

The Dawn of the Japanese “Post Petascale” Era

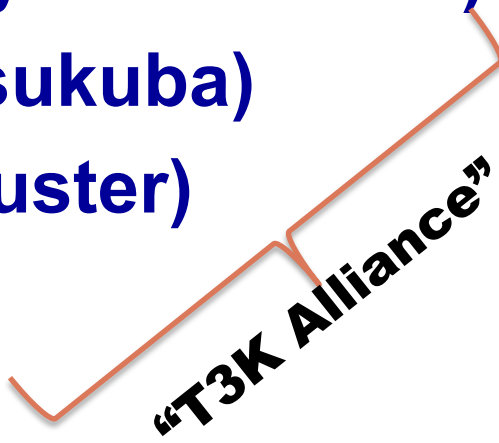
**Satoshi Matsuoka (Tokyo Inst. Tech.)
Yutaka Ishikawa (The Univ. Tokyo)**

Seemingly Sad State of Affairs for SCs in Japan circa 2010...

18 entries found.

Rank	Site	System	Cores	R _{max}	R _{peak}
22	Japan Atomic Energy Agency (JAEA) Japan	BX900 Xeon X5570 2.93GHz , Infiniband QDR Fujitsu	17072	191.4	200.08
37	Japan Agency for Marine -Earth Science and Technology Japan	SX-9/E/1280M160 NEC	1280	122.4	131.07
42	JAXA Japan	Fujitsu FX1, Quadcore SPARC64 VII 2.52 GHz, Infiniband DDR Fujitsu	12032	110.6	121.28
53	Information Technology Center, The University of Tokyo Japan	Hitachi opteron QC 2.3 GHz Myrinet 10G Hitachi	15104	101.74	138.96
56	Institute of Physical and Chemical Res. (RIKEN) Japan	RIKEN Intergrated Cluster of Clusters, Xeon X5570 2.93GHz, Infiniband DDR Fujitsu	9048	97.94	106.04
64	GSIC Center, Tokyo Institute of Technology Japan	Sun Fire x4600/x6250, Opteron 2.4/2.6 GHz, Xeon E5440 2.833 GHz, ClearSpeed CSX600, nVidia GT200; Voltaire Infiniband NEC/Sun	31024	87.01	163.19
71	Center for Computational Sciences, University of Tsukuba Japan	Appro Xtreme-X3 Server - Quad Opteron Quad Core 2.3 GHz, Infiniband Appro International	10368	77.28	95.39
91	National Institute for Fusion Science (NIFS) Japan	Hitachi SR16000 Model L2/128, Power6 4.7Ghz, Infiniband Hitachi	4096	56.65	77
96	University of Tokyo/Human Genome Center, IMS Japan	SunBlade x6250, Xeon E5450 3GHz, Infiniband Sun Microsystems	5760	54.21	69.12

Petascale Machines in Japan will be arriving fast circa 2012

- TSUBAME2.0 (2010Q4, 2.4PF Tokyo Inst. Tech.)
 - HA-PACS (2011Q4, ~1PF, Univ. Tsukuba)
 - Univ. Tokyo (2012Q1 ~1PF non-cluster)
 - Kyoto-U (2012Q2, 0.7PF?)
- 
- “T3K Alliance”
- “Kei” (2012Q2? 10PF, Kobe National Facility)
 - DDBJ (2012Q1 >20PBytes, N Genomics Inst.)
 - KEK (2012, ???, post BG/L)
 - Most facilities power constrained

TSUBAME2.0 Operational Nov.1, 2010



TSUBAME2.0: A GPU-centric Green 2.4 Petaflops Supercomputer

Tsubame 2.0: "Tiny" footprint, very power efficient

- Floorspace less than 200m² (2,100 ft²)
- Top-class power efficient machine on the Green 500

System

(42 Racks)

1408 GPU Compute Nodes,

34 Nehalem "Fat Memory" Nodes

Rack

(8 Node Chassis)



2.4 PFLOPS
80 TB

Node Chassis

(4 Compute Nodes)



6.7 TFLOPS
220 GB/412 GB

Compute Node

(2 CPUs, 3 GPUs)



1.6 TFLOPS
55 GB/103 GB

Chip

(CPU, GPU)



CPU(Westmere EP)
76.8 GFLOPS




GPUs(Tesla M2050)
515 GFLOPS
3 GB

Integrated by NEC Corporation

HA-PACS (Highly Accelerated Parallel Advanced system for Computational Sciences), U. Tsukuba

- GPU-base cluster with up to **1PFLOPS** peak
- High density of **multi-GPU** nodes
- Experimental technology for **GPU-GPU and node-node direct communication** based on **PCI-E active switch** by newly developed **PEACH** communication chip (partially)
- System delivery on **Dec. 2011**



Conceptual image of GPU-direct communication by PEACH (PCI-E Adaptive Communication Hub) utilizing PCI-E link for GPU-GPU, GPU-CPU and node-node communication

Evidence of Supercomputing Populism

“The Recent Japanese NextGen Kobe SC Project Cancellation Fiasco”

- May 2009 (precursor) – NEC&Hitachi announces pullout from Japanese NextGen Kobe SC Project
 - Result – public outcry denouncing NEC as “loser” “traitor”
- Sep. 2009 – new “Democratic Party” took over, slated to eliminate govt. waste
- Nov. 2009 – committees set-up, populated with “experts”(???) to review numerous projects and institutions out of the blue, just an hour each
- Nov. 16, 2009 – The NextGen project was reviewed, recommendation: “freeze (zero budget)” effectively killing the project(!)
 - “Why must we be #1? Why can’t we be just #2?”



Retaliation of Scientists and the Public

- Immediate reaction by numerous academic societies – Physics, Informatics, Mechanical Engineering, ...
- Press conference by famous Nobel laureates denouncing government's decision
 - Head of Riken “Do the reviewers have the guts to be stand the trial of history”?
- Public outcry blaming the government for sacrificing science and engineering, endangering Japan's core competence
 - “It is worthless to aim for #2 in Science and Engineering”
- Due to public pressure, the project was resurrected
 - “SUPACON = Supercomputer” becomes a household terminology
- But renamed to HPCI (High Performance Comp. Infrastructure)
 - HPCI = Kobe project + HPCI Consortium



Consortium and High-performance Computing Infrastructure (HPCI)

- **Background:**

- Goal NGS “reconsidered” by the new government for accountability for “taxpayers” : “Creation of the Innovative High-Performance Computing Infrastructure (HPCI)”.

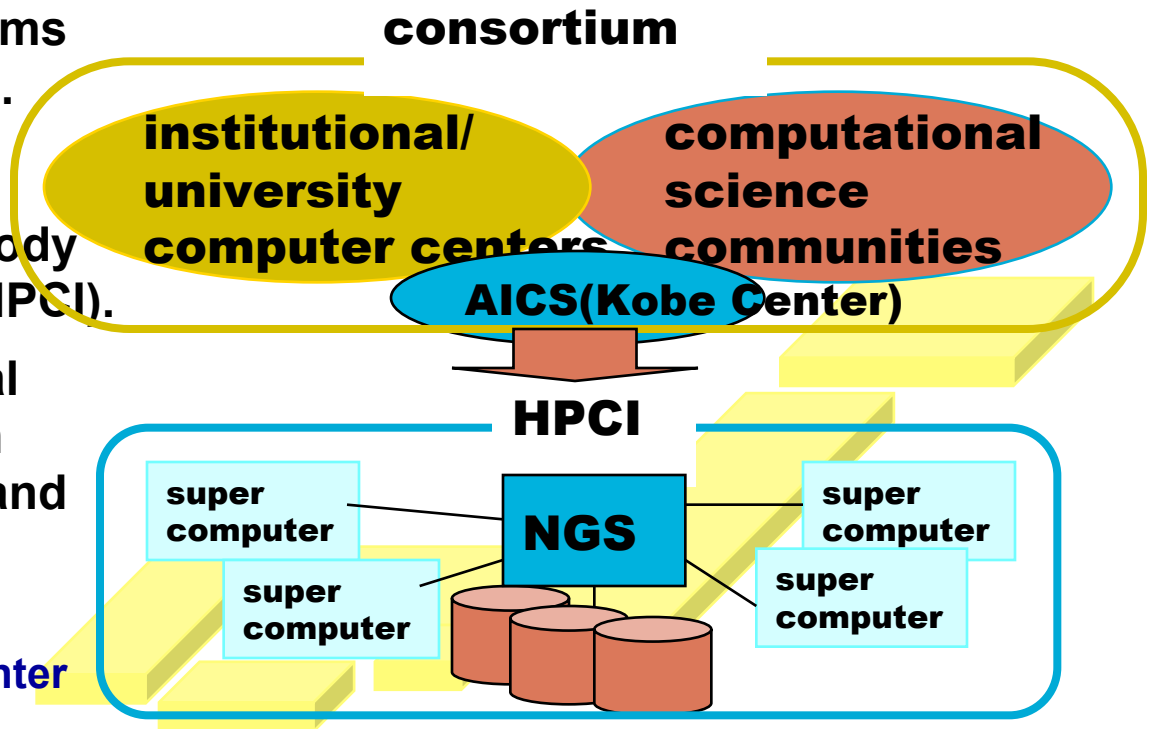
- **HPCI: High-Performance Computing Infrastructure**

- 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 the Kobe Center



HPCI Consortium

- HPCI = Kobe project + HPCI Consortium
- **Consortium Still in Planning Stage**
- **Expected to be a PRACE or TeraGrid-like organization, with the Kobe facility as a centerpiece**
- **Membership approval already done**
 - 25 “Resource Provider” Organizations (SCs)
 - 13 “User Community” Organizations

Still we got our work cut out...

Post Petascale Projects Just Started...

計算科学研究機構
Advanced Institute for Computational Science

TOP 機構概要 ライブラリ アクセス・お問合せ リンク

世界に誇る
拠点を目標して

計算科学研究機構 (AICS)。
それは、オールジャパン体制で臨む日本の計算科学の中心となる拠点。

10ペタフロップス級という世界最高水準の京速コンピュータ「京」を利用し、さまざまな分野における研究成果からブレイクスルーをもたらし、社会に貢献します。

また、次代を担う人材の育成にも積極的に取り組んでいきます。

トピックス バックナンバーはこちら>>>

2010. 08. 12 神戸新聞ジュニア記者が来訪。
2010. 07. 05 次世代スパコンの愛称が決定しました。
2010. 07. 01 計算科学研究機構が発足しました。

ご挨拶>>>

 計算科学研究機構 機構長
平尾 公彦

RIKEN WEB SITE
理化学研究所
ホームページ

次世代スパコン開発実施本部
ホームページ

Basic Research Programs

TOP Research Areas Overall Schedule How to Apply Contact Address JAPANESE

Japan Science and Technology Agency TOP » Development of System Software Technologies for post-Peta Scale High Performance Computing

“Outline of the Research Area” and “Research Supervisor’s Policy on Call for Application, Selection and Management of the Research Area”

[CREST]
Research area in the strategic sector:
“Creation of Basic Technologies for System Software Essential to Massive Parallel Processing (MPP) Computation with Manycore and other Processors”

Development of System Software Technologies for post-Peta Scale High Performance Computing

Research Supervisor: Akinori Yonezawa (Professor, Graduate School of Information Science and Technology, The University of Tokyo)

Outline of Research Area

The research area aims at developing system software technologies as well as related systems to be used for high performance computing in the post generations of the Japanese national supercomputer K.

More concretely, research and development will be conducted for system software enabling us to exploit maximum efficiency and reliability from supercomputers which will be composed of general purpose many-core processors as well as special purpose processors (so called GPGPU) in the second half of (and/or after) 2010's. In addition to the system software such as programming languages, compilers, runtime systems, operation systems, communication middleware, and file systems, application development support systems and ultra-large data processing systems are the targets for research and development. Also, the targets include system software in the overlapping layers of software stack, which encourages real usages of developed system software.

[TOP of this page](#)

Research Supervisor’s Policy on Call for Application, Selection and Management of the Research Area

Numerical simulation and data analysis utilizing super-scale computation and data processing are now regarded as the third methodology of science, which will play critical roles after the first and second methodologies, namely, theory and experiment/observation. Accordingly, Europe, US and China are engaged in severe competition in developing most advanced supercomputers. It is also one of the national project of developing the next-generation

**~\$30mil total,
\$3~5 mil 5 year projects x 10**

Expected funding and plans beyond NGS

- **Call for proposals for JST CREST (Core Research for Evolutional Science and Technology) for "new HPC technologies"**
 - 500-600M JPY(5M USD) for 5-7 years (4-5 B JPY in total)
~\$30 million total
 - System software-oriented, NOT hardware
 - Assume post-petascale architecture
 - Real Software deliverable pieces required, not just papers
 - Sato, Matsuoka, etc. to play proactive “advisory” role
- **Researches for the next of NGS (exascale?) will also be conducted in the AICS (Kobe Center).**
 - 2 teams currently by Ishikawa and Sato
- **In the 4th Science and Technology Basic Plan (FY2011-FY2015)**
 - Now under discussion toward¹¹ exaflops class HPC technology

International Fundings

- **France and Japan Fund: A collaborative call for proposals between "ANR-JST in ICT (Information and Communication Science and Technologies)" includes "Software and algorithm aspects of high performance computing (Axis 8)"**
 - **"Framework and Programming for Post Petascale Computing (FP3C)" has been accepted.**
 - **PI of Japan: M. Sato, PI of France: S. Petiton**
 - **Ishikawa, Boku, Nakashima, Sakurai, Matsuoka, ...**
 - **Key people from the French side: Cappello, Dayde, Namyst, Calvin, ...**
- **G8 calls**
 - **Many projects have been proposed.**
 - **Now in the second stage of the selection.**