Overview of Vector Supercomputer SX-ACE and Its Applications

Hiroaki Kobayashi
Tohoku University
koba@tohoku.ac.jp

Russian Supercomputing Days
Moscow, Russia
September 26-27, 2016

TOHOKU UNIVERSITY

WSSP

The 23rd Workshop on Sustained Simulation Performance
March 16-17, 2016 / Sendai, Japan

Contents

- Information
- Program
- > Day 1
- > Day 2
- Date & Venue
- Registration

Language: ja / en

About WSSP

Toward Future HPC Technologies

Tohoku University, Japan Agency for Marine-Earth Science and Technology (JAMSTEC), High Performance Computing Center Stuttgart (HLRS), and NEC Corporation are pleased to announce that the 23rd Workshop on Sustained Simulation Performance (WSSP) will be held on March 16th and 17th, 2016 in Sendai, Miyagi, Japan. The purpose of the workshop is to discuss future supercomputers, through the latest research efforts in large-scale computing with high performance and high efficiency.

We are looking forward to seeing you in the workshop.

Technical Program Overview

In the workshop, two keynote talks are scheduled.

Keynote Talk I

Parallel Algorithms: Theory, Practice and Education

Prof. Vladimir Voevodin
(Moscow State Univ.)

Date: March. 3. 2015
Signature

Susumu SATOMI
President, Tohoku University

Date: 3.03.2015
Signature

Victor SADOVNICHIY
Rector, Moscow State University,

TOHOKU UNIVERSITY

Missions of Cyberscience Center As a National Supercomputer Center

★ High-Performance Computing Center founded in 1969

- Offering leading-edge high-performance computing environments to academic users nationwide in Japan
  - 24/7 operations of large-scale vector-parallel and scalar-parallel systems
  - 1500 users registered in AY 2015
- User supports
  - Benchmarking, analyzing, and tuning users' programs
  - Holding seminars and lectures
- Supercomputing R&D, collaborating work with NEC
  - Designing next-generation high-performance computing systems and their applications for highly-productive supercomputing
  - 57-year history of collaboration between Tohoku University and NEC on High Performance Vector Computing
- Education
  - Teaching and supervising BS, MS and Ph.D. Students as a cooperative laboratory of graduate school of information sciences, Tohoku university

1969

1982

SX-1 in 1985

SX-2 in 1989

SX-3 in 1994

SX-4 in 1998

SX-7 in 2003

SX-9 in 2008

TOHOKU UNIVERSITY

Tohoku Univ.'s New Supercomputer System (2015.2.20~)

3D Tiled Display Wall System

1PB Lustre

10G IDS/IPS

HPCI

3PB Shared Disk

Infini band

Ethernet

IXS

Brand-new Vector System installed in 2015

68 nodes

Xeon IvyBridge Cluster installed in 2014

LX 406Re-2

24Cores

128GB

Scalar Parallel System

31.3TFlops, 8.5TB Memory

SX-ACE 2560 nodes

(5Clusters)

4Cores・64GB/node

Vector Parallel System

707TFlops,160TB Memory

HPCI: High-Performance Computing Infrastructure in JAPAN

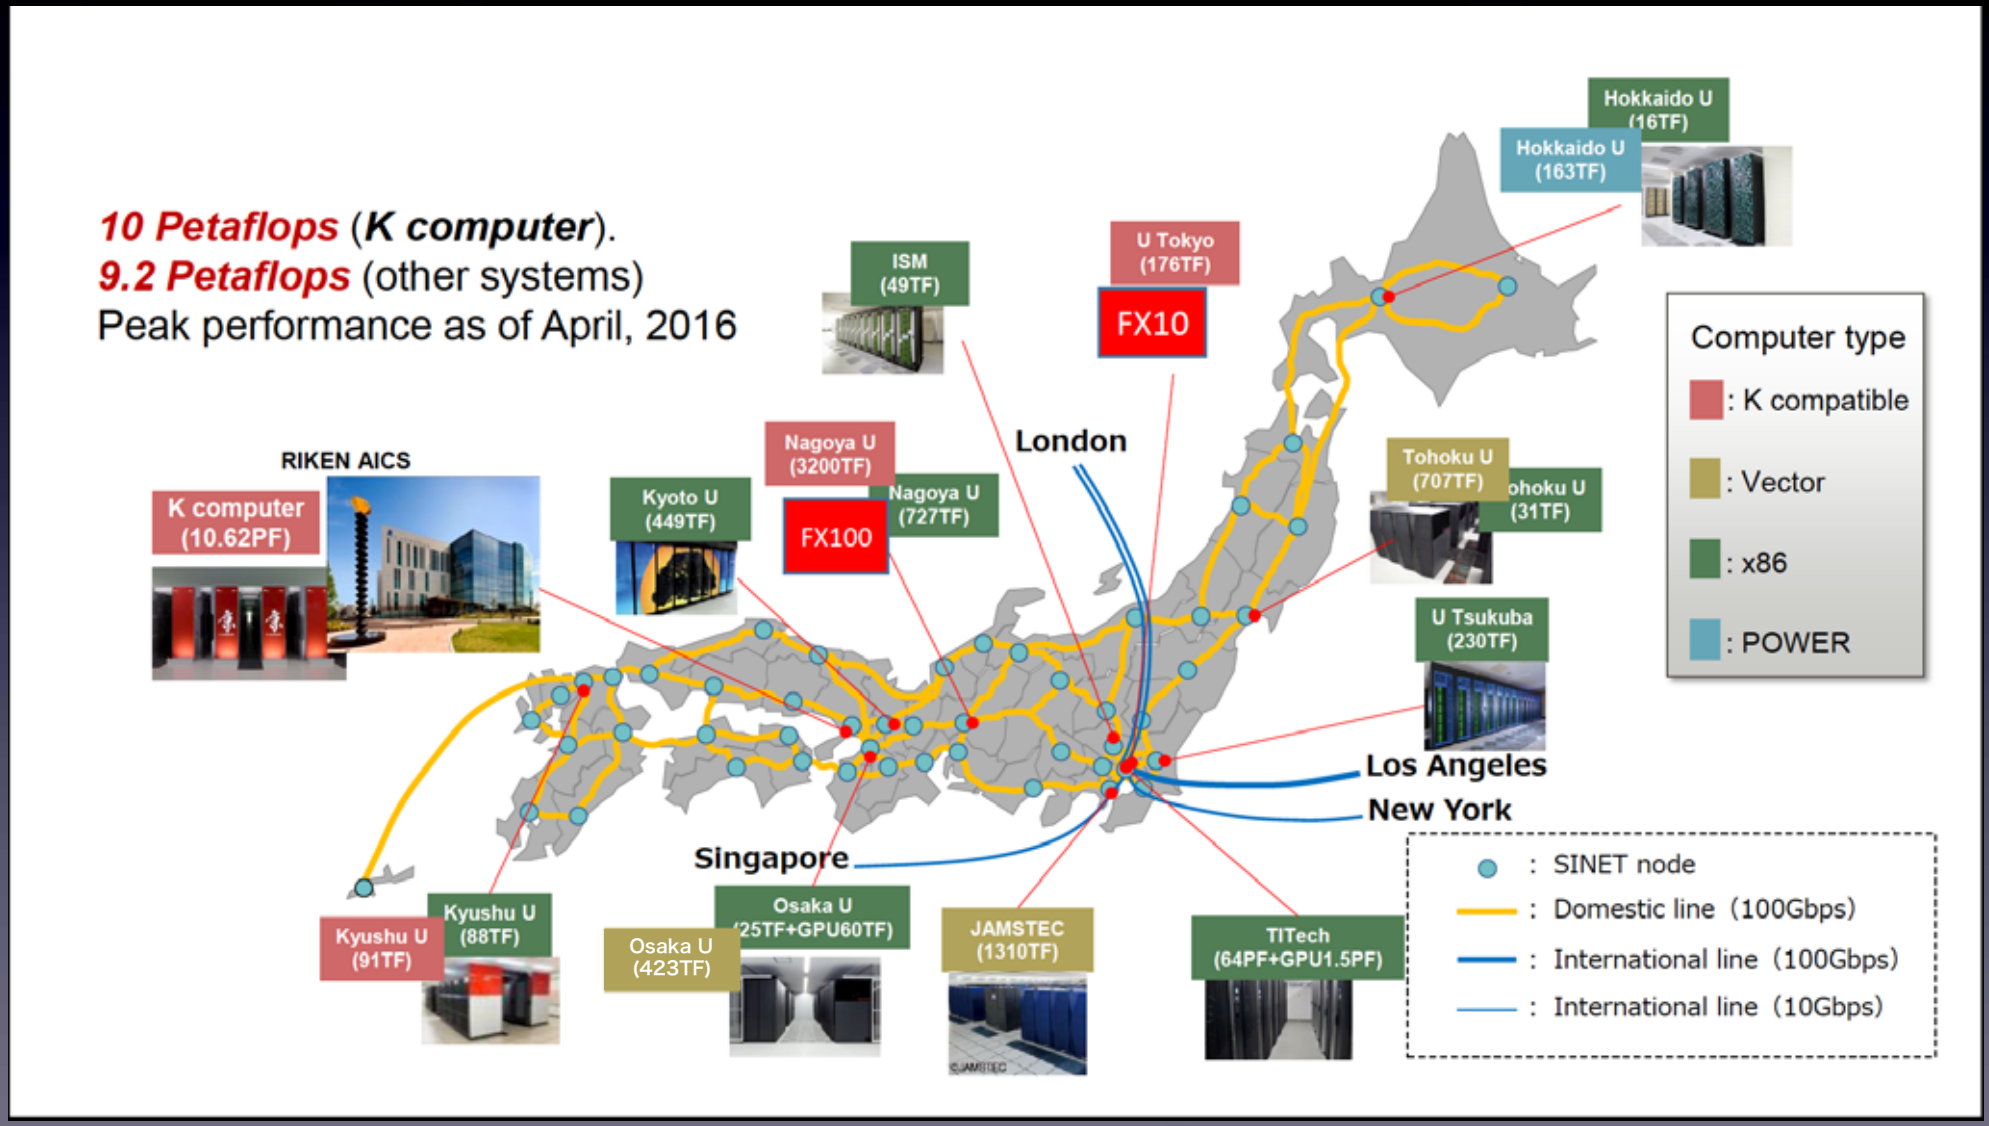

Organization of Tohoku Univ. SX-ACE System

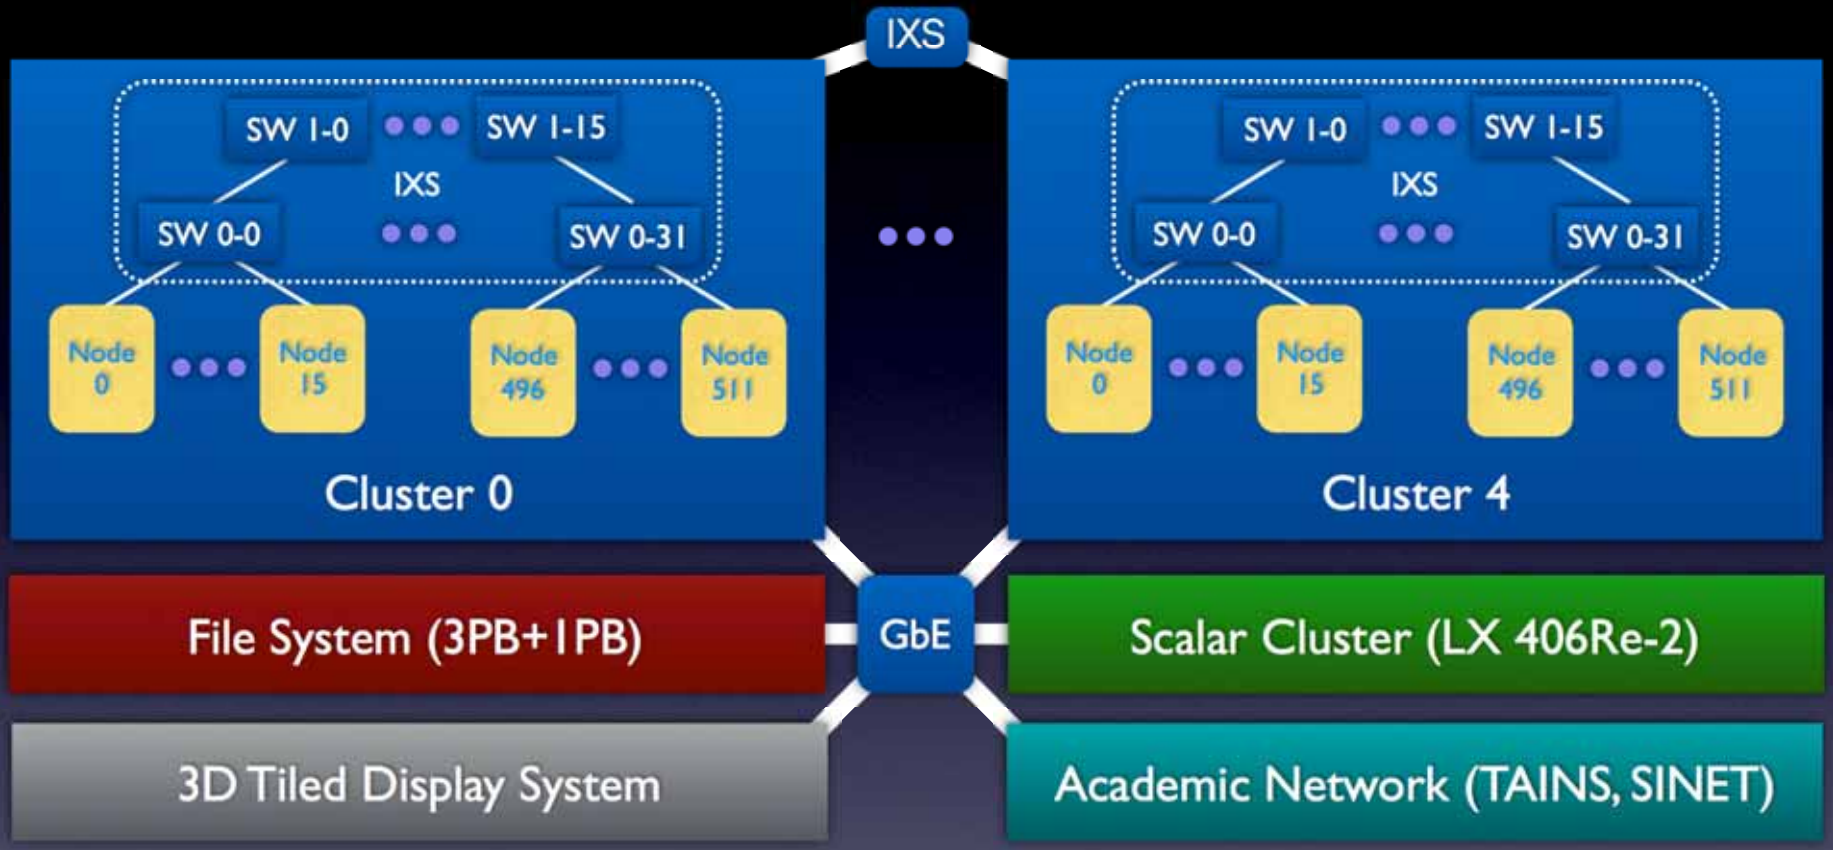

	Core	CPU(Socket)	Node	Cluster	Total System
Size	1	4 Cores	1CPU	512 Nodes	5 Clusters
Performance (VPU+SPU)	69GFlop/s (68GF+1GF)	276GFlop/s (272GF+4GF)		141Tflop/s (139TF+ 2TF)	707Tflop/s (697TF+10TF)
Mem. BW	256GB/s			131TB/s	655TB/s
Memory Cap.	64GB			32TB	160TB
IXS Node BW	-		4GB/s x2		-

TOHOKU UNIVERSITY

Features of Tohoku Univ. SX-ACE System

Significant Performance Improvement with Lower Power and Less Space

		SX-9 (2008)	SX-ACE (2014)	Improvement
CPU Performance	Number of Cores	1	4	4x
	Total Flop/s	118.4Gflop/s	276Gflop/s	2.3x
	Memory Bandwidth	256GB/sec	256GB/sec	1
	ADB Capacity	256KB	4MB	16x
Total Performance, Footprint, Power Consumption	Total Flop/s	34.1Tfop/s	706.6Tflop/s	20.7x
	Total Memory Bandwidth	73.7TB/s	655TB/s	8.9x
	Total Memory Capacity	18TB	160TB	8.9x
	Power Consumption (Max)	590kVA	1,080kVA	1.8x
	Footprint	**293m²**	**430m²**	1.5x

Powerful CPU/Node Performance and Higher B/F rate

		SX-ACE(2014)	K(2011)	Ratio
CPU (Node) Performance	Clock Frequency	1GHz	2GHz	0.5x
	Flop/s per Core	64Gflop/s	16Gflop/s	4x
	Cores per CPU	4	8	0.5x
	Flop/s per CPU	256Gflop/s	128Gflop/s	2x
	Bandwidth	256GB/s	64GB/s	4x
	Bytes per Flop (B/F)	1	0.5	2x
	Memory Capacity	64GB	16GB	4x

A Balanced System for High Sustained Performance, resulting in High Productivity in the Wide Area of Applications in Academia and Industry

Features of the SX-ACE Vector Processor

- 4 high-performance core Configuration, each with High-Performance Vector-Processing Unit and Scalar Processing Unit
  - 272Gflop/s of VPU + 4Gflop/s of SPU per socket
    - 68Gflop/s + 1Gflop/s per core
  - 1MB private ADB per core (4MB per socket)
    - Software-controlled on-chip memory for vector load/store
    - 4x compared with SX-9
    - 4-way set-associative
    - MSHR with 512 entries (address+data)
    - 256GB/s to/from Vec. Reg.
      - 4B/F for Multiply-Add operations
  - 256 GB/s memory bandwidth, Shared with 4 cores
    - 1B/F in 4-core Multiply-Add operations
      ~ 4B/F in 1-core Multiply-Add operations
    - 128 memory banks per socket
- Other improvement and new mechanisms to enhance vector processing capability, especially for efficient handling of short vectors operations and indirect memory accesses
  - Out of Order execution for vector load/store operations
  - Advanced data forwarding in vector pipes chaining
  - Shorter memory latency than SX-9

Source: NEC

SX-ACE Processor Architecture

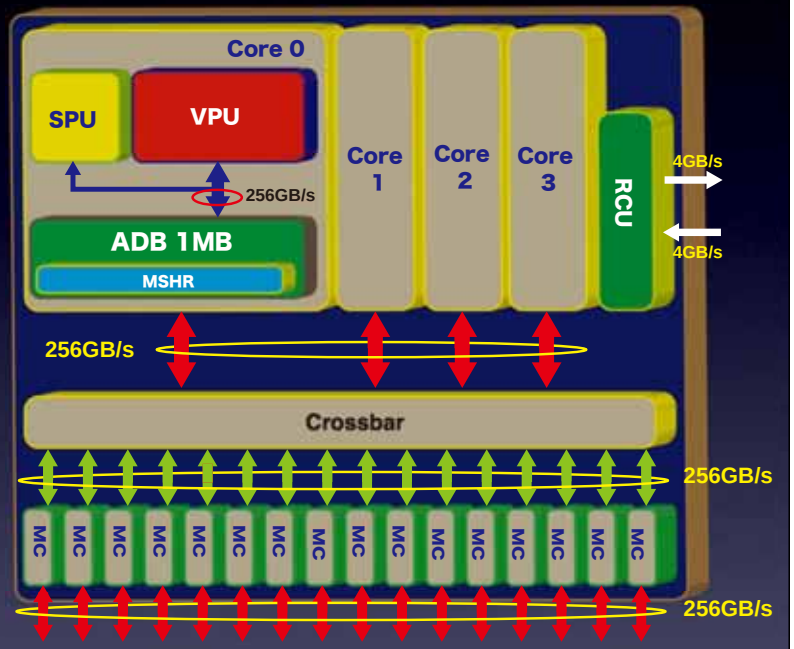

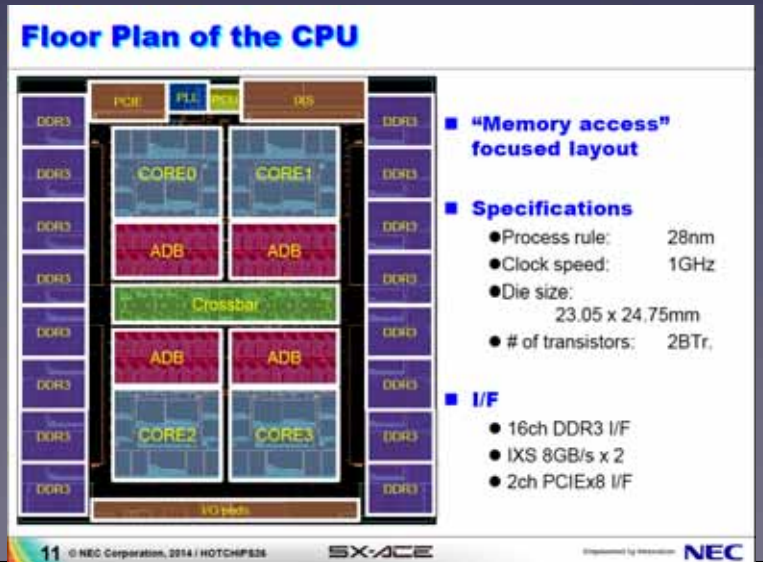

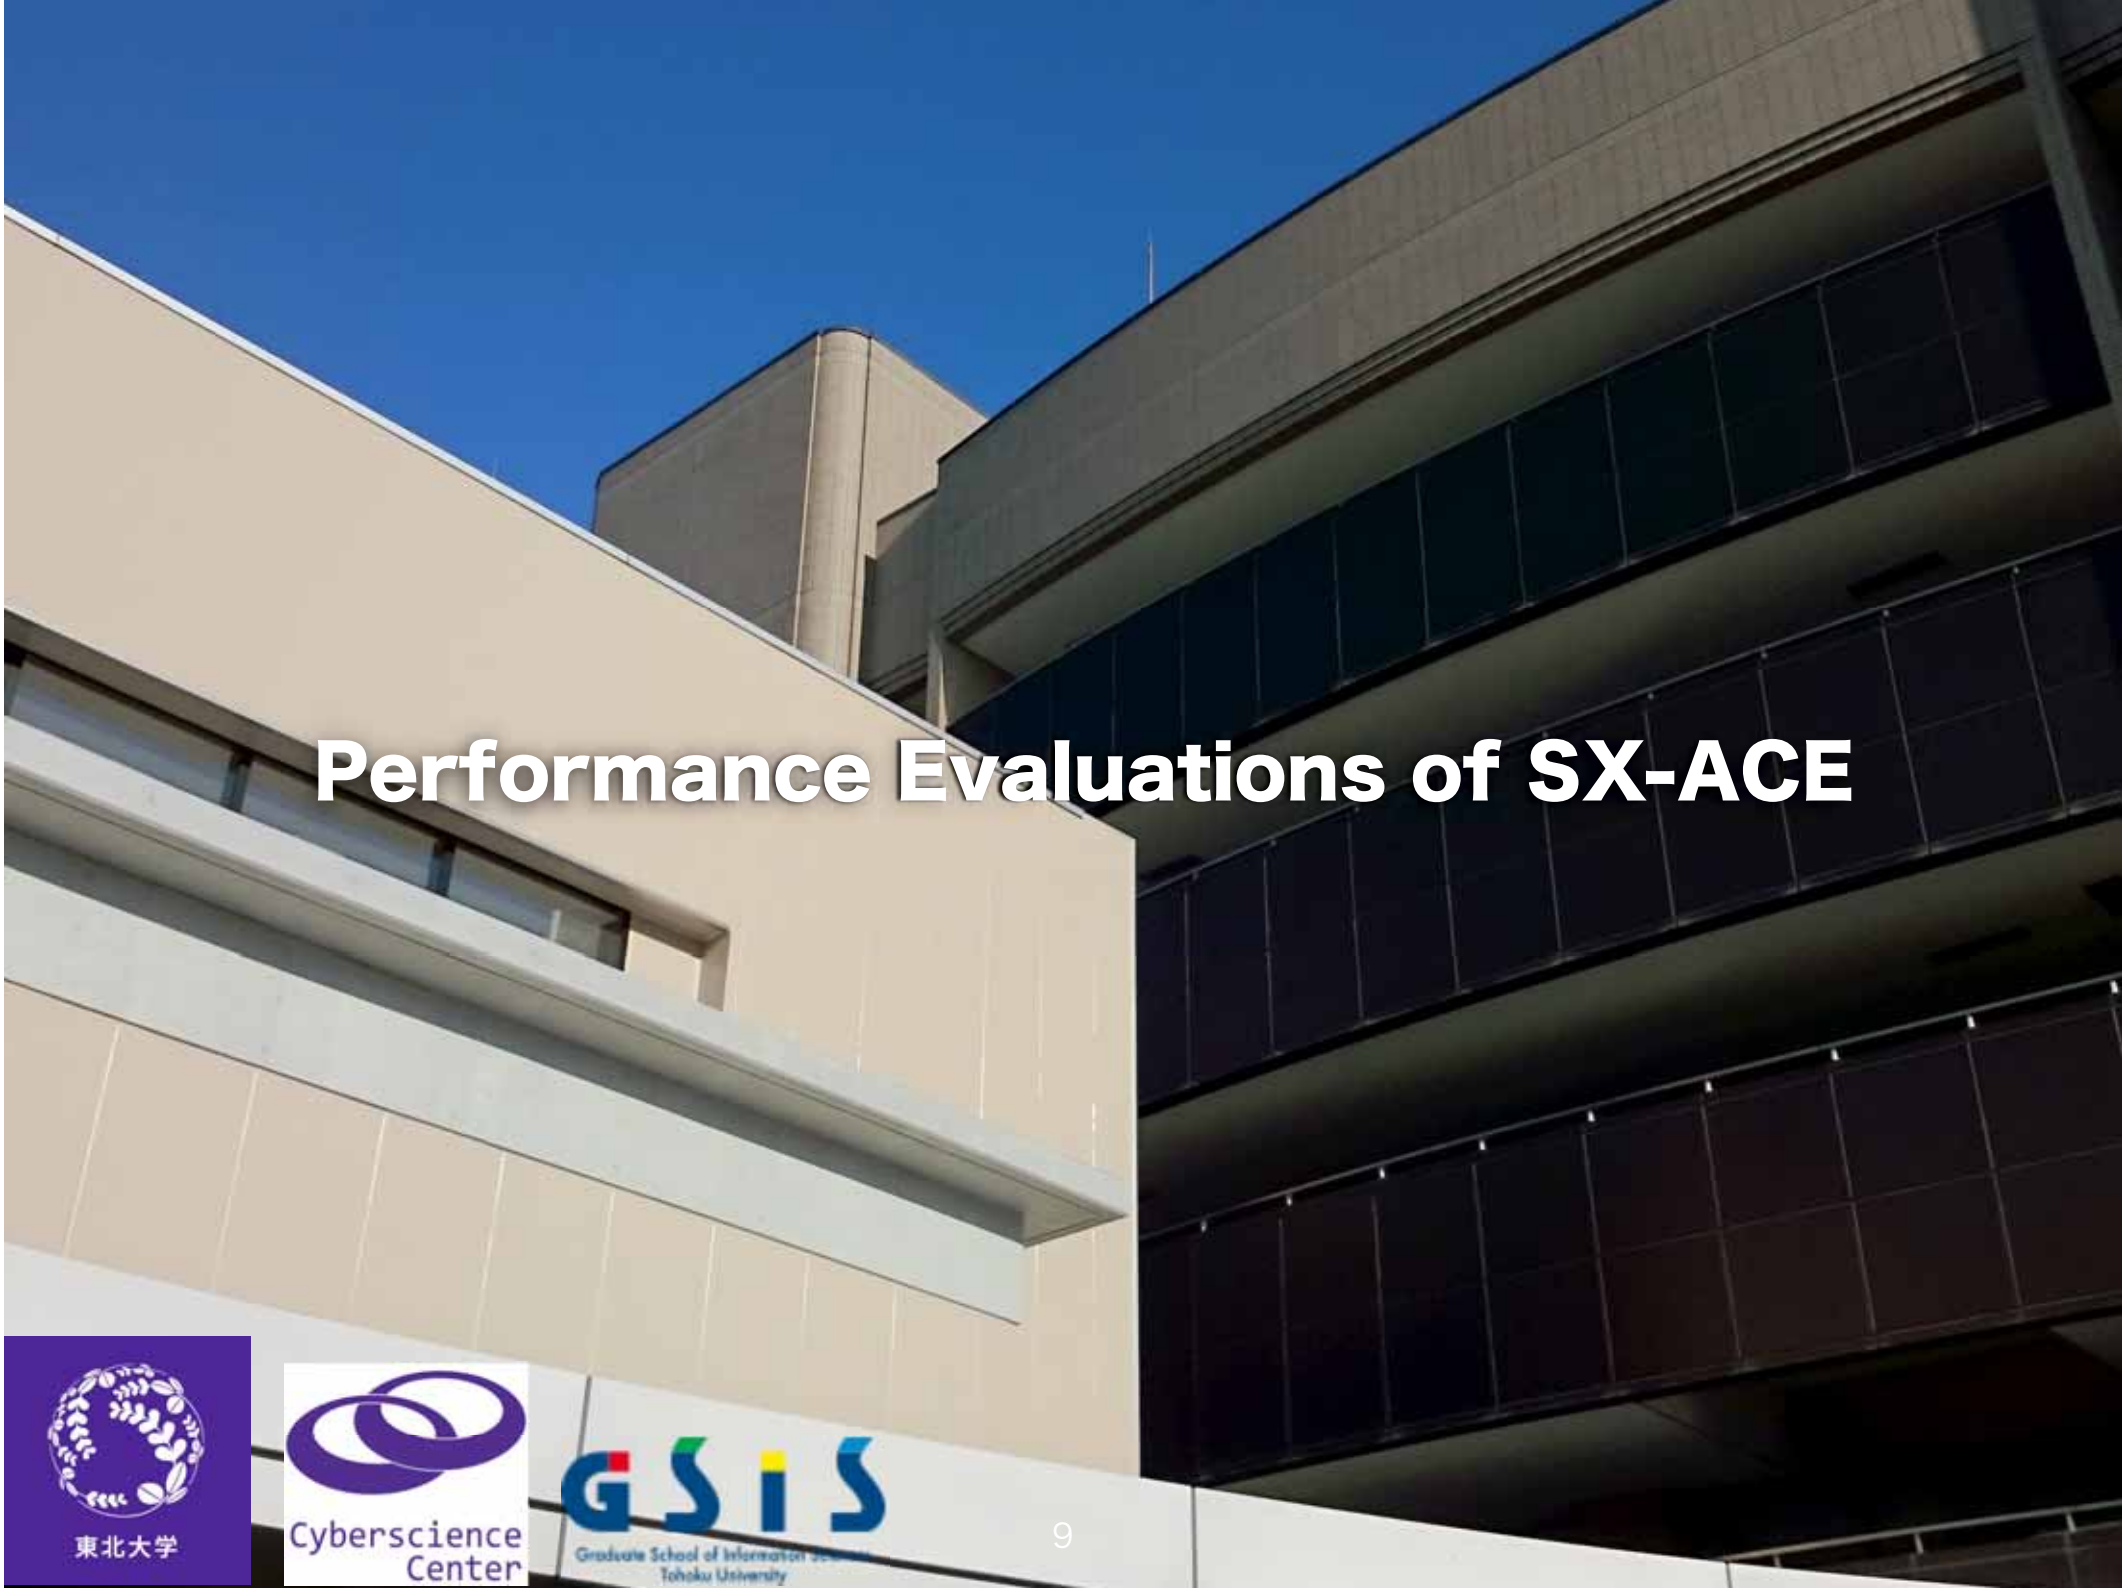
Performance Evaluations of SX-ACE
東北大学
Cyberscience Center
GSIS
Graduate School of Information Sciences
Tohoku University

TOHOKU UNIVERSITY

Performance Evaluation of SX-ACE by using HPCG

- HPCG (High Performance Conjugate Gradients) is designed
  - to exercise computational and data access patterns that more closely match a broad set of important applications, and
  - to give incentive to computer system designers to invest in capabilities that will have impact on the collective performance of these applications.
  - ✓ HPL for top500 is increasingly unreliable as a true measure of system performance for a growing collection of important science and engineering applications.
- HPCG is a complete, stand-alone code that measures the performance of basic operations in a unified code:
  - ✓ Driven by multigrid preconditioned conjugate gradient algorithm that exercises the key kernels on a nested set of coarse grids with Sparse matrix-vector multiplication.
    - Sparse triangular solve.
    - Vector updates.
    - Global dot products.
    - Local symmetric Gauss-Seidel smoother.
    - Reference implementation is written in C++ with MPI and OpenMP support.

Breakdown of the HPCG Benchmark

Benchmarking Flow

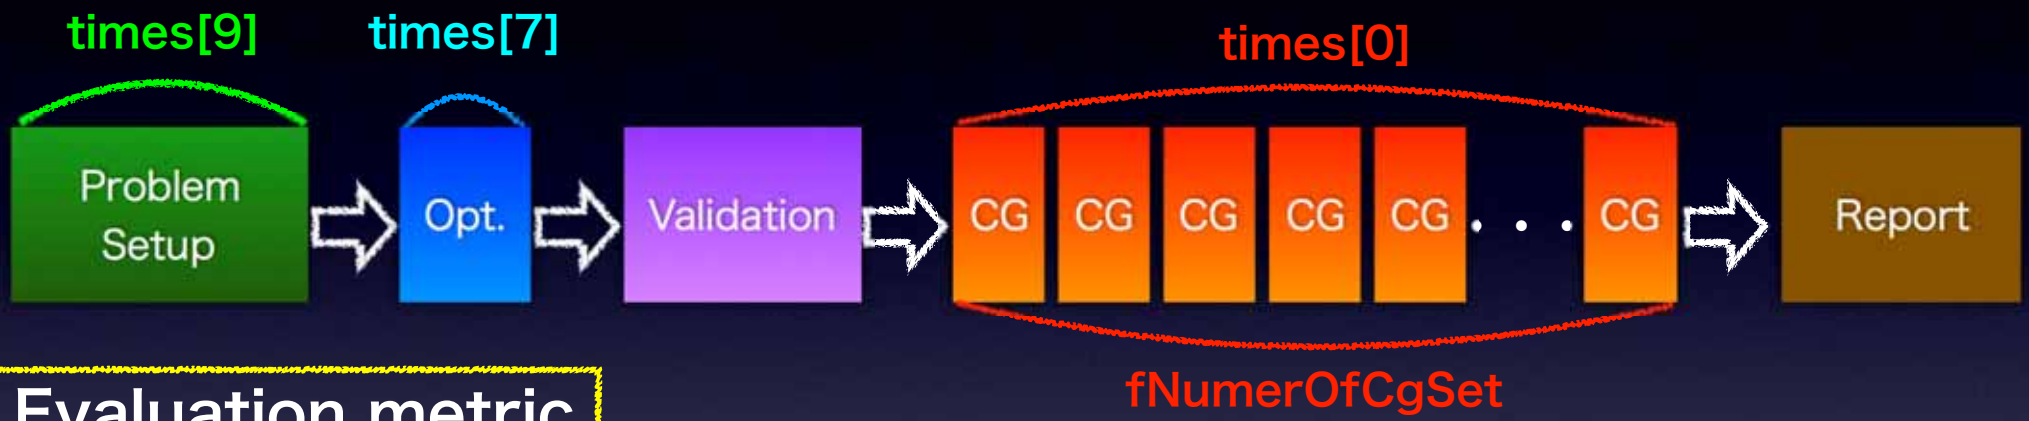

Evaluation metric

※ frefnops : total number of floating point operations for CG (# of iterations = 50)

Ver. 2.4 → Ver. 3.0: Setup overhead considered for individual CG iterations!

ver. 2.4

GFlop/s = frefnops/(times[0]+fNumberOfCgSets* times[7]/10.0)/1.0E9

ver. 3.0

GFlop/s = frefnops/(times[0]+fNumberOfCgSets*(times[7]/10.0+ times[9]/10.0))/1.0E9

Optimizations of the the HPCG Benchmark for SX-ACE

*Komatsu et al.@SC15

- Data packing for vector-friendly matrix memory allocation of sparse matrices
- Parallelization of 27-point stencil computation by using coloring and hyperplane methods
- Selective reusable-data caching and blocking for effective use of ADB

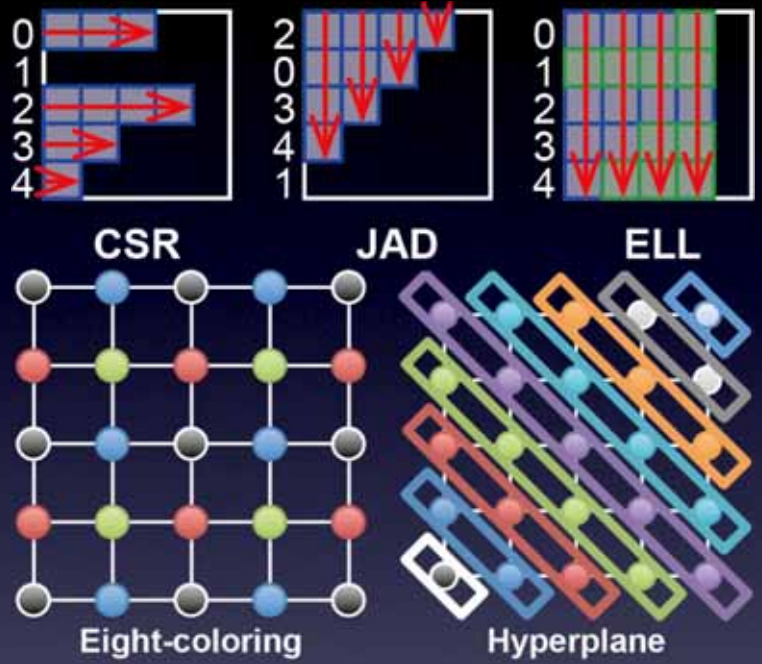

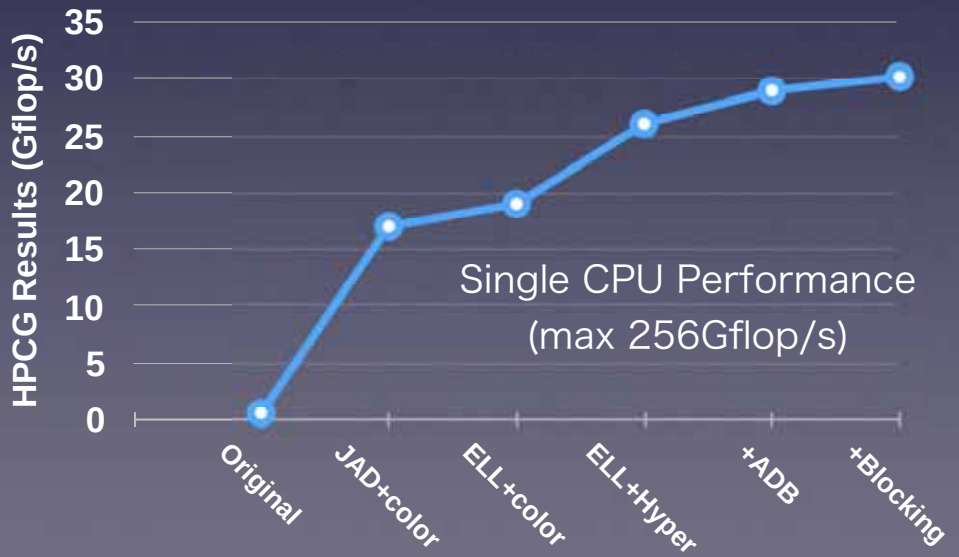

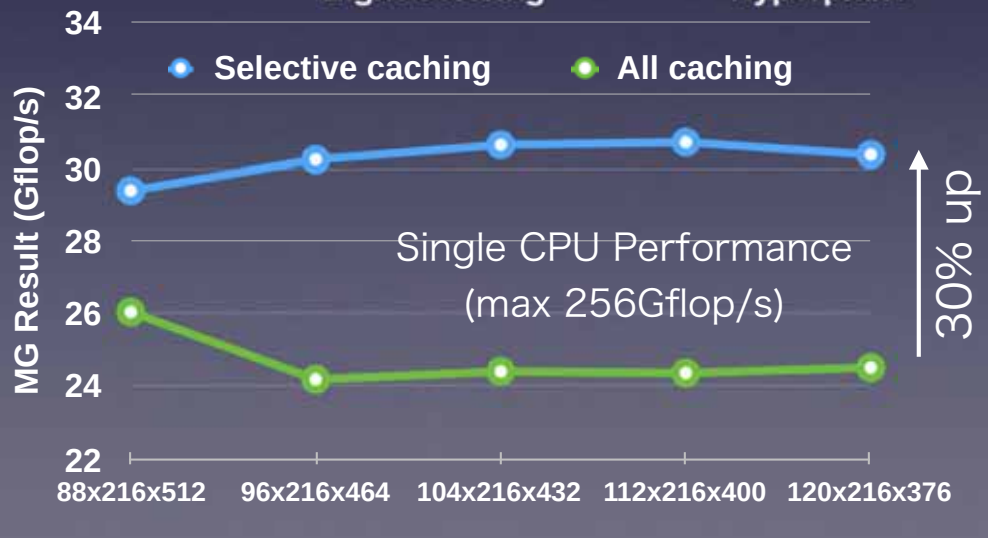

Scalability of the HPCG Benchmark

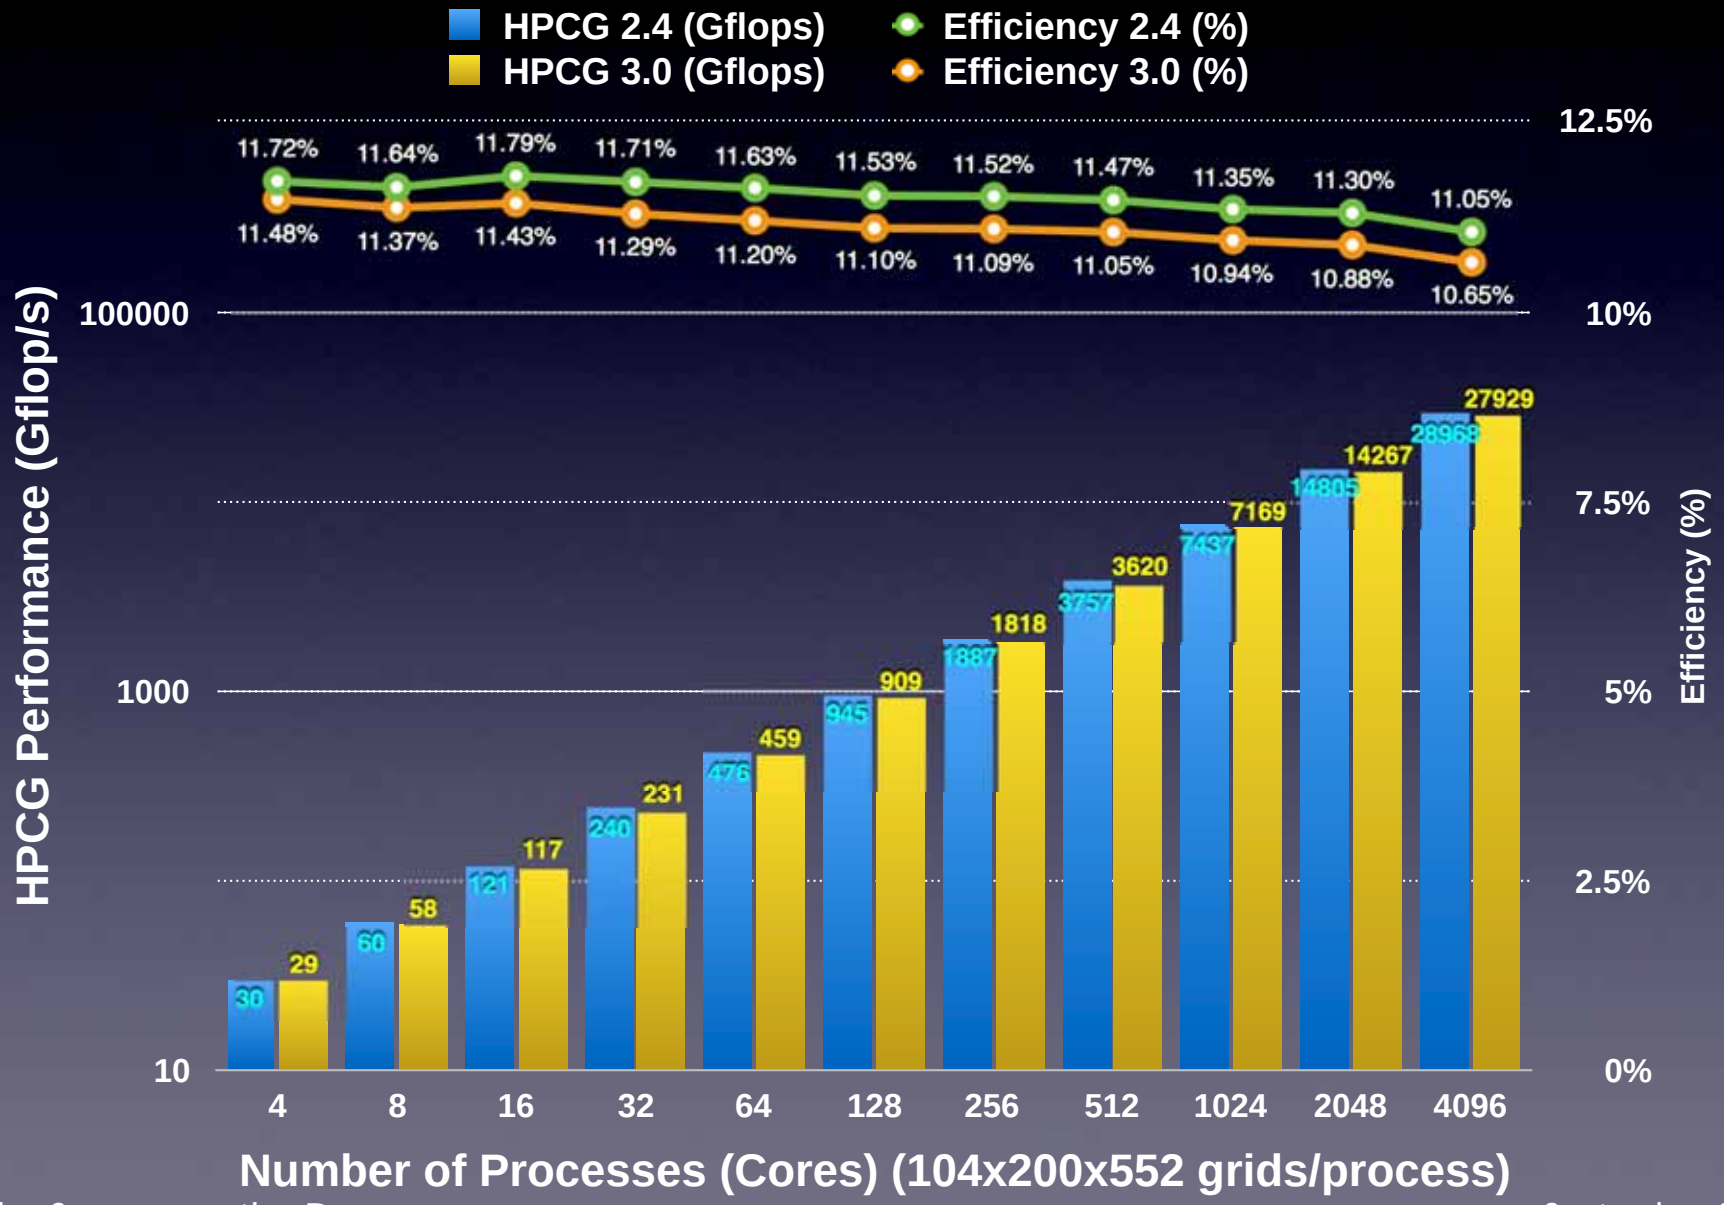

TOHOKU UNIVERSITY

Efficiency Evaluation of HPCG Performance (As of ISC16 data)

System	Eff.(%)	HPCG/W (MF/W)
SX-ACE	10.8	107.5
K/FX	3.7	49.2
Xeon	2.2	36.6
BlueGene	1.5	38.6
GPU	1.4	58.0
XeonPhi	1.0	25.8
Sunway TaihuLight	0.3	24.1

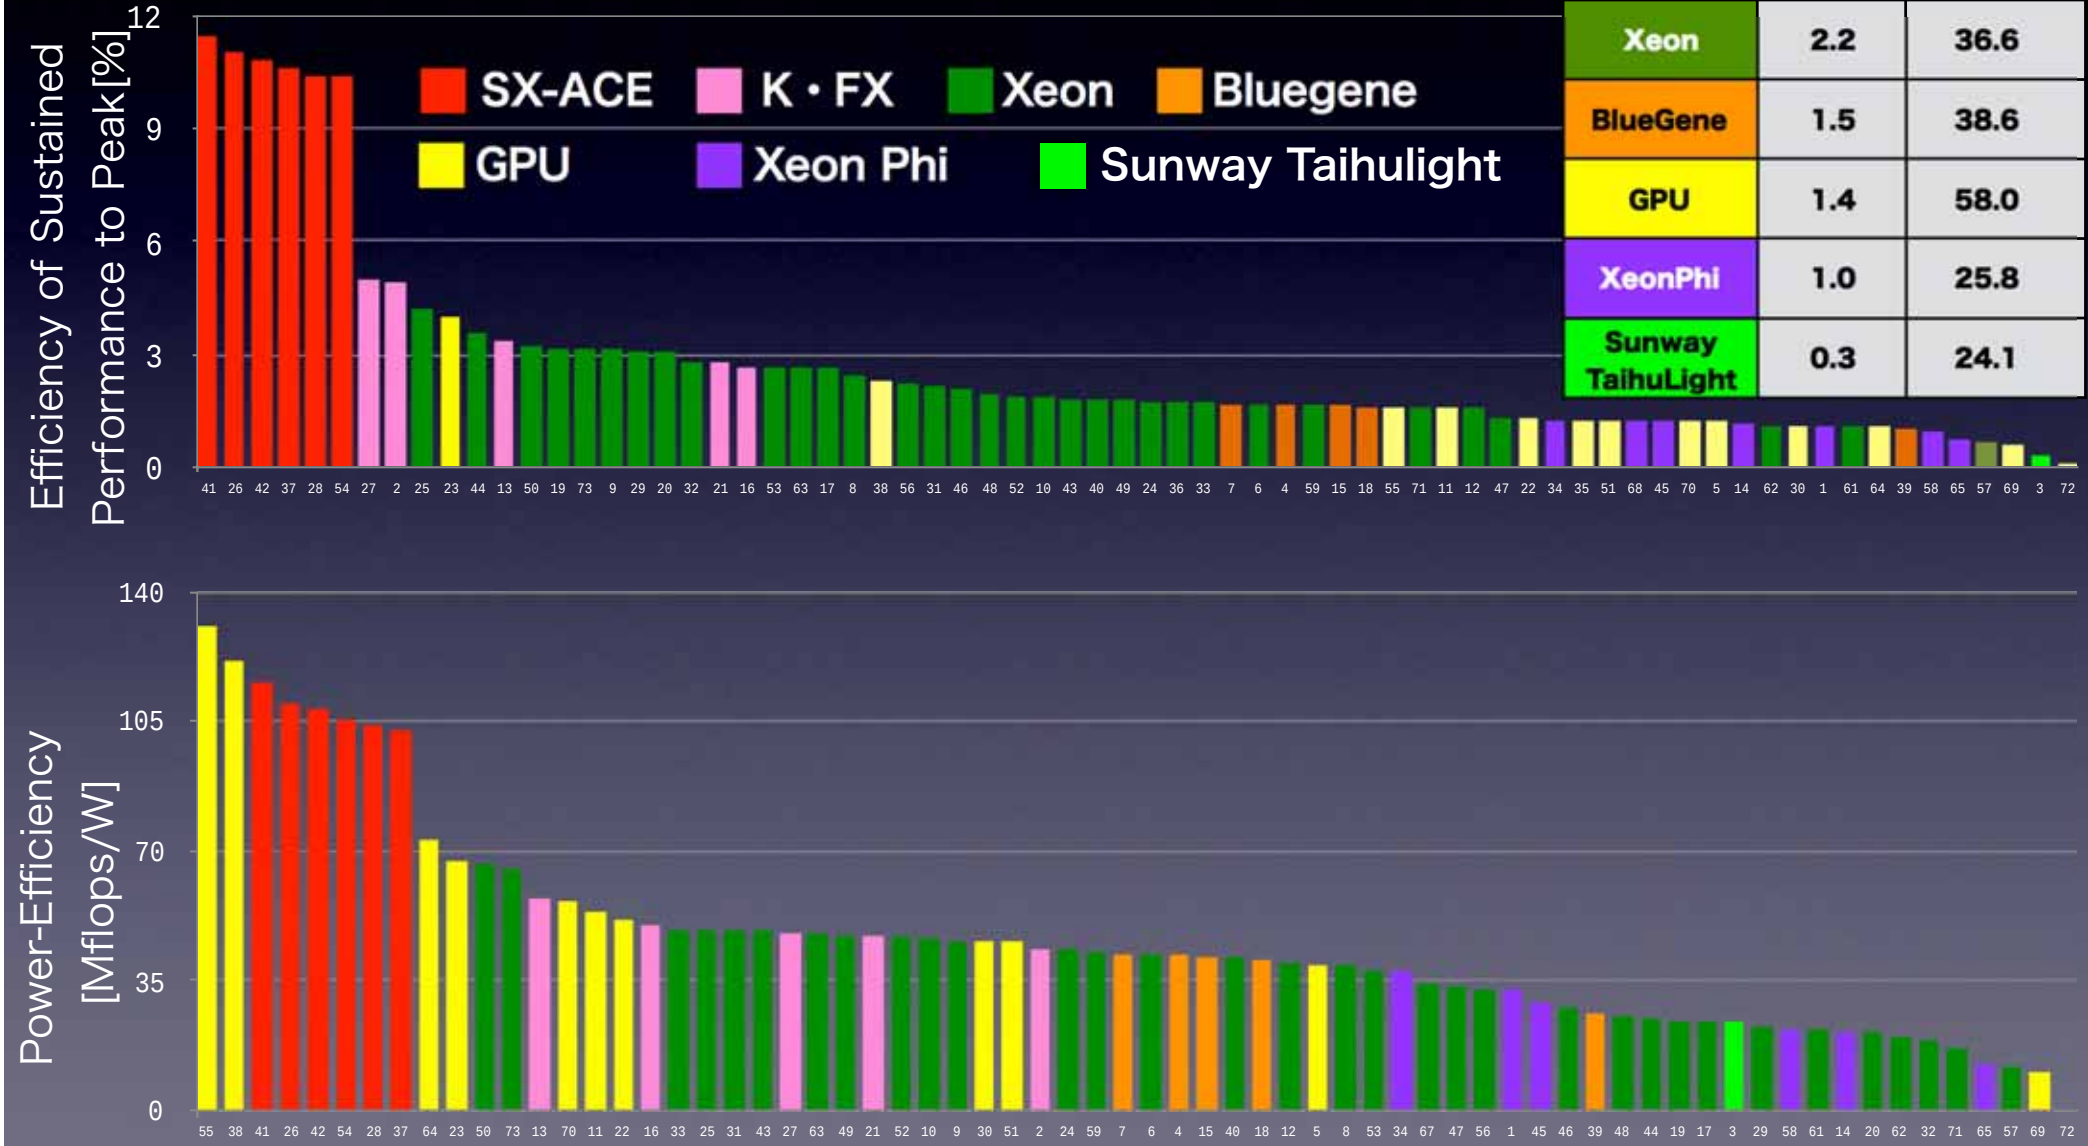

Necessary Peak Perf. to Obtain the Same Sustained Performance Normalized by SX-ACE (As of ISC 16 data)

System	Effi. (%)	HPCG/W (MF/W)	Req. Peak
SX-ACE	10.8	107.5	1.0x
K/FX	3.7	49.2	3.1x
Xeon	2.2	36.6	5.6x
BlueGene	1.5	38.6	7.3x
GPU	1.4	58.0	8.7x
XeonPhi	1.0	25.8	10.2x
Sunway TaihuLight	0.3	24.1	36.0x

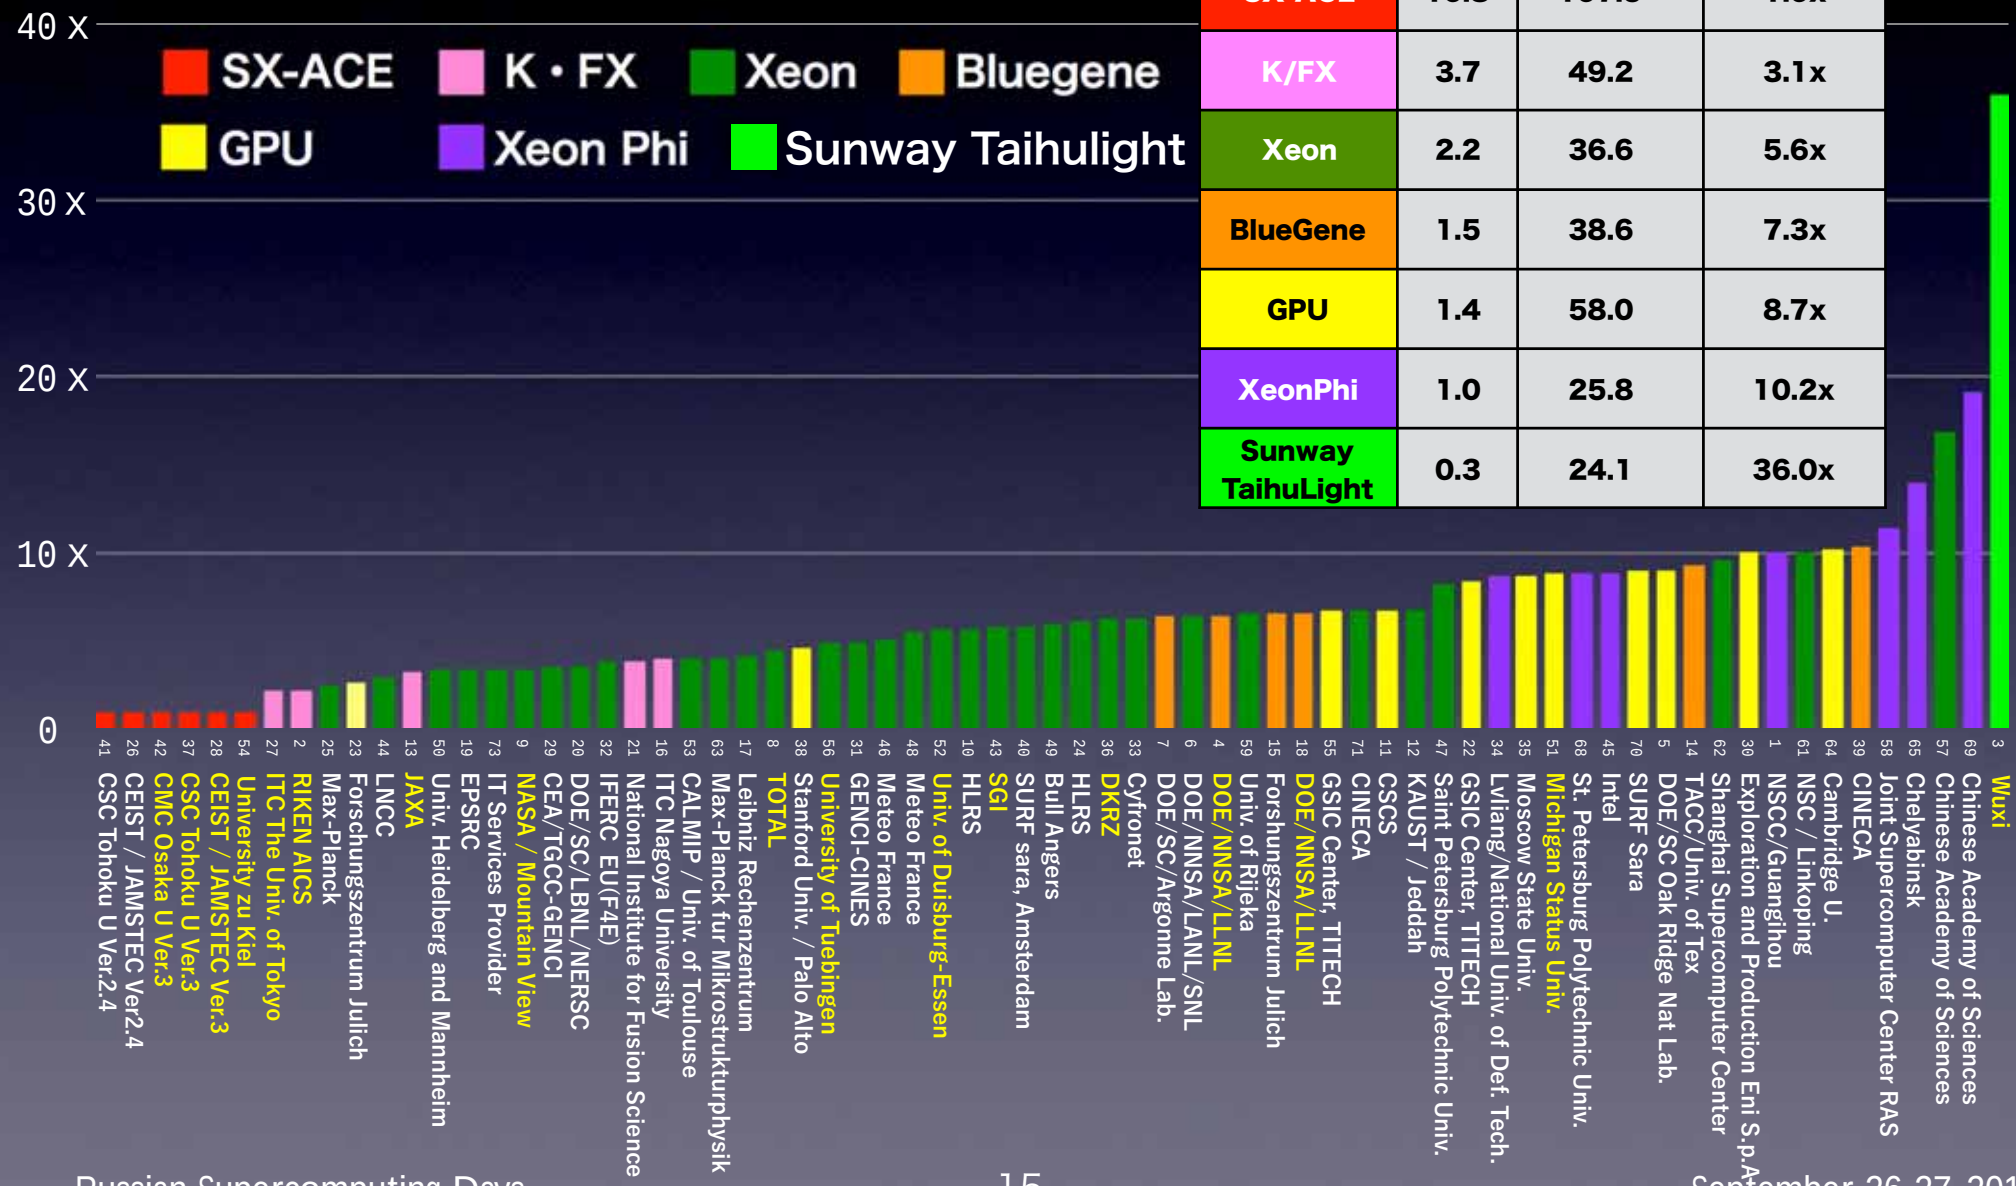

TOHOKU UNIVERSITY

Leading Science and Engineering Fields supported by the Supercomputer of Tohoku University

Next-Generation CFD Analysis

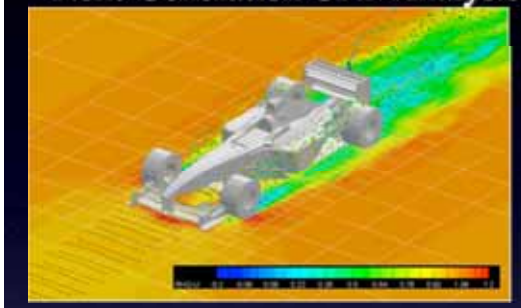

Turbine Design

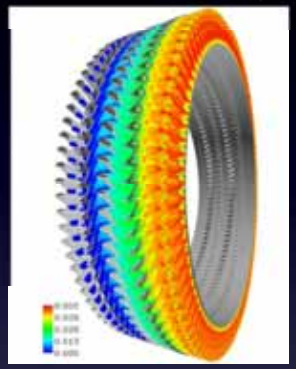

Perpendicular Magnetic Recording Medium Design

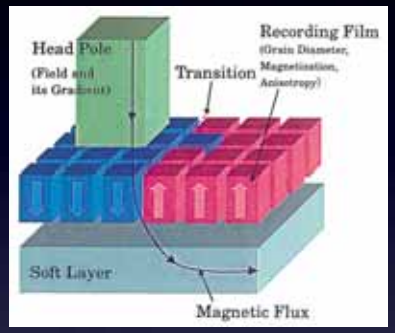

Nano Material Design

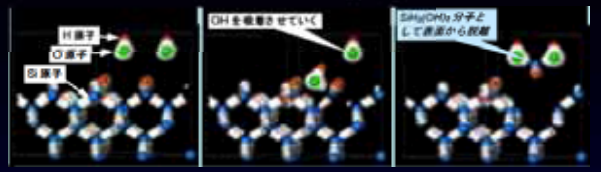

Tsunami Inundation Analysis

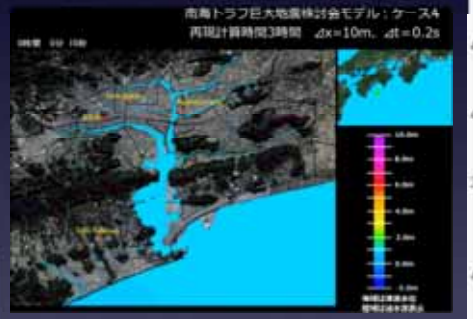

Earthquake Analysis

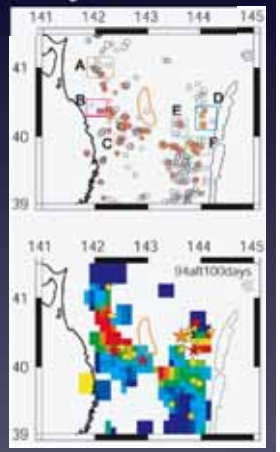

Industrial Use

Antenna Analysis

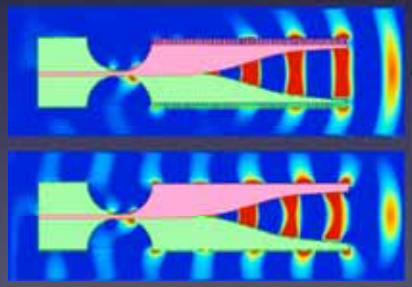

Heat Shock Analysis

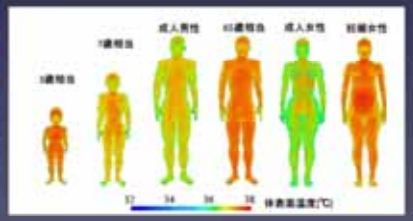

MRJ

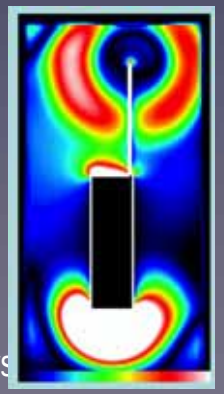

Combustion Flow Simulation

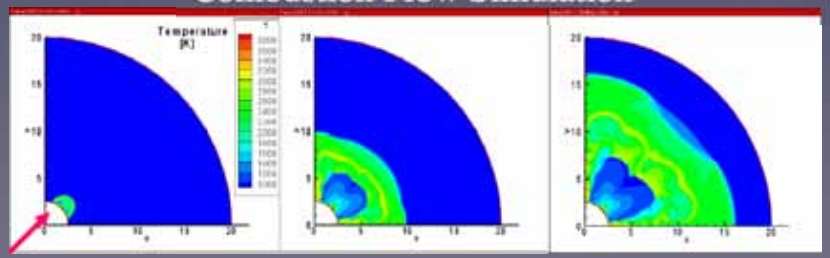

Ozone-hole Analysis

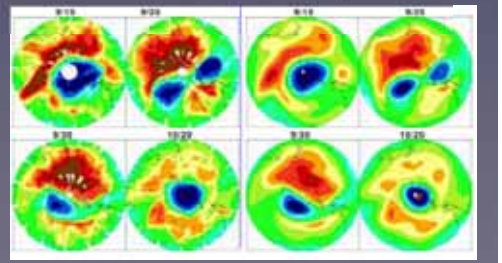

Heatstroke Risk Simulator on SX-ACE

- **The number of fatalities due to heat waves has increased in Europe, North America, and Asia.**
- **Heatstroke**
  - The number of people hospitalized suffering heatstroke is increasing in JAPAN
    - 58,000 patients in 2014
  - 12,000 fatalities from 1968 to 2014 in Japan
- **The changes of body temperature strongly depend on individual differences**
  - body size, age, male/female, tend to perspire a lot or not, difference in genders, etc.
- **Heatstroke Risk Simulator**
  - **Simulating the changes in body temperature**
    - Developed by Prof. Hirata (Nagoya Institute of Tech.
- **The body temperatures of children, elderly men/ women, pregnant are ease to increase.**

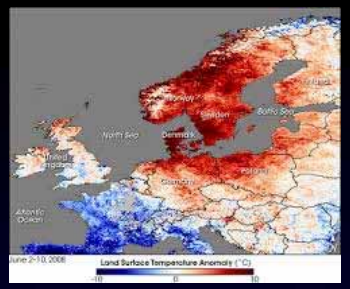

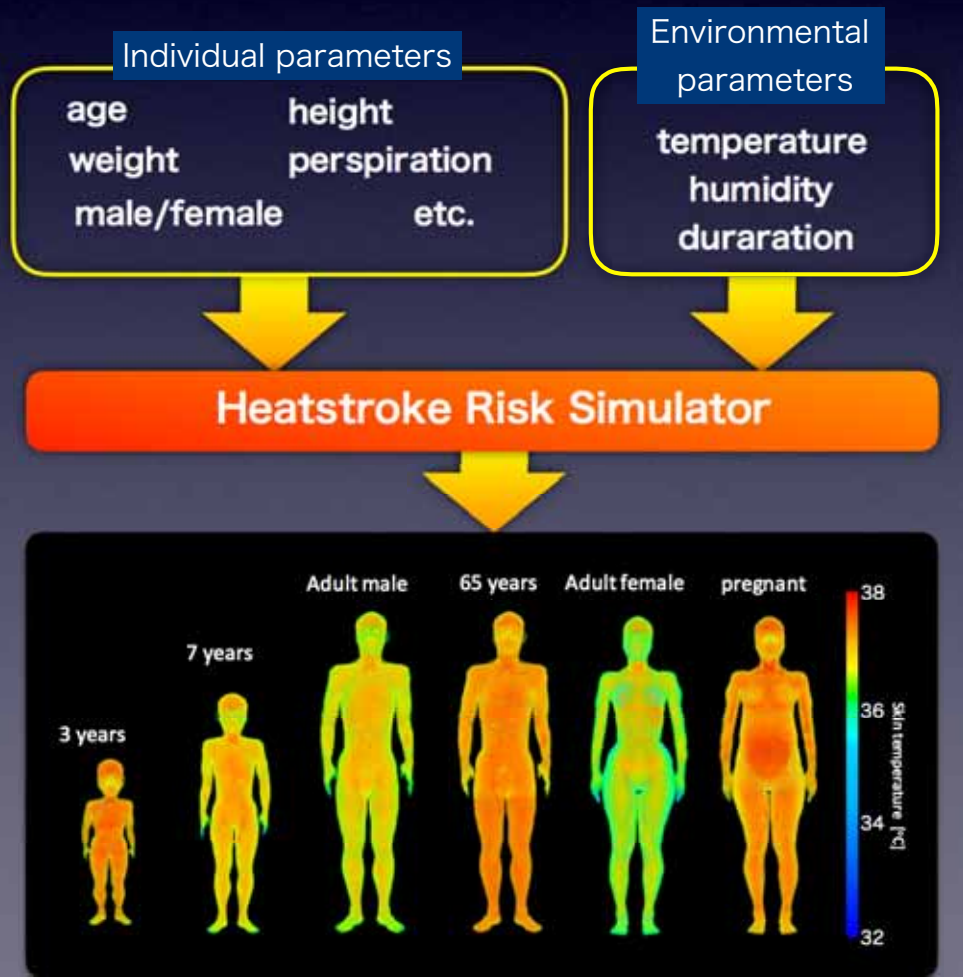

Scalability

- **Parallelize "temprise_k" sub-routine by MPI**
  - **866 x 320 x 160**
  - **5400 steps**

```
DO K=1,MODELZ
    DO J=1,MODELY
      DO I=1,MODELX
            temperature update calculation
      END DO
    END DO
END DO
```

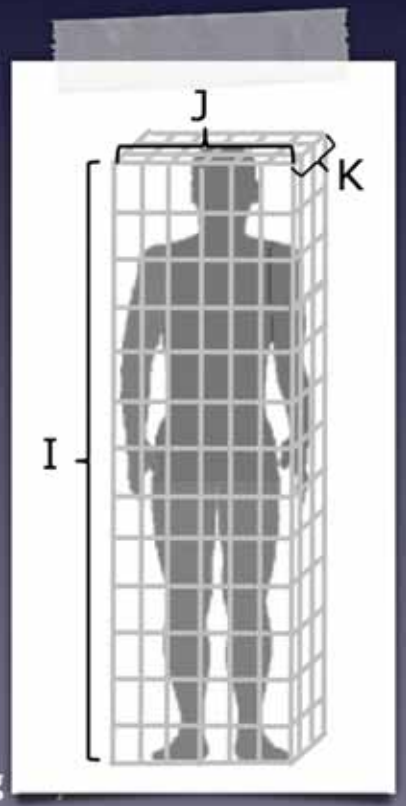

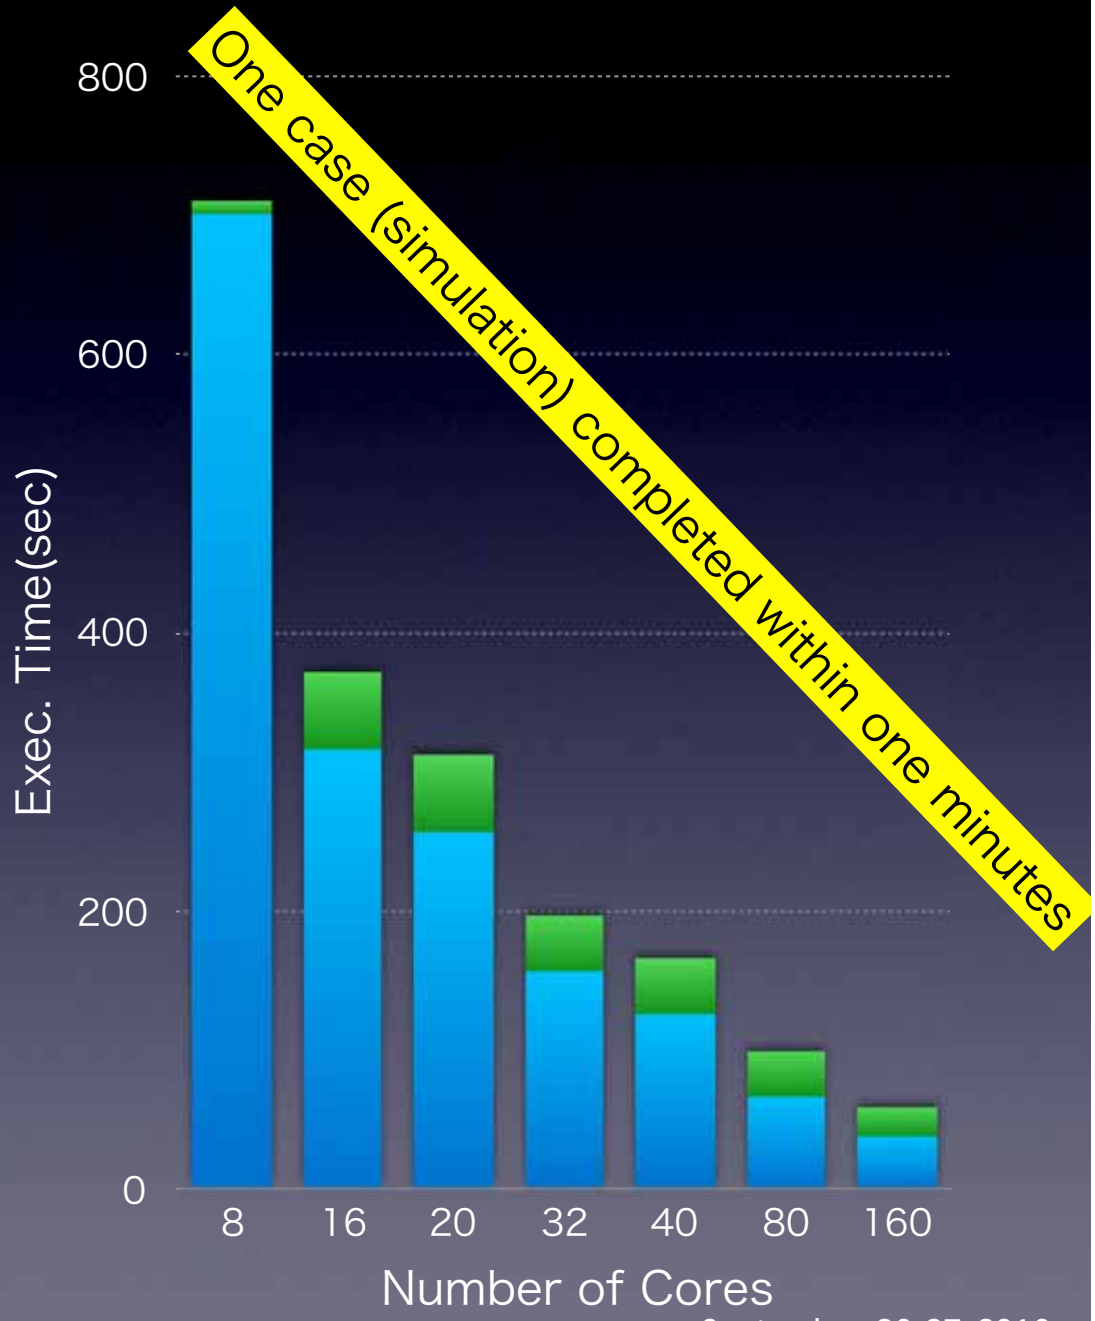

Performance of Magneto-Hydro-Dynamics Simulation on SX-ACE

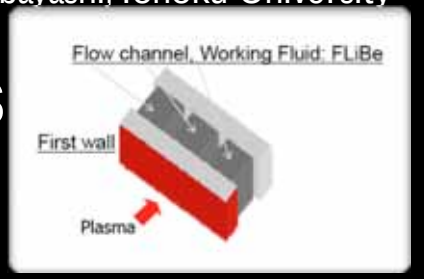

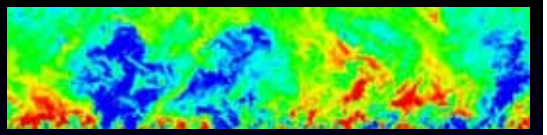

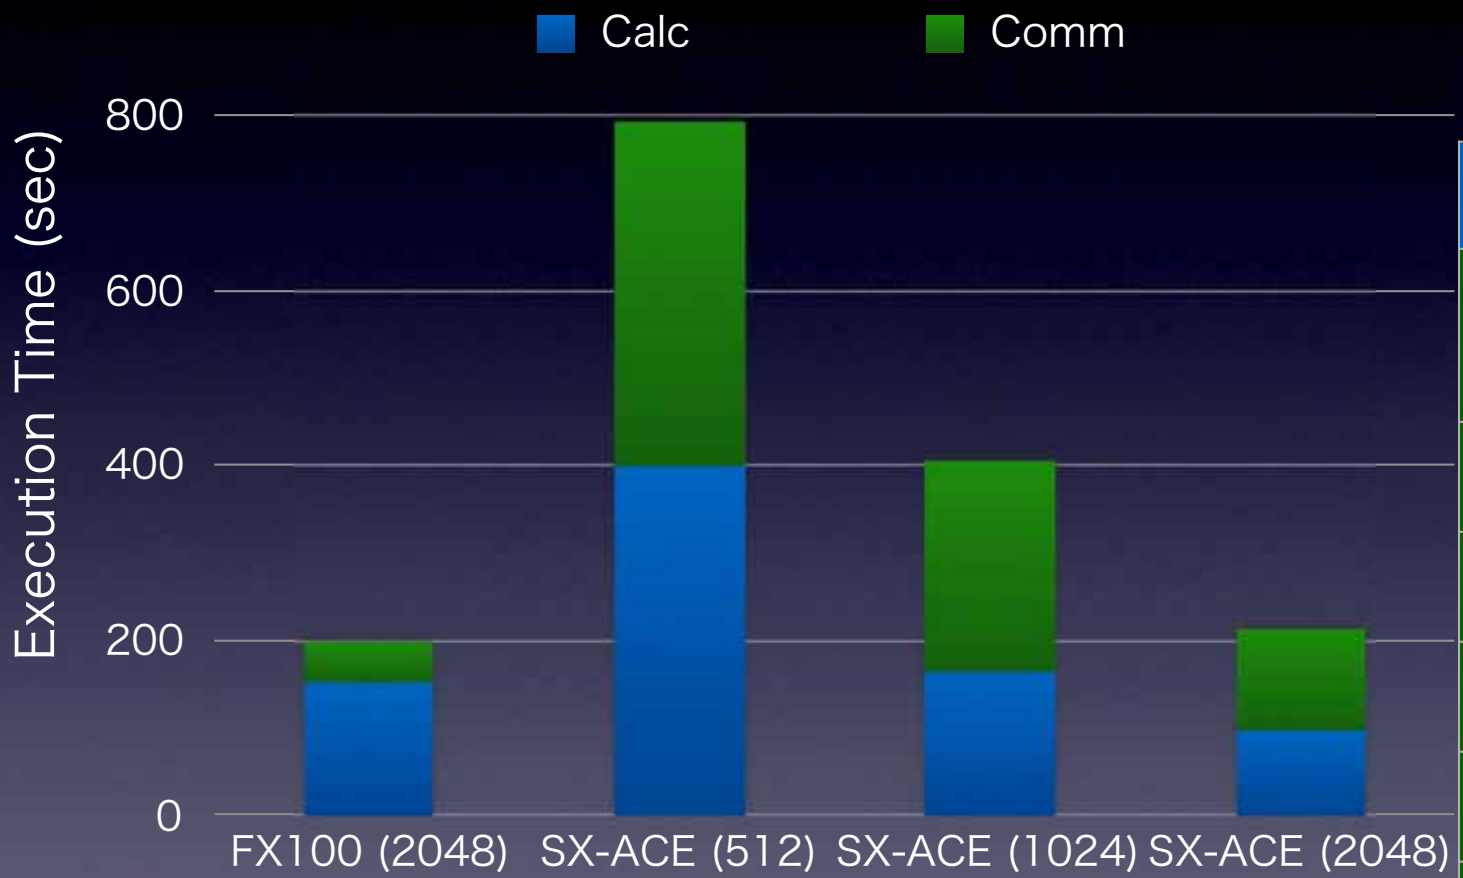

	FX100 (2048)	SX-ACE (512)	SX-ACE (1024)	SX-ACE (2048)
Sustained TFlop/s	50.3	15.4	27.6	42.7
Efficiency	2.5%	11.7%	10.5%	8.1%

	FX100	SX-ACE
Arch	SPARCv9 _HPC-ACE	SX
CPU/node	1	1
Cores/ CPU	32	4
Node Perf.	1.011 Tflop/s	256 Gflop/s
Mem cap./ node	32GB	64GB
mem BW	480 GB/s	256 GB/s
NW BW	12.5 GB/s	4GB/s

Ref : Y. Yamamoto, R. Egawa,Y. Isobe , and Y. Tsuji, "Performance evaluation of DNS code based on high-order accuracy finite difference methods, "Japan-Russia Workshop @ Nagoya, Dec 10, 2015.


Future Vector Systems R&D*
*This work is partially conducted with NEC, but the contents do not reflect any future products of NEC
東北大学
Cyberscience Center
GSIS
Graduate School of Information
Tohoku University

Timeline of the Cyberscience Center HPC System Development and R&D For the Future

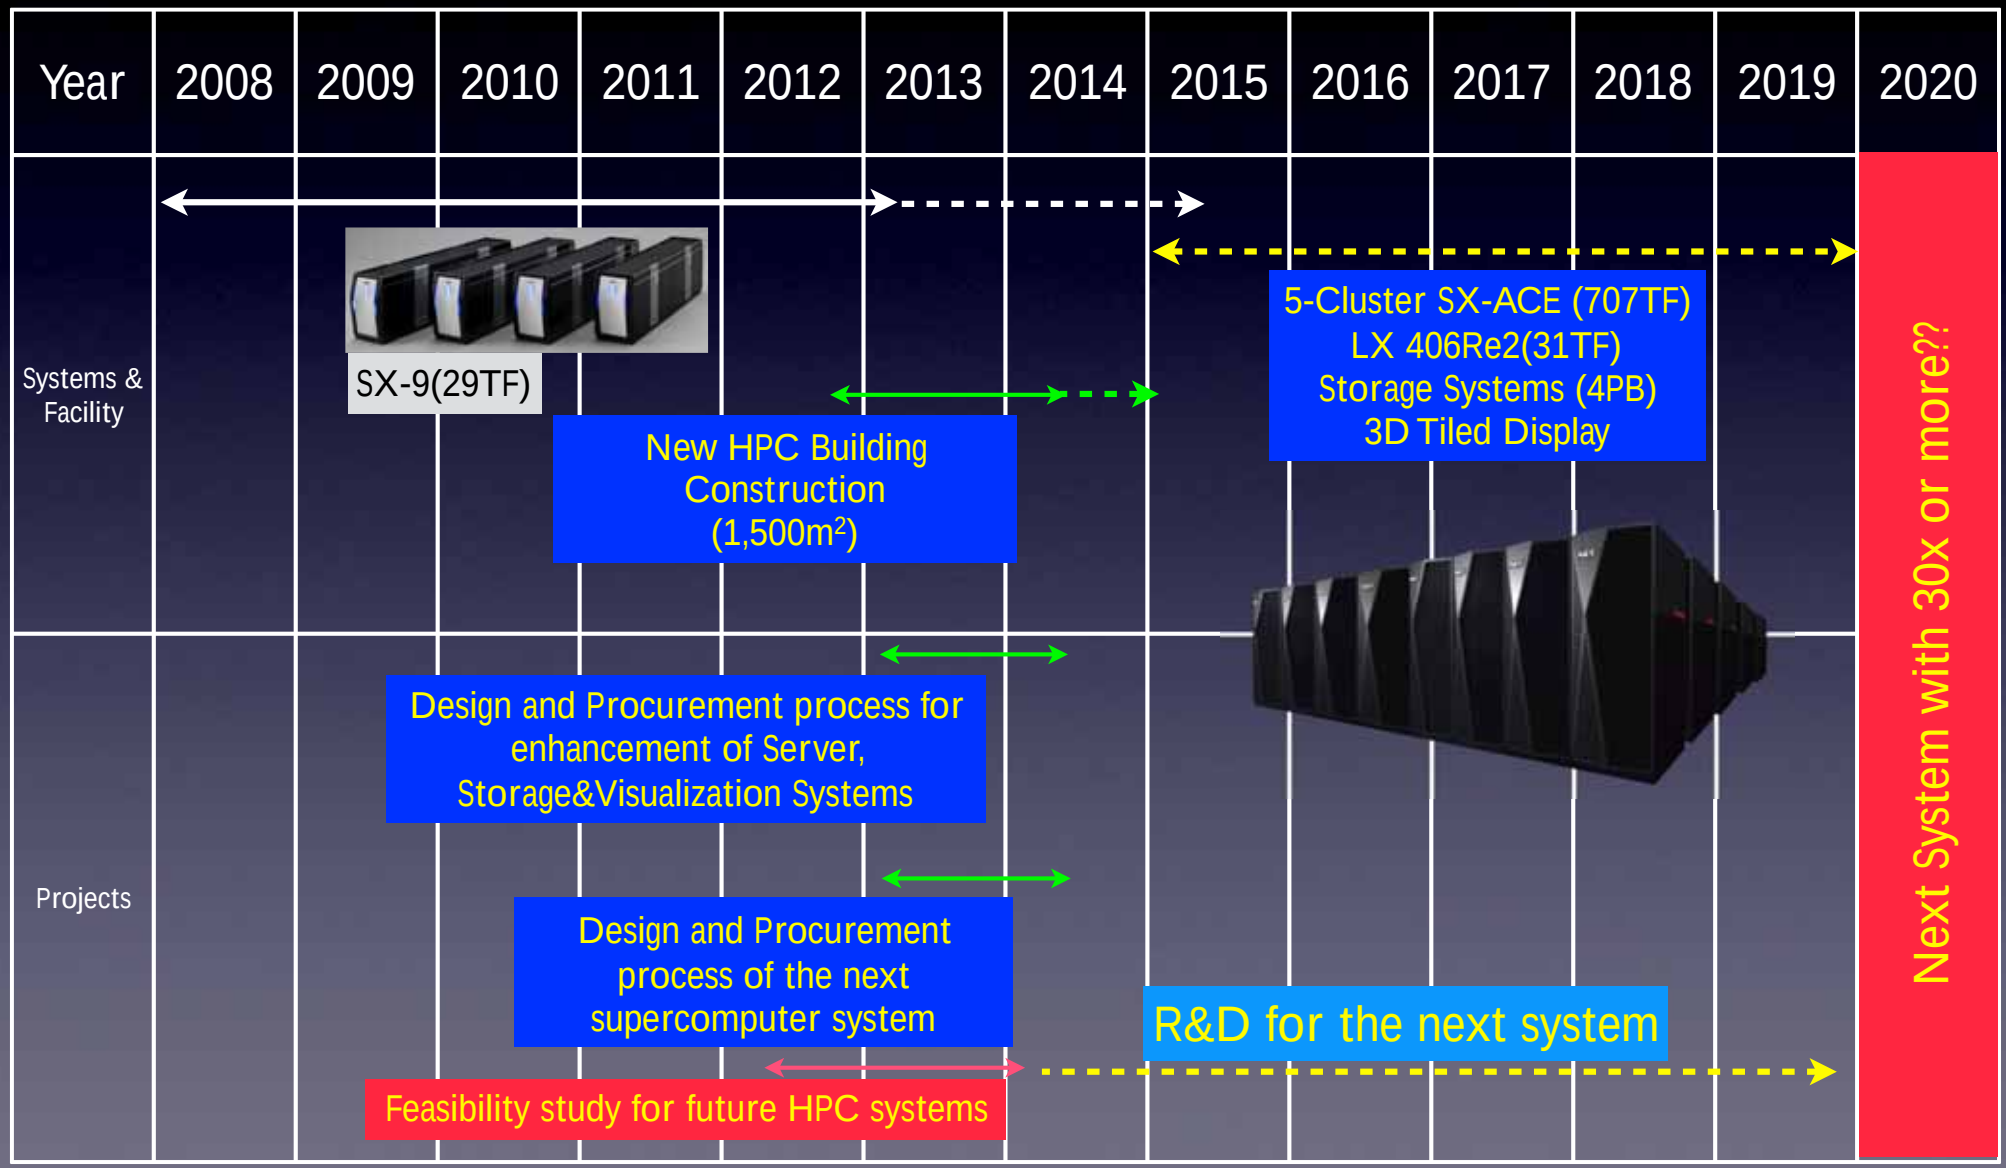

TOHOKU UNIVERSITY

The Time for Vector Computing has Come Again!

Modern and future microprocessors aggressively introduce vector computing mechanism for efficient processing of data-level parallelism

- ★ **Intel Many Cores with AVX (Advanced Vector eXtension)**
  - **Xeon with 256b-AVX2 and Xeon Phi with 72 AVX-512b**
- ★ **ARM with SVE (Scalable Vector Extension)**
  - **128b~2,48b-width Vector Extension**
  - **ARM64 with 512bit-vector for post-K computer in 2020?**
- ★ **Power9 with 256~512b-VSX (Vector Scaling eXtension)**
- ★ **GPU is the powerful vector computing platform**
  - **Pascal GP100 has 30 SMs(Streaming Multiprocessors), 2K-element vector processing capability**

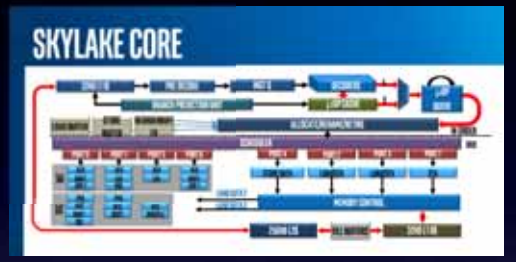

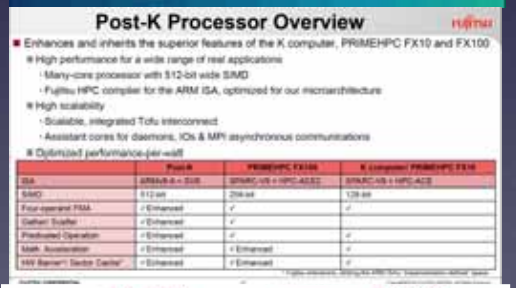

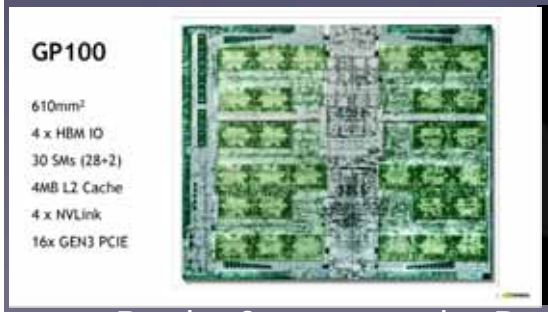

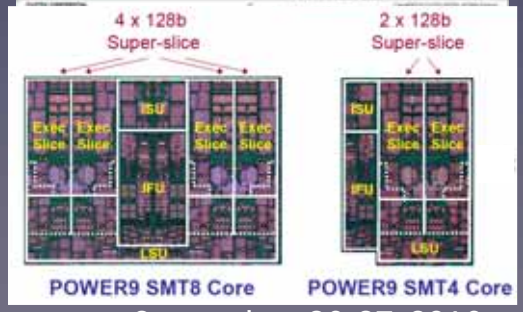

TOHOKU UNIVERSITY

The Ti ... Come Again!

Modern and future ... computing mechan ... ism

- ★ Intel Many ... n)
  - Xeon w ... 2b
- ★ ARM with ...
  - 128b~2 ...
  - ARM64 ... 020?
- ★ Power9 wi ...
- ★ GPU is the ...
  - Pascal ... s),
    2K-eler ...

RedLine Performance Solutions

HOME ABOUT COMPANY AREAS OF EXPERTISE SERVICE OFFERING

Blog

Vectors: How the Old Became New Again in Supercomputing

By Lynd Stringer in Uncategorized on August 17, 2016

http://www.redlineperf.com/vectors-how-the-old-became-new-again-in-supercomputing/

Vector instructions, once a powerful performance innovation of supercomputing in the 1970s and 1980s became an obsolete technology in the 1990s. But like the mythical phoenix bird, vector instructions have arisen from the ashes. Here is the history of a technology that went from new to old then back to new.

But first, a few definitions. A vector instruction is an SIMD instruction, Single Instruction Multiple Data. A vector instruction refers to vector registers where multiple data resides. For example, a Cray-1's vector register contained up to 64 64-bit double-precision floating point numbers. The Cray-1 had eight of these registers. Many operations, for example: add and multiply can be issued to add or multiply two vector registers and place the result in a third vector register.

Vectors Become New

In 1976 Cray Research and Seymour Cray created the Cray-1, the first commercially successful supercomputer with vector instructions. The first Cray-1 was delivered to Los Alamos National Laboratory

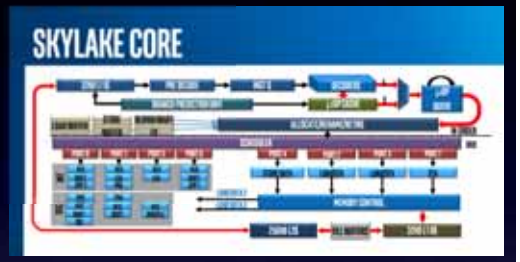

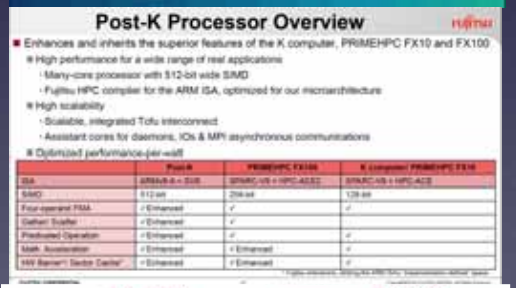

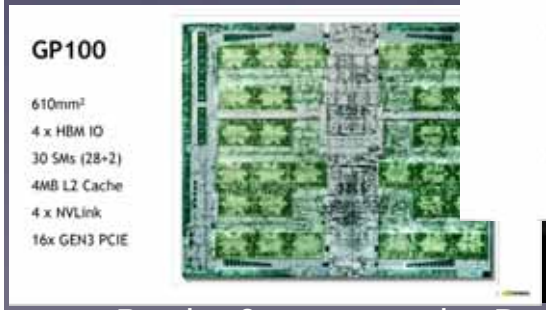

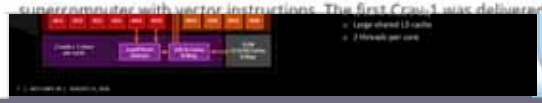

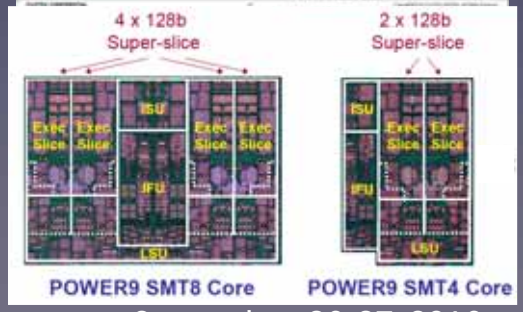

The time for Vector Computing has Come Again!

★ NEC's Next Generation Vector Supercomputer

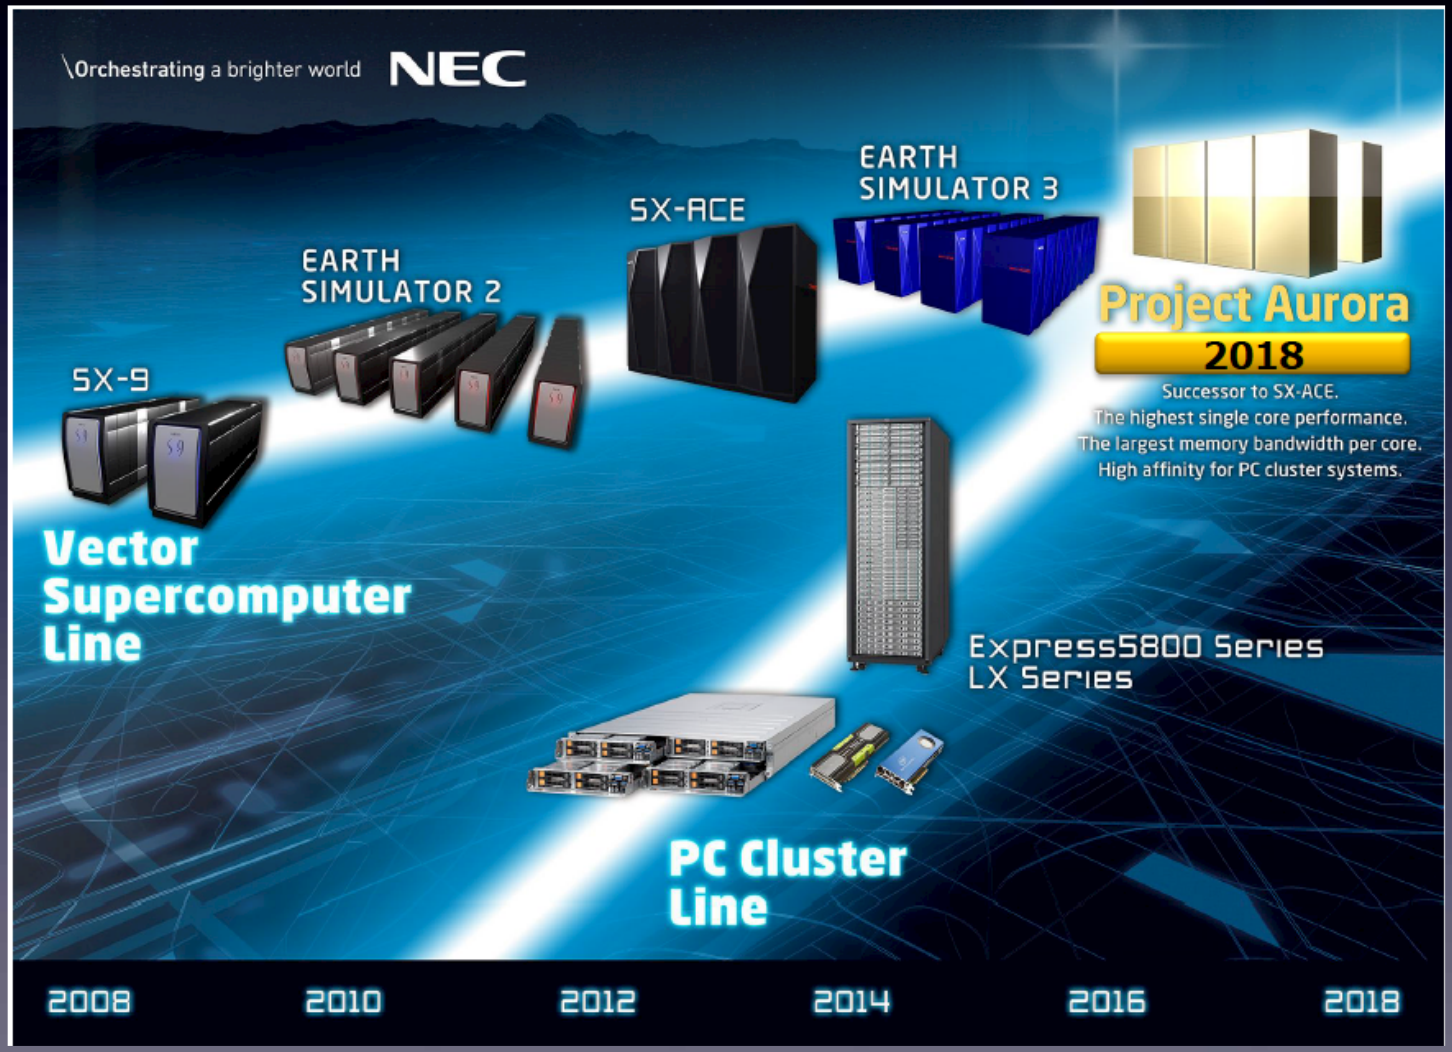

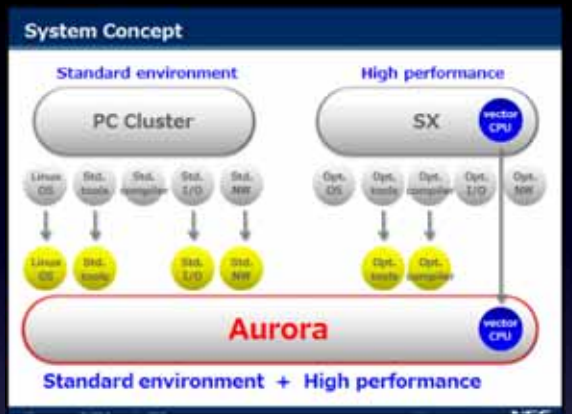

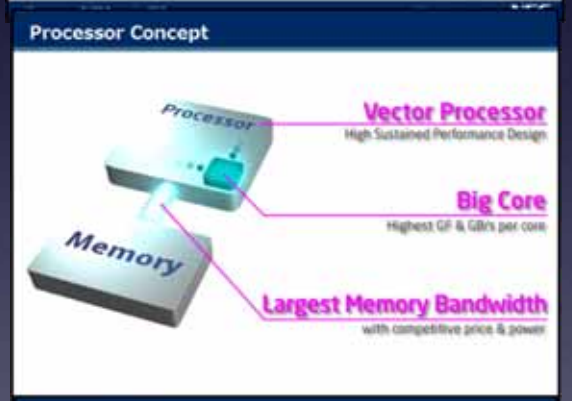

Source: NEC

Summary

★ SX-ACE, brand-new vector supercomputer of Tohoku University

- ✓ large-single core performance of 256-element vector processing
- ✓ a high-bandwidth memory subsystem
- ✓ No1. computing-efficiency and power-efficiency in the HPCG Benchmark ranking

★ The time for vector systems has come again

- ✓ Many modern processors employ vector-processing mechanism. and their vector lengths are increasing year by year.
- ✓ However, escalation of vector processing capability is not a only factor,
- ✓ Memory subsystem is now a key factor to increase sustained vector processing performance.

Acknowledgements to Members of Tohoku Univ-NEC Joint Research Division of HPC Technologies and Applications

★ Founded in June, 2014, 4-year period

★ Objectives

- R&D on HPC technologies to exploit high-sustained performance of science and engineering applications on current HPC Systems and to realize Future HPC Systems

- ✓ Evaluation and Improvement of the current HPC environments through migration of SX-9 applications to SX-ACE
- ✓ Detailed Evaluation and Analysis of Modern HPC Systems, not only Vector Systems but also Scalar-Parallel and Accelerator-Based Systems
- ✓ Feasibility study of a future highly balanced HPC system for high sustained performance of practical applications in the post-peta scale era

★ Faculty Members

- Hiroaki Kobayashi, Professor and division director
- Hiroyuki Takizawa, Associate Professor
- Ryusuke Egawa, Associate Professor
- Akihiko Musa (NEC), Visiting Professor
- Mistuo Yokokawa (Kobe Univ), Visiting Professor
- Shintaro Momose (NEC), Visiting Associate Professor
- Masayuki Sato, Assistant Professor

- ✓ In collaboration with visiting researchers from NEC and the technical staff of Cyberscience Center