



東北大学

# Experiences with NEC's New Vector System SX-Aurora TSUBASA and Its Extension for the Future

**Hiroaki Kobayashi**

Special Advisor to President (for ICT Innovation)  
Deputy Director for the HPC Strategy of Cyberscience Center  
Chair of Computer and Mathematics Sciences Department  
Professor of Graduate School of Information Sciences

Tohoku University  
koba@tohoku.ac.jp  
28th WSSP  
October 9-10, 2018



SX-Aurora TSUBASA  
A300-2 #00001

# Today's Agenda

## ★ Quick Introduction of NEC's New Vector System: SX-Aurora TSUBASA

- ✓ An X86-Attached SX Vector System Aiming at Standardization and Customization
- ✓ The New Execution Model of Scalar/OS Offloading

## ★ Early Evaluation of SX-Aurora TSUBASA

- ✓ Tohoku Univ.'s Application Kernels
- ✓ HPCG
- ✓ Vector Offloading Mechanism

## ★ On-going R&D

- ✓ Design consideration of SX-Aurora TSUBASA for the Next Generation
- ✓ R&D of a Quantum Computing-Assisted HPC Infrastructure



東北大学

# The First Impression of SX-Aurora TSUBASA



SX-Aurora TSUBASA  
A300-2 #00001

# NEC Brand-New Vector System: SX-Aurora-Tsubasa

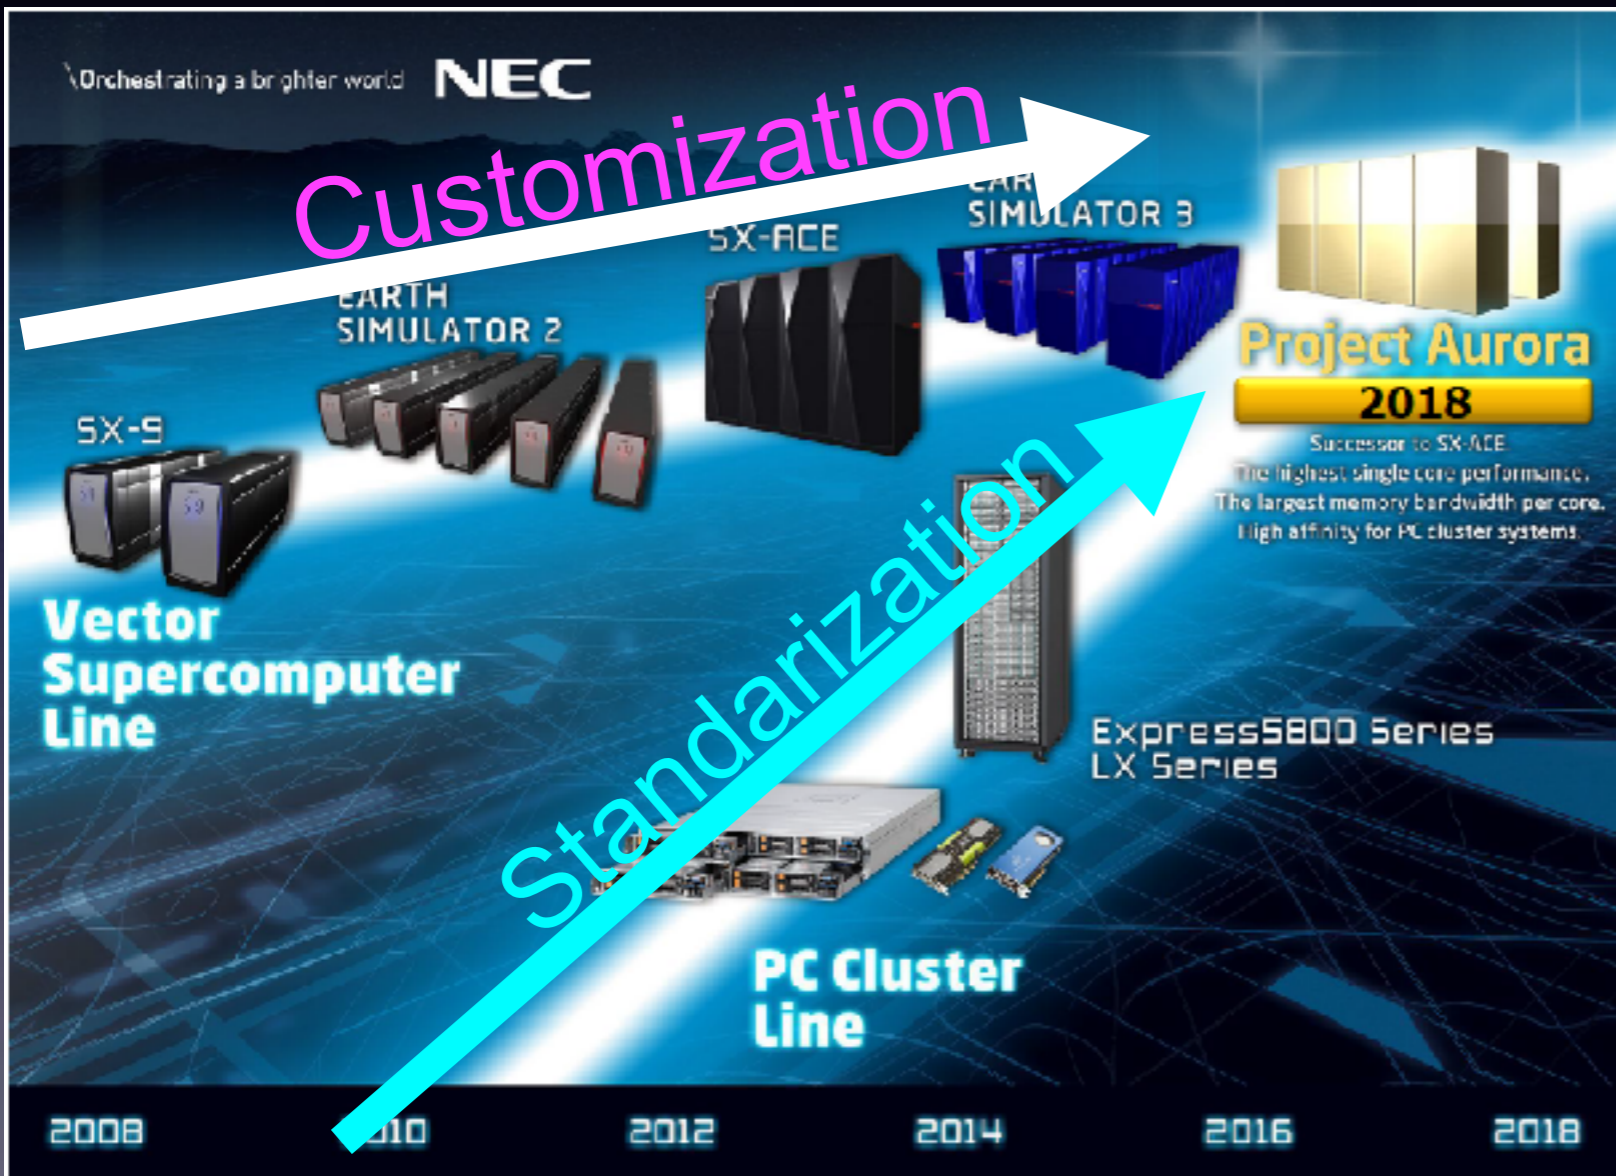
## Features

### ★ Customization

- ✓ Highest Mem. BW
- ✓ Largest Single Core Performance

### ★ Standardization

- ✓ Linux Environment
- ✓ New execution model centralized on vector computing



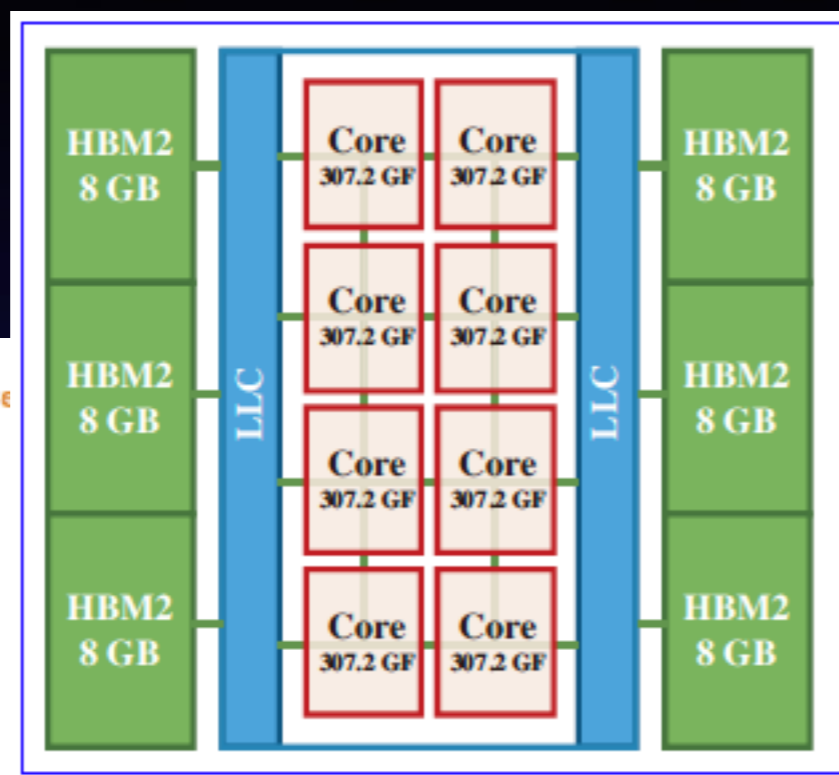
Source: NEC

# Hardware Specification of SX-Aurora TSUBASA

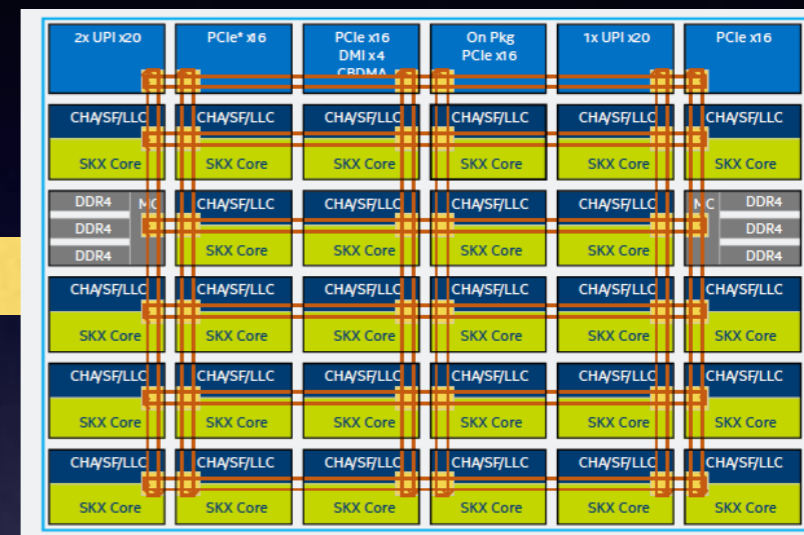
## SX Vector Processor




SX-Aurora TSUBASA  
A300-2 #00001



## X86 Processor(Xeon)



PCIe

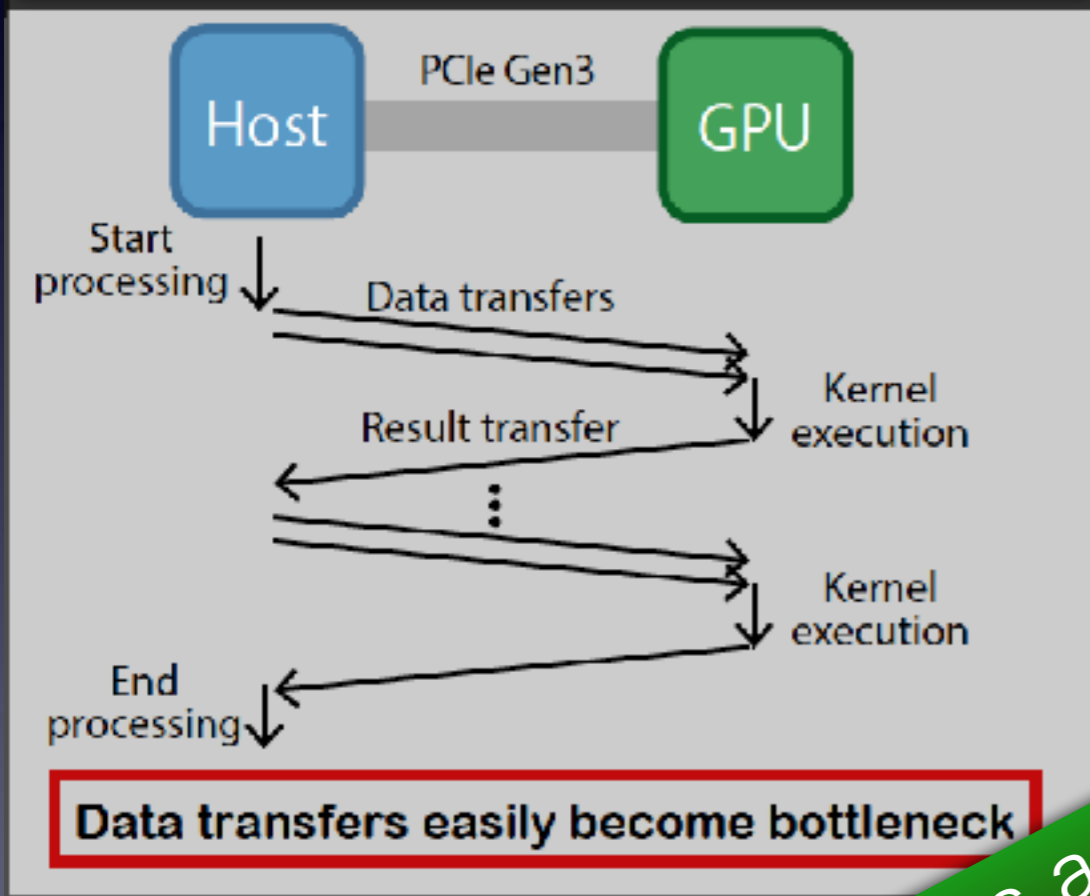


Vector Engine (VE)	Type 10B
Frequency	1.4 GHz
Performance/Core	537.6 GF(SP), 268.8 GF (DP)
No. of Cores	8
Performance/Socket	4.30 TFLOPS (SP) 2.15 TFLOPS (DP)
Memory Subsystem	HBM2 8GB x6
Memory Bandwidth	1.2 TB/s
Memory Capacity	48 GB

Vector Host (VH)	Intel Xeon Gold 6126
Frequency	2.60 GHz / 3.70 GHz (Turbo boost)
Performance/Core	166/236 GF(SP), 83/118 GF (DP)
No. of Cores	12
Performance/Socket	1,996/2,840 GF(SP) 998.4/1,420 GF(DP)
Memory Subsystem	DDR4-2666 DIMM 16GB x 6
Memory Bandwidth	128 GB/s
Memory Capacity	96 GB

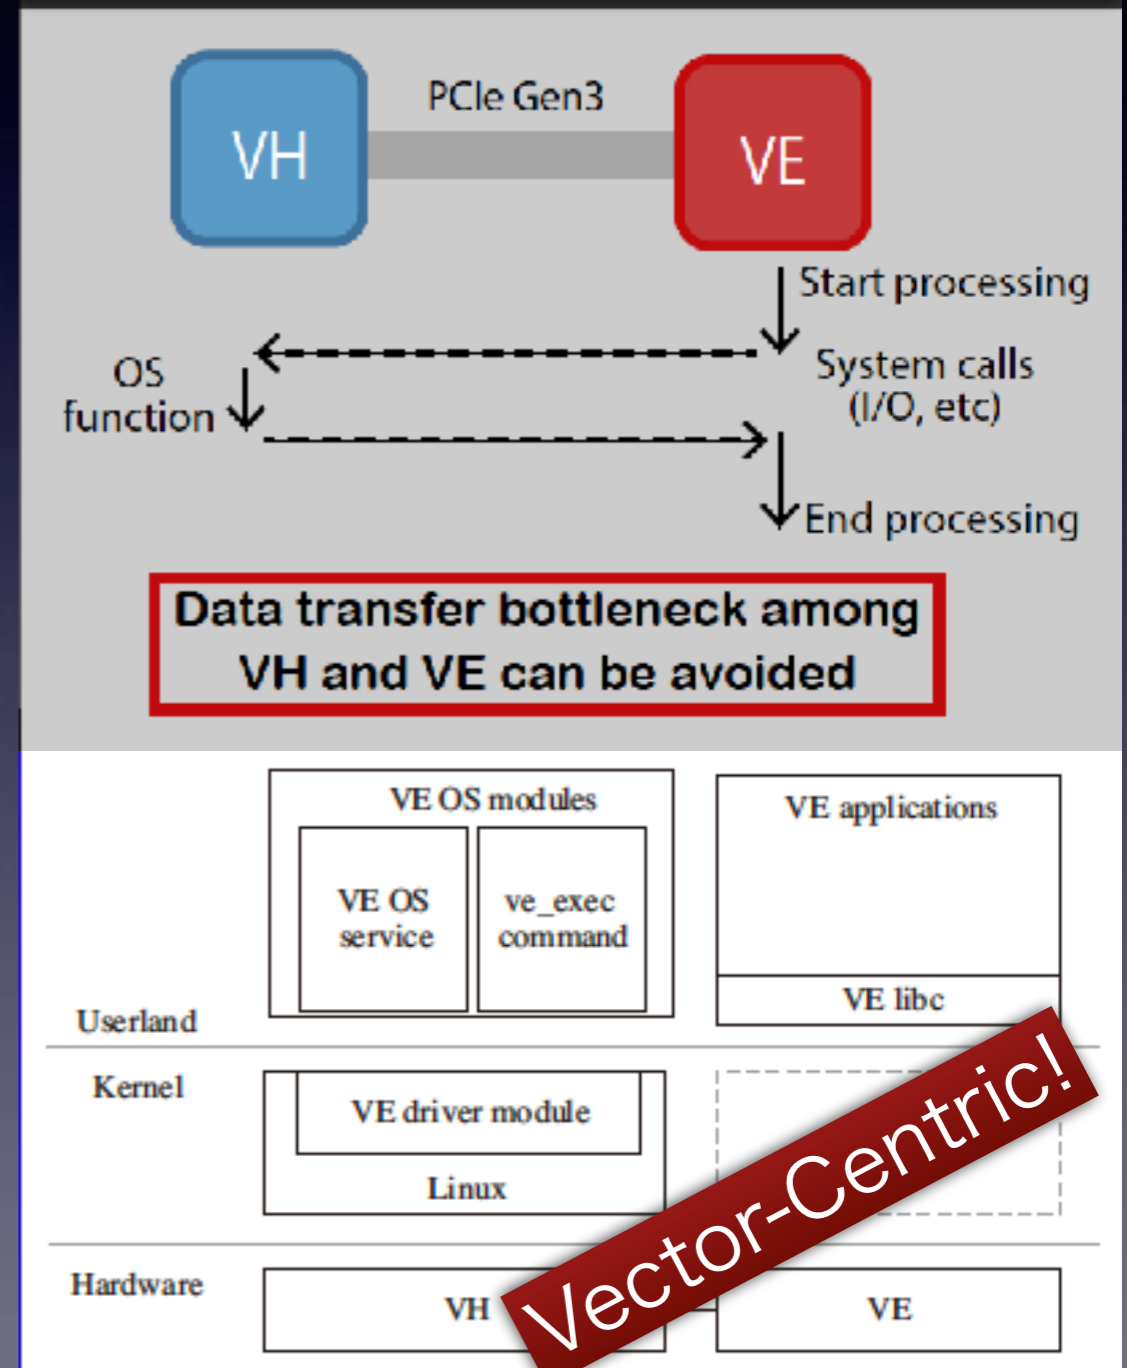
# A New Execution Model of SX-Aurora TSUBASA

## Conventional Execution Model of Accelerators



**Accelerator as a Slave!**

## SX-Aurora TSUBASA Execution Model






# Comparison between SX-ACE and SX-Aurora

		SX-Aurora (2018)	SX-ACE (2014)	Improvement
CPU Performance	Number of Cores	8	4	2x
	Total Flop/s in DP (Total Flop/s in SP)	2.15Tflop/s (4.3Tflop/s)	256Gflop/s	8.4x (16.8x)
	Memory Bandwidth	1.2TB/sec	256GB/sec	4.7x
	ADB Capacity	16MB(Shared)	4MB(Private)	16x
	B/F	0.55	1	0.55X
OS		Lunux	Super-UX	∞



# Comparison between Xeon Gold, SX-Aurora TSUBASA VE and V100

	Intel Xeon Gold 6126	NEC Vector Engine Type 10B	NVIDIA Tesla V100
<b>Frequency</b>	2.6 GHz / 3.7 GHz(Turbo)	1.4 GHz	1.245 GHz
<b>No. of cores</b>	12	8	5120
<b>Performance/socket</b>	1,996/2,840 GF (SP) 998.4/1,420 GF (DP)	4.3 TF (SP) 2.15 TF (DP)	14 TF (SP) 7 TF (DP)
<b>Memory subsystem</b>	DDR4-2666 DIMM 16GB x 6 channels	HBM2 8GB x 6 modules	HBM2 4GB x 4 modules
<b>Memory bandwidth</b>	128 GB/s	1.22 TB/s	900 GB/s
<b>Memory capacity</b>	96 GB	48 GB	16 GB
<b>Price?</b>		+	
		≈	



# You may be interested in Post-K Processor... ~Become available in 2021?~

## A64FX Chip Overview



### Architecture Features

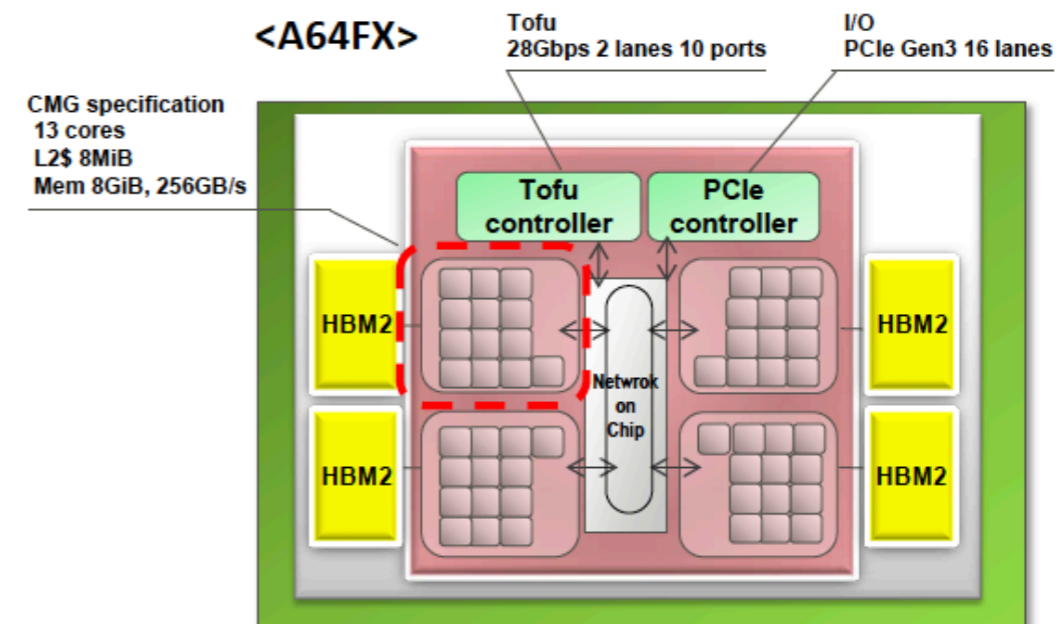
- Armv8.2-A (AArch64 only)
- SVE 512-bit wide SIMD
- 48 computing cores + 4 assistant cores\*  
\*All the cores are identical
- HBM2 32GiB
- Tofu 6D Mesh/Torus  
28Gbps x 2 lanes x 10 ports
- PCIe Gen3 16 lanes

### 7nm FinFET

- 8,786M transistors
- 594 package signal pins

### Peak Performance (Efficiency)

- >2.7TFLOPS (>90%@DGEMM)
- Memory B/W 1024GB/s (>80%@Stream Triad)



	A64FX (Post-K)	SPARC64 XIfx (PRIMEHPC FX100)
ISA (Base)	Armv8.2-A	SPARC-V9
ISA (Extension)	SVE	HPC-ACE2
Process Node	7nm	20nm
Peak Performance	>2.7TFLOPS	1.1TFLOPS
SIMD	512-bit	256-bit
# of Cores	48+4	32+2
Memory	HBM2	HMC
Memory Peak B/W	1024GB/s	240GB/s x2 (in/out)

6

All Rights Reserved. Copyright © FUJITSU LIMITED 2018

T. Yoshida, "Fujitsu High Performance CPU for the Post-K Computer," Hot Chips 30, 2018.


# The Similar Architecture with The Same Performance Available Right Now!

## Vector Engine Processor Overview

SX-Aurora TSUBASA

### Components

- 8 vector cores
- 16MB LLC
- 2D mesh network on chip
- DMA engine
- 6 HBM2 controllers and interfaces
- PCI Express Gen3 x16 interface




### Specs

Core frequency	1.6GHz
Core performance	307GF(DP) 614GF(SP)
CPU performance	2.45TF(DP) 4.91TF(SP)
Memory bandwidth	1.2TB/s
Memory capacity	24/48GB

### Technology


- 16nm FinFET process      4.8 billion transistors



# Benchmark Programs for Performance Evaluation

Kernels	Fields	Methods	Memory access	Grids	Code B/F	Vector Length	Vector Ratio	Actual B/F
Land Mine	Electromagnetic	FDTD	Sequential	100x750x750	6.22	250.9	99.2	5.14
Earthquake	Seismology	Dependent Friction Law	Sequential	2047x2047x256	4.00	255.9	99.4	4.00
Turbulent Flow	CFD	Navier-Stokes	Sequential	512x16384x512	8.00	255.8	99.1	1.47
Antenna	Electromagnetic	FDTD	Sequential	252755x9x97336	1.73	255.7	99.7	0.98
Plasma	Physics	Lax-Wendroff	Indirect	20M	0.82	256.0	70.9	0.11
Turbine	CFD	LU-SGS	Indirect	480x80x80x10	0.96	239.5	99.7	0.0084

# Tohoku Univ.'s Kernels Results



C B/F	6.22	4.00	8.00	1.73	0.82	0.96
S B/F	5.14	4.00	1.47	0.98	0.11	0.0084
Vec. Length	250.9	255.9	255.8	255.7	255.9	239.5
2) Vectrization	99.21	99.38	99.07	99.74	70.94	99.71

# Performance Evaluation of SX-Aurora TSUBASA by Using the HPCG Benchmark

★ HPCG (High Performance Conjugate Gradients) is designed to exercise computational and data access patterns that more closely match a broad set of important 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:

✓ Sparse matrix-vector multiplication.

✓ Sparse triangular solve.

✓ Vector updates.

✓ Global dot products.

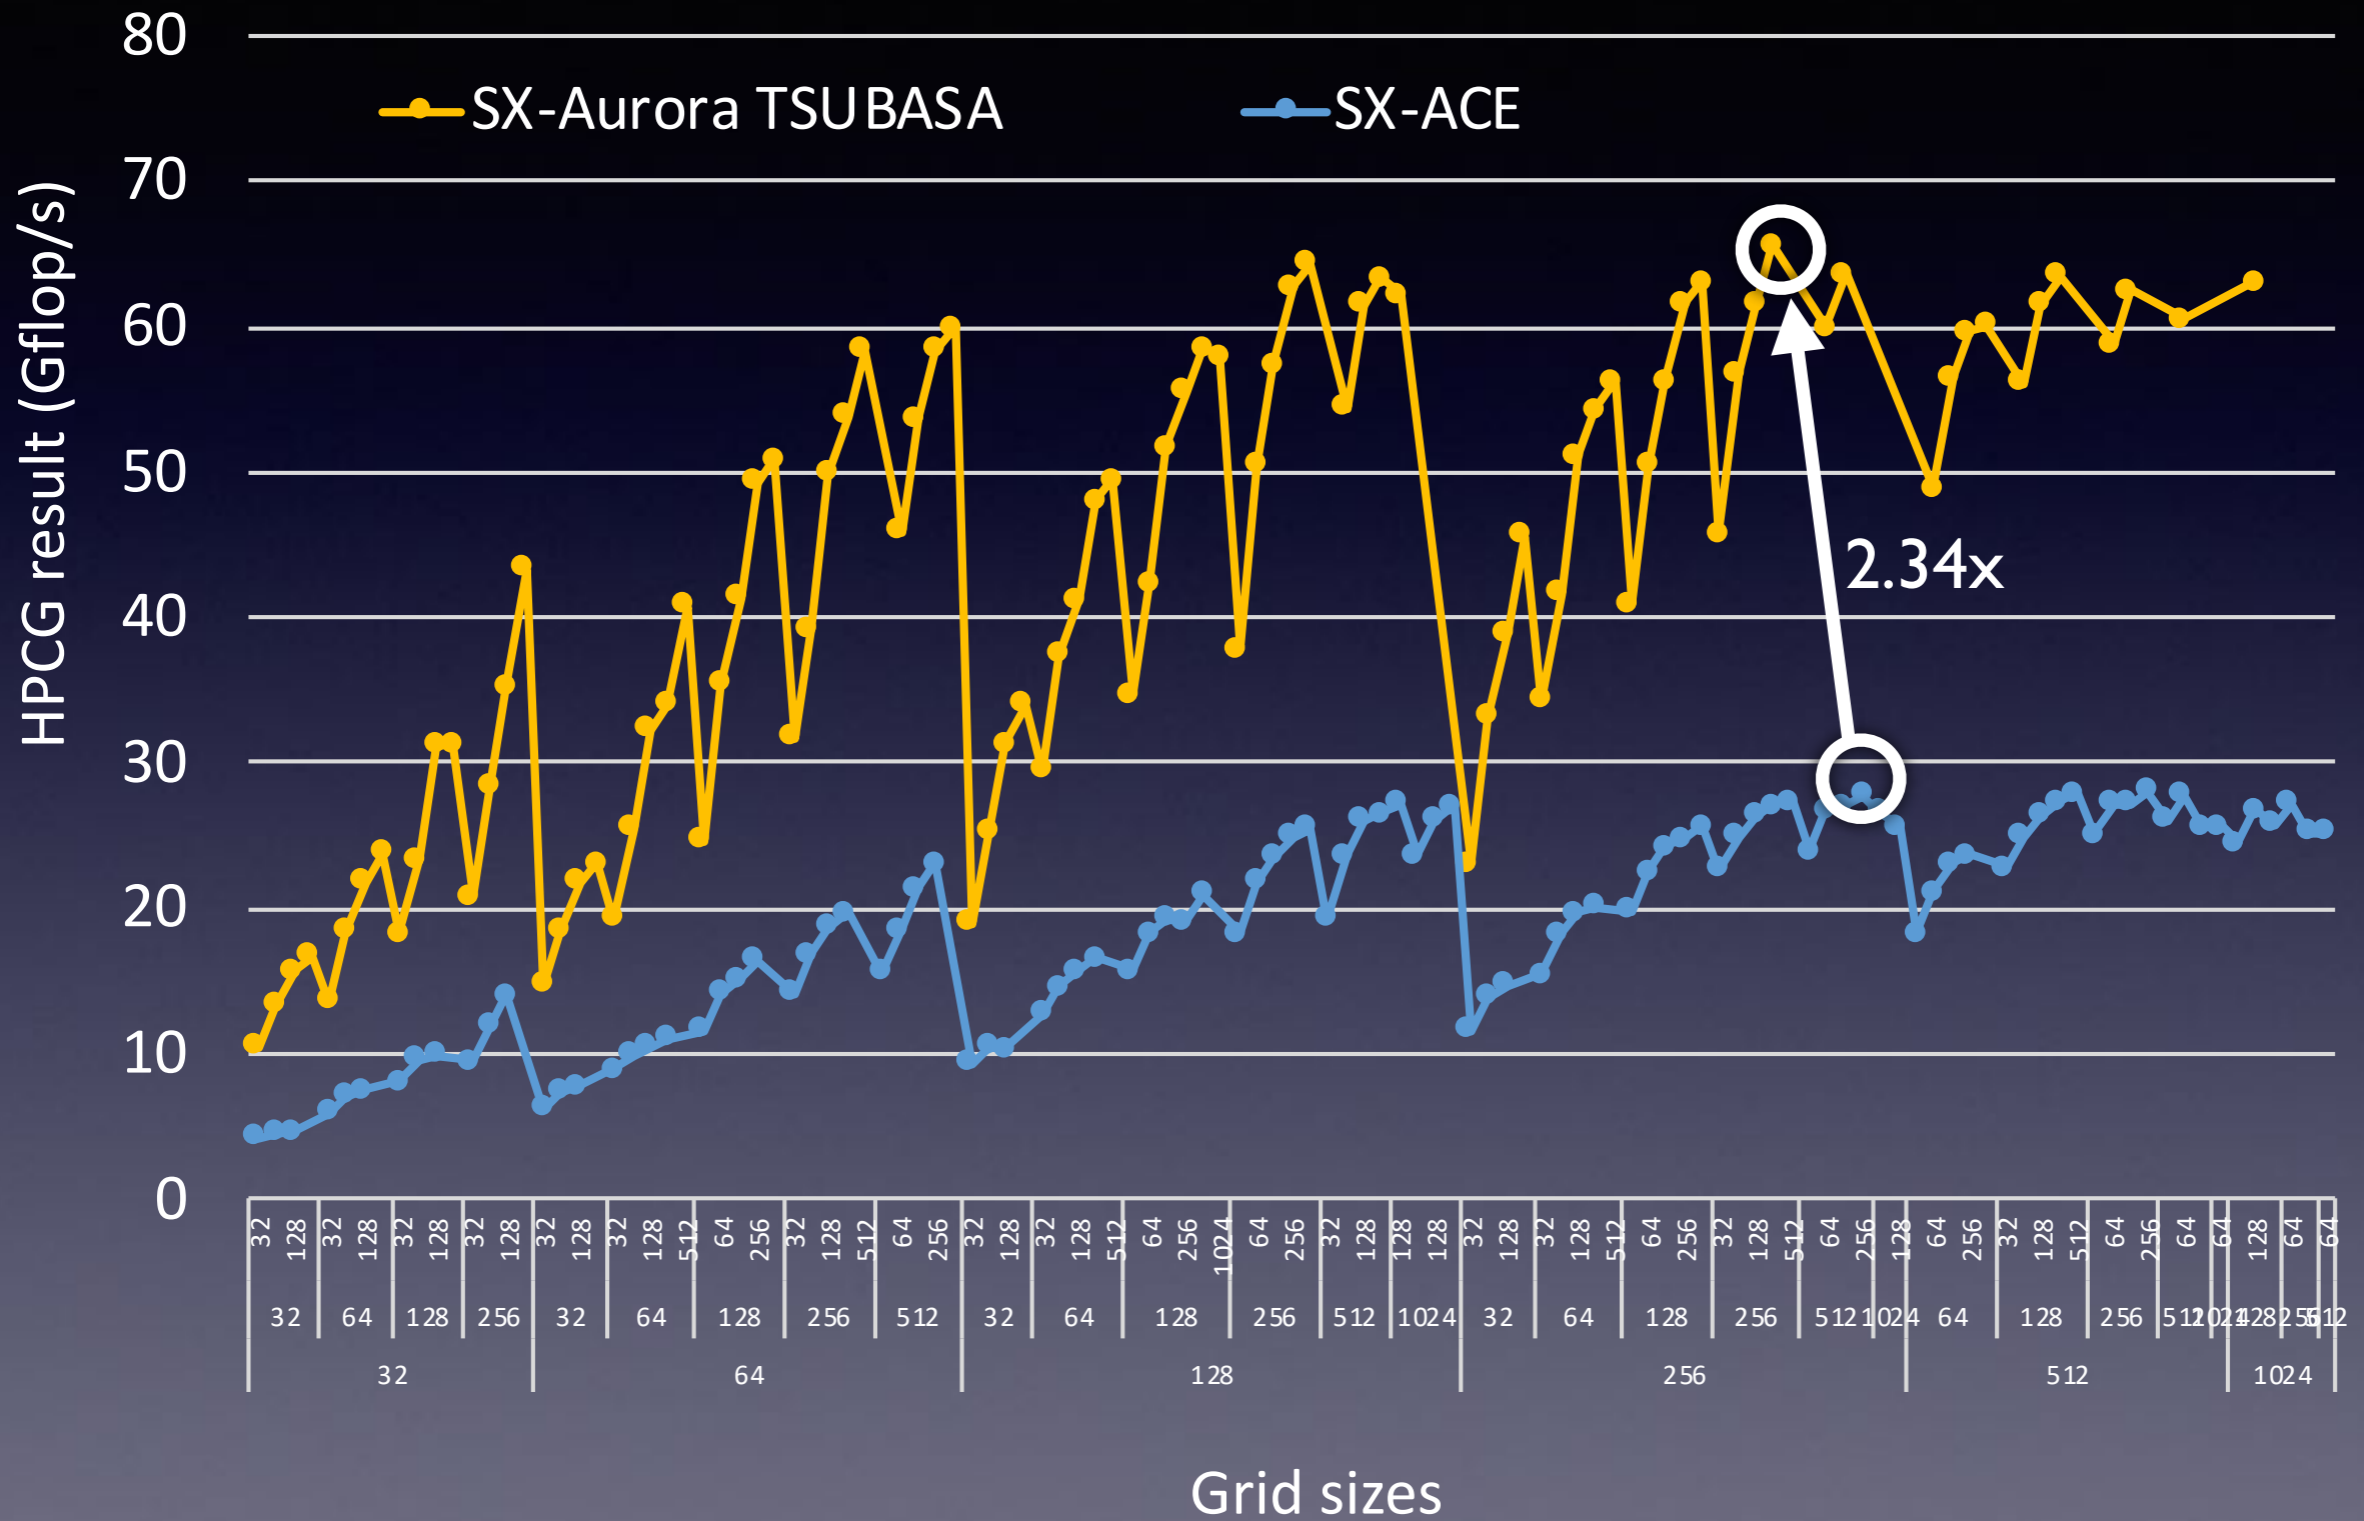
✓ Local symmetric Gauss-Seidel smoother.

✓ Driven by multigrid preconditioned conjugate gradient algorithm that exercises the key kernels on a nested set of coarse grids.

✓ Reference implementation is written in C++ with MPI and OpenMP support.

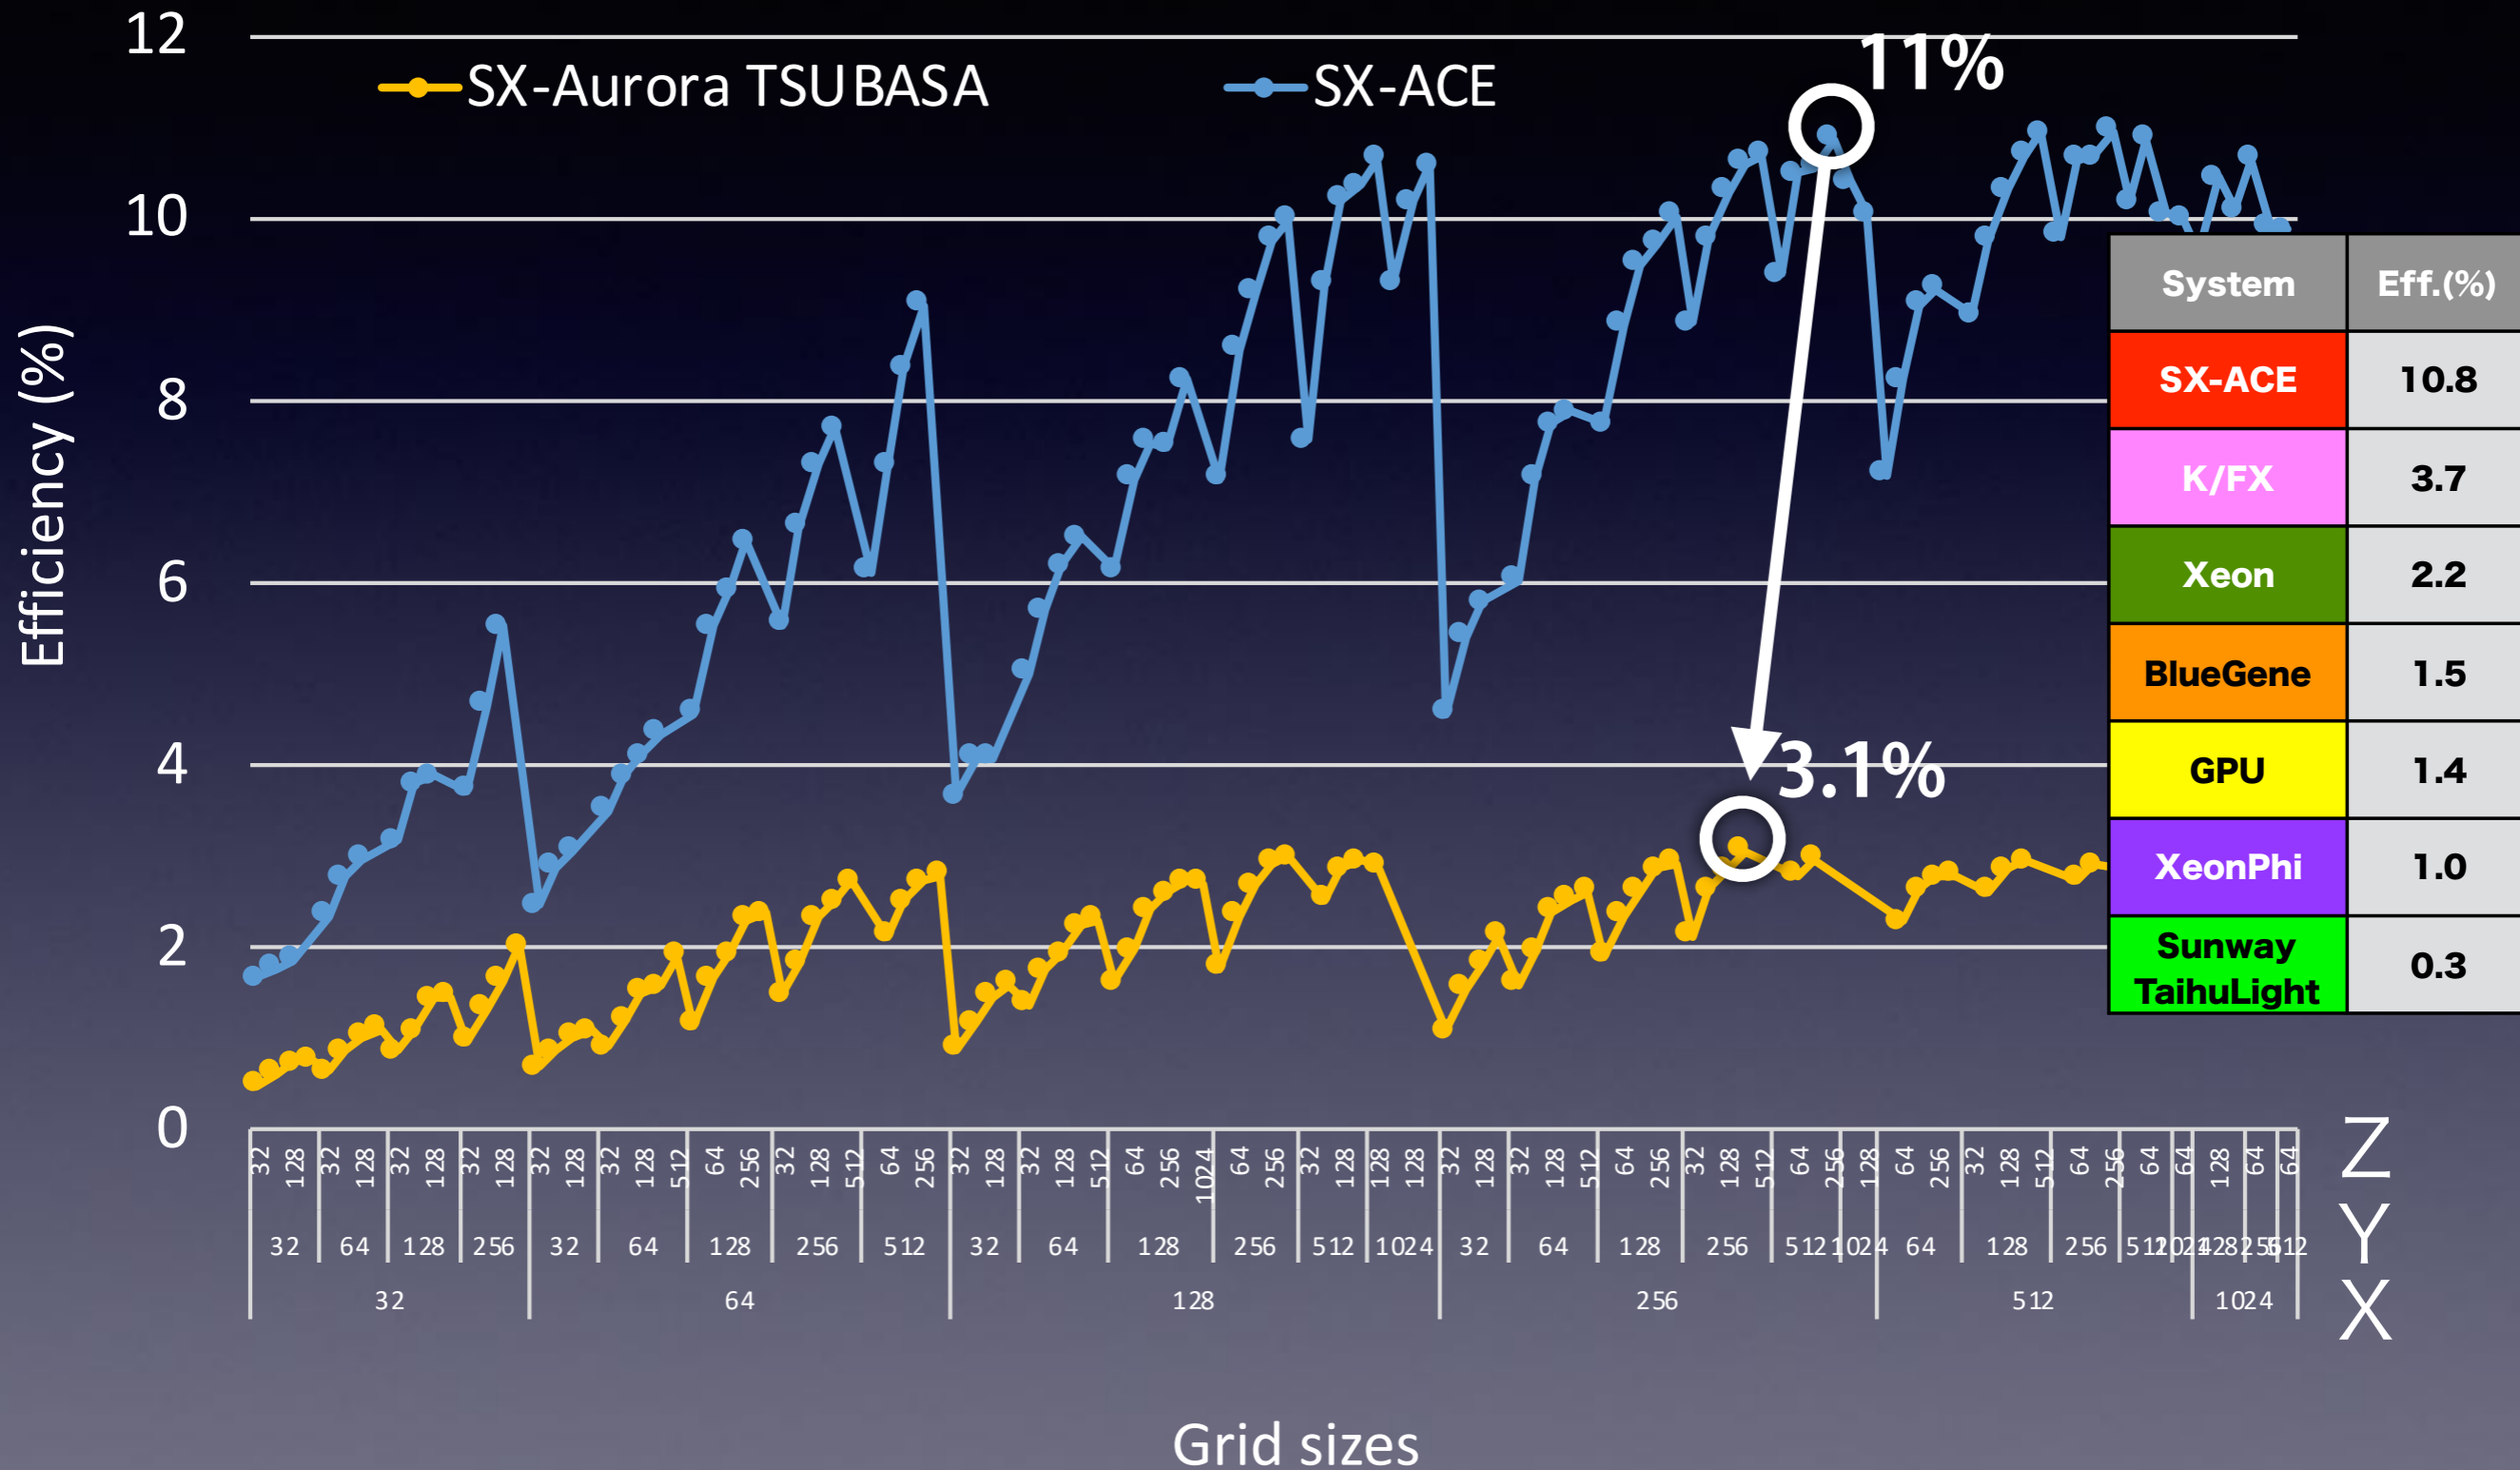
Benchmark	Kernel	Required B/F
HPL	DGEMM	< 0.1
HPGMG	GSRB	> 1
<b>HPCG</b>	<b>SpMV, SYMGS</b>	<b>&gt; 4</b>

# Sustained Performance of HPCG-Benchmark



N  
Y  
X

# HPCG-Benchmark Efficiency



# Evaluation of the New Execution Model: OS/Scalar Offloading from Vector Processing

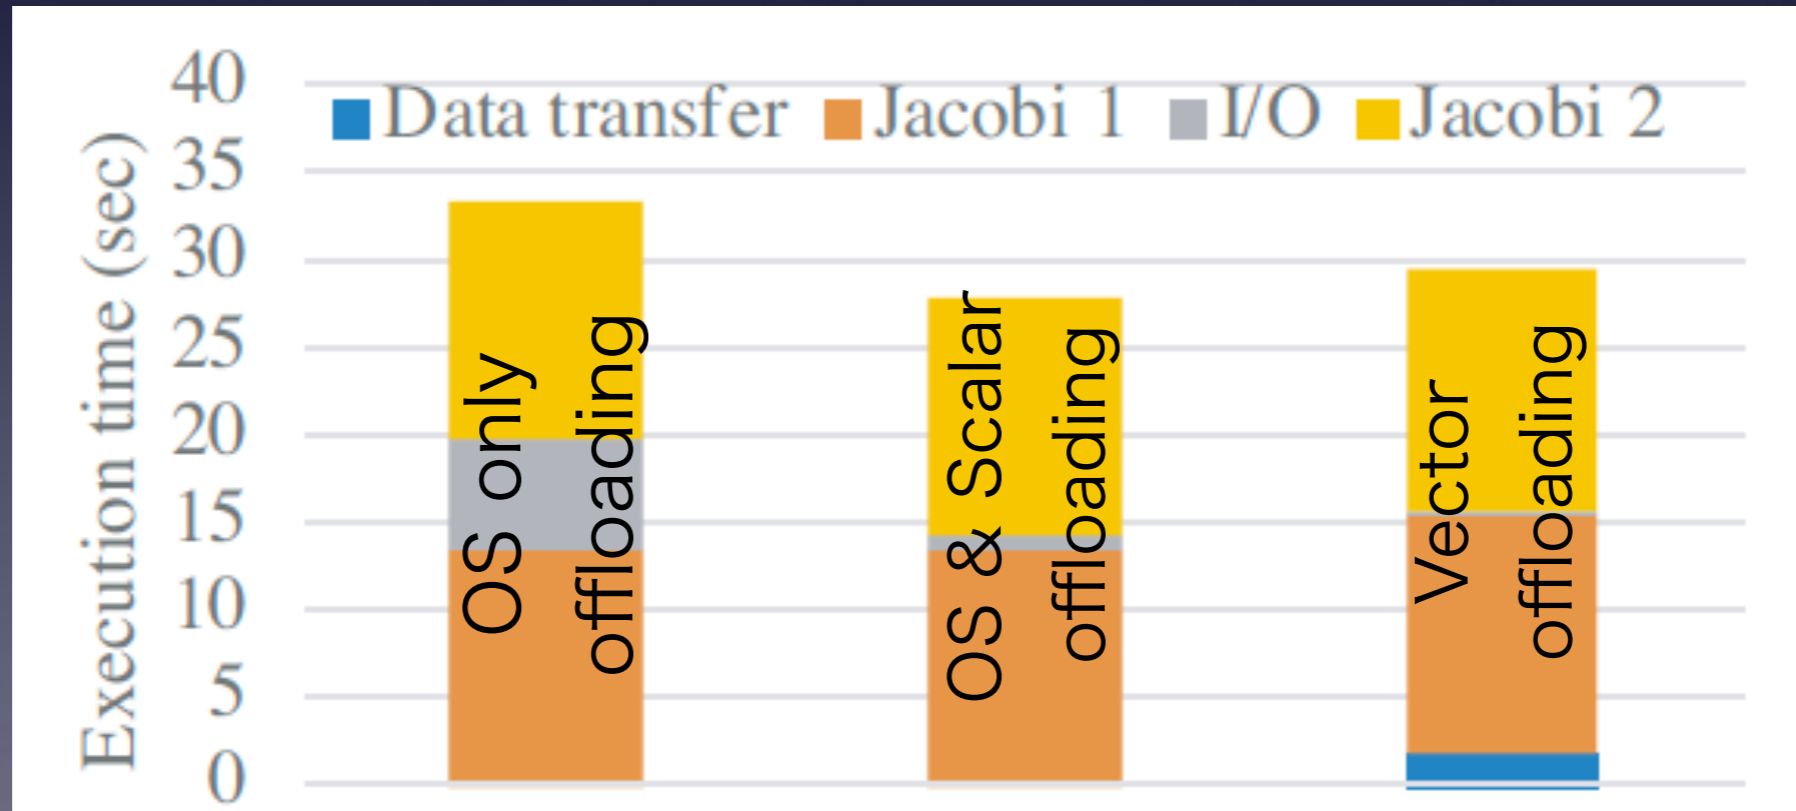


Offloading of OS and scalar operations

Offloading of vector operations

VH Offloading Mode

VE Offloading Mode



VH Offloading

VE Offloading



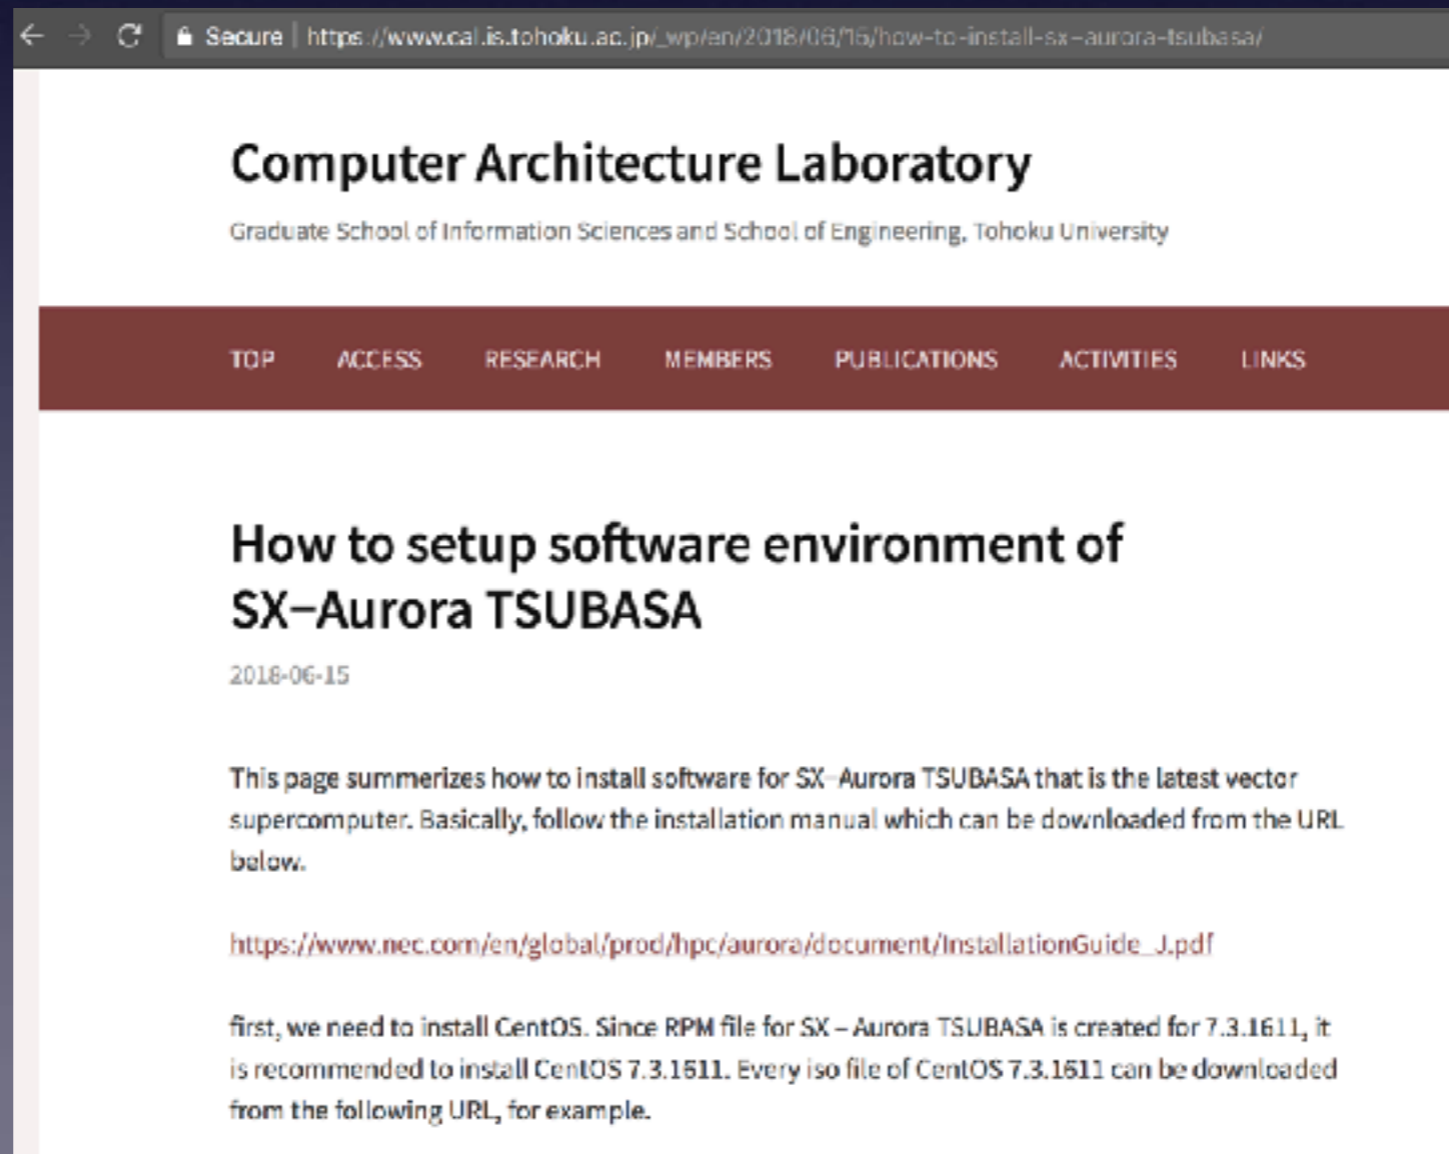
# Impressions of SX-Aurora TSUBASA

- ★ **SX-Aurora TSUBASA has a great potential to achieve a high sustained performance for memory-intensive applications, but...**
- ✓ Compiler development is still underway, limiting the sustained performance regarding auto-vectorization and auto-parallelization, anyway use the latest one for the best performance!
- ✓ Compiler is also not fully exploiting enlarged and core-shared capacity of LLC. Software controlled function is desired to make the best use of it for reducing off-chip memory transactions
- ✓ For some applications, the LLC bandwidth to cores becomes a bottleneck even with a high high rates
- Shared LCC of SX-Aurora, 2.66 against 1.2 of Mem. vs. Dedicated ADB of SX-ACE, 1 against 0.256 all in TB/s)

# Unofficial Web Site of SX-Aurora TSUBASA

[http://www.cal.is.tohoku.ac.jp/\\_wp/en/2018/06/15/how-to-install-sx-aurora-tsubasa/](http://www.cal.is.tohoku.ac.jp/_wp/en/2018/06/15/how-to-install-sx-aurora-tsubasa/)

- Our website provides the information about
  - How to setup software environments
  - How to update software environments
  - Events
  - etc



The screenshot shows a web browser window with the URL [https://www.cal.is.tohoku.ac.jp/\\_wp/en/2018/06/15/how-to-install-sx-aurora-tsubasa/](https://www.cal.is.tohoku.ac.jp/_wp/en/2018/06/15/how-to-install-sx-aurora-tsubasa/). The page header identifies the 'Computer Architecture Laboratory' at Tohoku University. A navigation menu includes links for TOP, ACCESS, RESEARCH, MEMBERS, PUBLICATIONS, ACTIVITIES, and LINKS. The main content area features the title 'How to setup software environment of SX-Aurora TSUBASA' with a date of 2018-06-15. The text explains that the page summarizes how to install software for SX-Aurora TSUBASA on the latest vector supercomputer, and provides a link to the installation manual: [https://www.nec.com/en/global/prod/hpc/aurora/document/InstallationGuide\\_J.pdf](https://www.nec.com/en/global/prod/hpc/aurora/document/InstallationGuide_J.pdf). It also notes that CentOS 7.3.1511 is recommended for installation.




東北大学

# Design Consideration of the Future Vector Systems \*

\*This work is partially conducted with NEC, but the contents do not reflect any future products of NEC



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



# Reenforce the academic and industry collaboration for the HPC R&D at Tohoku University

## ★ Tohoku-Univ NEC Joint Research Division of High-Performance Computing

★ Founded in June, 2014, 8-Year Period until 2022

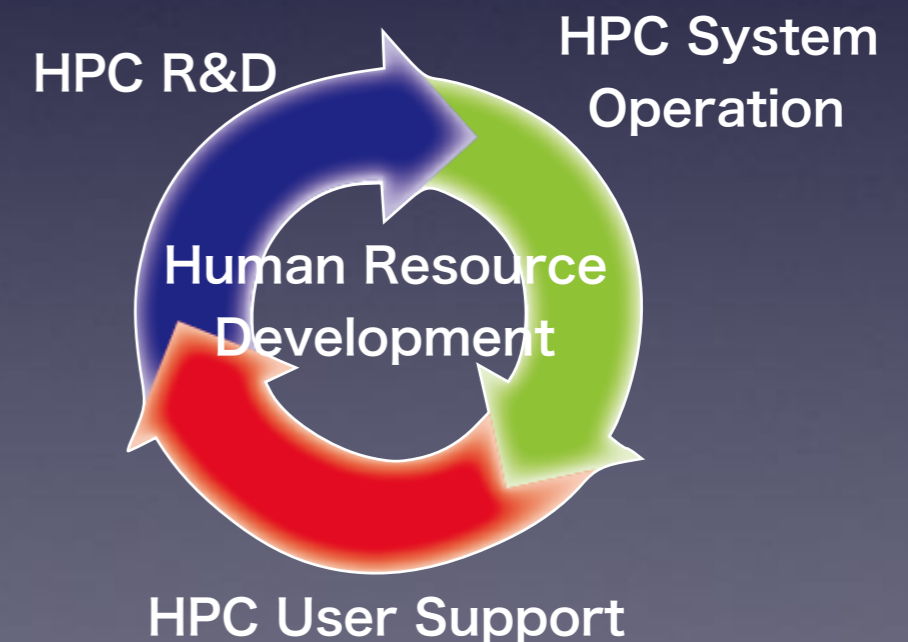
### ★ Objectives

- 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 targeting at 2021
- ✓ 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, Professor
- Ryusuke Egawa, Associate Professor
- Akihiko Musa (NEC), Visiting Professor
- Mitsuo Yokokawa (Kobe Univ), Visiting Professor
- Shintaro Momose (NEC), Visiting Associate Professor
- Masayuki Sato, Assistant Professor

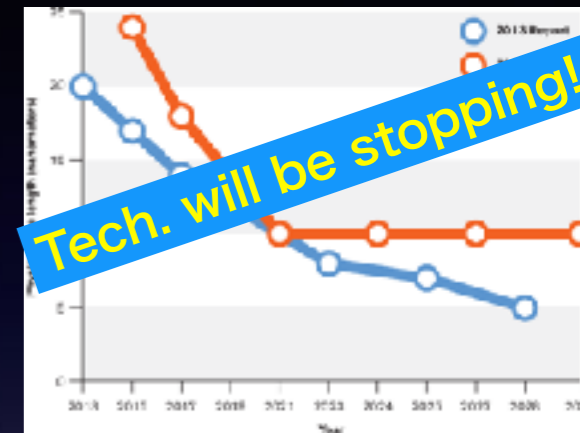
- ✓ In collaboration with visiting researchers from NEC and the technical staff of Cyberscience Center



# Scaling may be End, but Silicon is not End! And Use it Smart and Effective!

- ✓ We are facing the end of Moore's law due to the physical limitations, and the transistor cost is hard to reduce, however
- ✓ Silicon is still fundamental constructing material for computing platforms such as plastic, steel and concrete for automobiles, buildings and home appliances.
- ★ So, we have to become much more smart for design of Future HEC systems.
- ★ Use precious silicon budget (+ advanced device technologies) to effectively design mechanisms that can maximize the sustained performance and power-efficiency of individual applications domains.

Tech Scaling



Fabrication Cost



It's time to focus on **Domain-Specific Architectures(DSAs)** for computation-intensive, memory-intensive, I/O intensive, low-precision computing... etc applications to improve silicon/power efficiency!

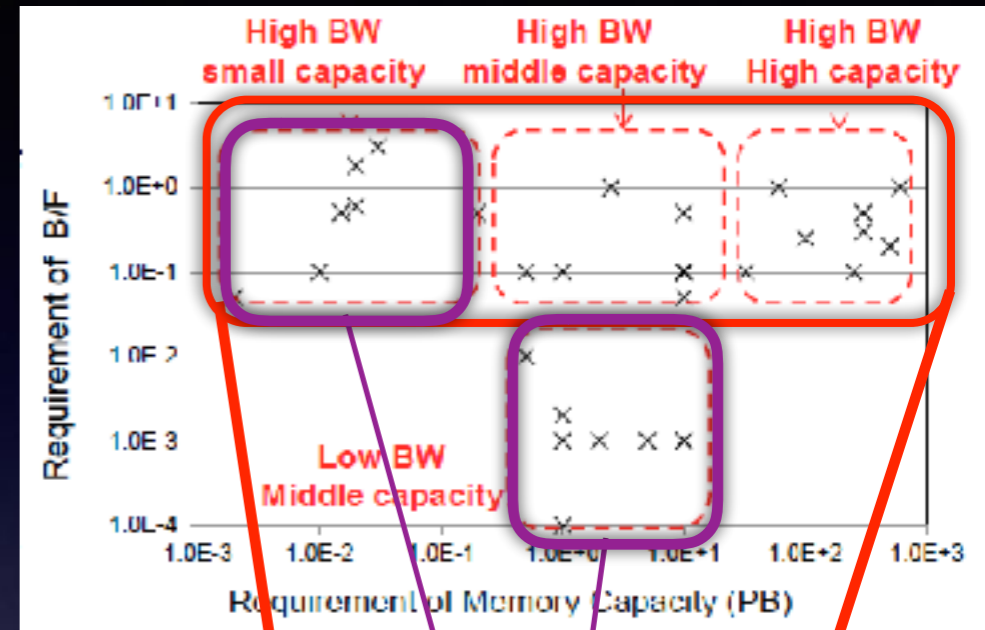
**New HPC System Architecture Design Concept of Ensemble Architecture:**

Make different DSAs combine and complementary work together to realize the general-purpose functionality as a single computing infrastructure

# Domain Specific Balanced Architecture Design Approach: Not Peak Performance, Turn Memory-BW into Sustained Performance!


For computation-intensive applications only!

Need balanced improvement both in flop/s and BW! for high-efficiency in wide application areas


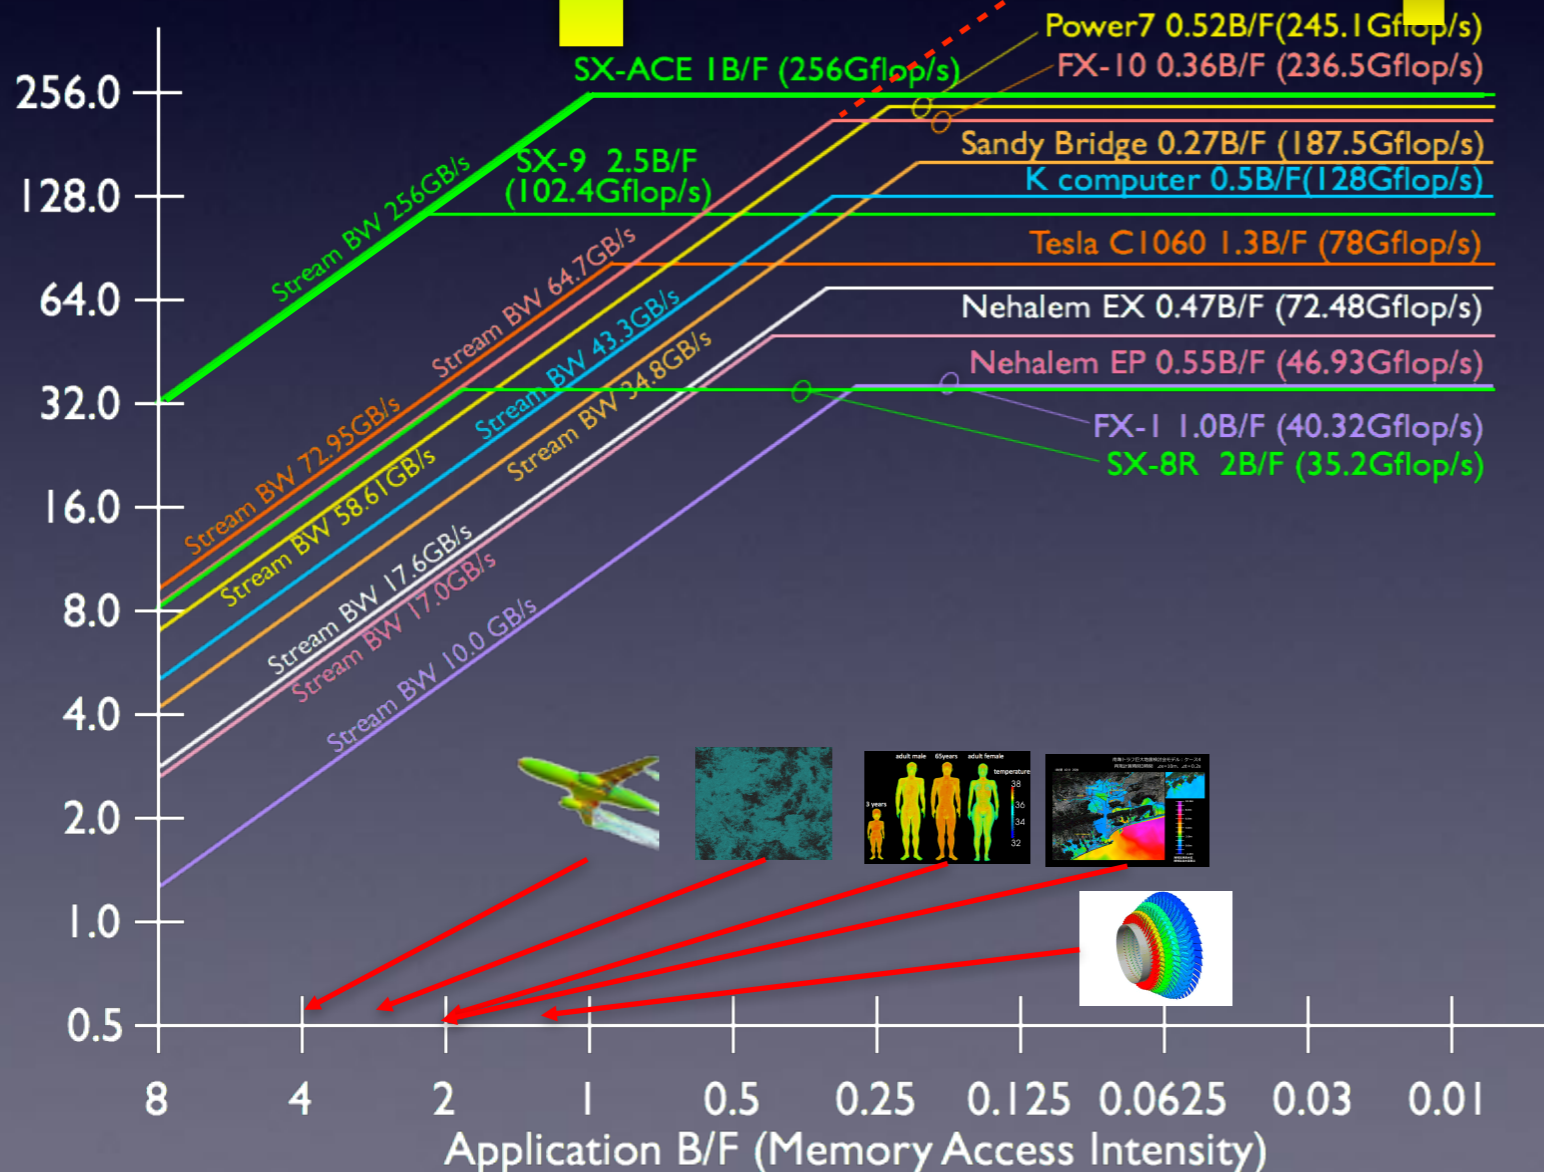


Flop/s-oriented, memory-limited design

Sustained Performance



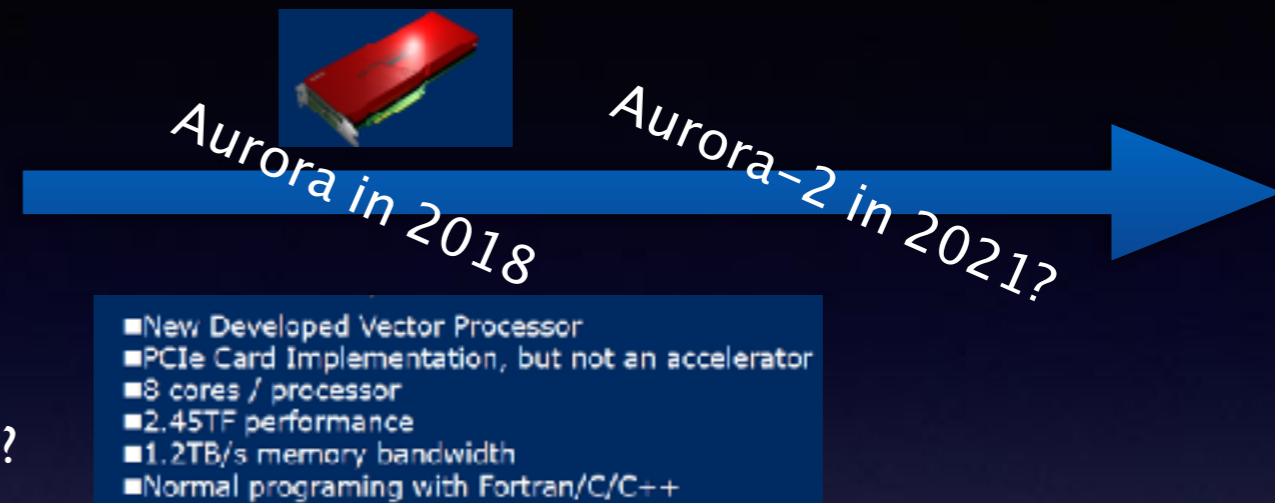
Application Spectrum



# What Does the Next Vector System Look Like in Year Around 2020-2021?

## ★ Vector Engine Spec.

- The 7nm Technology becomes available?
- 5X more transistors from 16nm tech?
- 5X in # of Cores, i.e. 50 VE cores feasible?
- up to 15TF, if the core performance is same, but should be lowered due to power/thermal limitation of the chip.



## ★ Memory Subsystem

- 2x in Memory BW, and 1.5X in Memory Capacity when using HBM 3 under the assumption of the same chip size of Aurora-TSUBASA
- ~3TB/s and ~96GB??

## ★ Design targets of 0.5BF (20 cores of 6TF for memory-intensive applications) to 0.25 BF (40 Cores of 12TF for compute-intensive applications)


- be competitive with contemporary HEC systems at that time, such as Post-K (JP), A2I (US), NERSC-9 (USA), Crossroads (US), EU Exa-System (FR/GE), NUDT2020 (Ch)...




# What Does the Next Vector System Look Like in Year Around 2020-2021? (Cont'd)

★ How 20~40 cores are integrated and connected.

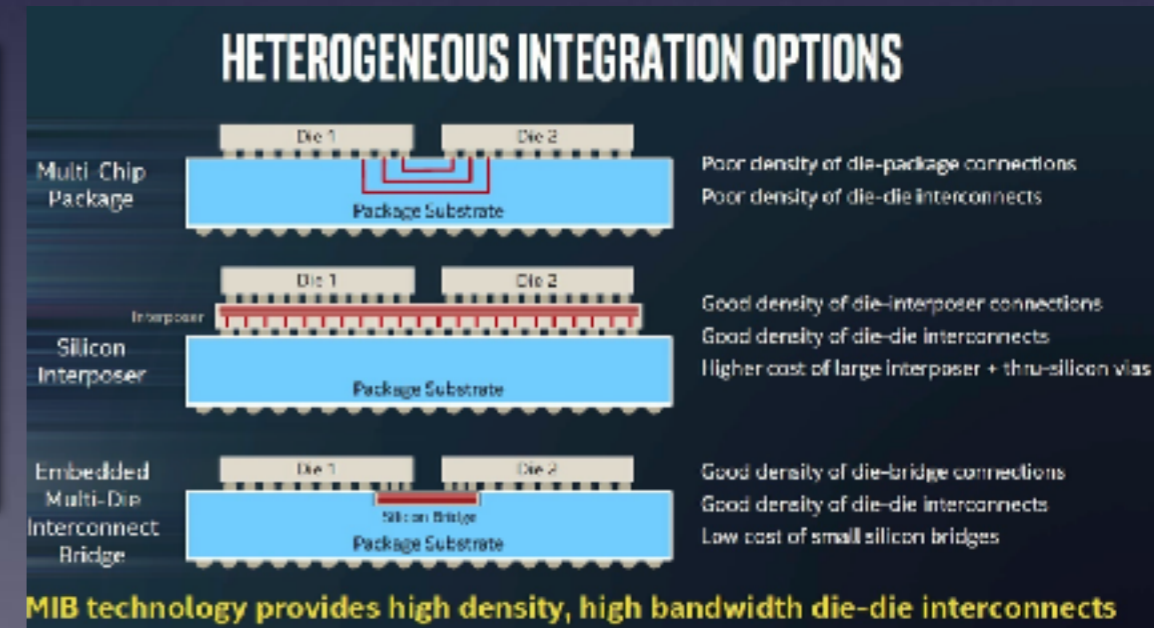
- Single chip or multi-chip (SIP) ?
- If SIP is employed, how multiple chips are connected?
  - ✓ If EMIB available, BW could be increased?
  - ✓ Silicon photonics with WDM becomes available?
- Single SMP or clustered SMP
- crossbar, mesh, ring, etc or their hierarchical and hybrid?
- coherency protocol of ADB (Snoopy or Directory)



Source by IBM



32 Cores, 9.6TF, 3TB/s, 0.3BF




Source by Intel

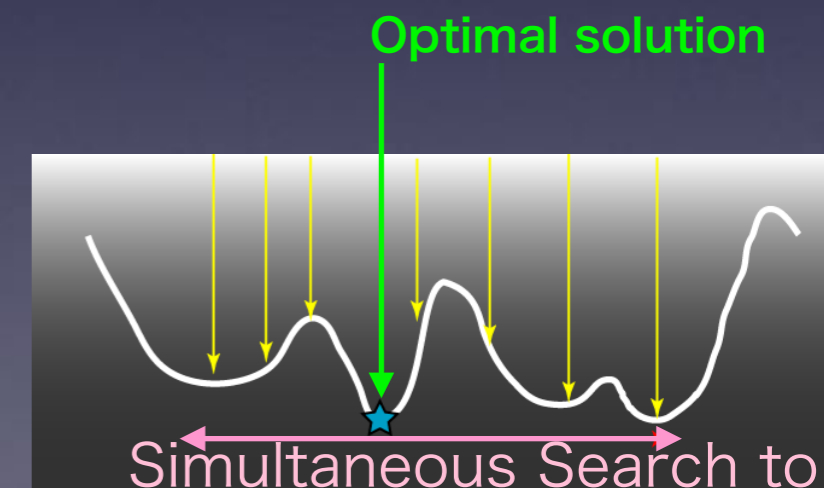
# Quantum Computer: Emerging Domain Specific Architecture

- ★ Quantum computing is drawing much attention recently as an emerging technology in the era of post-Moore
- ✓ In particular, quantum computers for quantum annealing are commercialized by the D-wave systems, and their applications are developed world-widely.
  - ✓ Google, NASA, Volkswagen, Lockheed, Denso...
  - ✓ The base model named the Ising model to design and implement the D-wave machines has been proposed by Prof. Nishimori et al of Tokyo Inst. Tech. In 1998.
- ★ The quantum annealing is a metaheuristic for finding the global minimum of a given objective function over a given set of candidate solutions (candidate states), by a process using quantum fluctuations

**An ideal solver for combinatorial problems!**



Transverse magnetic field type quantum annealing  
Chip and System (D-Wave)



reach optimal one by  
Quantum Fluctuation

# Toward Realization of Quantum Computing-Assisted HPC Infrastructure

- ★ Tohoku University has established an interdisciplinary priority research institute, named Q-HPC, for Quantum Computing-Accelerated HPC in 2018
- ★ As Q-HPC members, we start a new 5-year research program named “R&D of Quantum Annealing-Assisted HPC Infrastructure”, supported by MEXT
  - ✓ Becomes an innovative infrastructure to develop next-generation applications in the fields of computational science, data sciences and their fusions
  - ✓ provides transparent accesses to not only classical HPC resources but also Quantum Computing one in a unified fashion.

Innovative Applications in the fields of Computational Science, Data Science and their Fusion

QA-Assisted HPC infrastructure

Inductive Processing  
(for Data Science)

Deductive Processing  
(for Computational Science)

Coupling

QA  
Platform

QA-Assisted  
AI-ML Platform

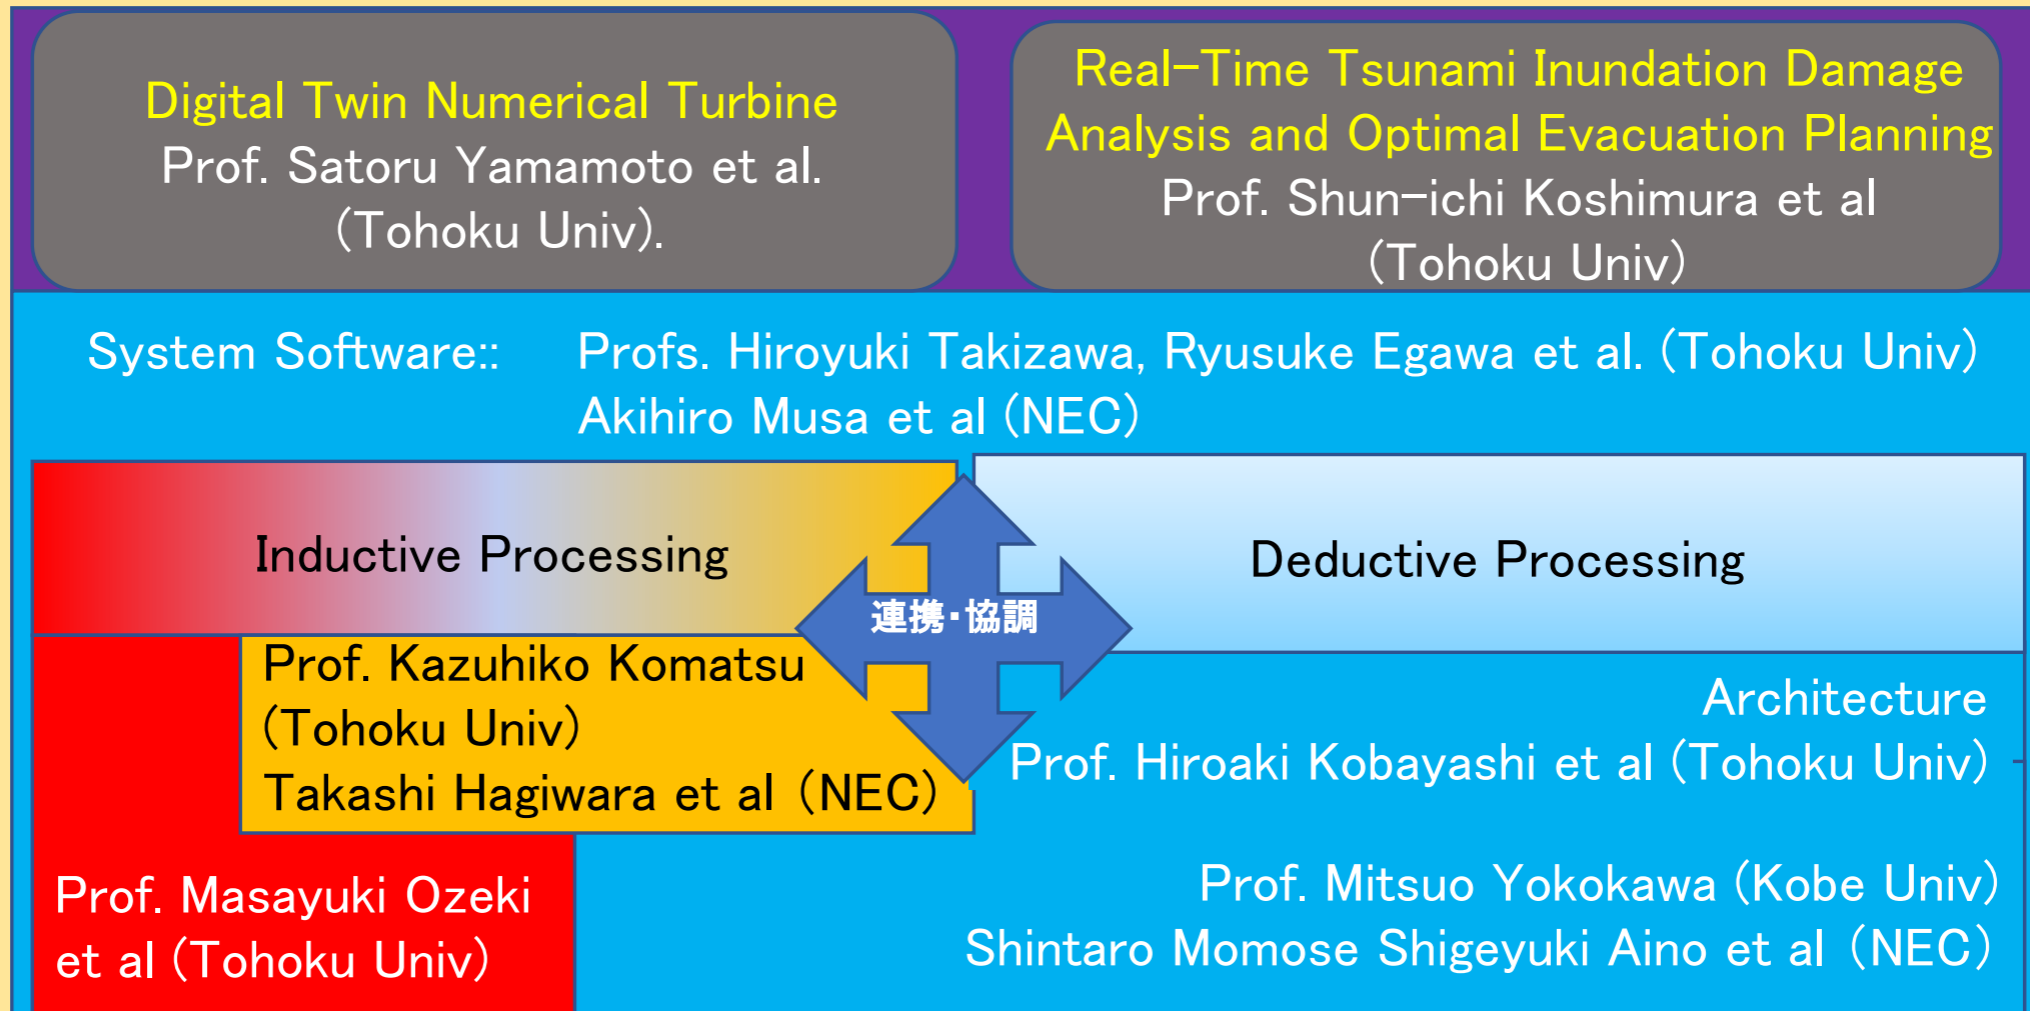

X86-Assisted Vector Computing Platform

D-Wave Machines

SX-Aurora TSUBASA and its Successors

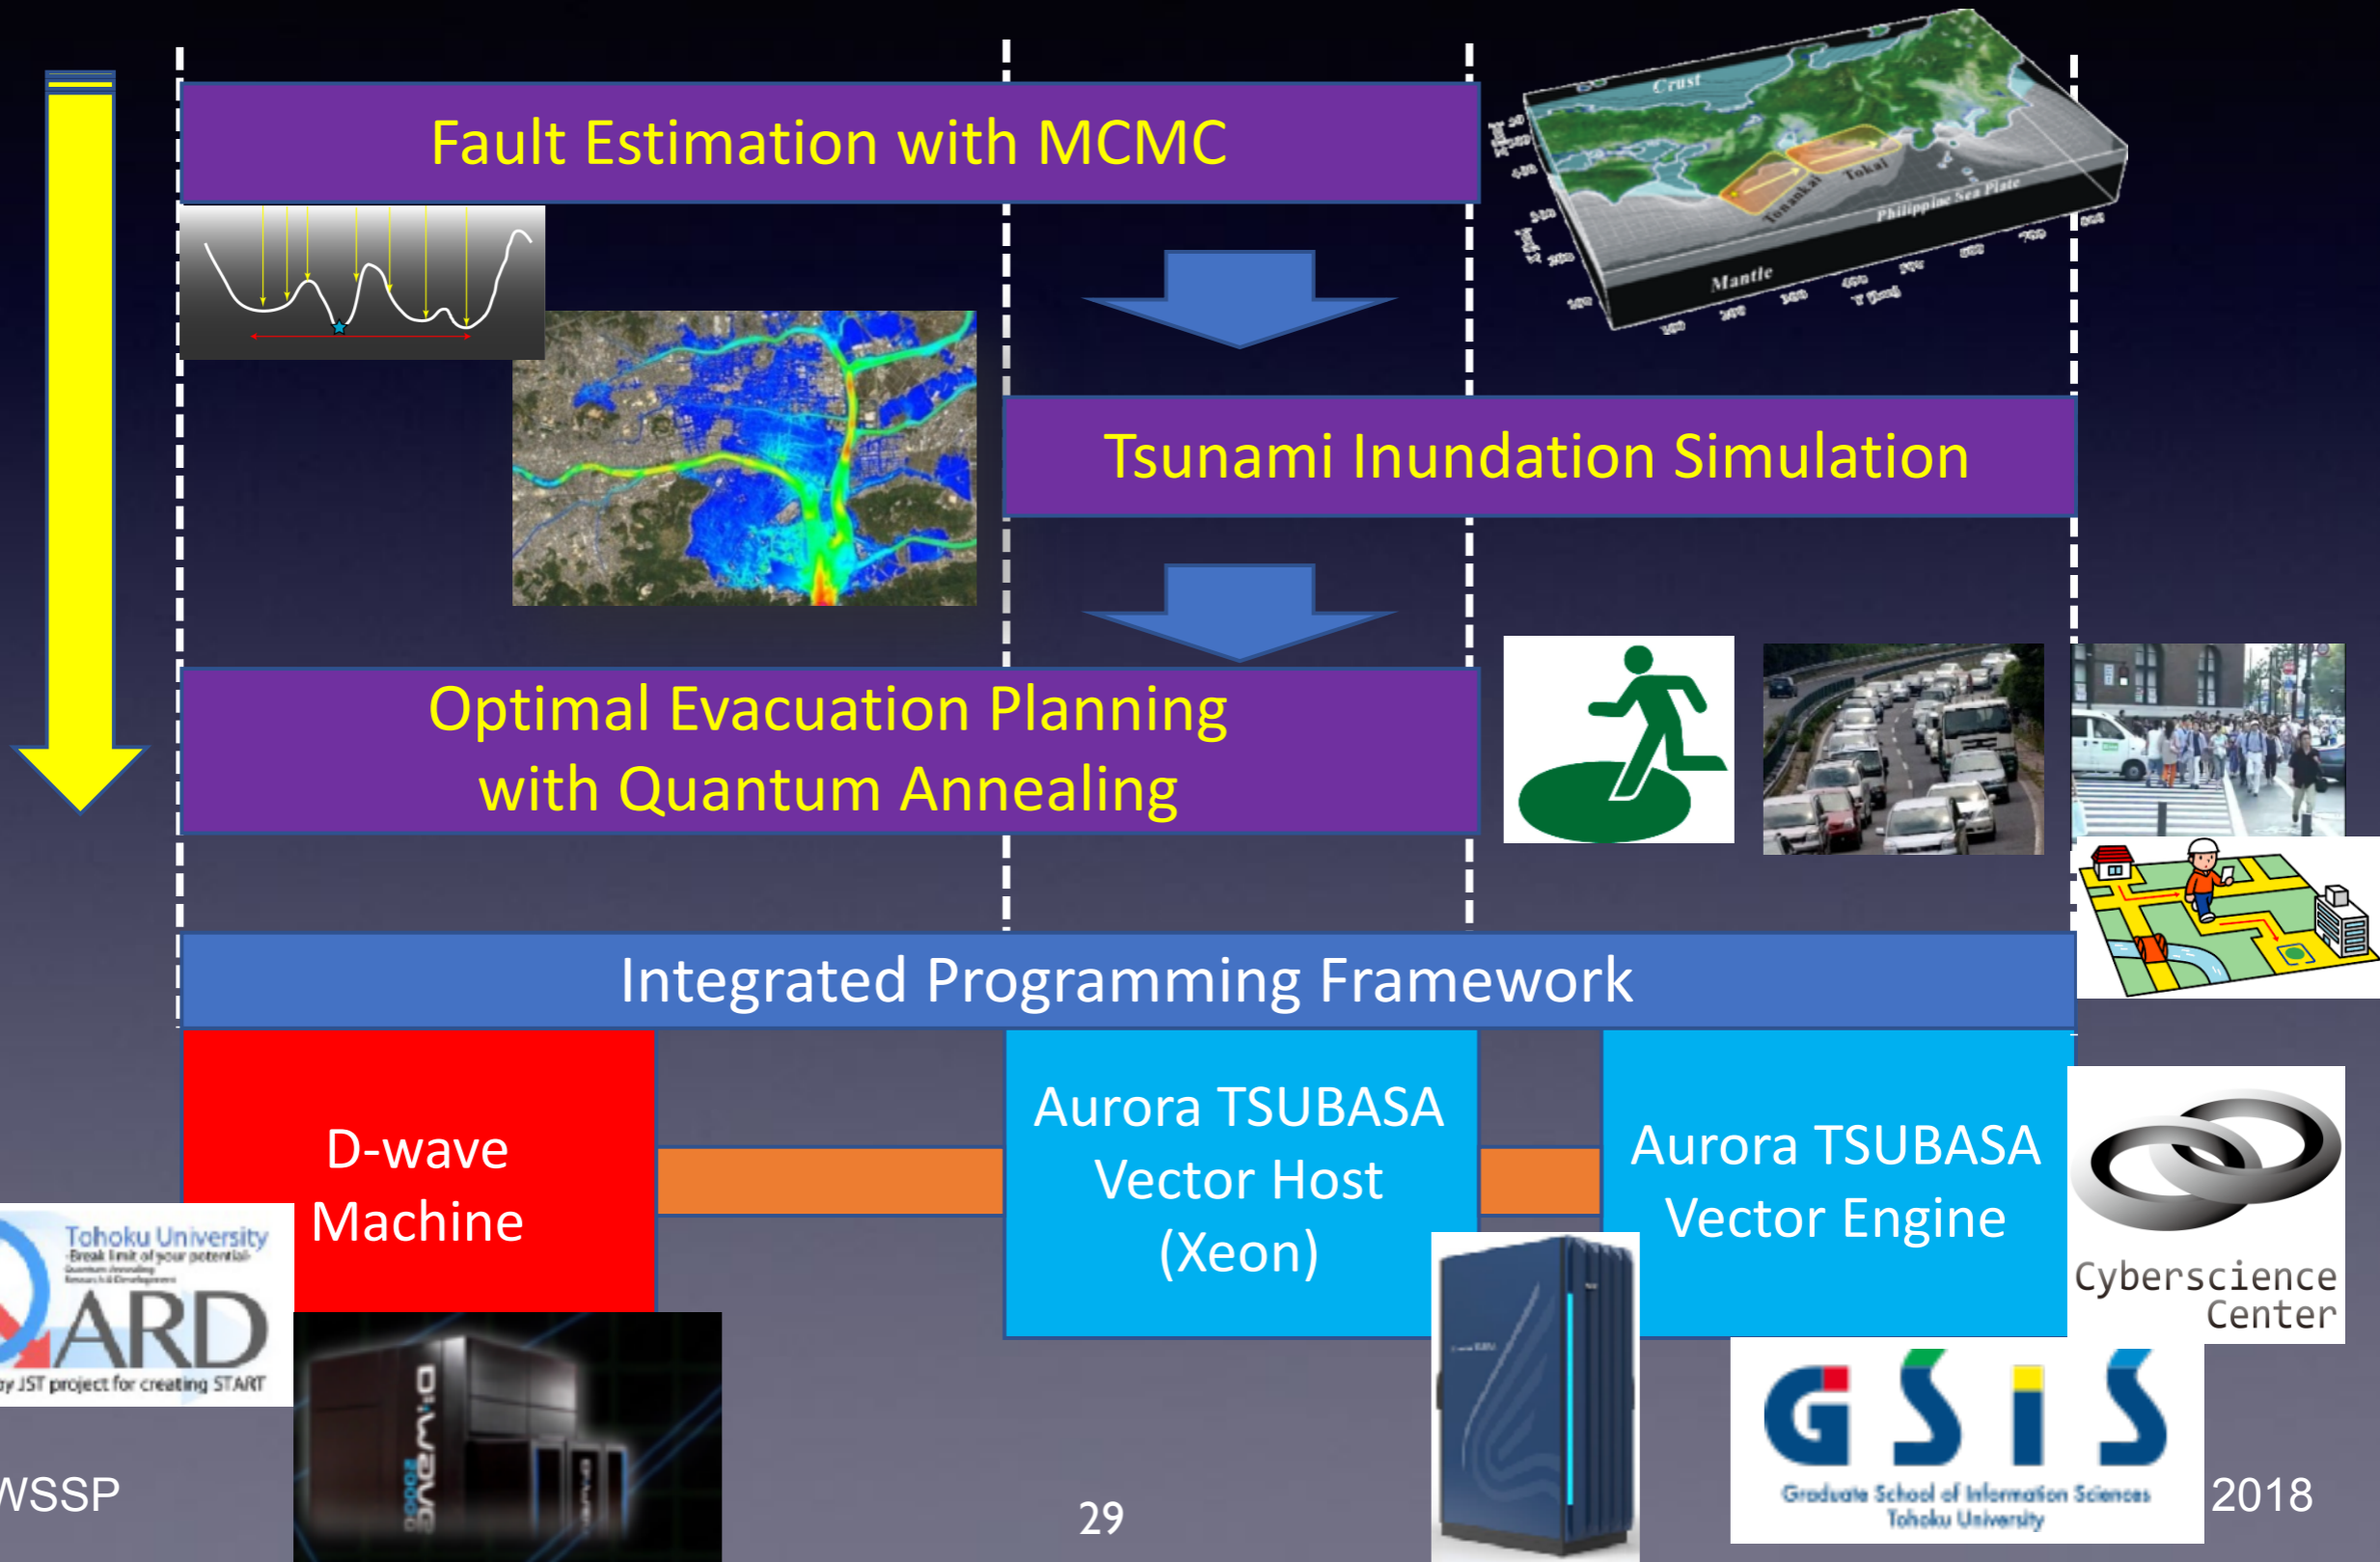


# Team Organization

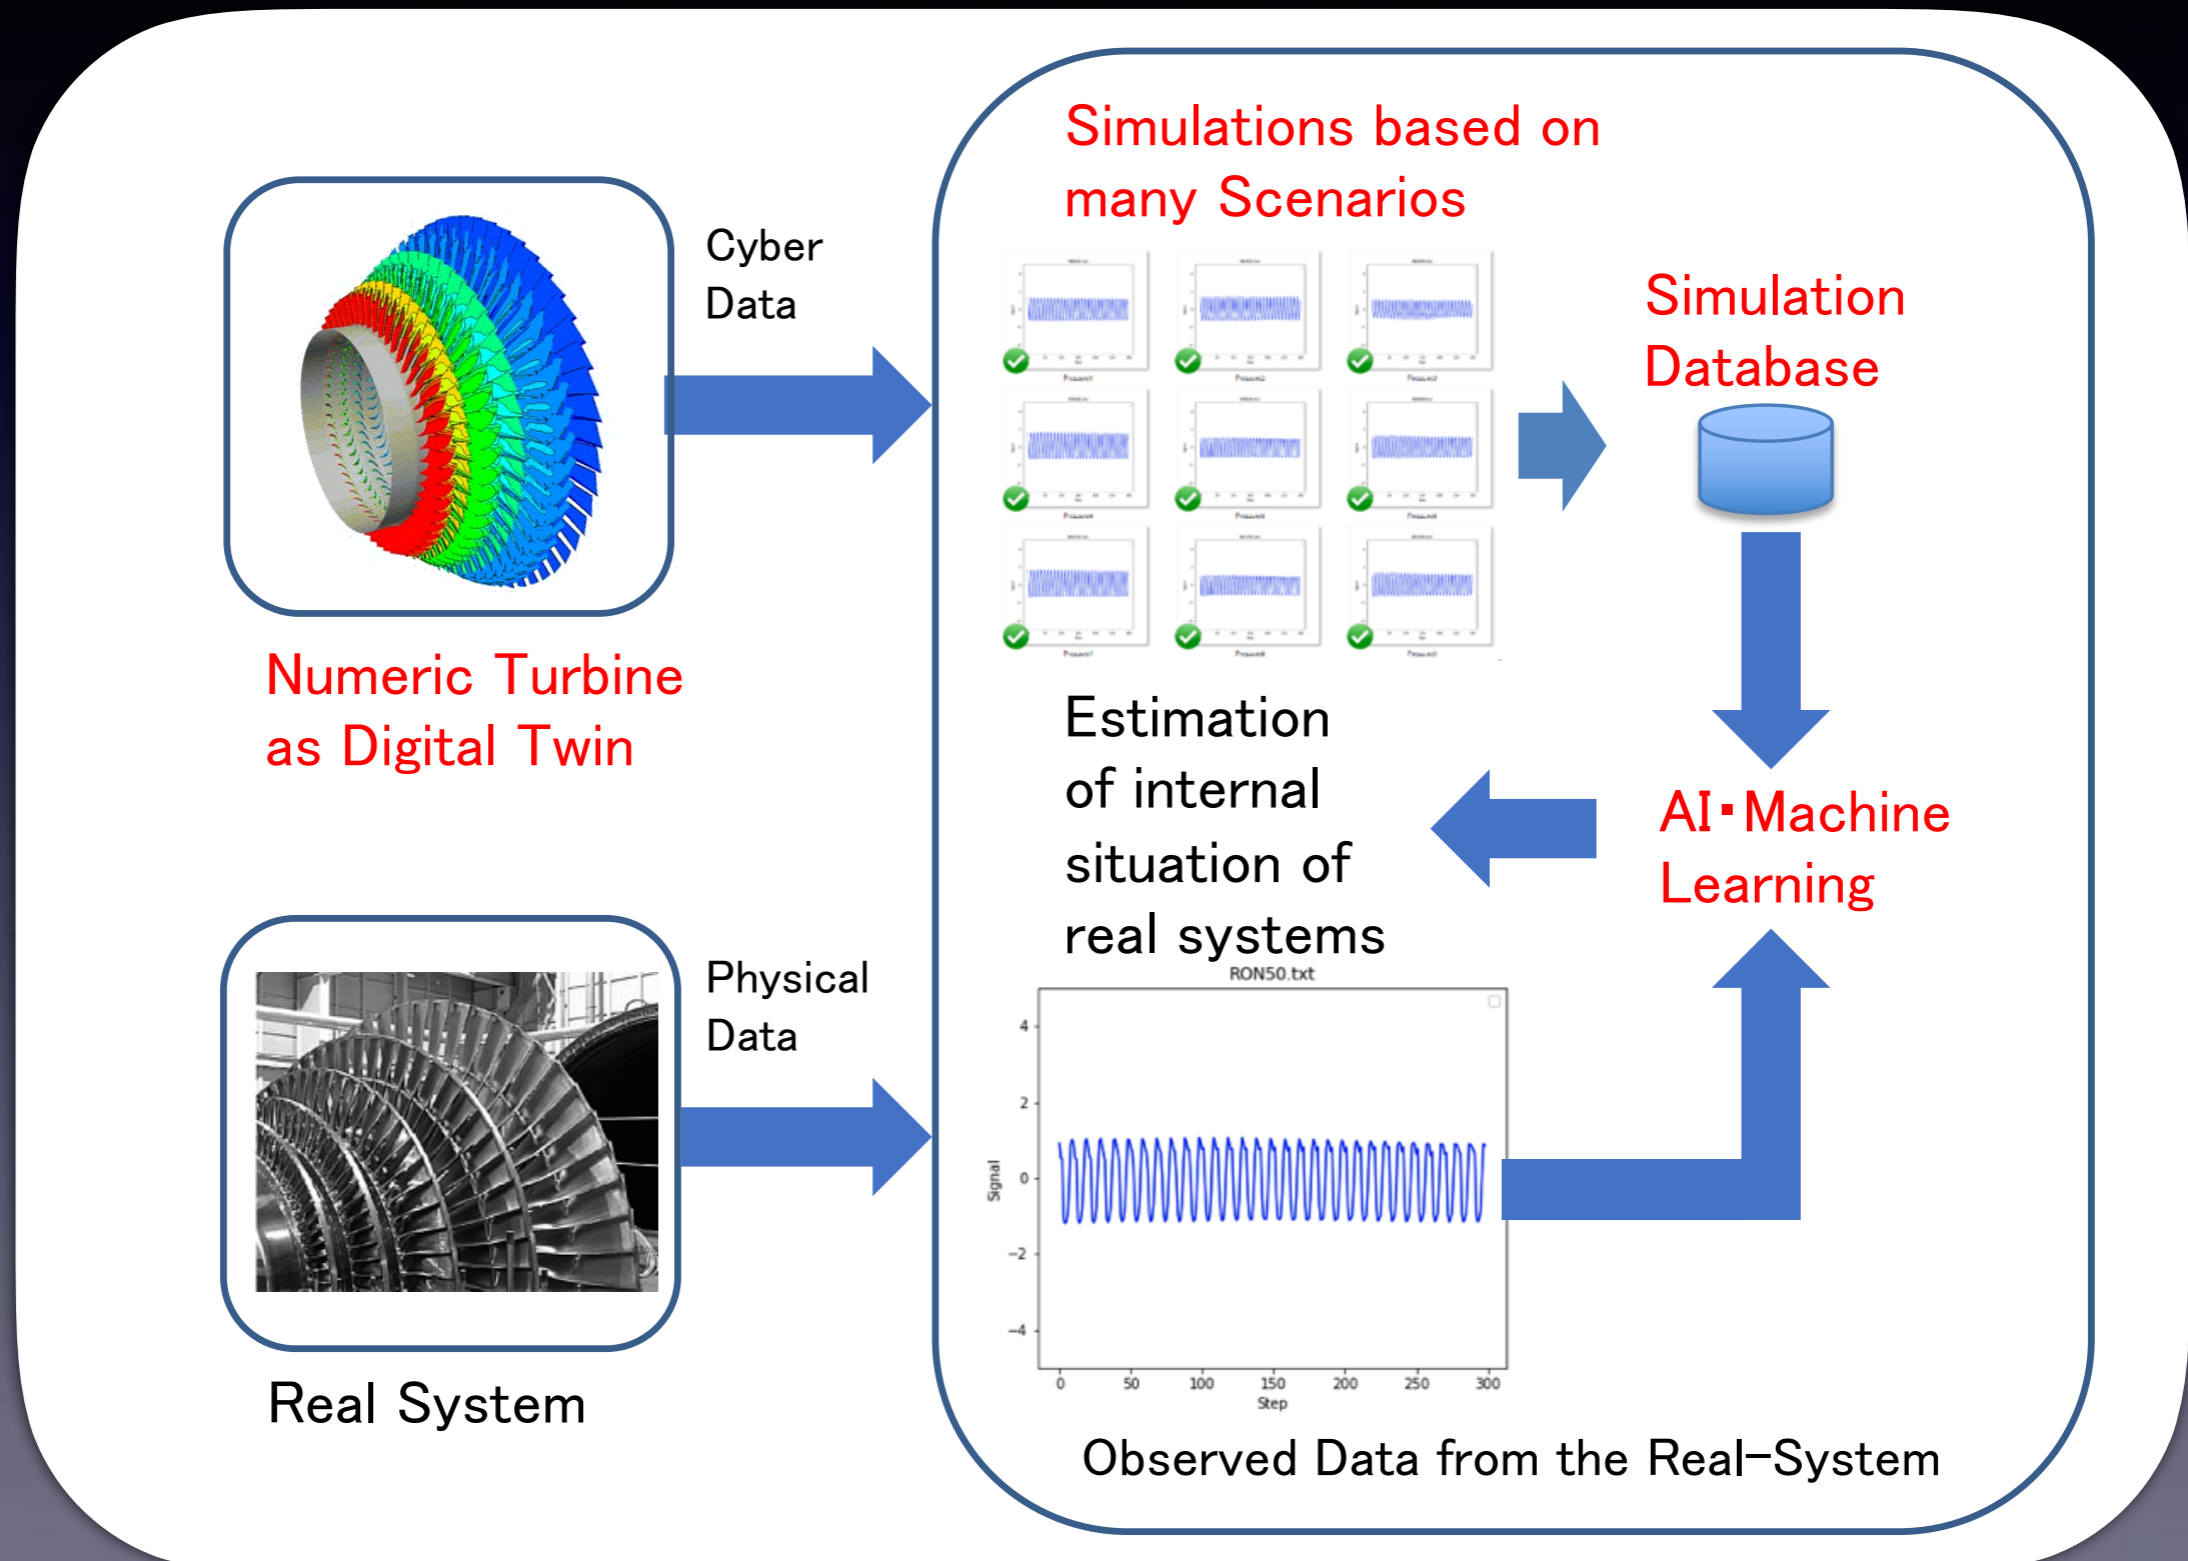


**International Advisory Committee:**  
 Mateo Valero (BSC,ES) Michael Resch (HLRS,DE), Vladimir Voevodin (MSU,RU),  
 Bo Ewald(D-wave, CA), Hans Peter Graf(NEC Lab. America, US)

# An Example of Target Application: QA-Enhanced Real-Time Tsunami Inundation Forecasting and Optimal Evacuation Planning



# An Example of Target Applications: Digital Twin Numerical Turbine



Let's Meet together again at the next WSSP at Tohoku Univ.

★ 29th Workshop on Sustained Simulation Performance

- Date: March 19-20, 2019
- Place: Tohoku University

