 4.0;
#pragma xmp device loop (x, y) on t(x, y) threads
    for (y = 1; y < N-1; y++)
      for (x = 1; x < N-1; x++)
        u[y][x] = uu[y][x];
  }
#pragma xmp device replicate_sync out (u)
} // #pragma xmp device replicate
```

XMP-dev: XcalableMP acceleration device extension [U of Tsukuba]

- Offloading a set of distributed array and operation to a cluster of GPU

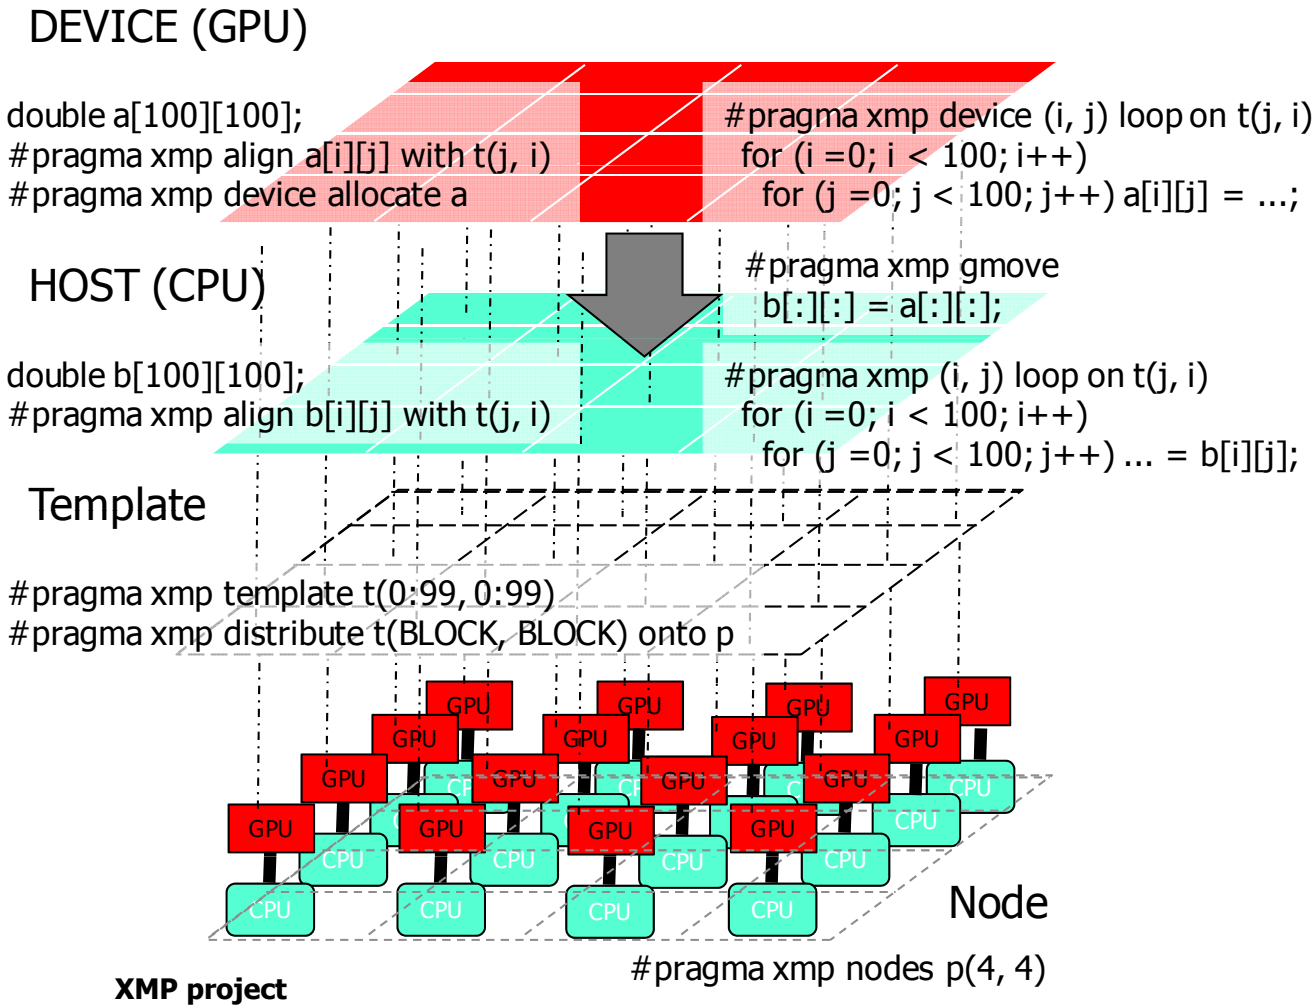

History and Plan of CCS computing facilities

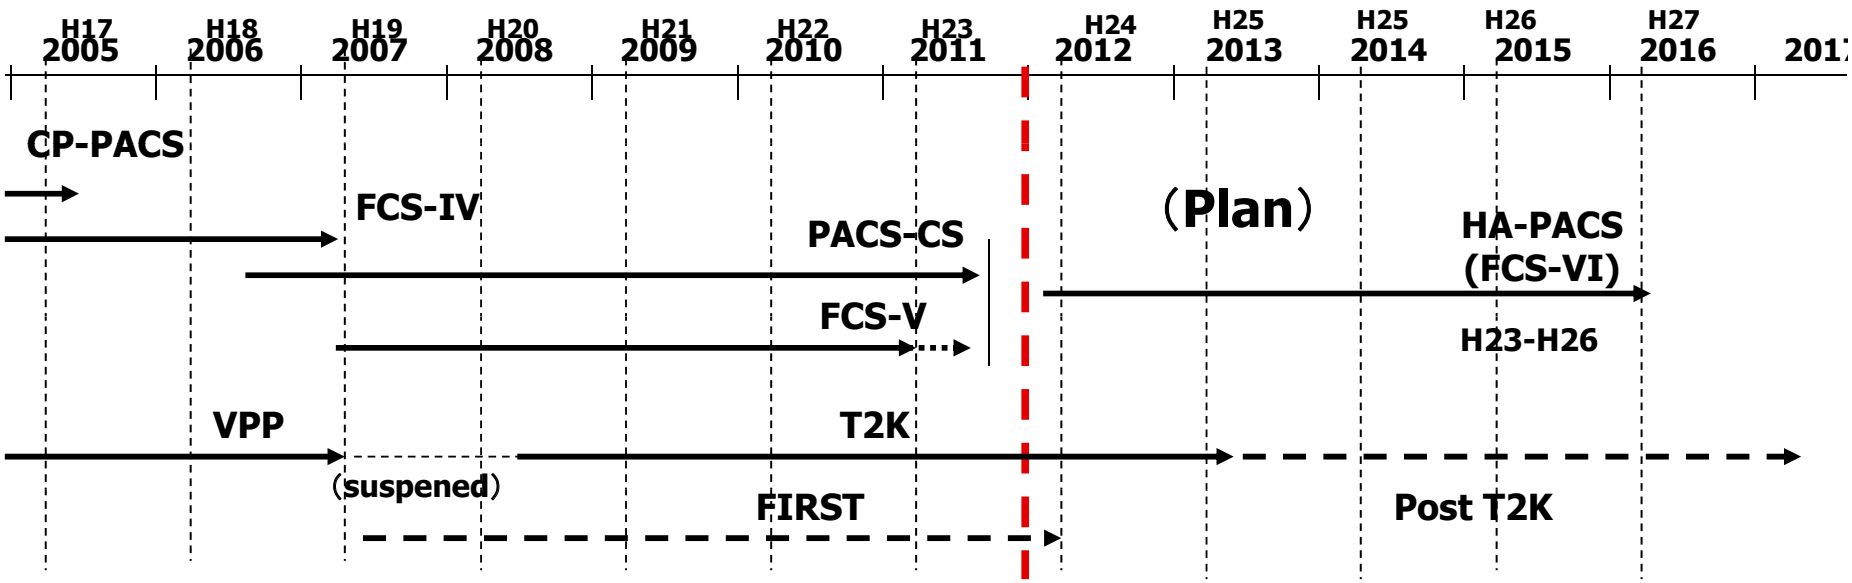

- **Agenda: Post-T2K system**
  - Extension of HA-PACS?
  - Intel MIC based?
  - Targets may be more than 4PF?