Green Computing Involves a lot of Innovations but also a lot of Sweat

Satoshi Matsuoka
Director, Riken Center for Computational Science /
Professor, Tokyo Institute of Technology

SC18 EE Workshop 20181111
Dallas TX USA

Tokyo Tech. TSUBAME Supercomputing History

World's Leading Supercomputer. ***x100,000*** *speedup in 17 years developing world-leading use of* ***massively parallel, many-core technology***

2000
Matsuoka
GSIC
Appointment

2000
128 Gigaflops
Custom
Supercomputer
32 cores

2002 "TSUBAME0"
1.3 TeraFlops
First "TeraScale"
JP Univ. Supercomputer
800 cores

2006 TSUBAME1.0
80 TeraFlops
No.1 Asia, No.7 World
10,000 cores

2008 TSUBAME1.2
170 TeraFlops
Word's first GPU
Supercomputer

2010 TSUBAME2.0
2.4 Petarlops No1 World
No.1 Production Green
ACM Gordon Bell Prize

General Purpose CPU & Many Core Processor (GPU), Advaned Optical Networks, Non-Volatile Memory, Efficient Power Control and Cooling

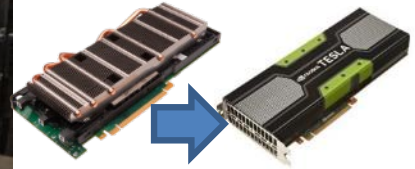

2013 TSUBAME2.5
4118 GPUs Upgraded
5.7 Petaflps, No.2 Japan
AI Flops 17.1 Petaflops

2013 TSUBAME-KFC
TSUBAME3 Prototype
Oil Immersive Cooling
Green World No.1
2015 AI Prototype Upgrade (KFC/DL)

2017 TSUBAME3.0, **> 10 million cores**
12.1 Petaflops (AI Flops 47.2 Petaflops)
Green World No1
HPC and Big Data / AI Convergence

Towards 2.0: Biggest Problem is Power…

Machine	CPU Cores	Watts	Peak GFLOPS	Peak MFLOPS/ Watt	Watts/ CPU Core	Ratio c.f. TSUBAME
TSUBAME(Opteron)	10480	800,000	50,400	63.00	76.34	
TSUBAME2006 (w/360CSs)	11,200	810,000	79,430	98.06	72.32	
TSUBAME2007 (w/648CSs)	11,776	820,000	102,200	124.63	69.63	1.00
Earth Simulator	5120	6,000,000	40,000	6.67	1171.88	0.05
ASCI Purple (LLNL)	12240	6,000,000	77,824	12.97	490.20	0.10
AIST Supercluster (Opteron)	3188	522,240	14400	27.57	163.81	0.22
LLNL BG/L (rack)	2048	25,000	5734.4	229.38	12.21	1.84
Next Gen BG/P (rack)	4096	30,000	16384	546.13	7.32	4.38
TSUBAME 2.0 (2010Q3/4)	160,000	810,000	1,024,000	1264.20	5.06	10.14

MegaProto [SC2005]

- "MegaScale" HPC power efficiency project w/H.Nakashima, H.Nakamura, T.Boku, M.Sato… (2002-2006)
- Custom node w/IBM, custom PCI-e backplane MB
- 17 nodes in 1U, power, HDD, …
- Embedded Dual Rail GigE switch, direct chassis-chassis
- 10W / node, < 300W chassis
- Faster than 1U Xeon servers in many apps
- ***But largely ignored… (though predating HPE Moonshot by 10 years)***

- Effecion TM8800
- 2.4GFlops@5W
- ⇒ 480MFlops/W
- 512MB DDR-SDRAM

- 256MB SDRAM
  - (512M DDR in V2)
- 512KB flash

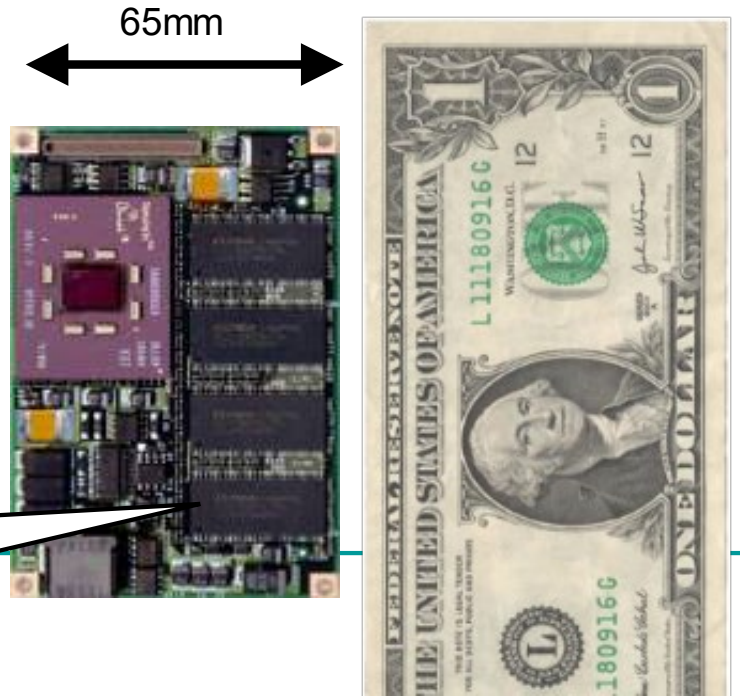

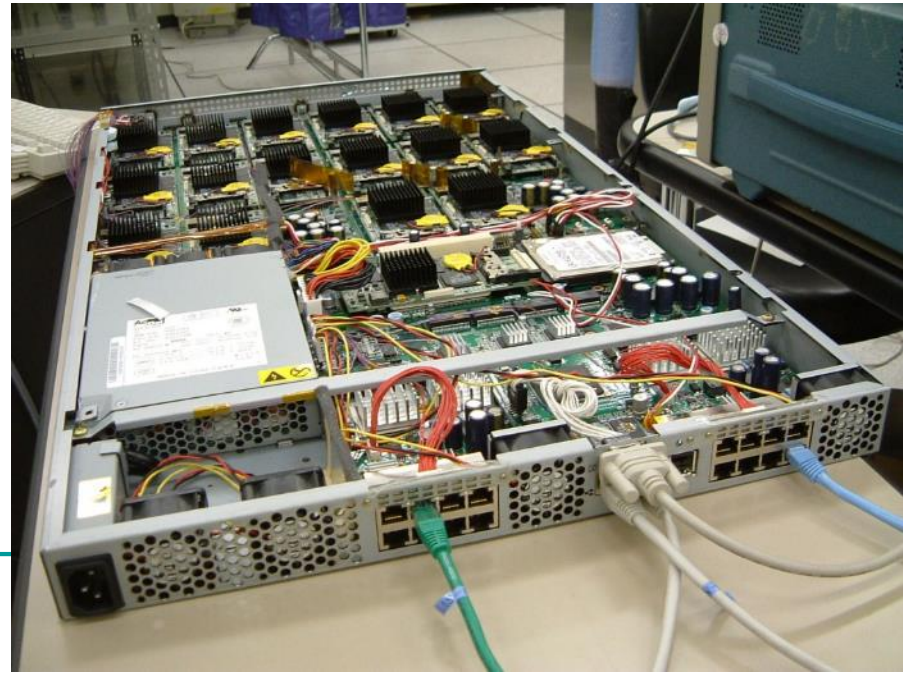

JST-CREST "MegaScale" [2003-2007]

- *"Embedded tech. meets HPC"*
- Increased importance w/emphasis on mobile/ubiquitous computing
- Attain ~x2 with effective **DVS**

=> 2004 CTWatch J. (Matsuoka ed.)

- Can't attain x10 tho…
  - Low hanging fruit achieved with simple DVS
  - Memory, NW, I/O… s tarts to dominate
  - Simple slow cores limited to weak scaling
  - Conclusion: new devices, new architectures, cooling, apps need to be exploited and optimized as a whole

JST-CREST MegaProto (SC2005)
17Sockets Effecion/1U
64GFlops/300W

Aggressive Power Saving in HPC

Methodologies	Enterprise/Business Clouds	HPC
Server Consolidation	Good	NG!
DVFS (Dynamic Voltage/Frequency Scaling)	Good	Poor
New Devices	Poor (Cost & Continuity)	Good
New Architecture	Poor (Cost & Continuity)	Good
Novel Cooling	Limited (Cost & Continuity)	Good (high thermal density)

JST-CREST "Ultra Low Power (ULP)-HPC" Project 2007-2012

ULP-HPC SIMD-Vector (GPGPU, etc.)

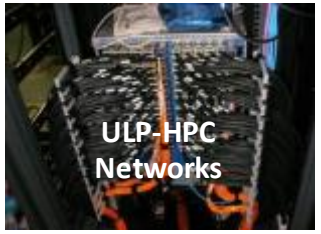

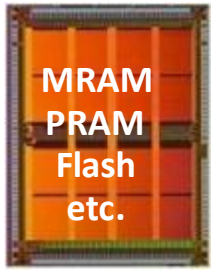

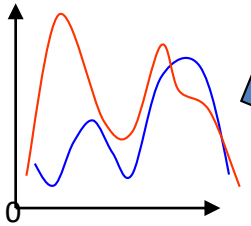

Low Power
High Perf Model

Power Optimize using Novel Components in HPC

Auto-Tuning for Perf. & Power

モデルと実測の Bayes 的融合

- Bayes モデルと事前分布

$$y_i \sim N(\mu_i, \sigma_i^2)$$

$$\mu_i \mid \beta, \sigma_i^2 \sim N(x_i^T \beta, \sigma_i^2 / \kappa_0)$$

$$\sigma_i^2 \sim \mathrm{Inv}\text{-}\chi^2(\nu_0, \sigma_0^2)$$

モデルによる所要時間の推定

- n 回実測後の事後予測分布

$$y_i \mid (y_{i1}, y_{i2}, \cdots, y_{in}) \sim t_{\nu_n}(\mu_{in}, \sigma_{in}^2 \kappa_{n+1} / \kappa_n)$$

$$\nu_n = \nu_0 + n, \quad \kappa_n = \kappa_0 + n, \quad \mu_n = (\kappa_0 x_i^T \beta + n\bar{y}_i) / \kappa_n$$

$$\nu_n \sigma_n^2 = \nu_0 \sigma_0^2 + \sum (y_m - \bar{y}_i)^2 + \kappa_0 n (\bar{y}_i - x_i^T \beta)^2 / \kappa_n$$

所要時間の実測データ

ABCLibScript: アルゴリズム選択

```
!ABCLib$  static  select region start
!ABCLib$  parameter (in CacheS, in NB, in NPrc)
!ABCLib$      select sub region start
!ABCLib$      according estimated
!ABCLib$        (2.0d0*CacheS*NB)/(3.0d0*NPrc)
          対象1（アルゴリズム1）
!ABCLib$      select sub region end
!ABCLib$      select sub region start
!ABCLib$      according estimated
!ABCLib$        (4.0d0*ChcheS*dlog(NB))/(2.0d0*NPrc)
          対象2（アルゴリズム2）
!ABCLib$      select sub region end
!ABCLib$  static select  region  end
```

実行起動前自動チューニング指定、アルゴリズム選択処理の指定

コスト定義関数で使われる入力変数

コスト定義関数

対象領域1、2

Optimization Point

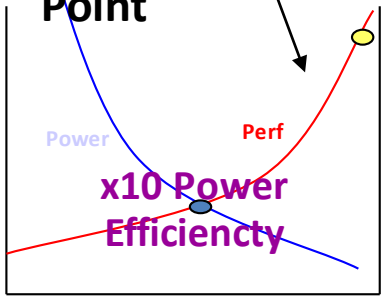

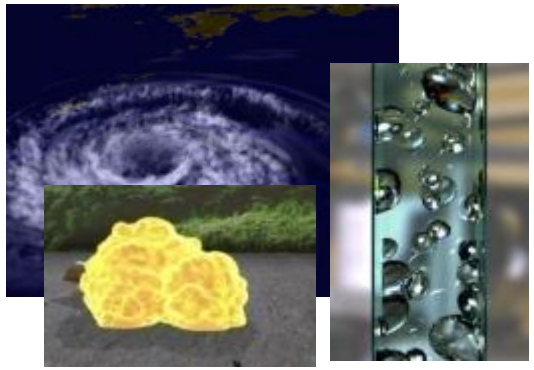

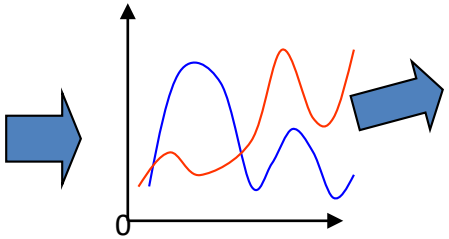

Perf. Model
Algorithms

Power-Aware and Optimizable Applications

x1000
Improvement in 10 years

Slide from Circa 2004 My Prediction for a Petaflops Machine (as TSUBAME 1 was being designed)

"Future AV Vector-Parallel"

PlayStation3
(Nov. 2006)

IBM/Toshiba/SONY Cell (Chip Vector) + SMT/CMT 256GF-1TFlops (@0.065-0.03)

OR

"Future PC Vector-Graphics"

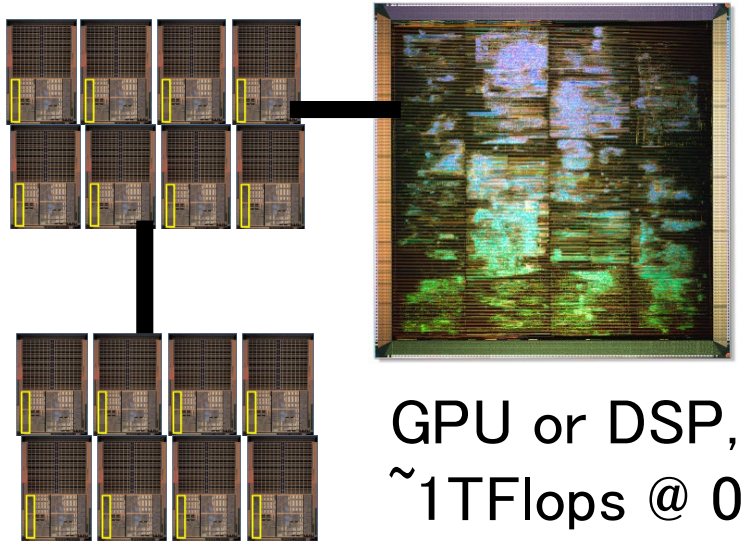

GPU or DSP, ~1TFlops @ 0.045

Multicore SMT/CMT, ~50Gflops@0.065 μ

NVIDIA GTX8800 CUDA GPU (Nov. 2006)

100CPUs/Rack => A 100 Teraflops per Rack

Power efficient GPUs for HPC? (What? Gaming stuff for HPC? – at DoE meeting)

2007 GPU-Heavy Proto Cluster ($$$ by Microsoft)
128 GTX9800 GPUs, 32 nodes

Fantastic perf/W Proteomics docking app [Y.Akiyama] /w GPU-optimized FFT [SC08]
~= 4-rack BG/L, x4 perf/W

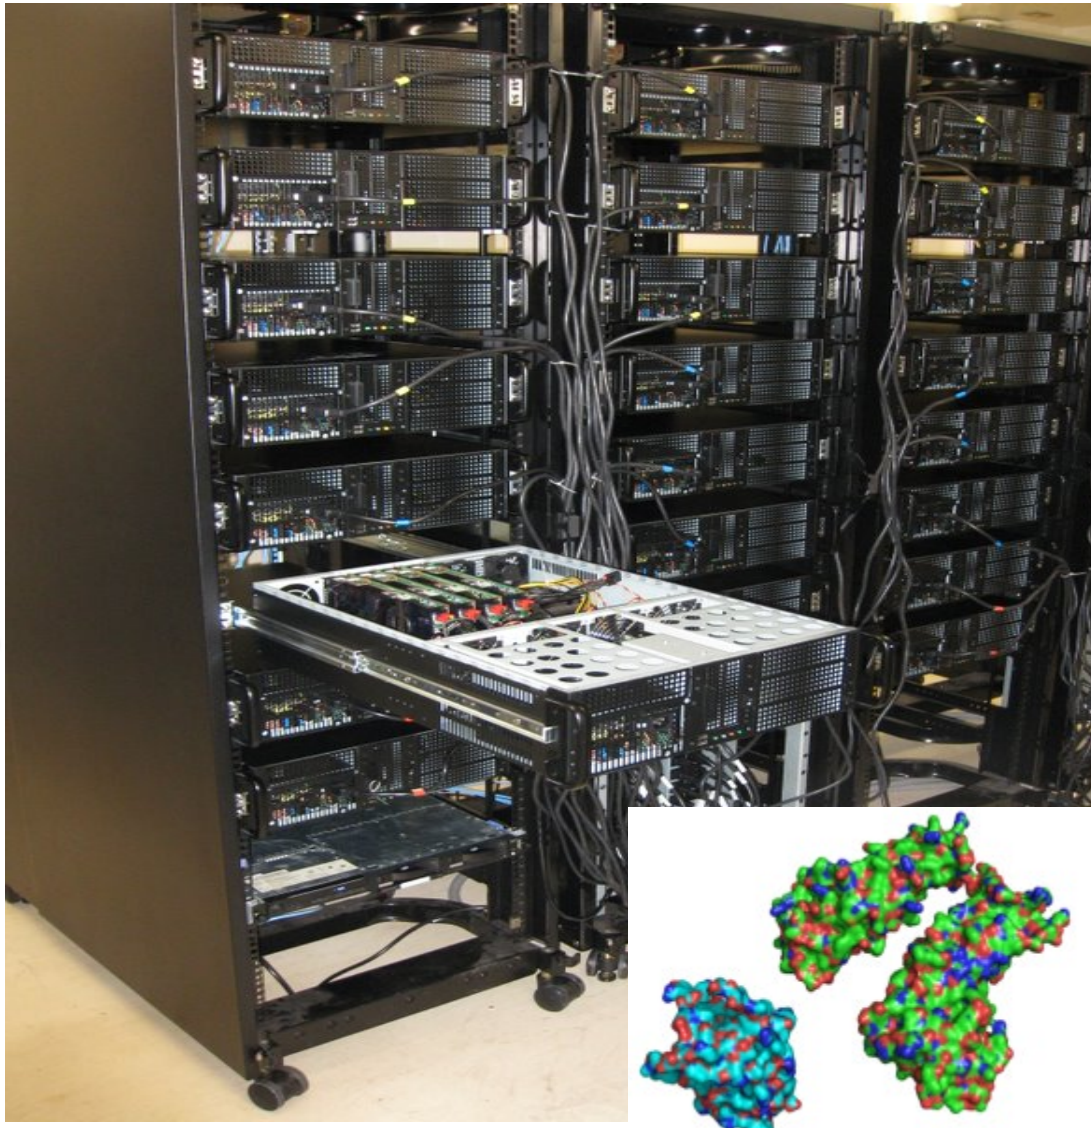

TSUBAME 1.2 Experimental Evolution (Oct. 2008)

"Retrofitting a production machine (live!)"

Voltaire ISR9288 Infiniband x8
10Gbps x2 ~1310+50 Ports
~13.5Terabits/s
(3Tbits bisection)
10Gbps+External NW

NEC SX-8i

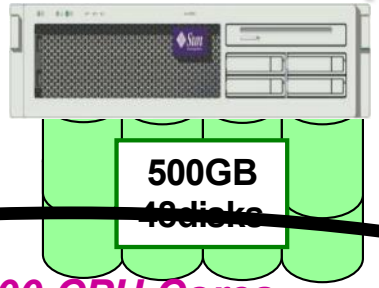

Unified Infiniband network

10,000 CPU Cores
300,000 SIMD Cores
> 20 Million Threads
~900TFlops-SFP, ~170TFlops-DFP
80TB/s Mem BW (1/2 ES)

Storage
1.5 Petabyte (Sun x4500 x 60)
0.1Petabyte (NEC iStore)
Lustre FS, NFS, CIF, WebDAV (over I
60GB/s aggregate I/O BW

Sun x4600 (16 Opteron Cores)
32~128 GBytes/Node
10480core/655Nodes
21.4TeraBytes
50.4TeraFlops
OS Linux (SuSE 9, 10)
NAREGI Grid MW

PCI-e

ClearSpeed CSX600
SIMD accelerator
648 boards,
52.2TeraFlops
SFP/DFP

170 Nvidia Tesla 1070, ~680 Tesla cards
High Performance in Many BW-Intensive Apps
ASUCA Weather app [Shimokawabe&Aoki, SC10]

TSUBAME2.0 Nov. 1, 2010, world 4th fastest Nov. 2010
"The Greenest Production Supercomputer in the World"

- GPU-centric (>4000) high performance & low power
- Small footprint (~200m2 or 2000 sq.ft), low TCO
- High bandwidth memory, optical network, SSD storage…

TSUBAME 2.0
New Development

Chip
(CPU ,GPU)

CPU(Westmere EP)
76.8 GFLOPS
32nm

GPUs(Tesla M2050)
515 GFLOPS
3 GB
40nm

Compute Node
(2 CPUs,3 GPUs)

1.6 TFLOPS
55 GB/103 GB
>400GB/s Mem BW
80Gbps NW BW
~1KW max

Node Chassis
(4 Compute Nodes)

6.7 TFLOPS
220 GB/412 GB
>1.6TB/s Mem BW

Rack
(8 Node Chassis)

53.6 TFLOPS
1.7 TB/3.2 TB
>12TB/s Mem BW
35KW Max

System
(42 Racks)
1408 GPU Compute Nodes,
34 Nehalem "Fat Memory" Nodes

2.4 PFLOPS
80 TB
4224 GPUs
>600TB/s Mem BW
220Tbps NW
Bisecion BW
1.4MW Max

Integrated by NEC Corporation

TSUBAME2.0=>2.5 Power Improvement (GPU Update Fermi->Kepler)

The Green500 List

Listed below are the November 2013 The Green500's energy-efficient supercomputers ranked from 1 to 10.

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	4,503.17	GSIC Center, Tokyo Institute of Technology	TSUBAME-KFC - LX 1U-4GPU/104Re-1G Cluster, Intel Xeon E5-2620v2 6C 2.100GHz, Infiniband FDR, NVIDIA K20x	27.78
2	3,631.86	Cambridge University	Wilkes - Dell T620 Cluster, Intel Xeon E5-2630v2 6C 2.600GHz, Infiniband FDR, NVIDIA K20	52.62
3	3,517.84	Center for Computational Sciences, University of Tsukuba	HA-PACS TCA - Cray 3623G4-SM Cluster, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband QDR, NVIDIA K20x	78.77
4	3,185.91	Swiss National Supercomputing Centre (CSCS)	Piz Daint - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Level 3 measurement data available	1,753.66
5	3,130.95	ROMEO HPC Center - Champagne-Ardenne	romeo - Bull R421-E3 Cluster, Intel Xeon E5-2650v2 8C 2.600GHz, Infiniband FDR, NVIDIA K20x	81.41
6	3,068.71	GSIC Center, Tokyo Institute of Technology	TSUBAME 2.5 - Cluster Platform SL390s G7, Xeon X5670 6C 2.930GHz, Infiniband QDR, NVIDIA K20x	922.54
7	2,702.16	University of Arizona	iDataPlex DX360M4, Intel Xeon E5-2650v2 8C 2.600GHz, Infiniband FDR14, NVIDIA K20x	53.62
8	2,629.10	Max-Planck-Gesellschaft MPI/IPP	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband, NVIDIA K20x	269.94
9	2,629.10	Financial Institution	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband, NVIDIA K20x	55.62
10	2,358.69	CSIRO	CSIRO GPU Cluster - Nitro G16 3GPU, Xeon E5-2650 8C 2.000GHz, Infiniband FDR, Nvidia K20m	71

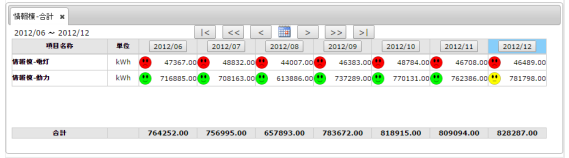

18% Power Reduction inc. cooling

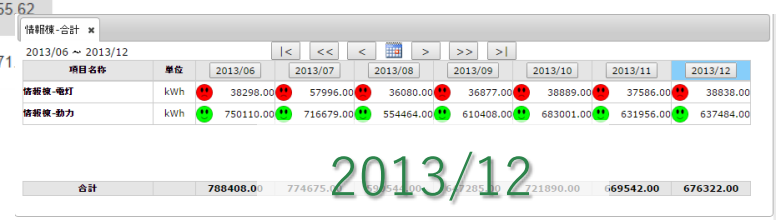

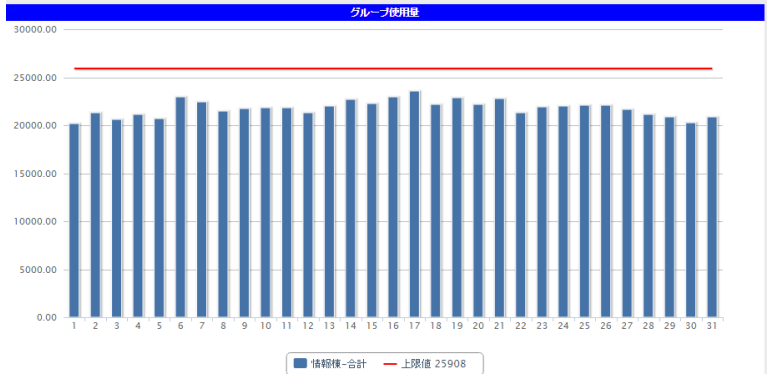

2013/11 Green 500 #6 in the world

- Along with TSUBAME-KFC (#1)
- 2014/6 #9

TSUBAME2.0=>2.5
System-Wide Power Capping via Scheduling

- 2.0: After the Fukushima accident, turn off 1/3 nodes during summer daytime
- 2.5: Dynamic power capping via power monitoring and resource scheduling

TSUBAME2.5 Compute Resources

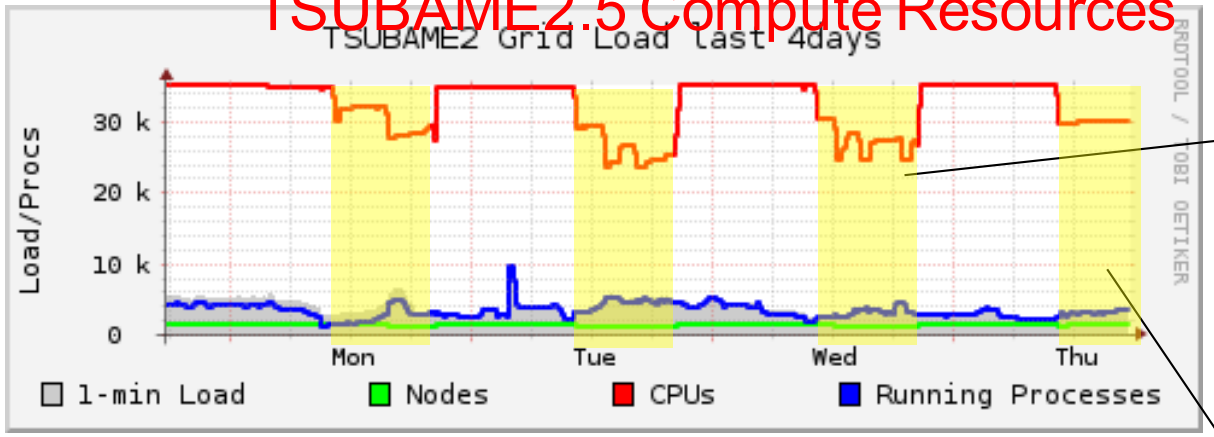

Dynamically control # active nodes according to power cap

Yellow=Summer Daytime

TSUBAME2.5 Power Profile

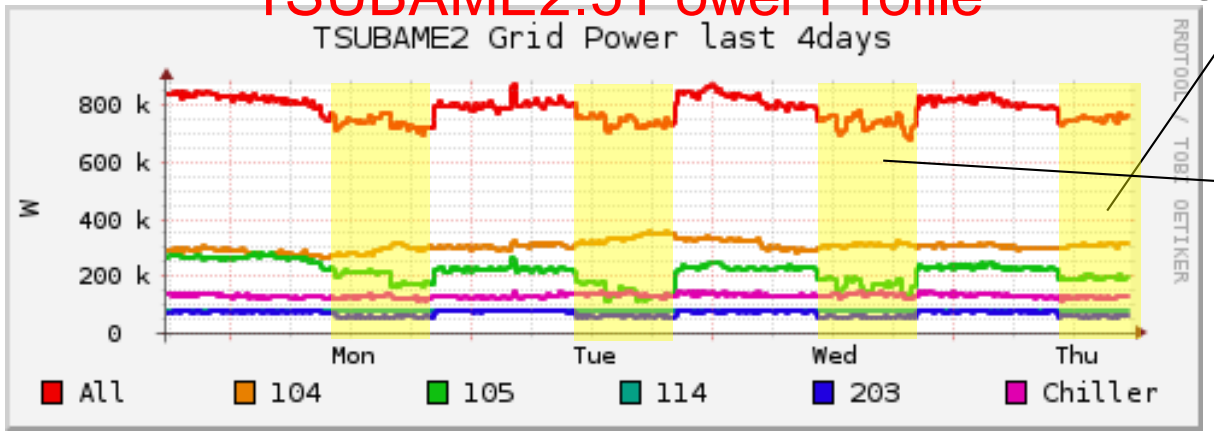

Power capped to 800kW during daytime via dynamic power-capped scheduling we developed

TSUBAME-KFC/DL: TSUBAME3 Prototype [ICPADS2014]

Oil Immersive Cooling＋ Hot Water Cooling + High Density Packaging + Fine-Grained Power Monitoring and Control, *upgrade to /DL Oct. 2015*

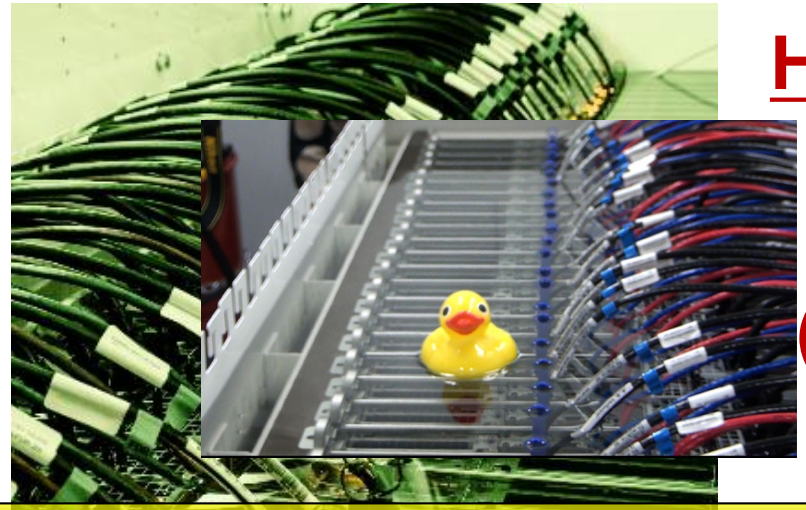

High Temperature Cooling

Oil Loop 35~45℃

⇒ Water Loop 25~35℃

(c.f. TSUBAME2: 7~17℃)

Cooling Tower:

Water 25~35℃

⇒ To Ambient Air

Single Rack High Density Oil Immersion

168 NVIDIA K80 GPUs + Xeon

413+TFlops (DFP)

1.5PFlops (SFP)

~60KW/rack

Container Facility

20 feet container (16m²)

Fully Unmanned Operation

Modern AI is enabled by Supercomputing

- 25 years of AI winter after failure of symbolic logic based methods (e.g., Prolog, ICOT) -> resurrection by DNN, basic algorithms in the 1980s but too expensive -> HPC made machines 10 million times faster in 30 years -> expensive training now possible
- Recent trends require more supercomputing power
  - Deeper, more complex networks (Capsule Networks)
  - Complex, multidimensional data (e.g., 3-D Hi-Res images)
  - Increasing training sets (incl. GANs)
  - Coupling with high-fidelity simulations
  - Etc.

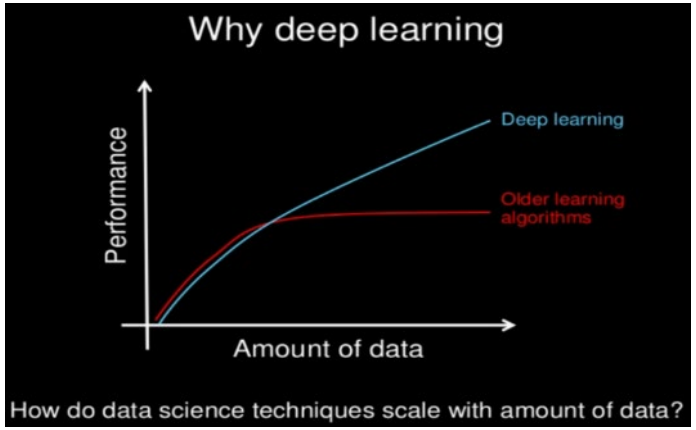

Fig. 2: Andrew Ng (Baidu) "What Data Scientists Should Know about Deep Learning"

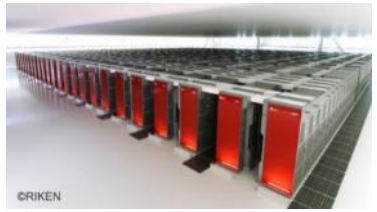

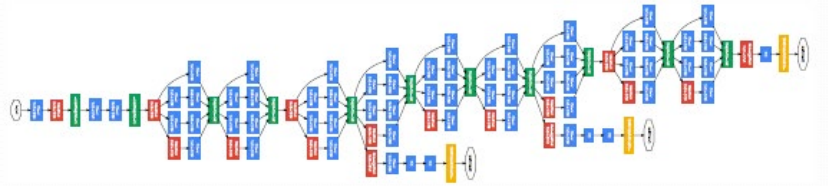

What is meant by Convergence of HPC & AI?

- **Acceleration of Simulation (first principles methods) with AI (empirical method)**
  - Interpolation & Extrapolation of long trajectory MD
  - Reducing parameter space on Paretho optimization of results
  - Adjusting convergence parameters for iterative methods etc.
- **AI *replacing* simulation**
  - When exact physical models are unclear, or excessively costly to compute
- **Acceleration of AI with HPC**
  - HPC Processing of training data -data cleansing
  - Acceleration of (Parallel) Training: Deeper networks, bigger training sets, complicated networks, high dimensional data…
  - Acceleration of Inference: above + real time streaming data
  - Various modern training algorithms: Reinforcement learning, GAN, Dilated Convolution, etc.

4 Layers of Parallelism in DNN Training

- Hyper Parameter Search
  - Searching optimal network configs & parameters
  - Parallel search, massive parallelism required
- Data Parallelism
  - Copy the network to compute nodes, feed different batch data, average => network reduction bound
  - TOFU: Extremely strong reduction, x6 EDR Infiniband

Inter-Node

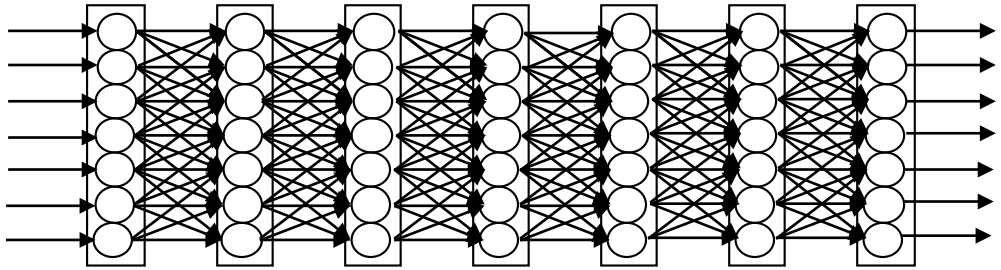

- Model Parallelism (domain decomposition)
  - Split and parallelize the layer calculations in propagation
  - Low latency required (bad for GPU) -> strong latency tolerant cores + low latency TOFU network

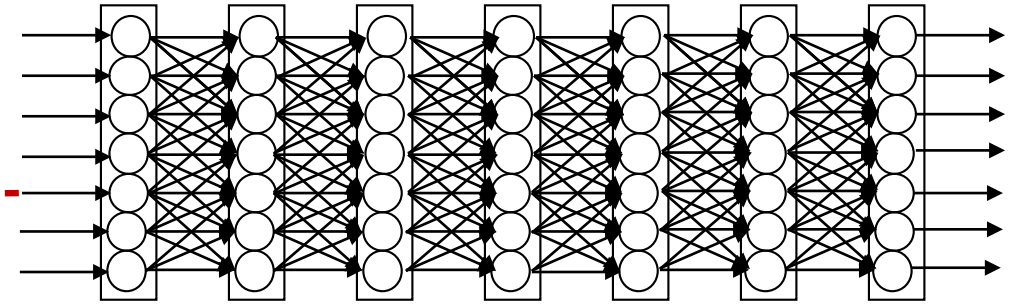

- Intra-Chip ILP, Vector and other low level Parallelism
  - Parallelize the convolution operations etc.
  - SVE FP16+INT8 vectorization support + extremely high memory bandwidth w/HBM2

Intra-Node

- Post-K could become world's biggest & fastest platform for DNN training!

Massive amount of total parallelism, only possible via supercomputing

Deep Learning is "All about Scale" Massive Parallelization is the key

- **Data-parallel training with (Asynchronous) Stochastic Gradient Descent**
  - Replicate network to all the nodes, feed different data, average the gradients periodically
  - Network All-Reduce Reduction in Megabytes~Gigabytes becomes the bottleneck at scale
  - NVIDIA: NVLink Hardware + NICL library (up to 8 GPUs on DGX-1, 16 on DGX-2 w/ NVL Switch)

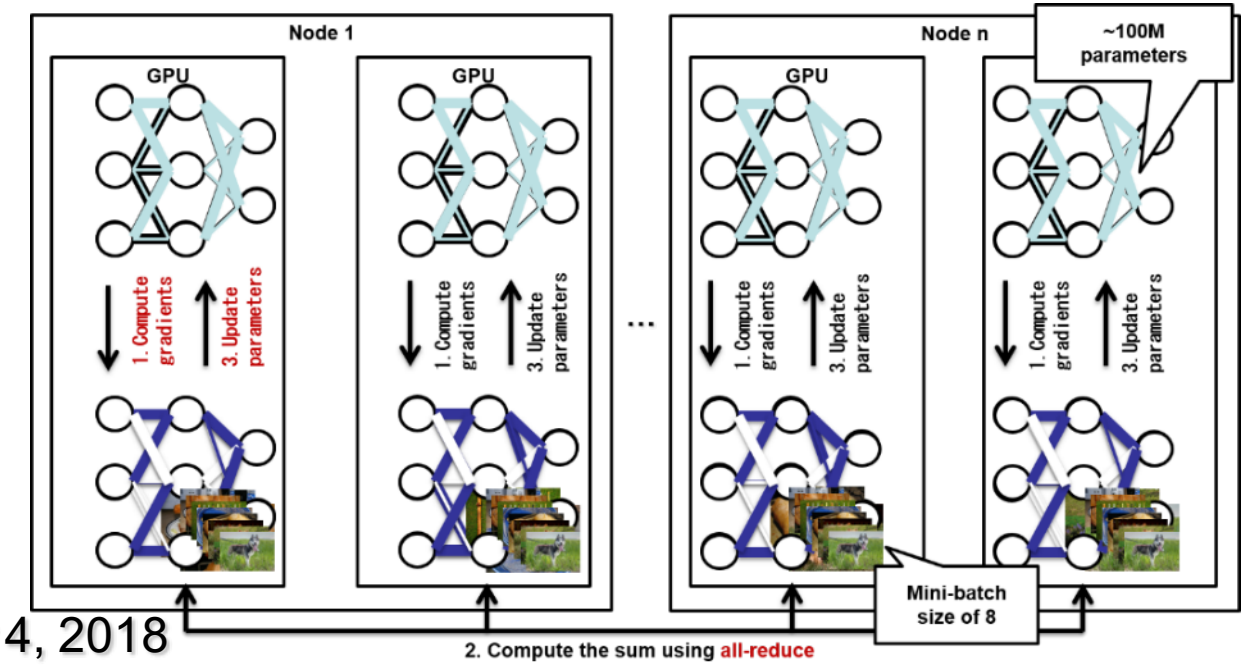

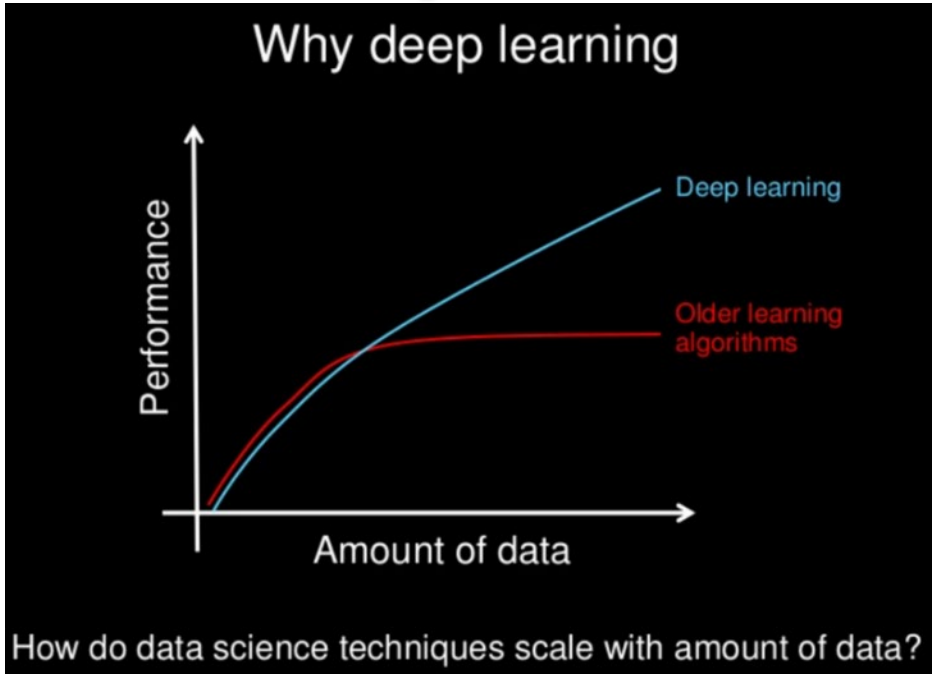

Fig. 2: Andrew Ng (Baidu) "What Data Scientists Should Know about Deep Learning"

Fig. 3: Simplified DL workflow with ASGD per iteration:
1. Compute gradient
2. Exchange gradients via all-reduce; and
3. Update network parameters

Predicting Statistics of Asynchronous SGD Parameters for a Large-Scale Distributed Deep Learning System on GPU Supercomputers

Background

- In large-scale Asynchronous Stochastic Gradient Descent (ASGD), mini-batch size and gradient staleness tend to be large and unpredictable, which increase the error of trained DNN

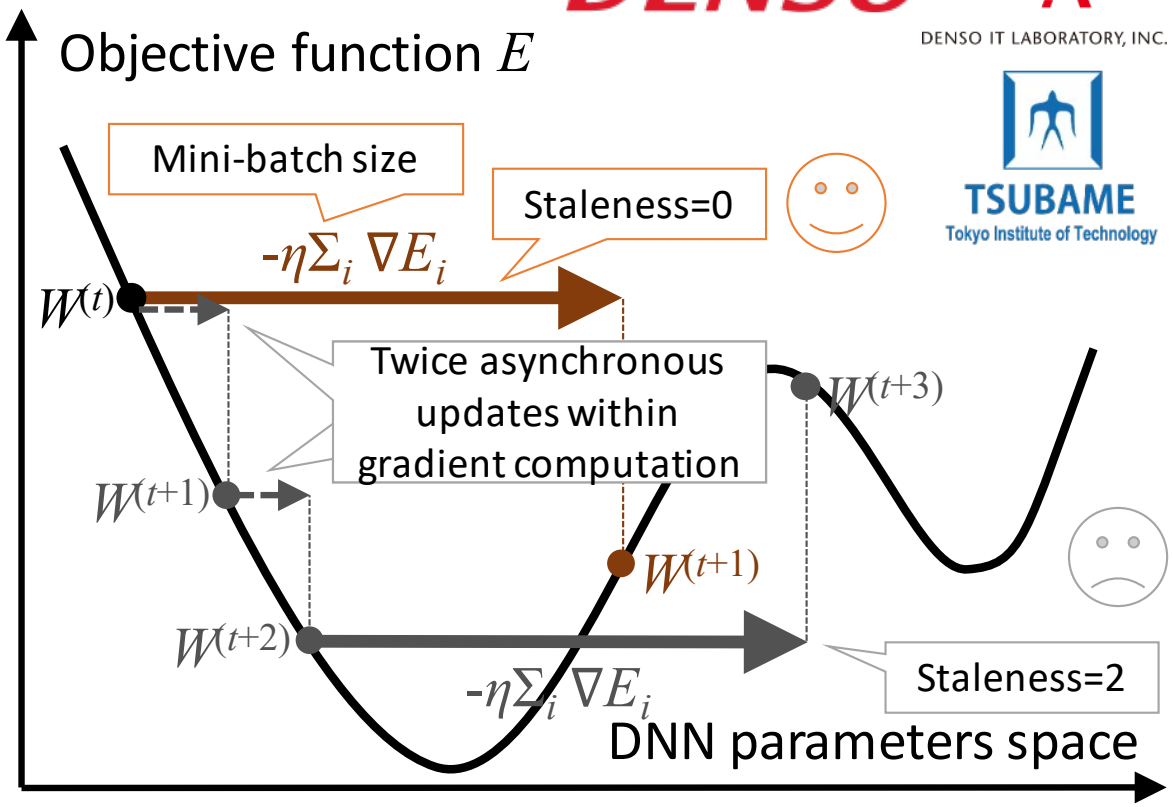

Proposal

- We propose a empirical performance model for an ASGD deep learning system SPRINT which considers probability distribution of mini-batch size and staleness

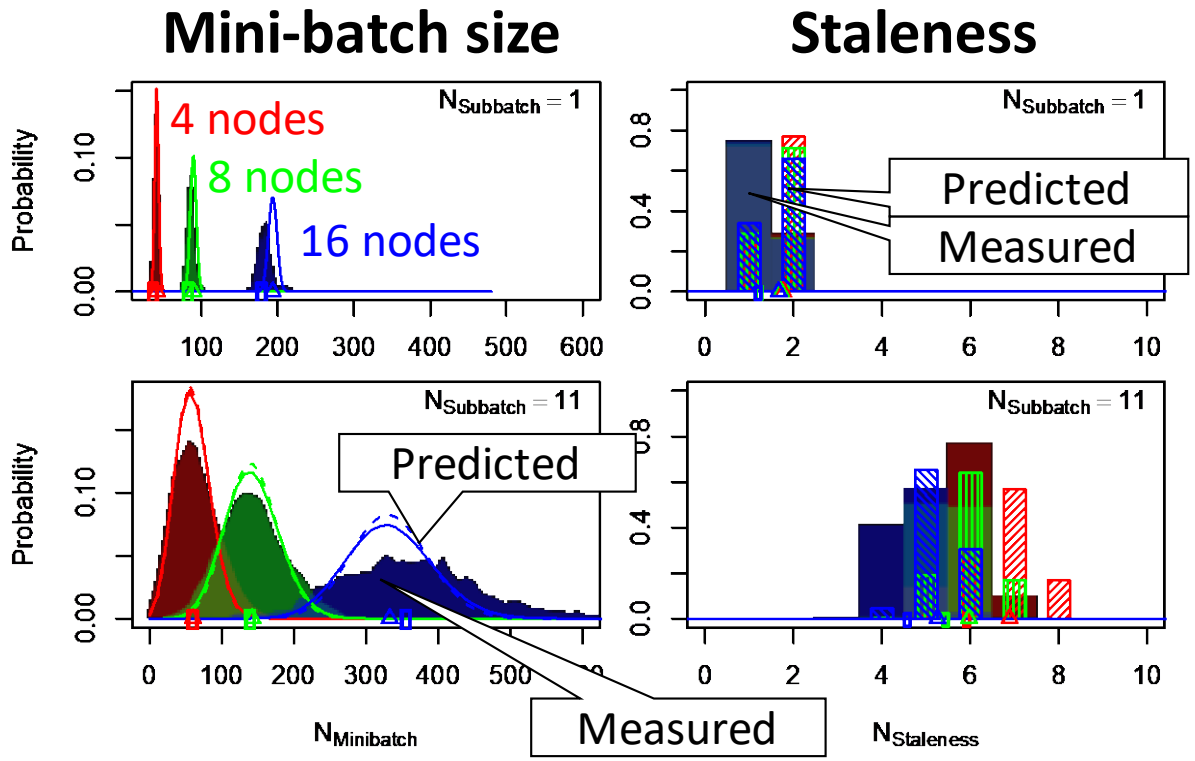

($N_{Subbatch}$: # of samples per one GPU iteration)

- Yosuke Oyama, Akihiro Nomura, Ikuro Sato, Hiroki Nishimura, Yukimasa Tamatsu, and Satoshi Matsuoka, "**Predicting Statistics of Asynchronous SGD Parameters for a Large-Scale Distributed Deep Learning System on GPU Supercomputers**", in proceedings of 2016 IEEE International Conference on Big Data (IEEE BigData 2016), Washington D.C., Dec. 5-8, 2016

Interconnect Performance as important as GPU Performance to accelerate DL

- **ASGD DL system SPRINT (by DENSO IT Lab) and DL speedup prediction with performance model**

$$T_{Epoch} = \frac{N_{File} \times T_{GPU}}{N_{Node} \times N_{GPU} \times N_{Subbatch}}$$

  - Data measured on T2 and KFC (both FDR) fitted to formulas
  - Allreduce time ($\in T_{GPU}$) dep. on #nodes and #DL_parameters

$$T_{Barrier} + (\alpha \log_2(N_{Node}) + \beta) \times N_{Param}$$

The Optimal Predicted Configurations of CNN-A on TSUBAME-KFC/DL

	N_{Node}	$N_{Subbatch}$	Average mini-batch size	Epoch time[s]	Speedup
Baseline K80+ FDR	8	8	165.1	1779	-
FP16	7	22	170.1	1462	**1.22**
EDR IB	12	11	166.6	1245	**1.43**
FP16 + EDR IB	8	15	171.5	1128	**1.58**

Fig. 4: Oyama et al. "Predicting Statistics of Asynchronous SGD Parameters for a Large-Scale Distributed Deep Learning System on GPU Supercomputers

- **Other approaches == similar improvements:**
  - Cuda-Aware CNTK optimizes communication pipeline ➔ 15%—23% speedup
    (Banerjee et al. "Re-designing CNTK Deep Learning Framework on Modern GPU Enabled Clusters")
  - Reduced precision (FP[16|8|1]) to minimize msg. size w/ no or minor accuracy loss

Overview of TSUBAME3.0
BYTES-centric Architecture, Scalaibility to all 2160 GPUs, all nodes, the entire memory hiearchy

Full Operations
Aug. 2017

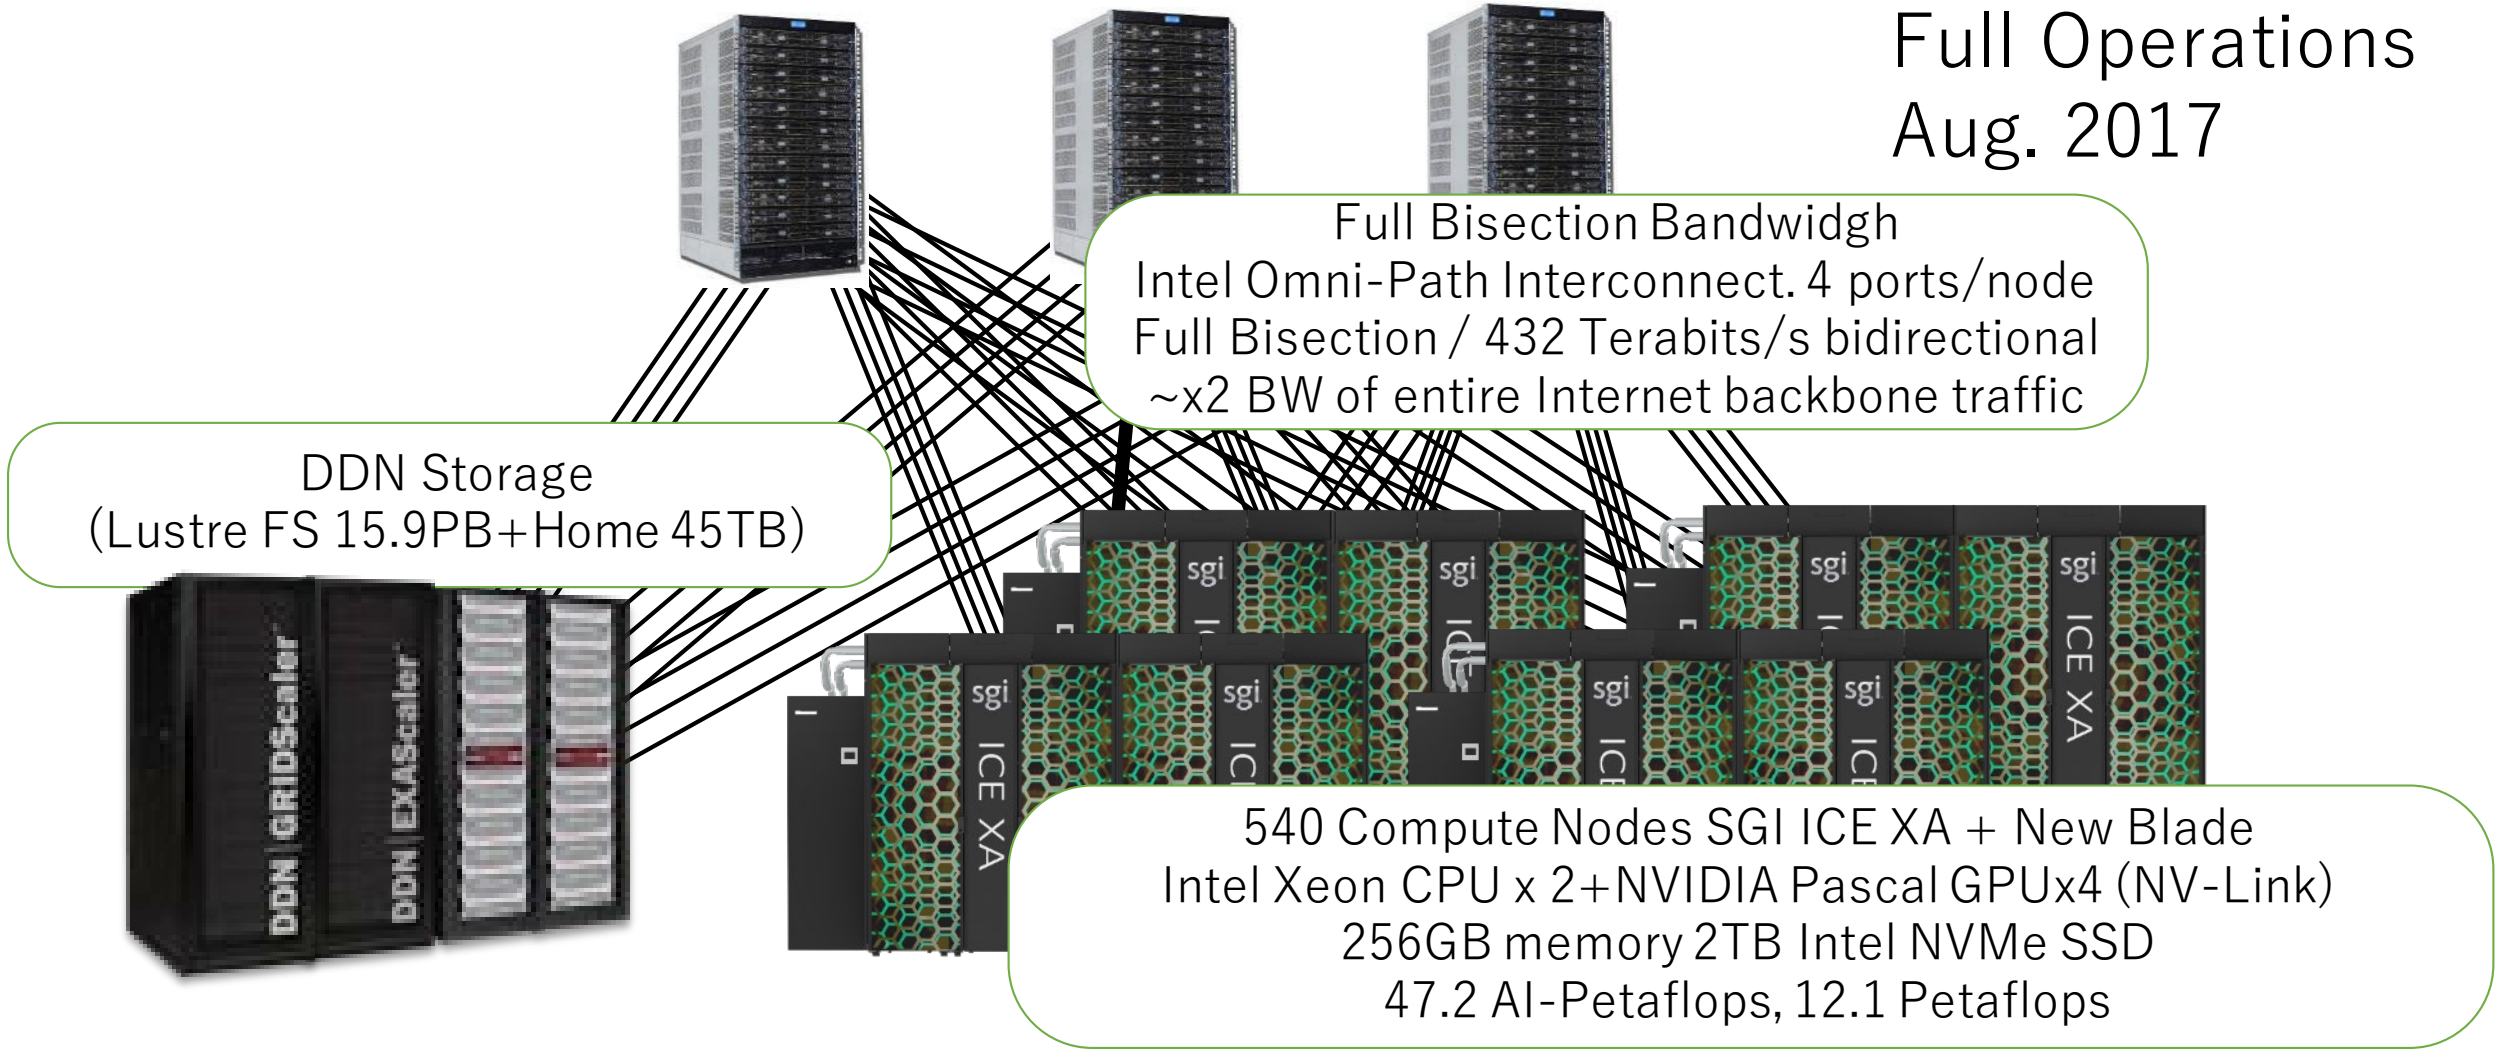

TSUBAME3.0 Co-Designed SGI ICE-XA Blade (new)

- No exterior cable mess (power, NW, water)
- Plan to become a future HPE product

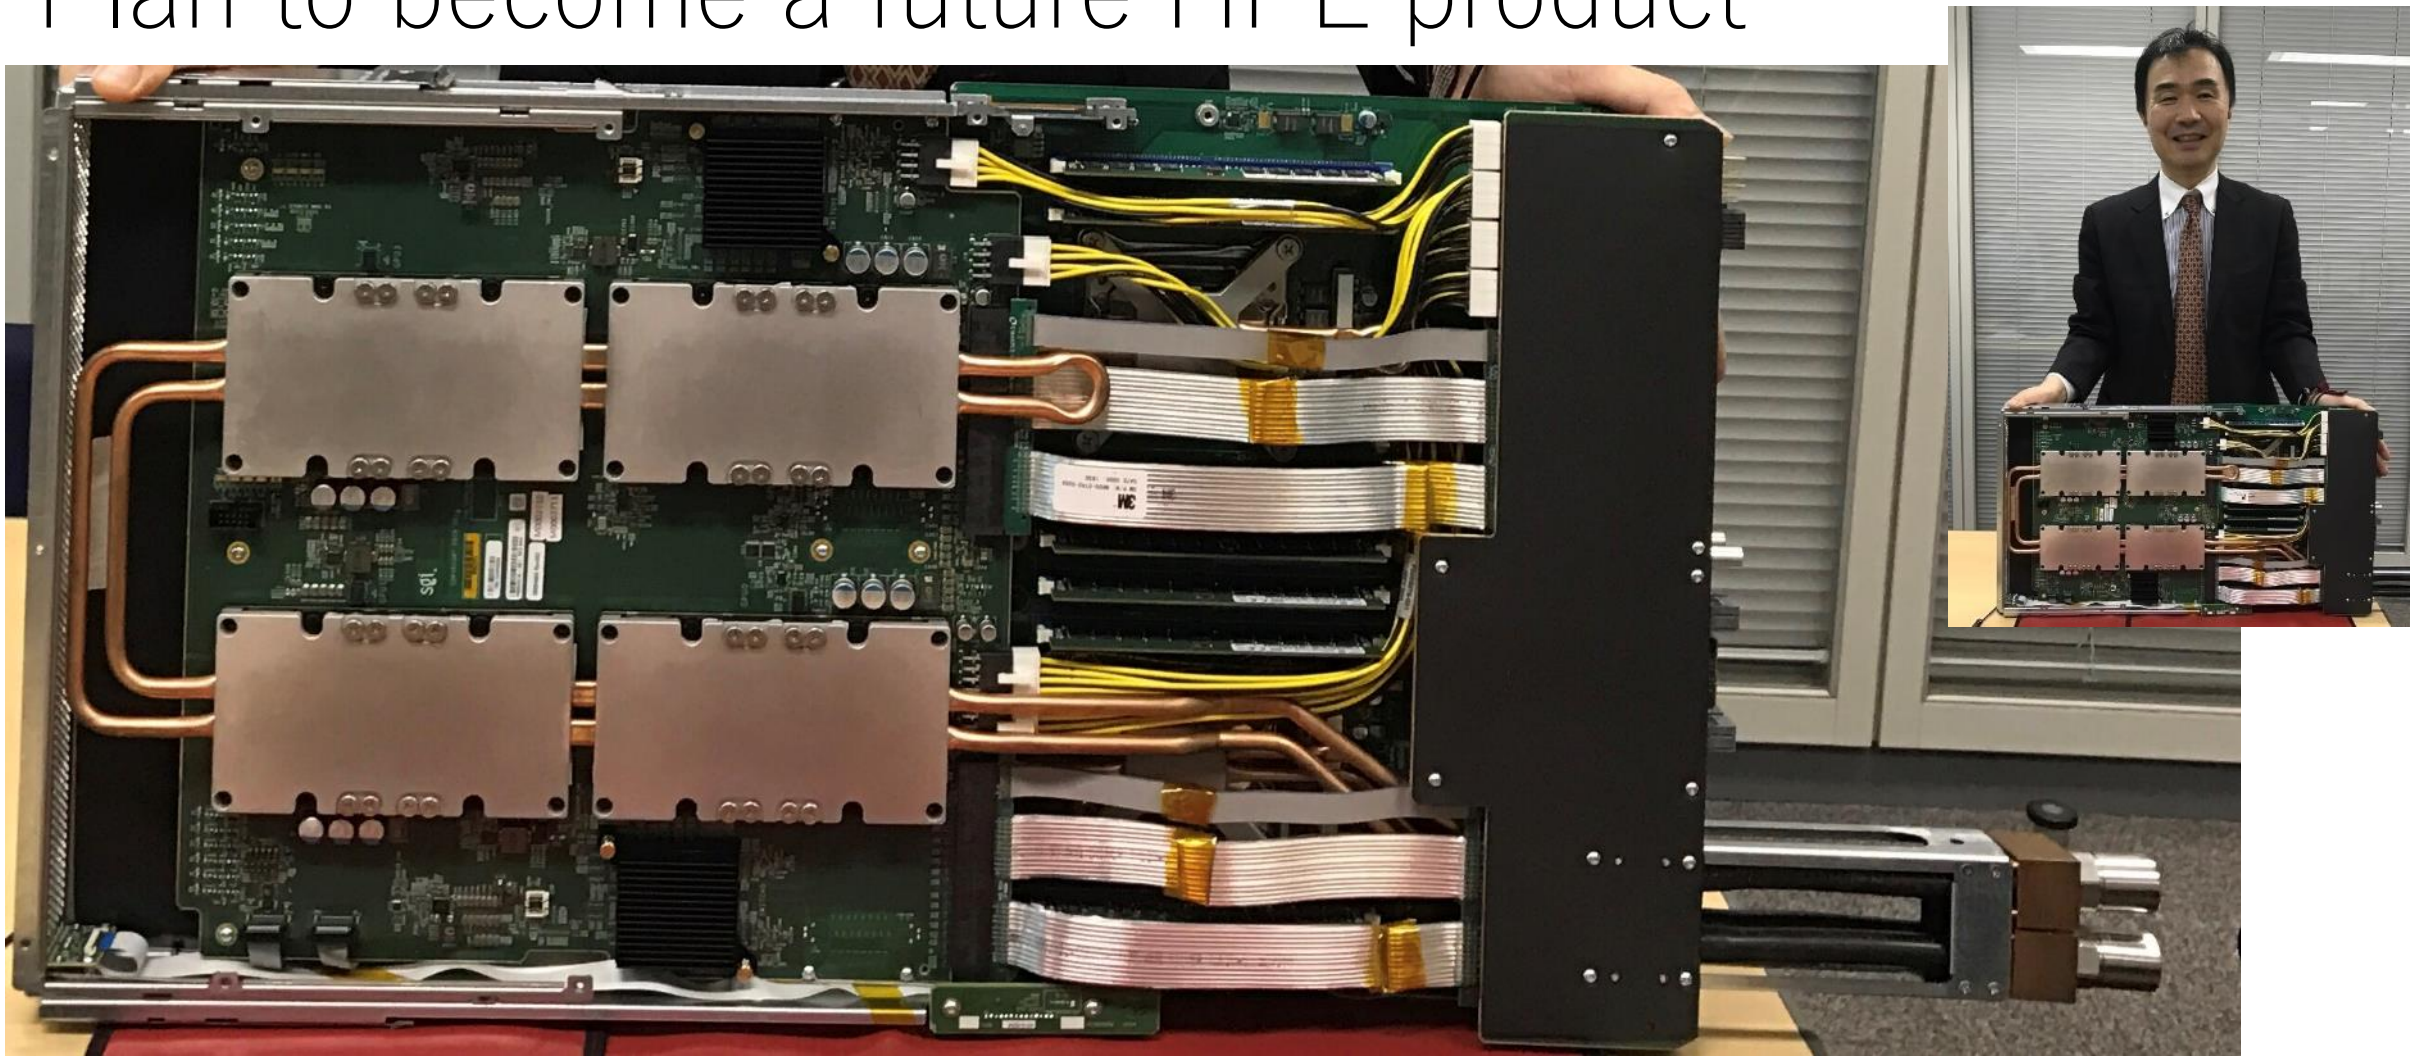

Liquid Cooled
"Hot Pluggable" ICE-XA Blade
Smaller than 1U server, no cables or pipes

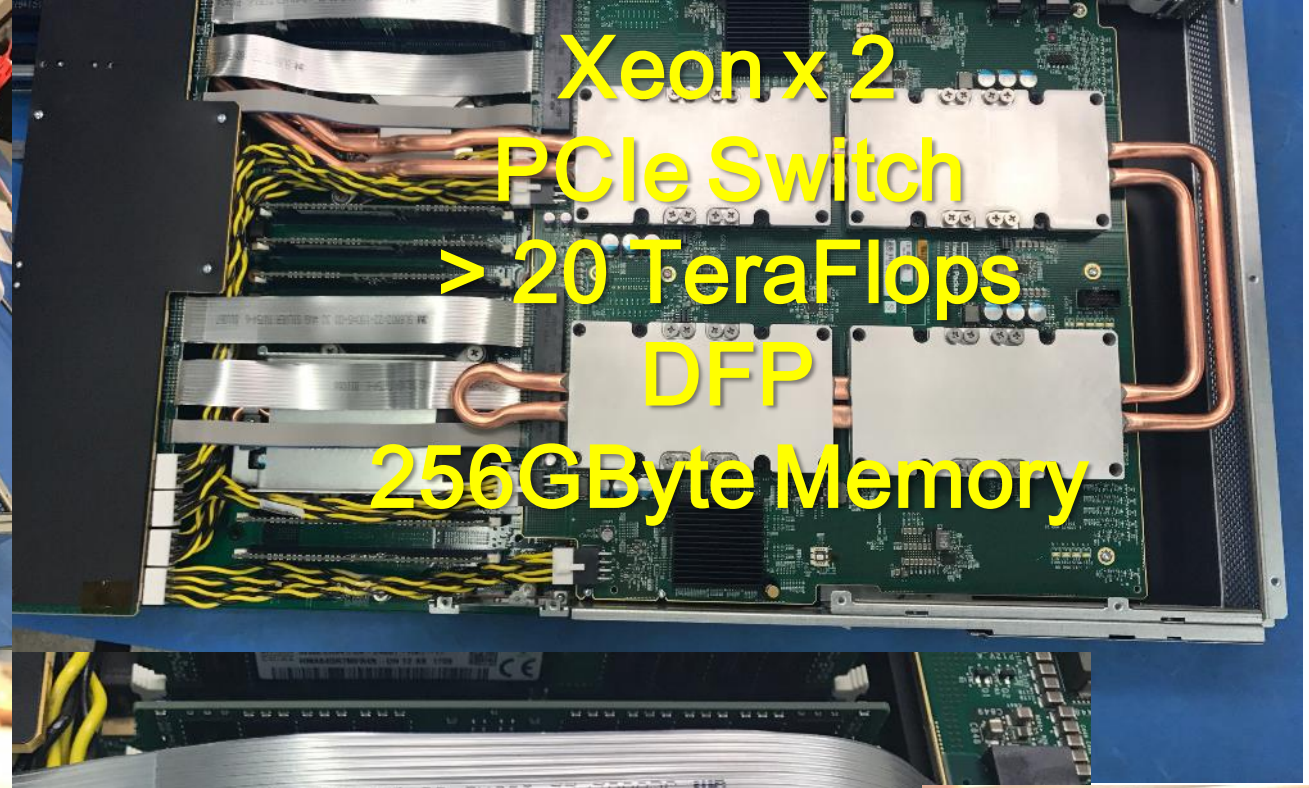
Xeon x 2
PCIe Switch
> 20 TeraFlops DFP
256GByte Memory

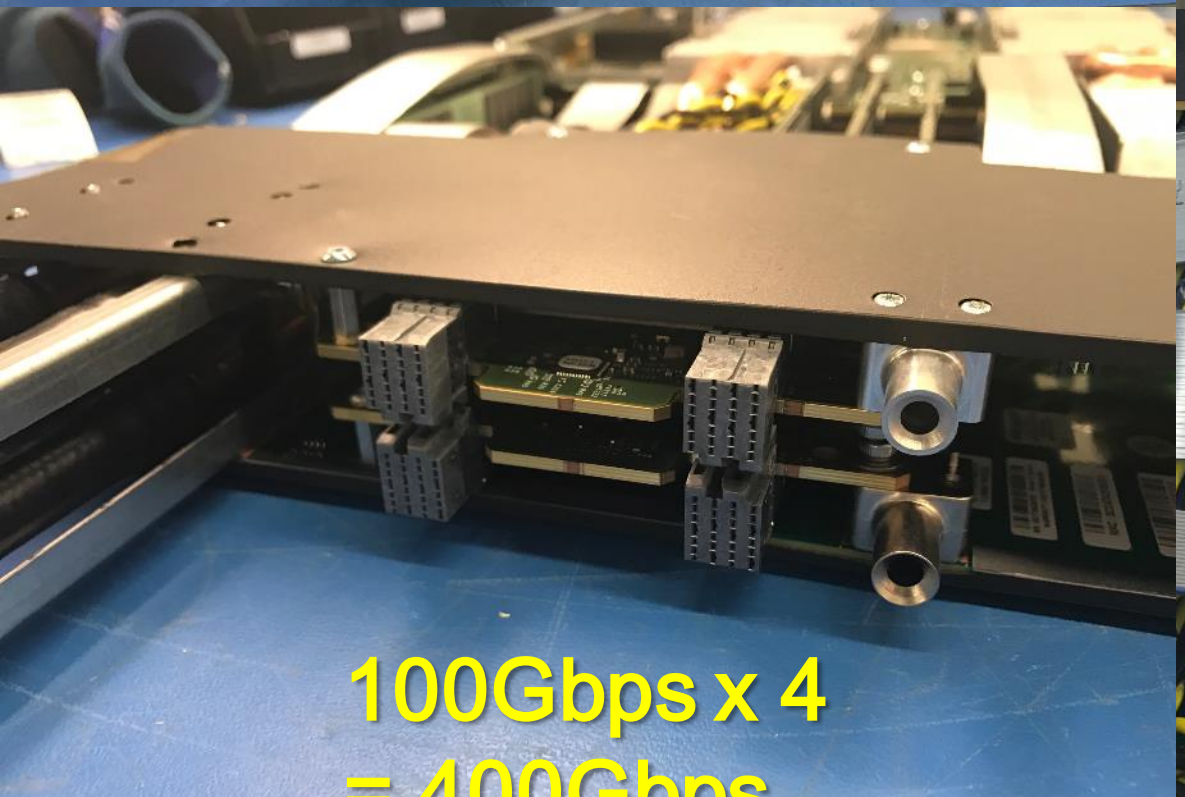
100Gbps x 4
= 400Gbps

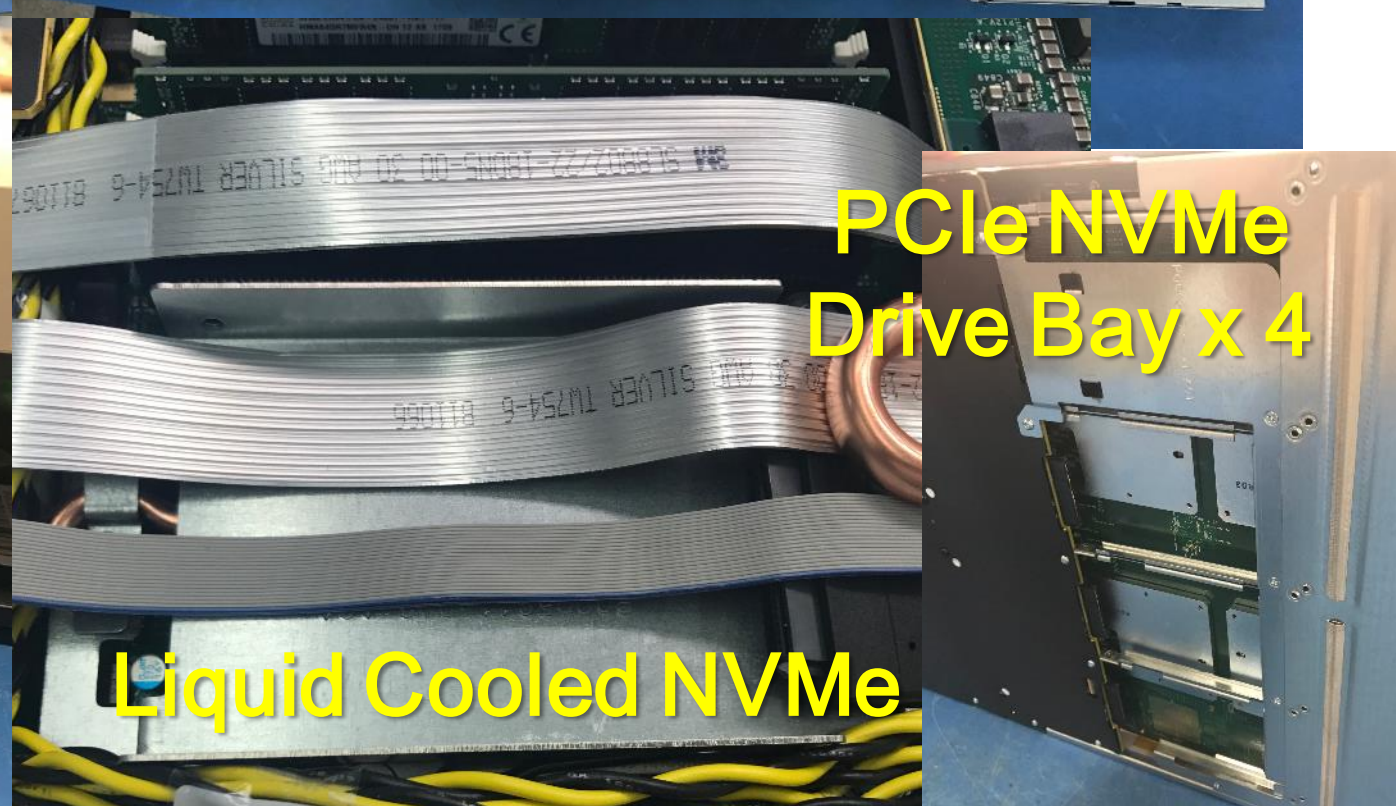
PCIe NVMe
Drive Bay x 4
Liquid Cooled NVMe

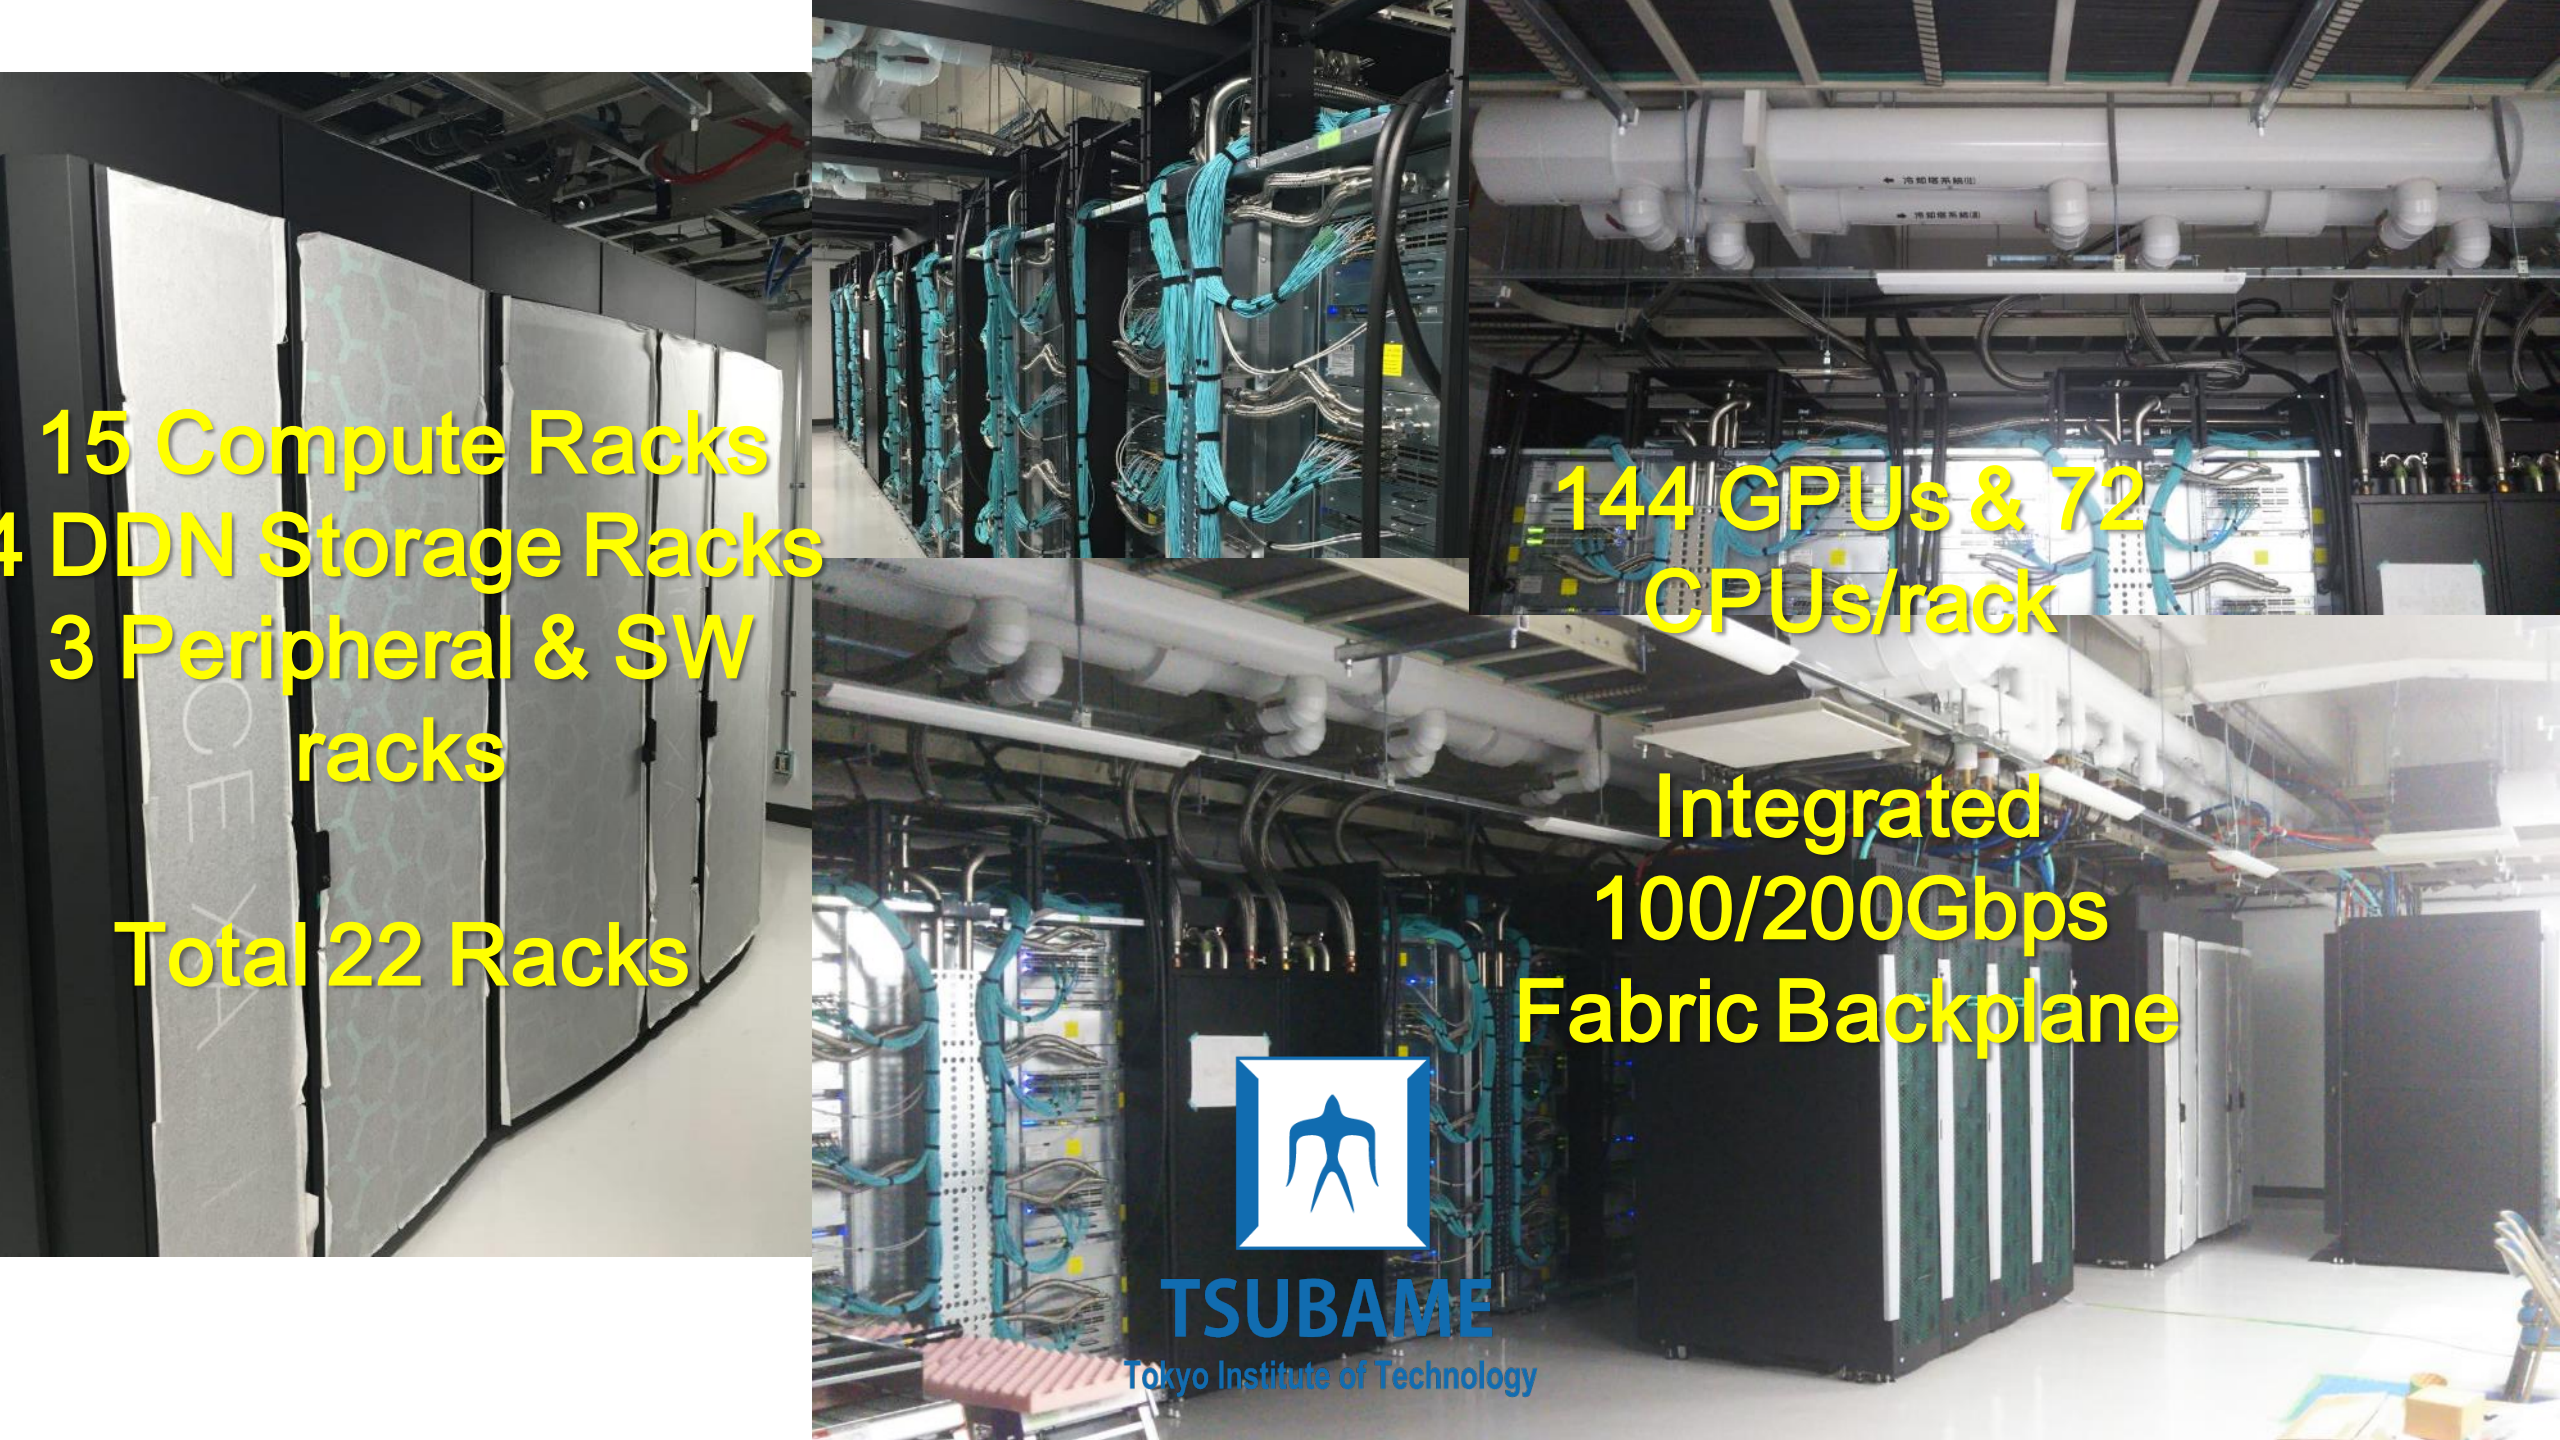
15 Compute Racks
4 DDN Storage Racks
3 Peripheral & SW racks
Total 22 Racks
144 GPUs & 72 CPUs/rack
Integrated 100/200Gbps Fabric Backplane
TSUBAME
Tokyo Institute of Technology

TSUBAME3: *A Massively BYTES Centric Architecture* for Converged BD/AI and HPC

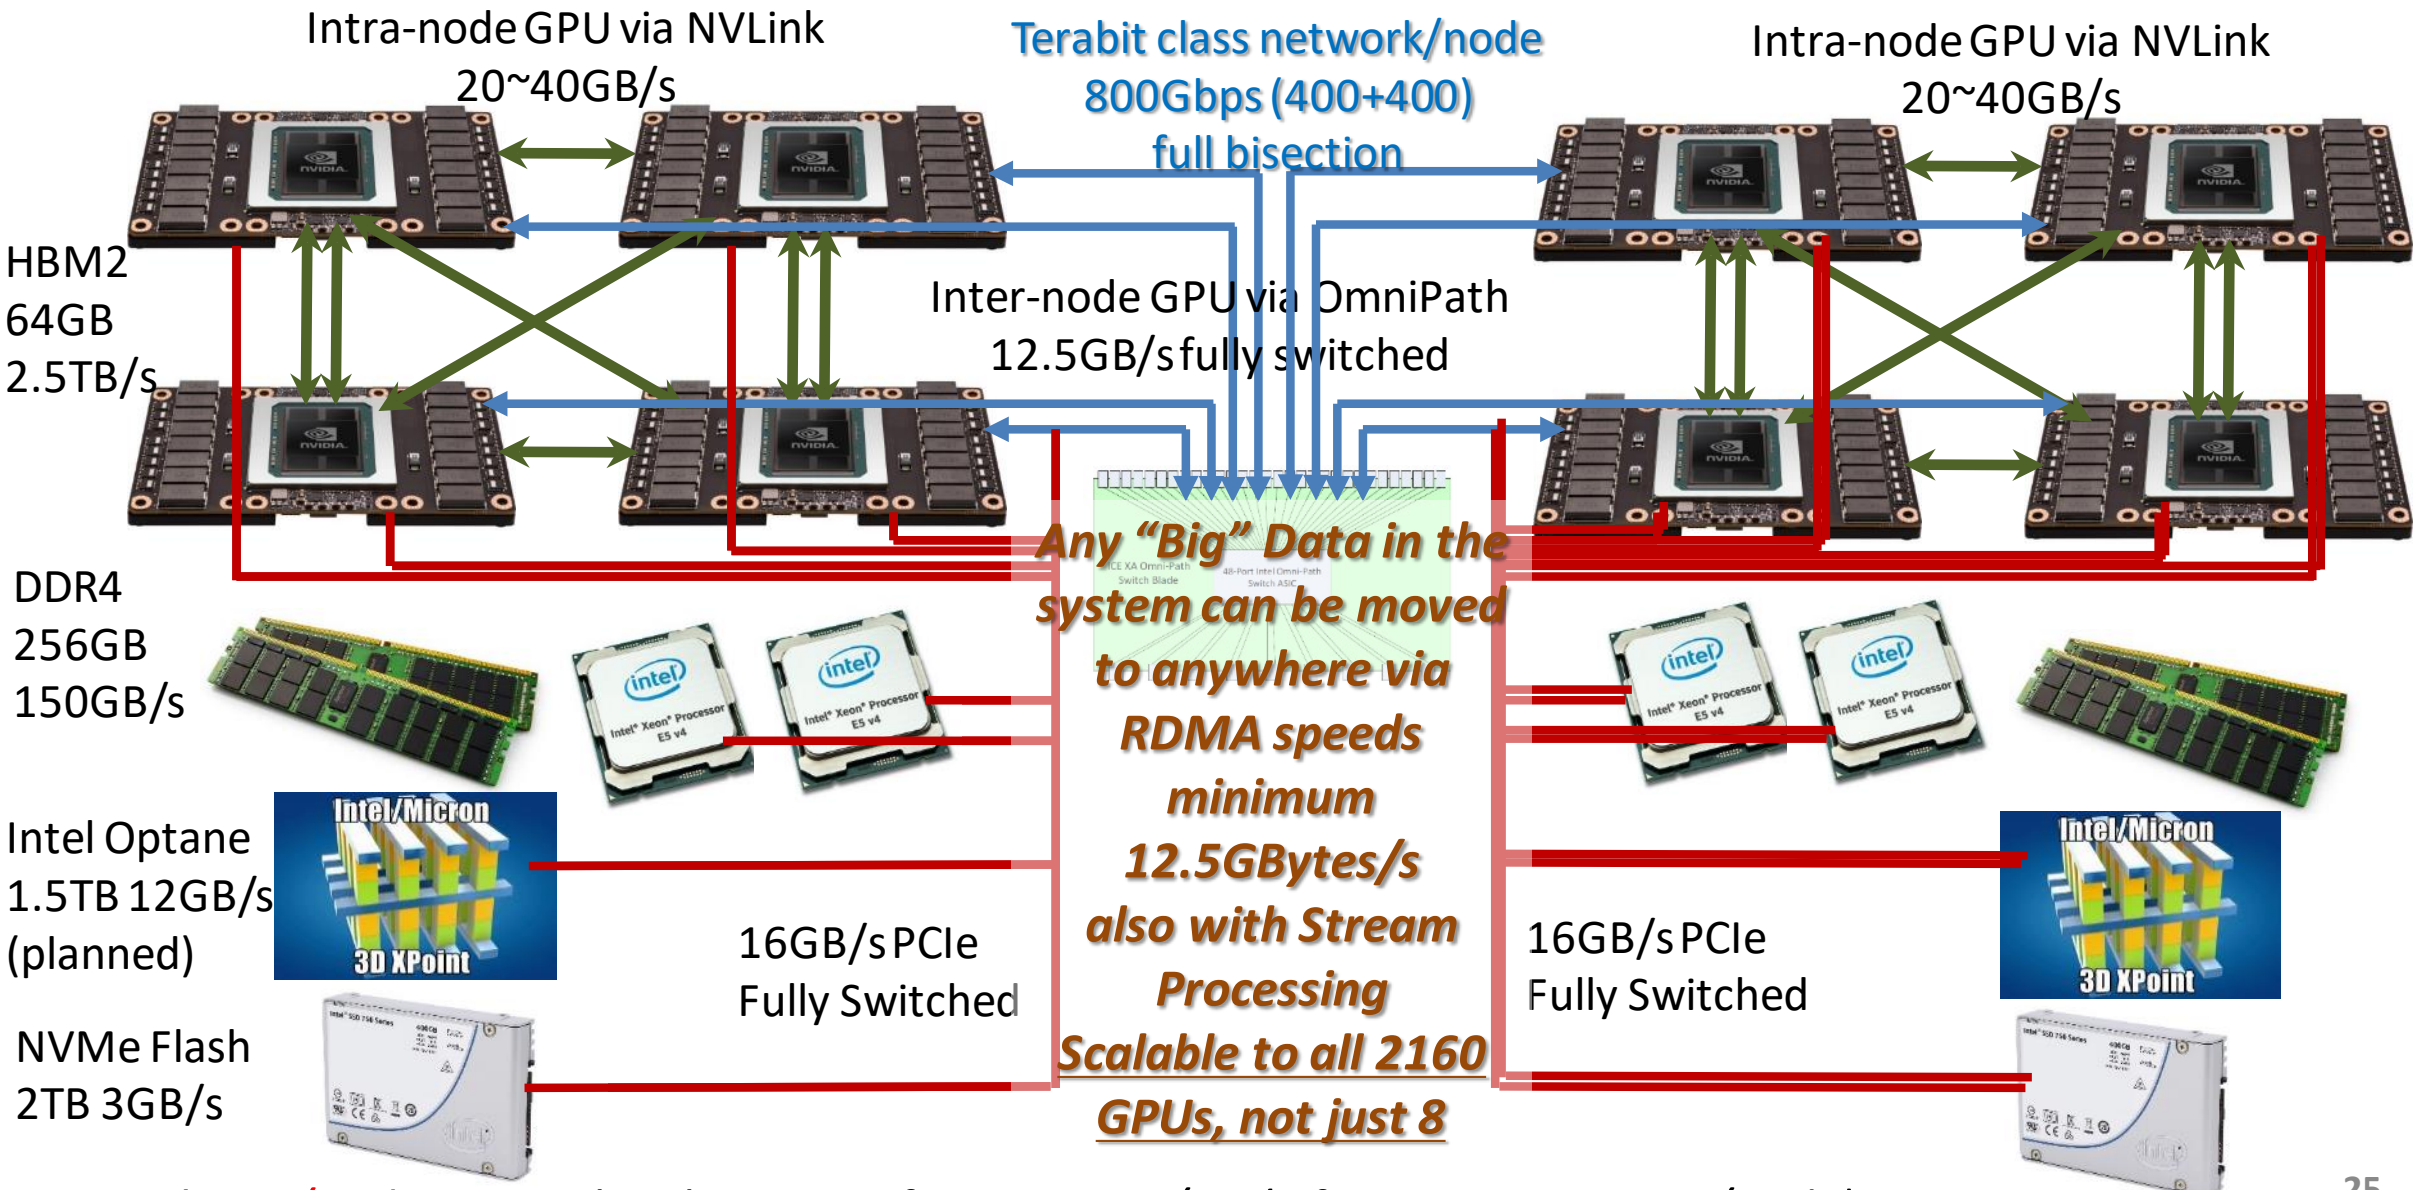

~4 Terabytes/node Hierarchical Memory for Big Data / AI (c.f. K-compuer 16GB/node)

➔ *Over 2 Petabytes in TSUBAME3, Can be moved at 54 Terabyte/s or 1.7 Zetabytes / year*

TSUBAME3: *A Massively BYTES Centric Architecture* for Converged BD/AI and HPC

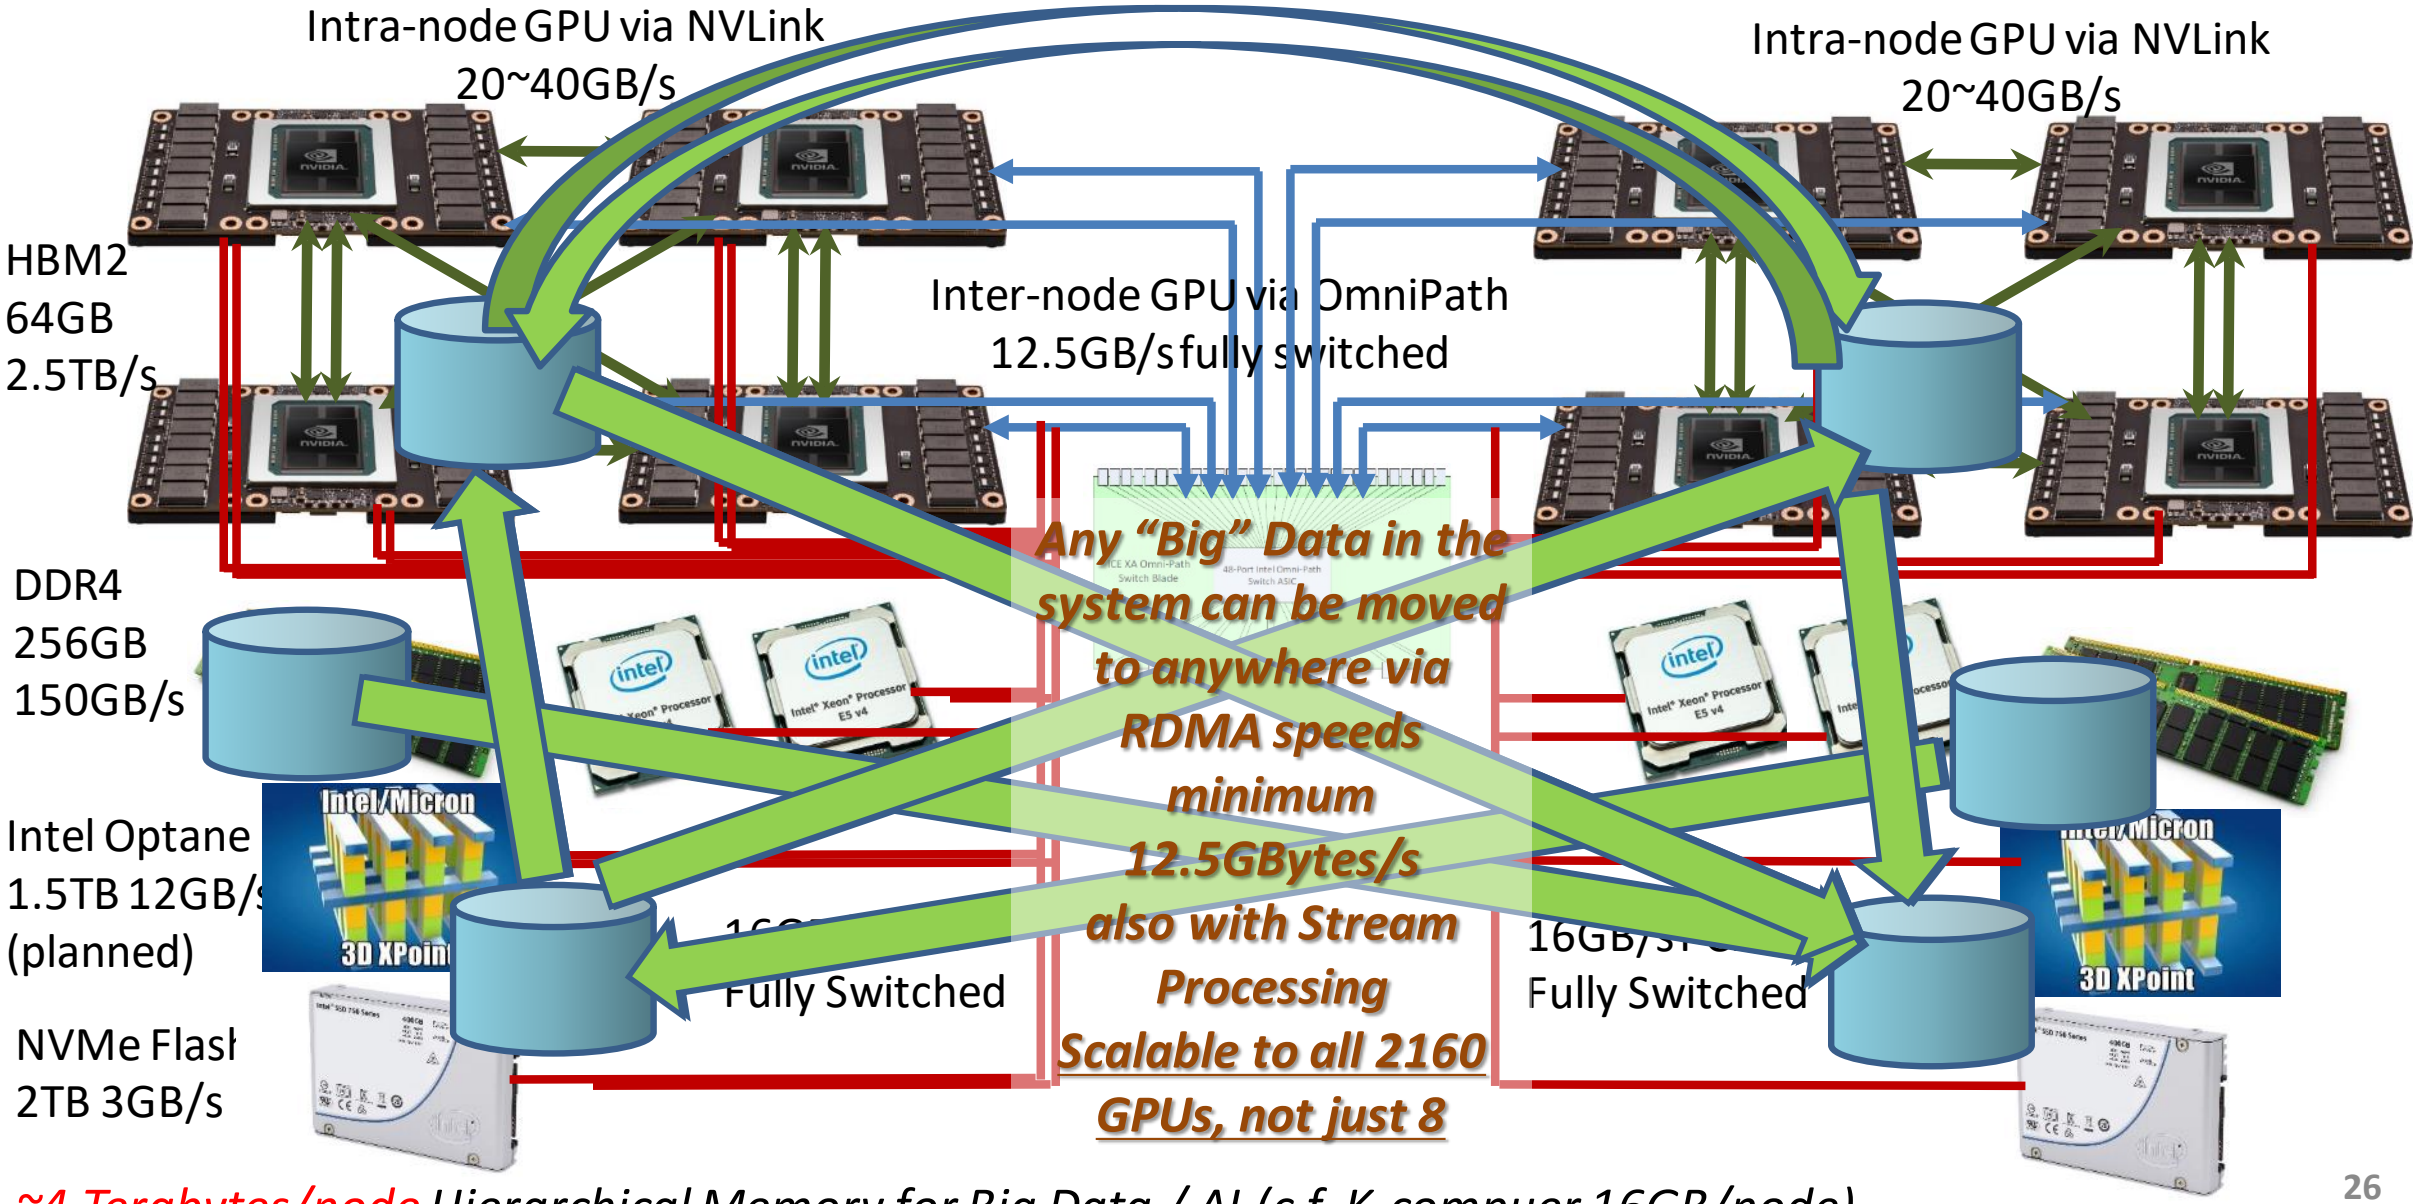

~4 Terabytes/node Hierarchical Memory for Big Data / AI (c.f. K-compuer 16GB/node)

➔ *Over 2 Petabytes in TSUBAME3, Can be moved at 54 Terabyte/s or 1.7 Zetabytes / year*

Tsubame3 Highly Efficient Datacenter
Machine PUE ~= 1.03 (~1.1 w/storage)

Over 100t total floor load

Reinforced "Slab-Like" flat floor surface.
Over 1t/m^2 floor load

Ultra High Density

Space Efficient

Power, water, and cabling are all above with ceiling support, for space efficiency and freedom of layout
IDC space ~= 130 m2

Piping and cabling hang from the ceiling

Max 32 degrees Celsius water

Year-round free "warm-water" cooling with cooling tower, PUE = 1.03, machine power ~= facility power

Efficient Warm-Water Cooling

Low Electrical Distribution Loss

420V High Voltage to minimize electrical distribution loss and cheaper cabling

420V Tri-Phase AC Power

Warm Water Cooling Distribution in T3

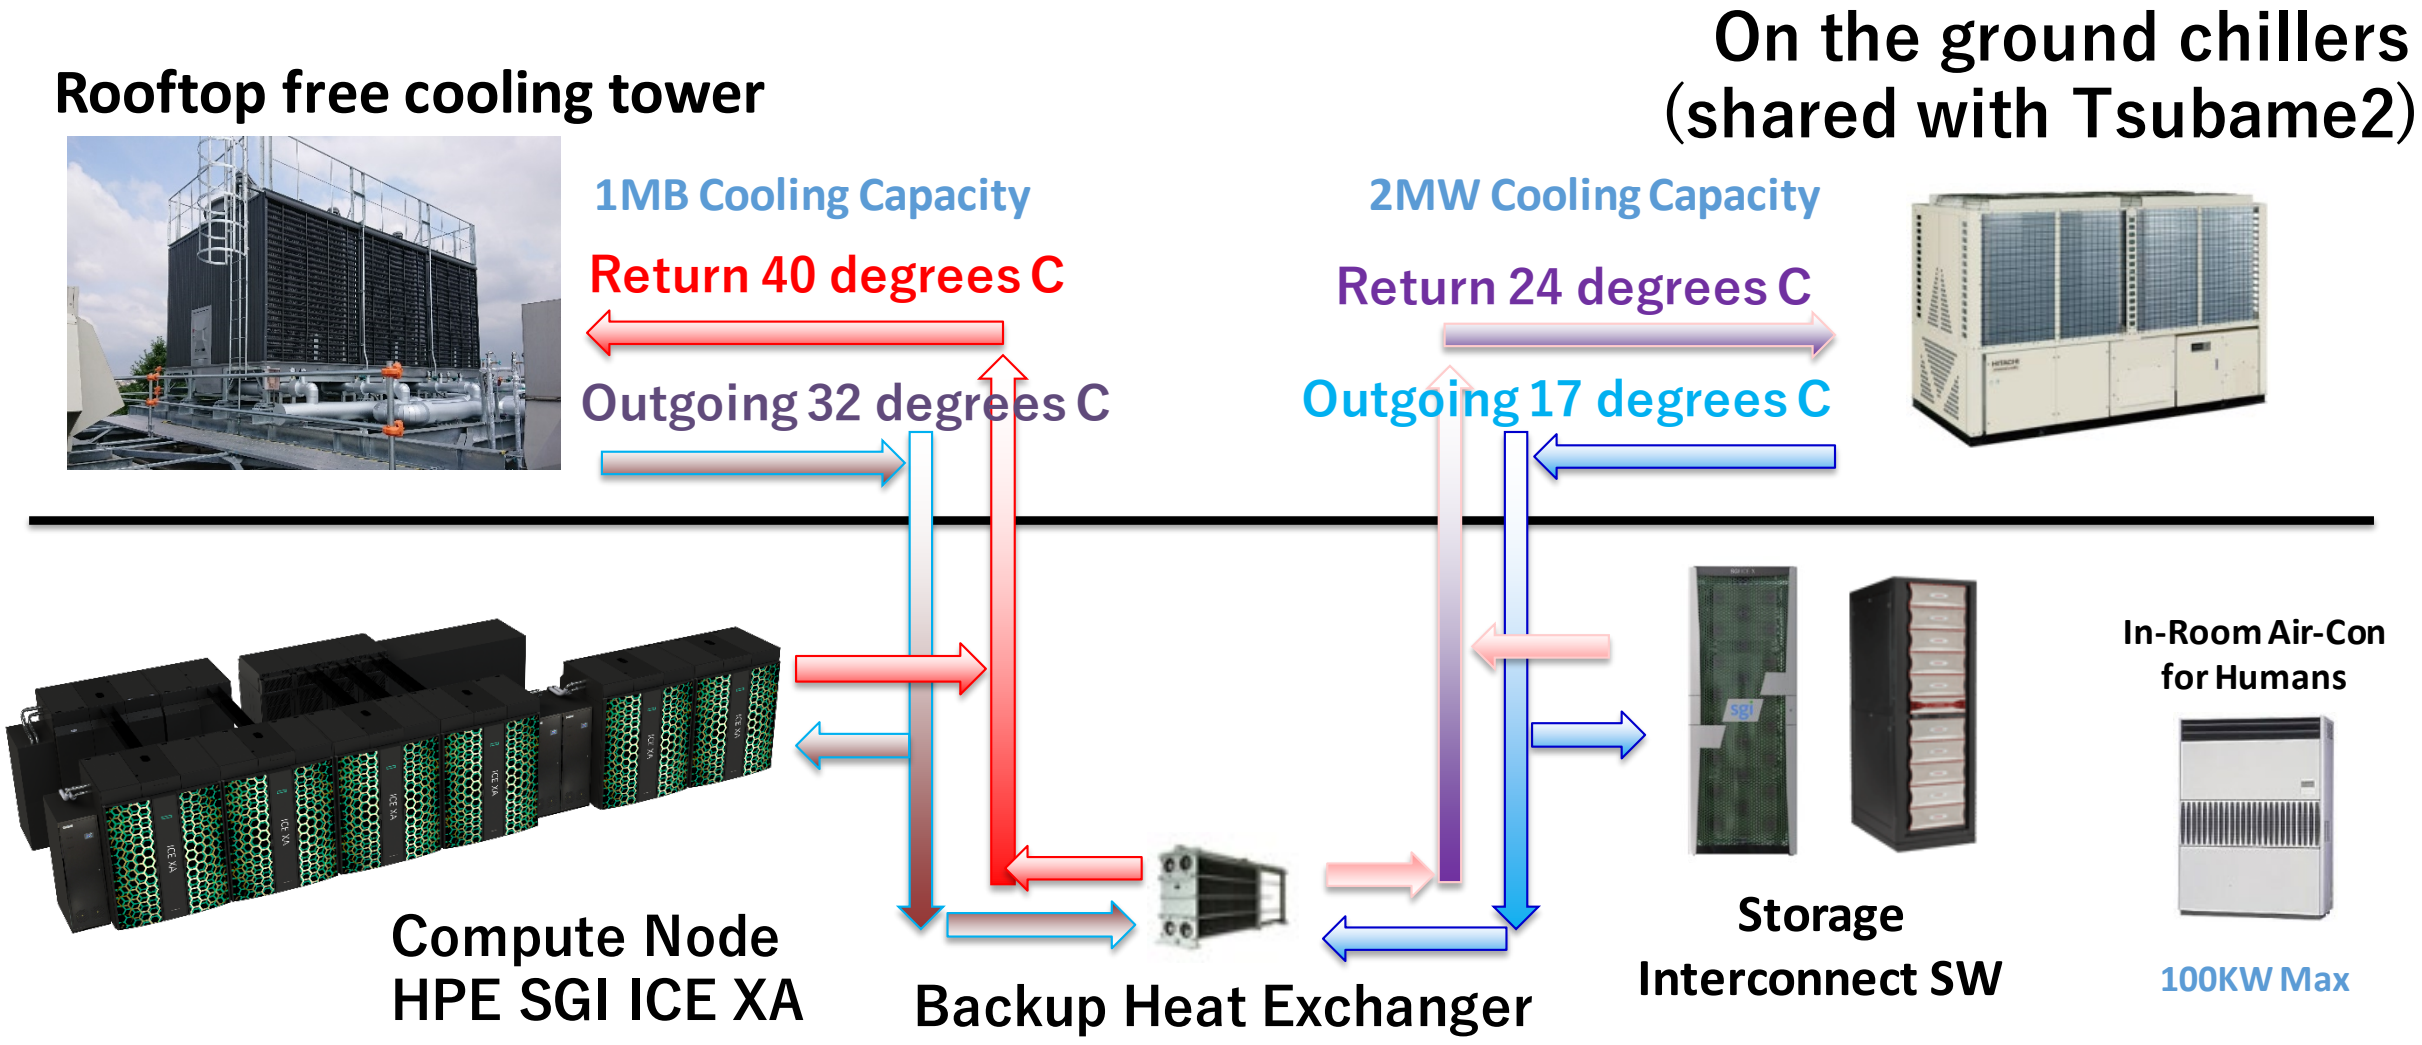

TSUBAME3.0 became the first large production petaflops-scale supercomputer in the world to be #1 on the "Green500" power efficiency W world ranking of supercomputers

14.1 Gigaflops/W is more than x10 more efficient than PCs and Smartphones!

The GREEN 500 CERTIFICATE

Tsubame 3.0, a modified HPE ICE XA System at the GSIC Center, Tokyo Institute of Technology

is ranked

No. 1 in the Green500

among the World's TOP500 Supercomputers
with 14.1 GFlops/Watt Linpack Power-Efficiency
on the Green500 List published at ISC High Performance, June 19, 2017

Congratulations from the Green500 Editors

Wu Feng
Virginia Tech

Kirk Cameron
Virginia Tech

The current status of AI & Big Data in Japan

We need the triage of advanced algorithms/infrastructure/data but we lack the cutting edge infrastructure dedicated to AI & Big Data (c.f. HPC)

Tokyo Tech

AIRC

Joint RWBC Open Innov. Lab (OIL) (Director: Matsuoka)

AIST NATIONAL INSTITUTE OF ADVANCED INDUSTRIAL SCIENCE AND TECHNOLOGY (AIST)

AIST-AIRC

東京工業大学 Tokyo Institute of Technology

AI Venture Startups

XC

UEI

Preferred Networks

R&D ML Algorithms & SW

Big Companies AI/BD R&D (also Science)

IT LAB DENSO IT LABORATORY, INC.

MIZUHO みずほ情報総研

Panasonic

ABEJA

DeNA

ZMP

Seeking Innovative Application of AI & Data

Riken -AIP

NICT National Institute of Information and Communications Technology

UCRI

NICT-UCRI

AI/BD Centers & Labs in National Labs & Universities

Over $1B Govt. AI investment over 10 years

Massive Rise in Computing Requirements (1 AI-PF/person?)

SAKURA Internet

Google Cloud Platform

Windows Azure

amazon web services

In HPC, Cloud continues to be insufficient for cutting edge research => dedicated SCs dominate & racing to Exascale

AI&Data Infrastructures

Massive "Big" Data in Training

"Big" Data

Use of Massive Scale Data now Wasted

DENSO Petabytes of Drive Recording Video

FA&ロボット&ロボマシン FANUC

FA&Robots

TOYOTA

JARI 一般財団法人 日本自動車研究所

Web access and merchandice

Y! YAHOO! JAPAN

SoftBank

NTT

IoT Communication, location & other data

ABCI: AI Bridging Cloud Infrastructure "Architecting Ultra Efficient AI-HPC IDC"

ABCI
AI Bridging Cloud Infrastructure

- **Open, Public, and Dedicated** infrastructure for AI & Big Data Algorithms, Software, and Applications
- **Open Innovation Platform** to accelerate joint academic-industry R&D for AI, international collaborations are also welcome
- Top-level compute capability: **0.55 EFLOPS (DL), 37 PFLOPS (DP)**
- Top-level energy efficiency: lower PUE
- All design and implementations will be open-sourced

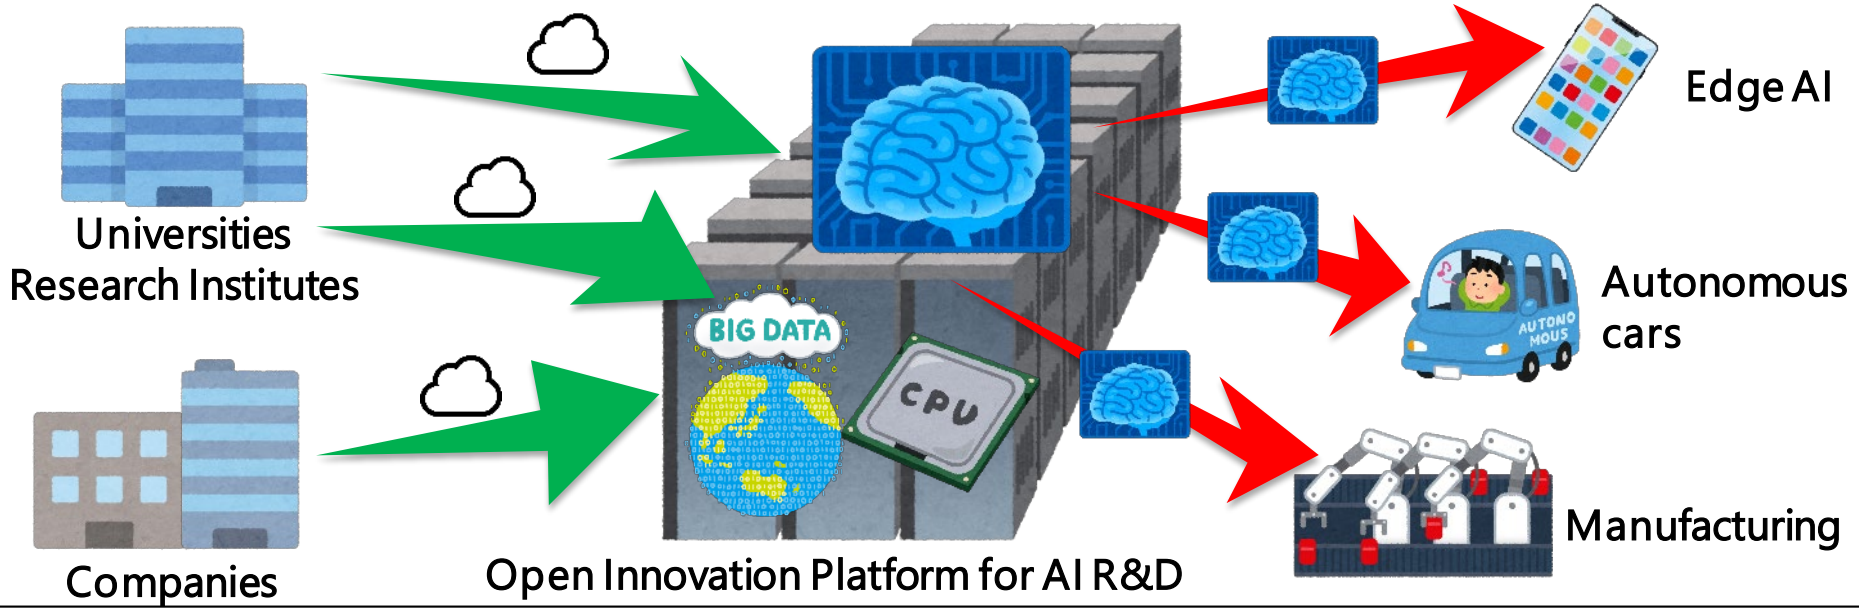

AIST ABCI Hardware Overview

High-Performance Computing System

550 PFlops(FP16), 37.2 PFlops(FP64)
476 TiB Memory, 1.74 PB NVMe SSD

Computing Nodes (w/ GPU) x 1088

GPU	NVIDIA Tesla V100 SXM2 x 4
CPU	Intel Xeon Gold 6148 (2.4GHz/20cores) x 2
Memory	384GiB
Local Storage	Intel SSD DC P4600 (NVMe) 1.6TB x 1
Interconnect	InfiniBand EDR x 2

Multi-platform Nodes (w/o GPU) x 10

- Intel Xeon Gold 6132 (2.6GHz/14cores) x 2
- 768GiB Memory, 3.8TB NVMe SSD

Interactive Nodes x 4

Management and Gateway Nodes x 15

Large-scale Storage System

22 PB GPFS

DDN SFA14K (w/ SS8462 Enclosure x 10) x 3

- 12TB 7.2Krpm NL-SAS HDD x 2400
- 3.84TB SAS SSD x 216
- NSD Server x 12

Protocol Nodes, etc x 8

Interconnect (Infiniband EDR)

- Mellanox CS7500 x 2
- Mellanox SB7890 x 229

Gateway and Firewall

- Nexsus 3232C x2
- FortiGate 1500D x2
- FortiAnalyzer 400E x1

Service Network (10GbE)

100Gbps

SINET5

ABCI High-Performance Computing System

0.550 EFlops(FP16), 37.2 PFlops(FP64)
19.88 PFlops(Peak), Ranked #5 Top500 June 2018

Chips
(GPU, CPU)

Tesla V100

Xeon Skylake-SP

Compute Node
(4GPUs, 2CPUs)

PRIMERGY CX2570 M4

Node Chassis
(2 Compute Nodes)

PRIMERGY CX400 M4

Rack
(17 Chassis)

System
(32 Racks)

1088 Compute Nodes
4352 GPUs

GPU:	**CPU:**				
7.8 TFlops(FP64)	**1.53 TFlops(FP64)**	**34.2 TFlops(FP64)**	**68.5 PFlops(FP64)**	**1.16 PFlops(FP64)**	**37.2 PFlops(FP64)**
125 TFlops(FP16)	**3.07 TFlops(FP32)**	**506 TFlops(FP16)**	**1.01 PFlops(FP16)**	**17.2 PFlops(FP16)**	**0.55 EFlops(FP16)**
		～3.72 TB/s MEM BW		～131TB/s MEM BW	～4.19 PB/s MEM BW
NVIDIA Tesla V100 (16GB SMX2)		384 GiB MEM 200 Gbps NW BW 1.6TB NVMe SSD		Full Bisection BW within Rack 70kW Max	1/3 of Oversubscription BW 2.3MW

ABCI Computing Node

FUJITSU PRIMERGY Server (2 servers in 2U)	
CPU	Xeon Gold 6148 (27.5M Cache, 2.40 GHz, 20 Core) x2
GPU	NVIDIA Tesla V100 (SXM2) x4
Memory	384GiB DDR4 2666MHz RDIMM
Local Storage	1.6TB NVMe SSD (Intel SSD DC P4600 u.2) x1
Interconnect	InfiniBand EDR x2

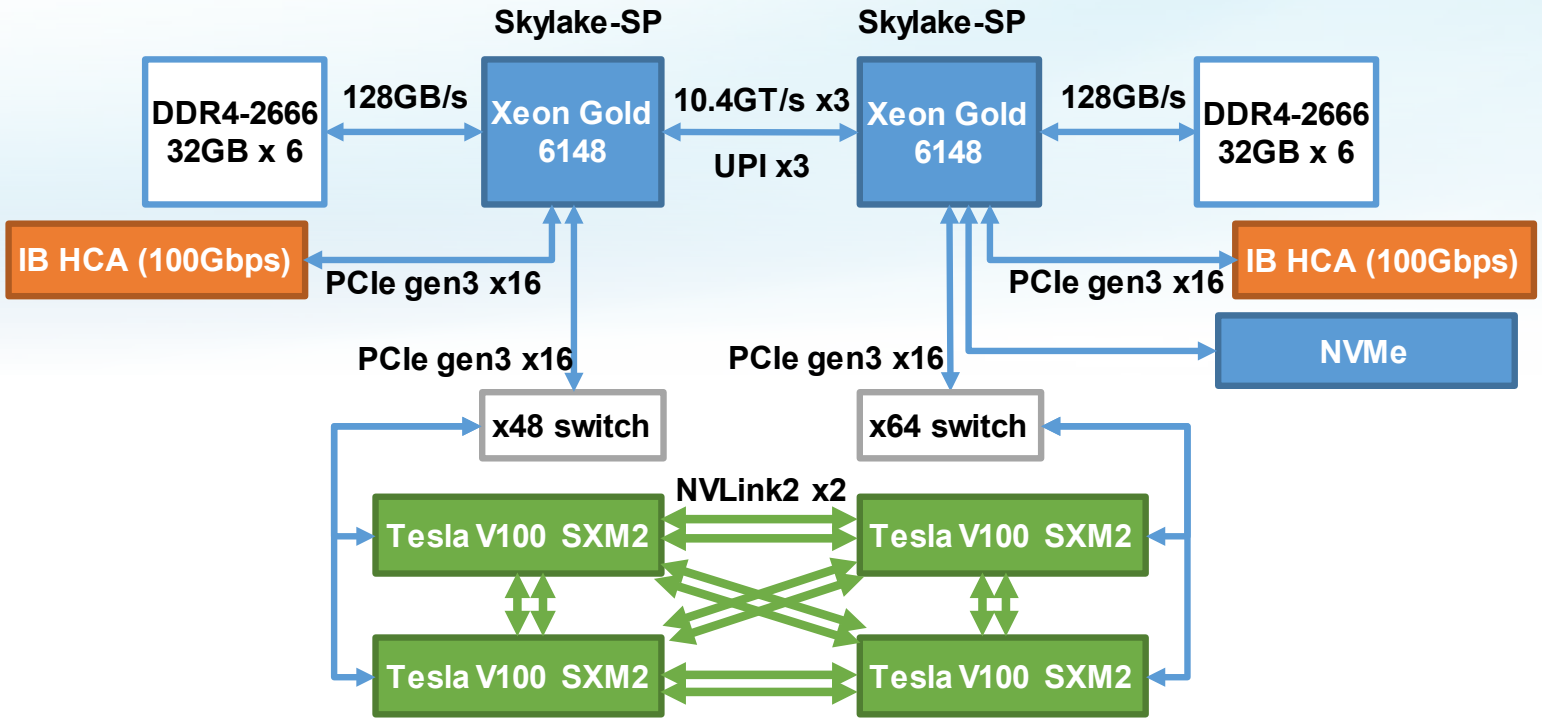

CPU blade

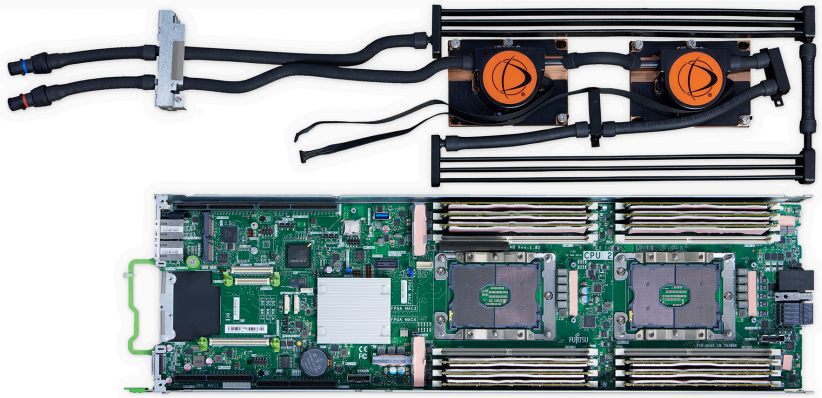

GPU blade

ABCI Computing Rack / Interconnect

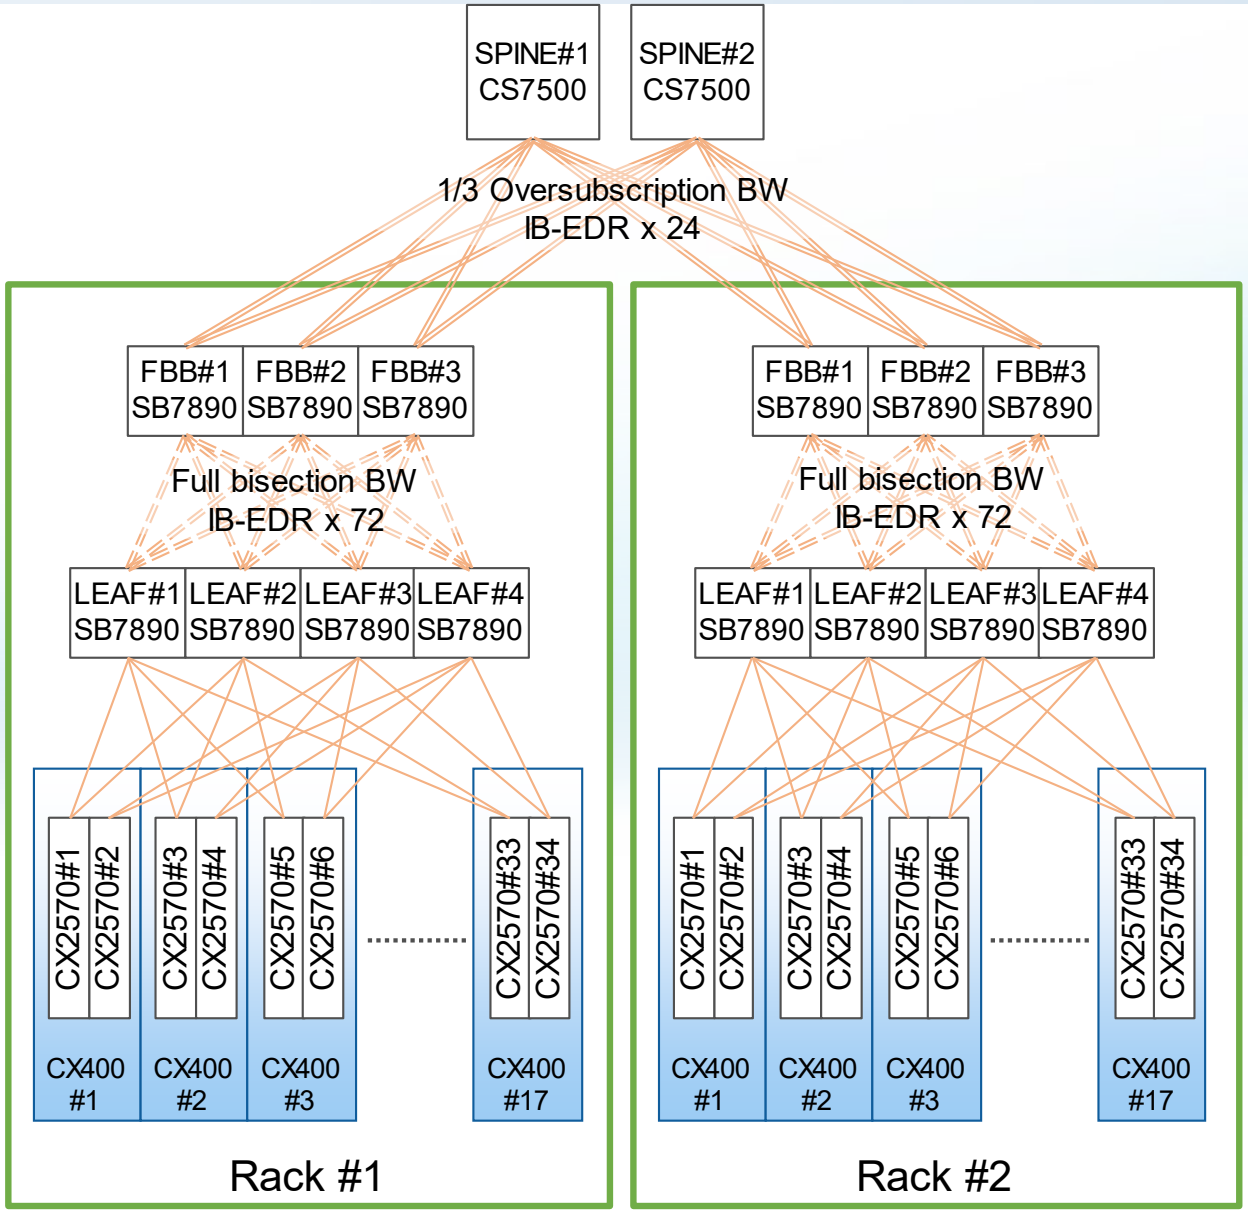

- **Dense-packaged rack: 34 nodes, 136 Tesla V100**
  - Theoretical peak performance per rack : 1.16 PFlops (FP64), 17 PFlops (FP16) c.f. Google TPU 3.0 Pod (>100PFlops)
  - Power consumption per rack: 67.33 kW
- **Interconnect**
  - Fat-tree topology
  - Intra-rack: full bisection BW
  - Inter-rack : 1/3 over-subscription (2400/6800)
  - Large-scale storage system: full bi-section BW

.....

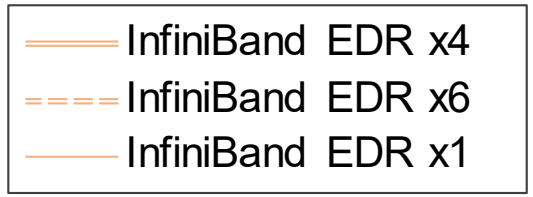

ABCI Procurement Benchmarks

- **Big Data Benchmarks**
  - (SPEC CPU Rate)
  - Graph 500
  - MinuteSort
  - Node Local Storage I/O
  - Parallel FS I/O

No traditional HPC Simulation Benchmarks except SPEC CPU.
Plan on “open-sourcing”

- **AI/ML Benchmarks**
  - Low precision GEMM
    - CNN Kernel, defines “AI-Flops”
  - Single Node CNN
    - AlexNet and GoogLeNet
    - ILSVRC2012 Dataset
  - Multi-Node Scalable CNN
    - Caffe+MPI
  - Large Memory CNN
    - Convnet on Chainer
  - RNN / LSTM
    - Neural Machine Translation on Torch

Container is critical for AI R&D

- lang2program (referred in ACL2017)
  - https://github.com/kelvinguu/lang2program
  - Provided as a Dockerfile
  - Bunch of software needed to be run
    - Tensorflow, PostgresQL, Python Pip Packages, etc.
- It is not possible to run it on traditional large-scale HPC systems
  - Arbitrary/Voluntary installation of software => chaos
  - Docker => security reason
- We're developing an easy-to-manage and flexible-to-use platform for deploying AI apps as "modules"

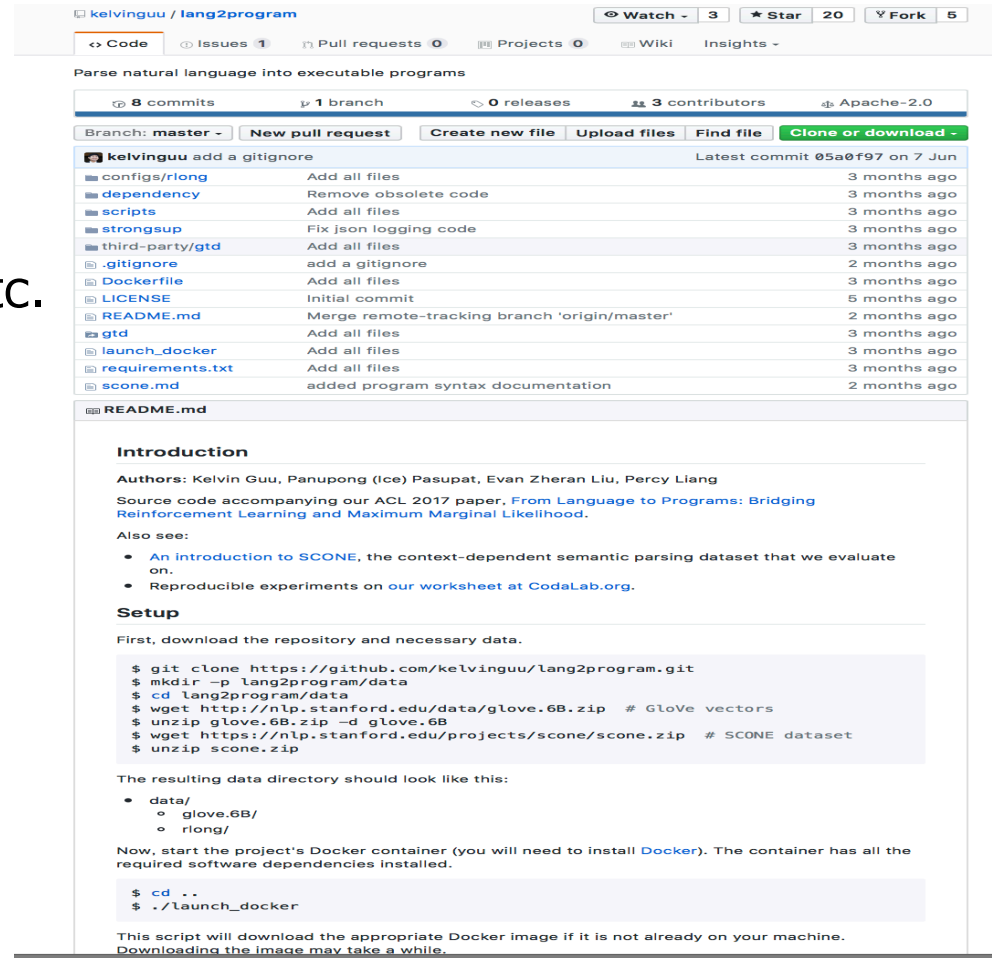

AI Frameworks & Modules

- Host OS only provides the minimum set of software including:
  - Base drivers, libraries
- "Base" modules provide customized OS images including:
  - CUDA, cuDNN, NCCL, MPI, etc.
- "DL" modules provide deep learning frameworks and apps, which extend "Base" images.
- We mainly employ Singularity as the basis of our AI platform.

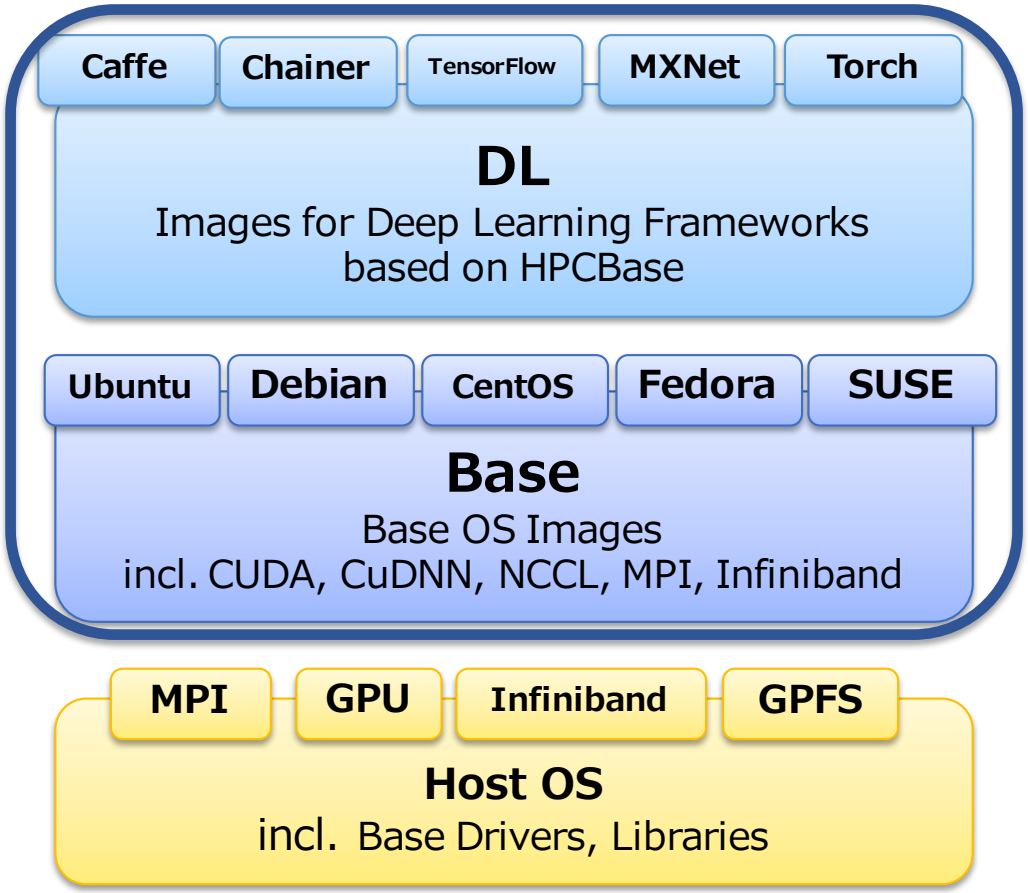

ABCI Software Stack

Software

Operating System	CentOS, RHEL
Job Scheduler	Univa Grid Engine
Container Engine	Docker, Singularity
MPI	OpenMPI, MVAPICH2, Intel MPI
Development tools	Intel Parallel Studio XE Cluster Edition, PGI Professional Edition, NVIDIA CUDA SDK, GCC, Python, Ruby, R, Java, Scala, Perl
Deep Learning	Caffe, Caffe2, TensorFlow, Theano, Torch, PyTorch, CNTK, MXnet, Chainer, Keras, etc.
Big Data Processing	Hadoop, Spark

Container support

- Containers enable users to instantly try the state-of-the-art software developed in AI community
- ABCI supports two container technologies
  - **Docker**, having a large user community
  - **Singularity**, recently accepted HPC community
- ABCI provides various single-node/distributed deep learning framework container images optimized to achieve high performance on ABCI

docker

The image part with relationship ID rId5 wa

ABCI Cloud Services

Service Types

Service Type	Description	#Nodes (Min./Max.)
Spot	Batch job service	1 / 512
On-demand	Interactive job service	1 / 32
Reserved	Advanced reservation service	1 / 32
Group Storage	Shared storage service	N/A

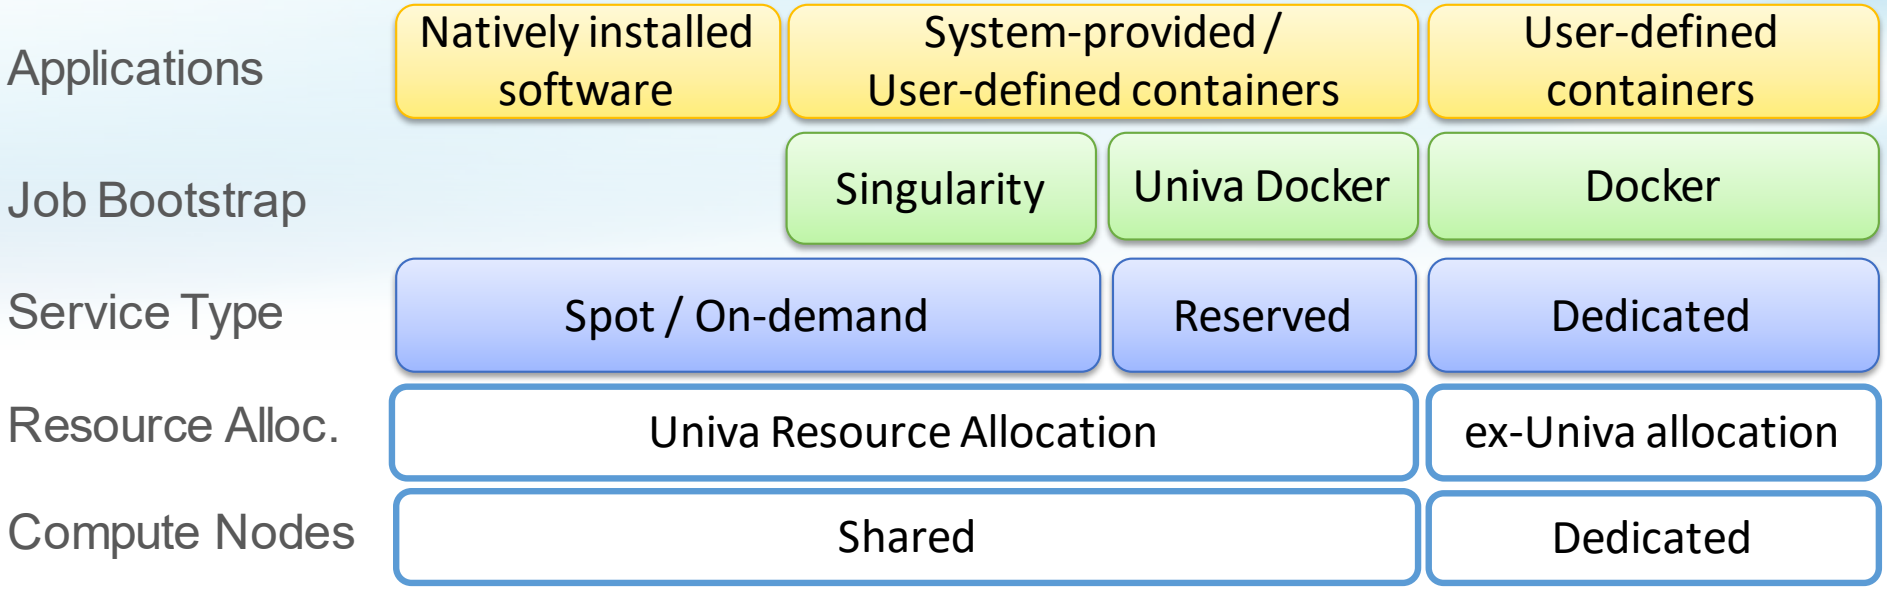

※ABCI provides batch and interactive-type job execution service for maximizing throughput, advanced reservation service for dedicated use of nodes, IDE, and storage services.

ABCI Cloud Services

Resource Types

※Users can choose the most suitable computing instance from five resource types.

Types	#CPU core Assign / Total	#GPU Assign / Total	Memory (GB) Assign / Total	Storage (TB) Assign / Total
F (Full Node)	40 / 40	4 / 4	360 / 384	1.4 / 1.6
G.large	20 / 40	4 / 4	240 / 384	0.7 / 1.6
G.small	5 / 40	1 / 4	60 / 384	0.175 / 1.6
C.large	20 / 40	0 / 4	120 / 384	0.7 / 1.6
C.small	5 / 40	0 / 4	30 / 384	0.175 / 1.6

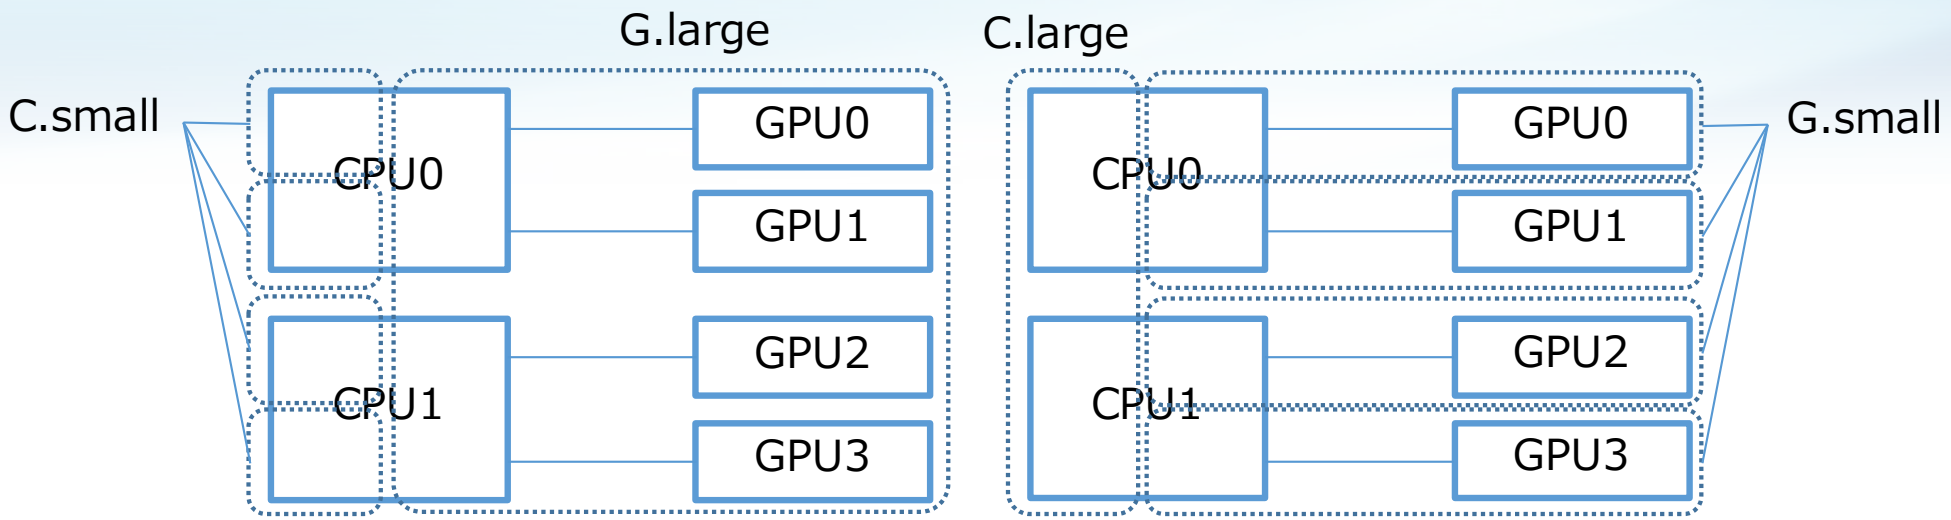

100% Free Cooling over the Entire Year

- Free cooling using a passive cooling tower
  - ○ Low OPEX. No active chiller (mechanical refrigeration)
  - ✕ Generated water temperature is depended on the external weather (temp. and humidity)
- Hybrid water/air cooling with high-temperature cooling water (32℃)

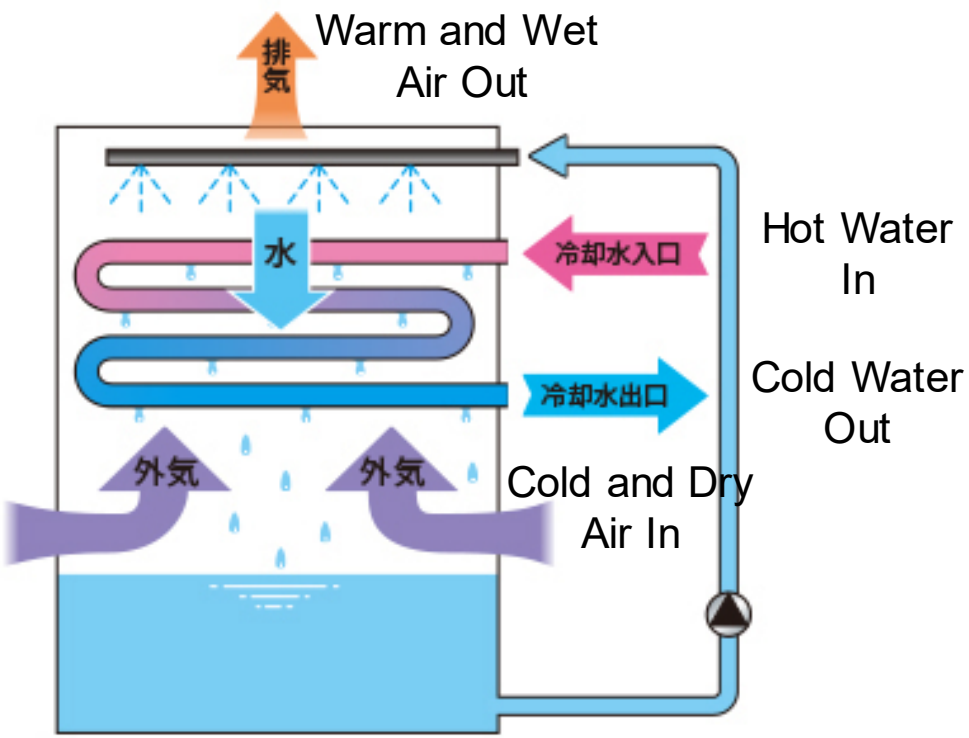

https://www.sinko.co.jp/product/technical-column/09/

AI Datacenter
"Commoditizing supercomputer cooling technologies to Cloud (70kW/rack)"

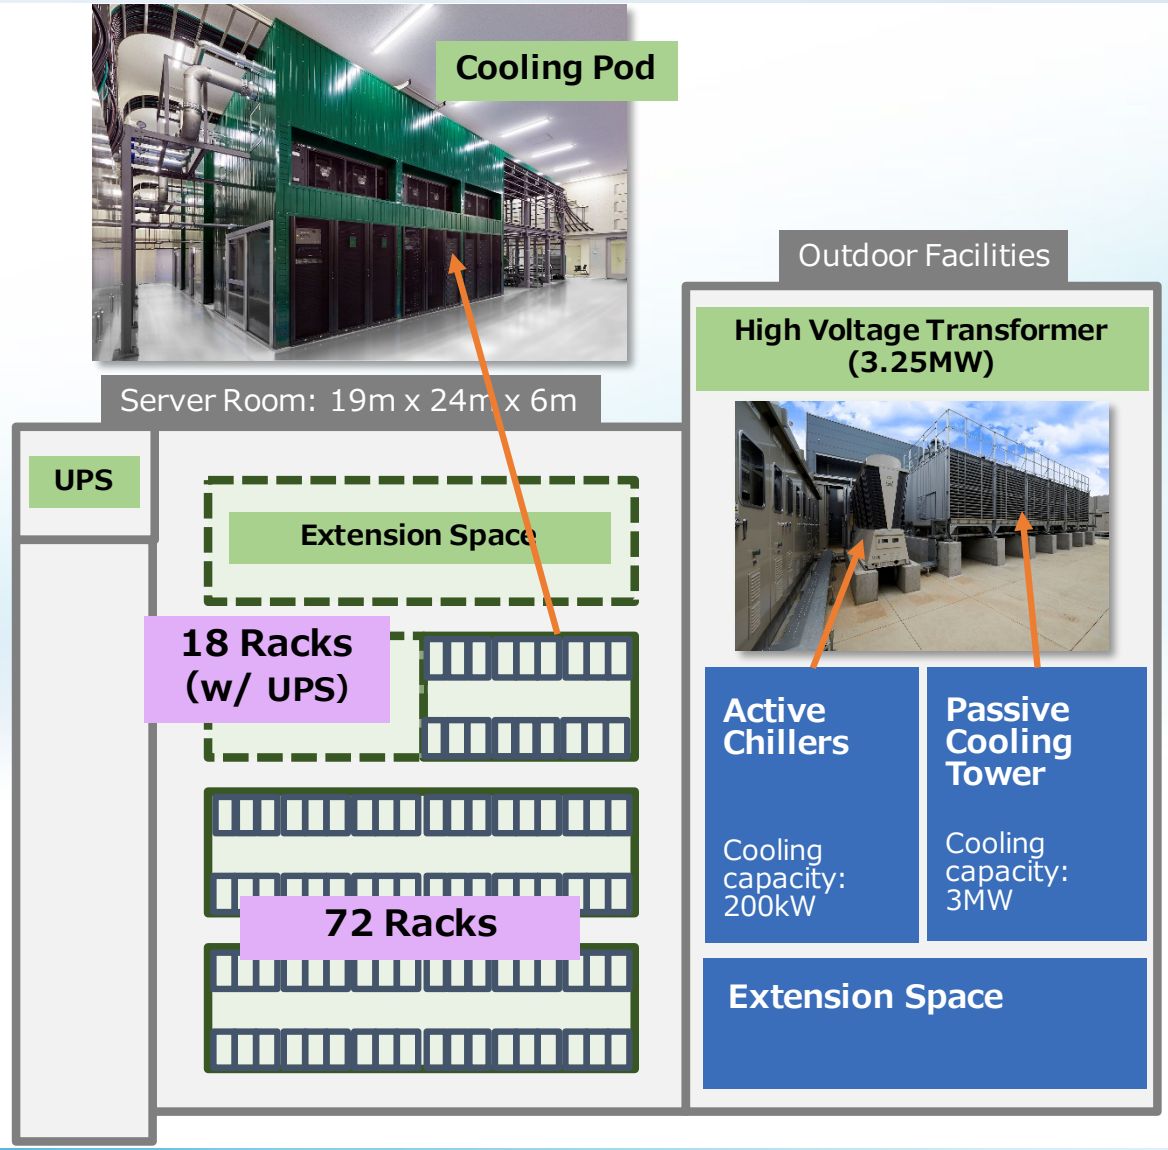

- **Single floor, cost effective building**
- **Hard concrete floor 2t/m² weight tolerance for racks and cooling pods**
- **Number of Racks**
  - Initial: 90 (ABCI uses 41 racks)
  - Max: 144
- **Power capacity: 3.25 MW**
  - ABCI uses 2.3MW max
- **Cooling capacity: 3.2MW**
  - 70kW/rack: 60kW water+10kW air

January 16, 2017

October 30, 2017

Jan 2018

Cooling Pod

- Modularization of DC facilities: rack space, water/air cooling and power equipment
- No raised floor and skeleton frame built on concrete slab
- Effective air cooling by hot aisle containment
- Ease of maintenance

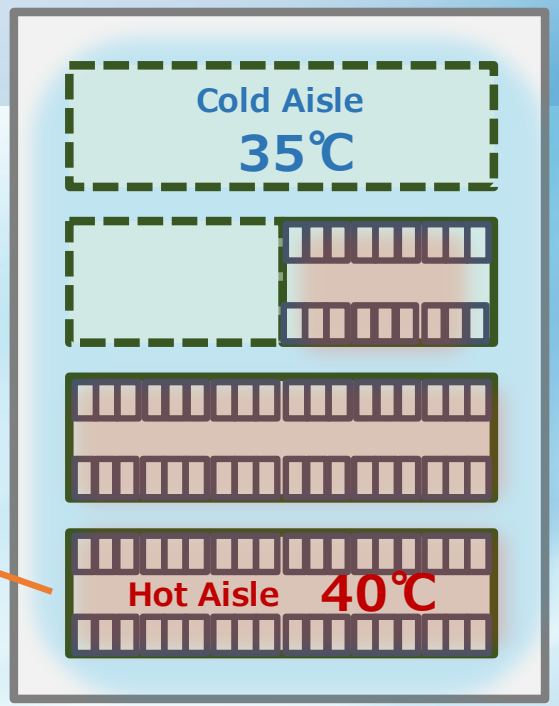

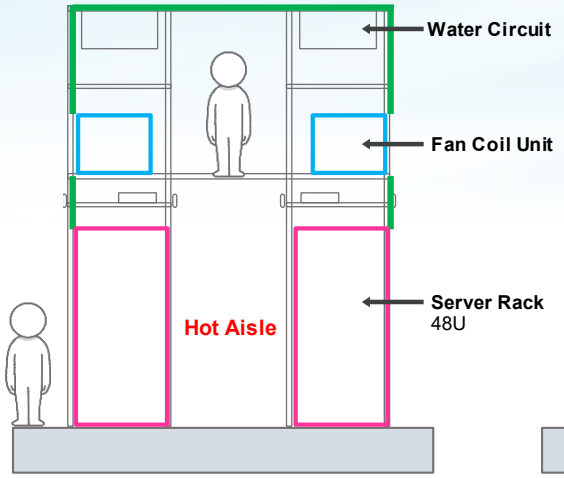

Front view

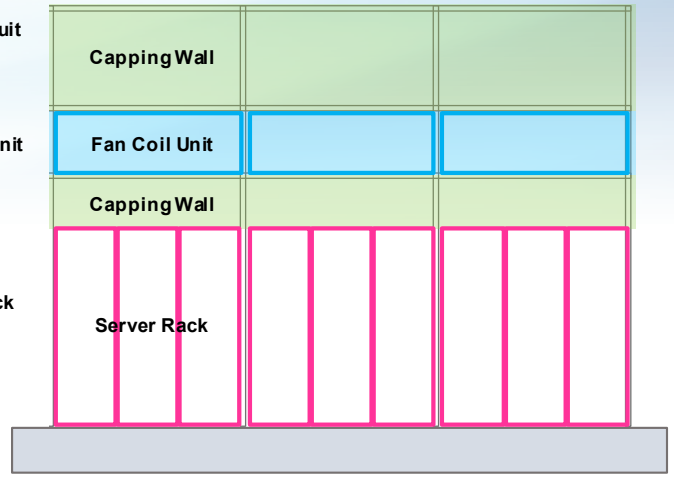

Side view

Inside Cooling Pod

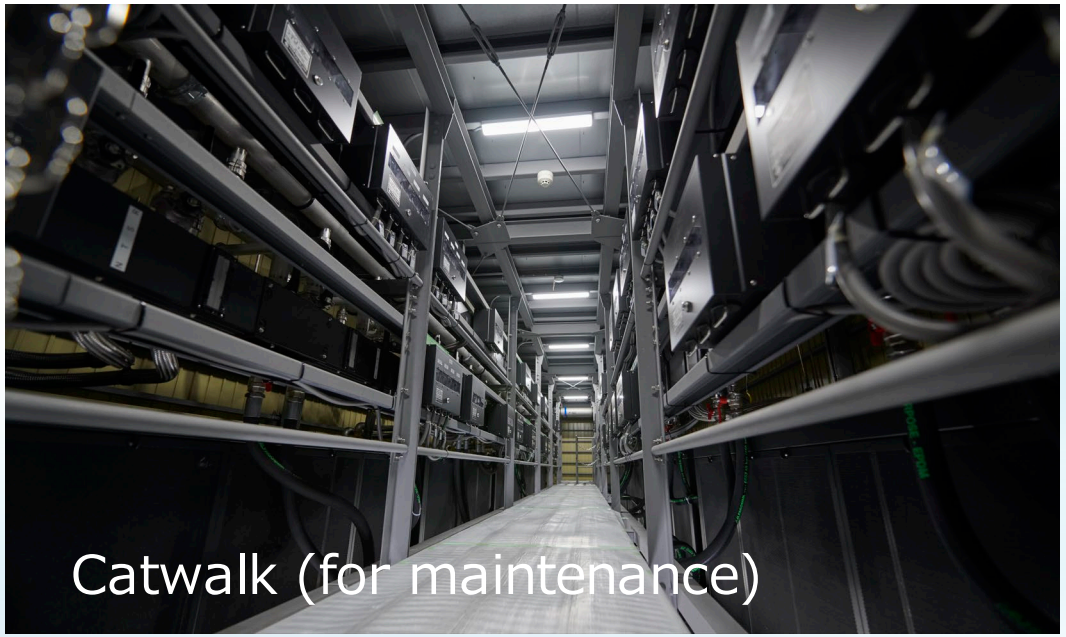

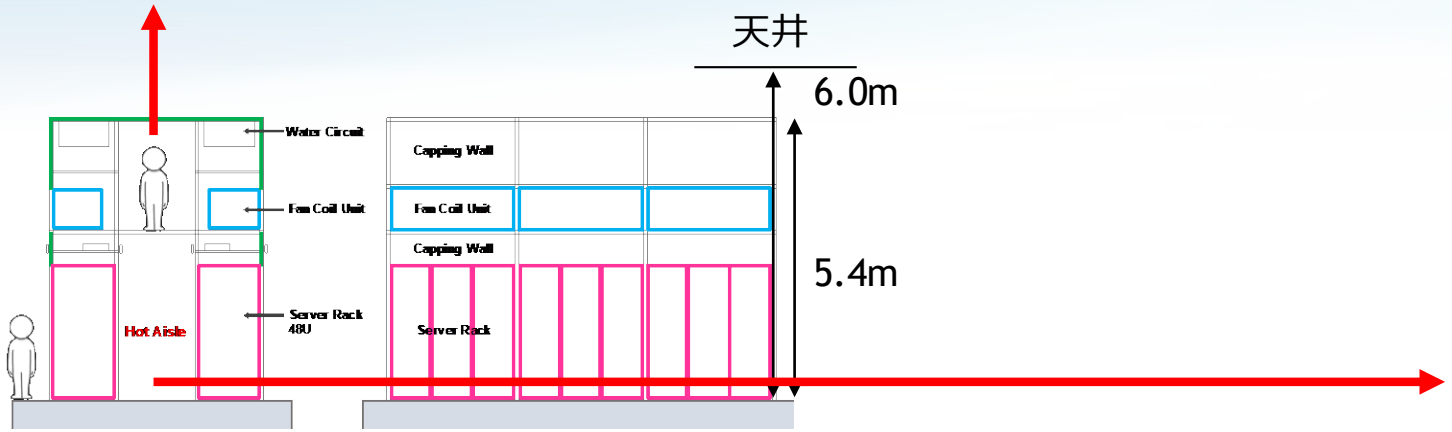

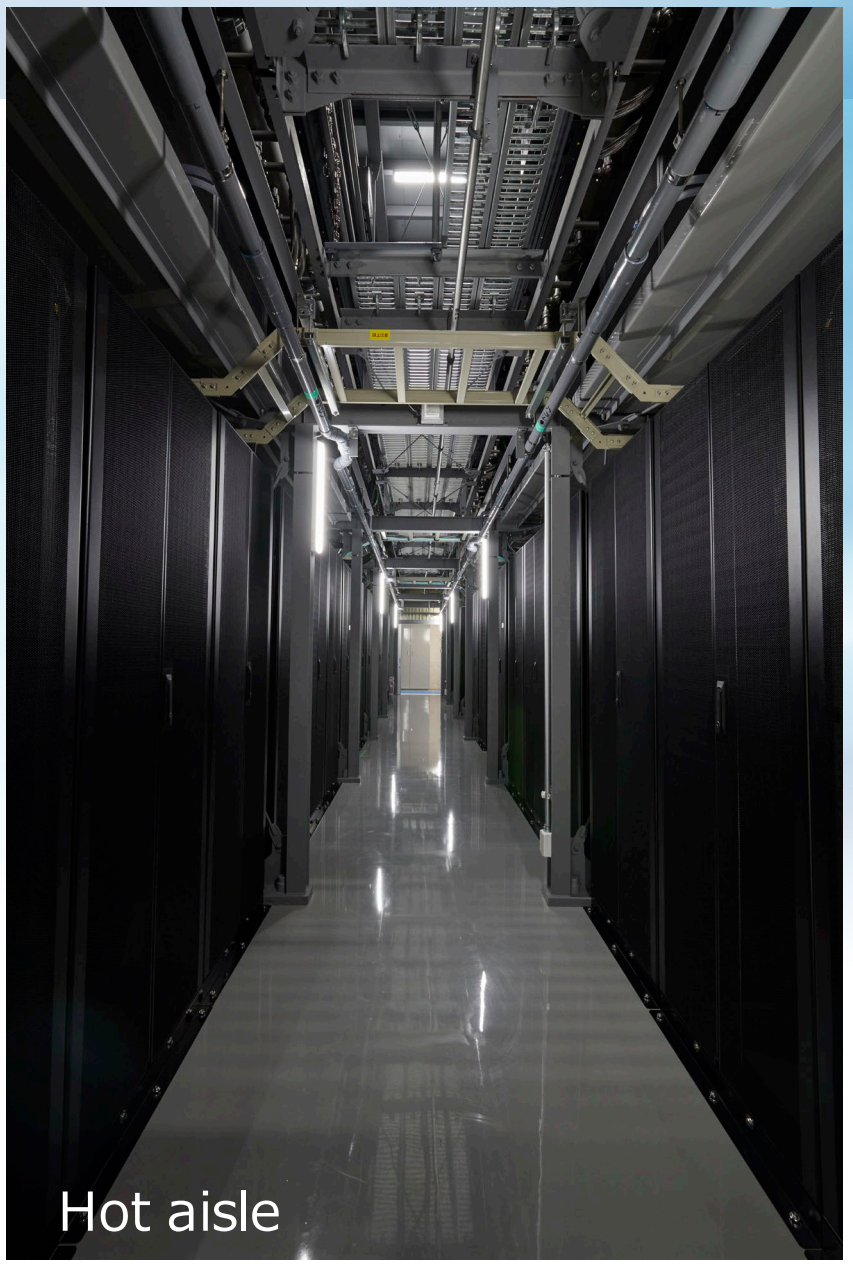

Jan 2018

Feb 2018

Apr 2018

Hybrid Water/Air Cooling System

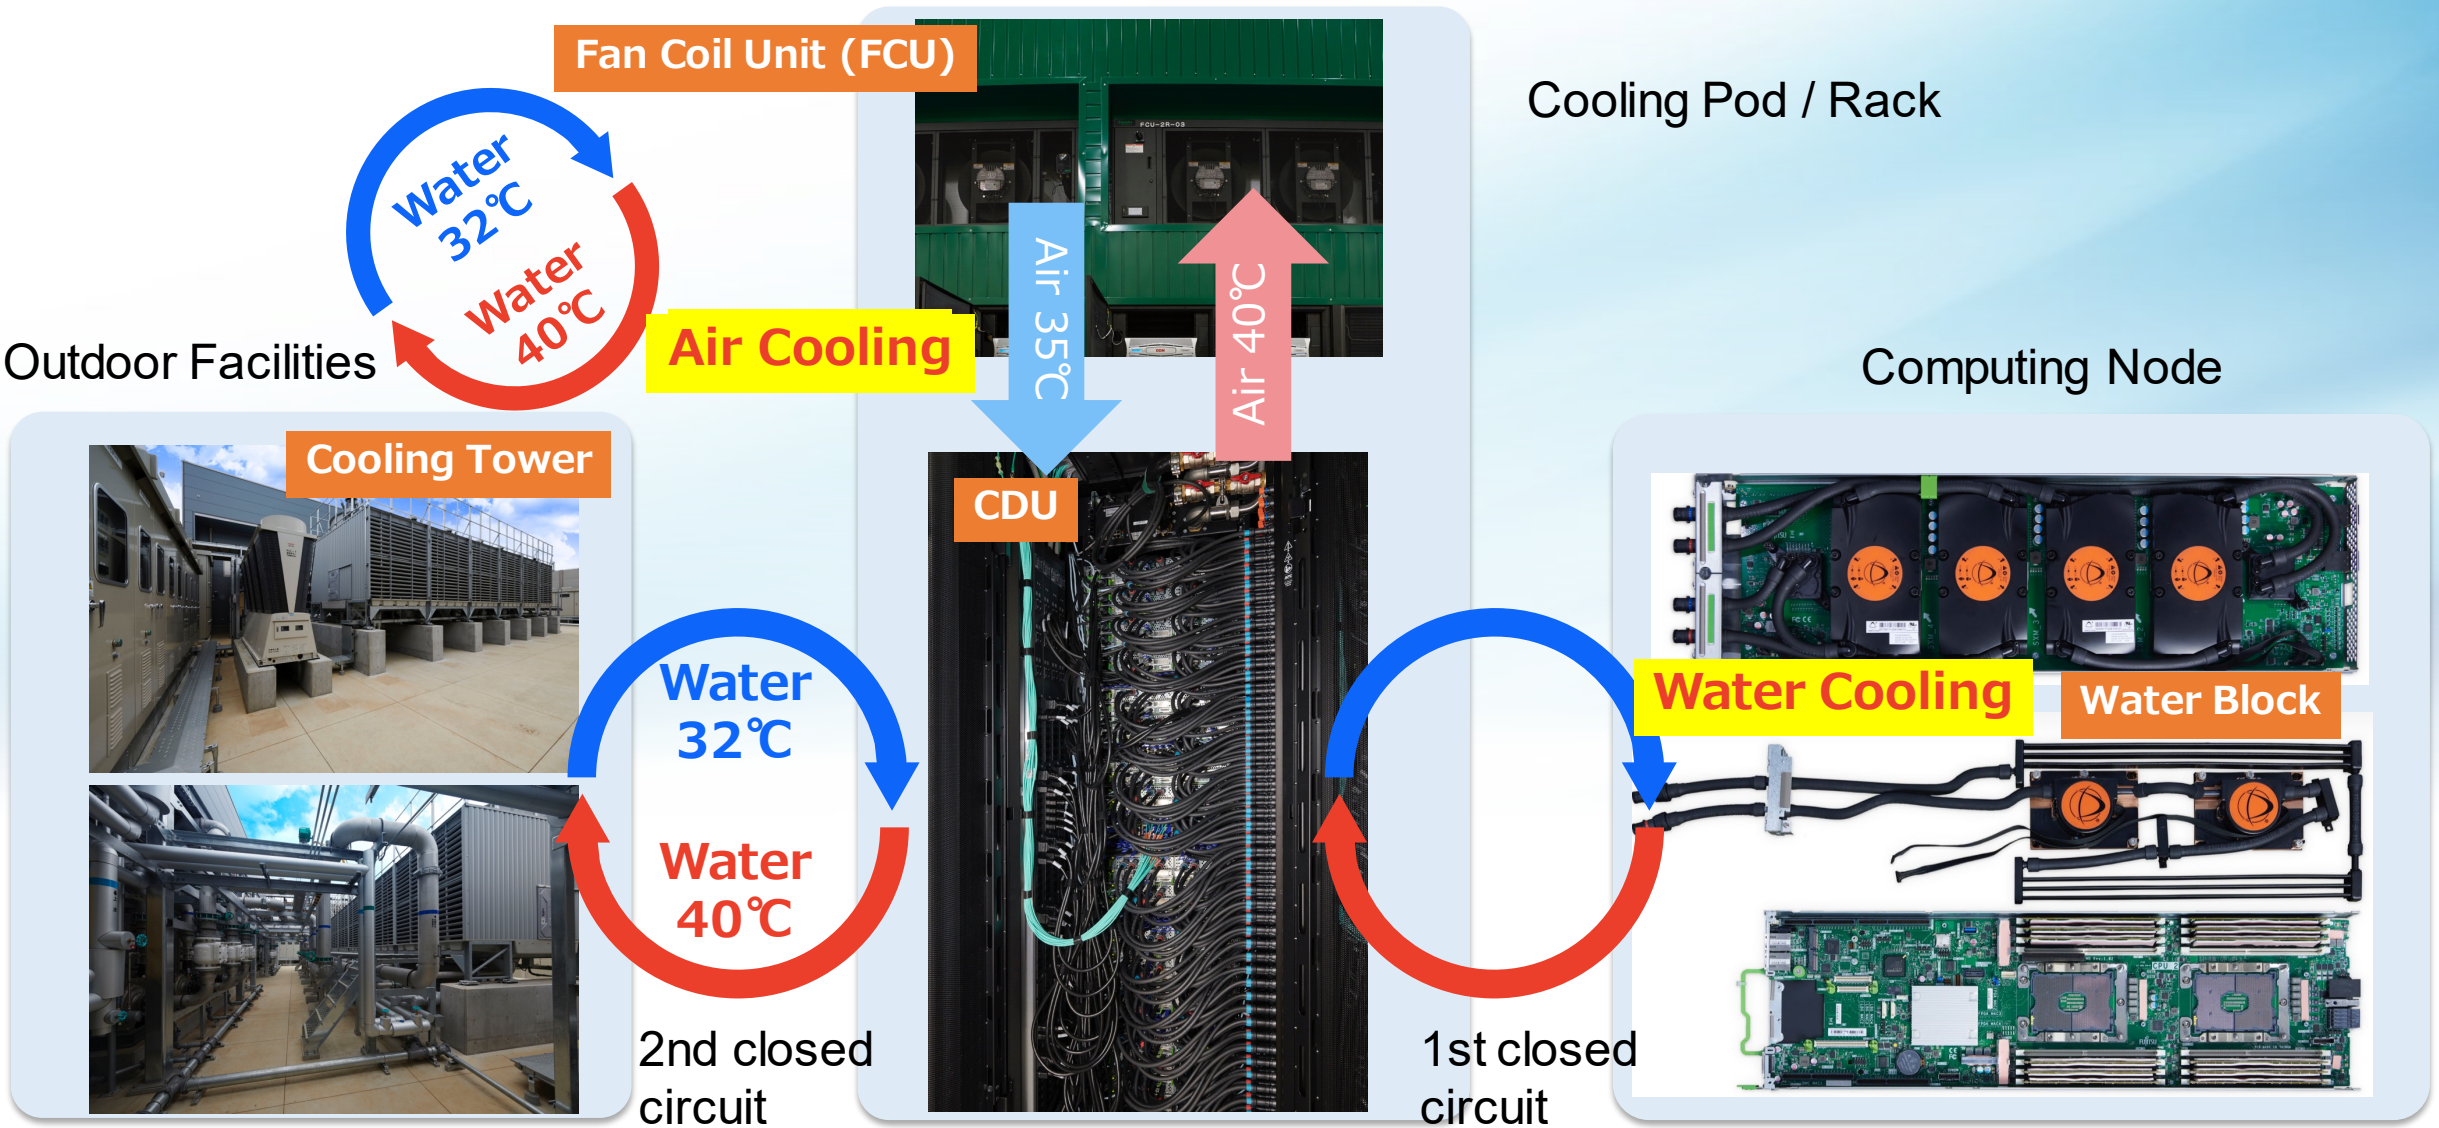

Survived the extreme heat this summer!

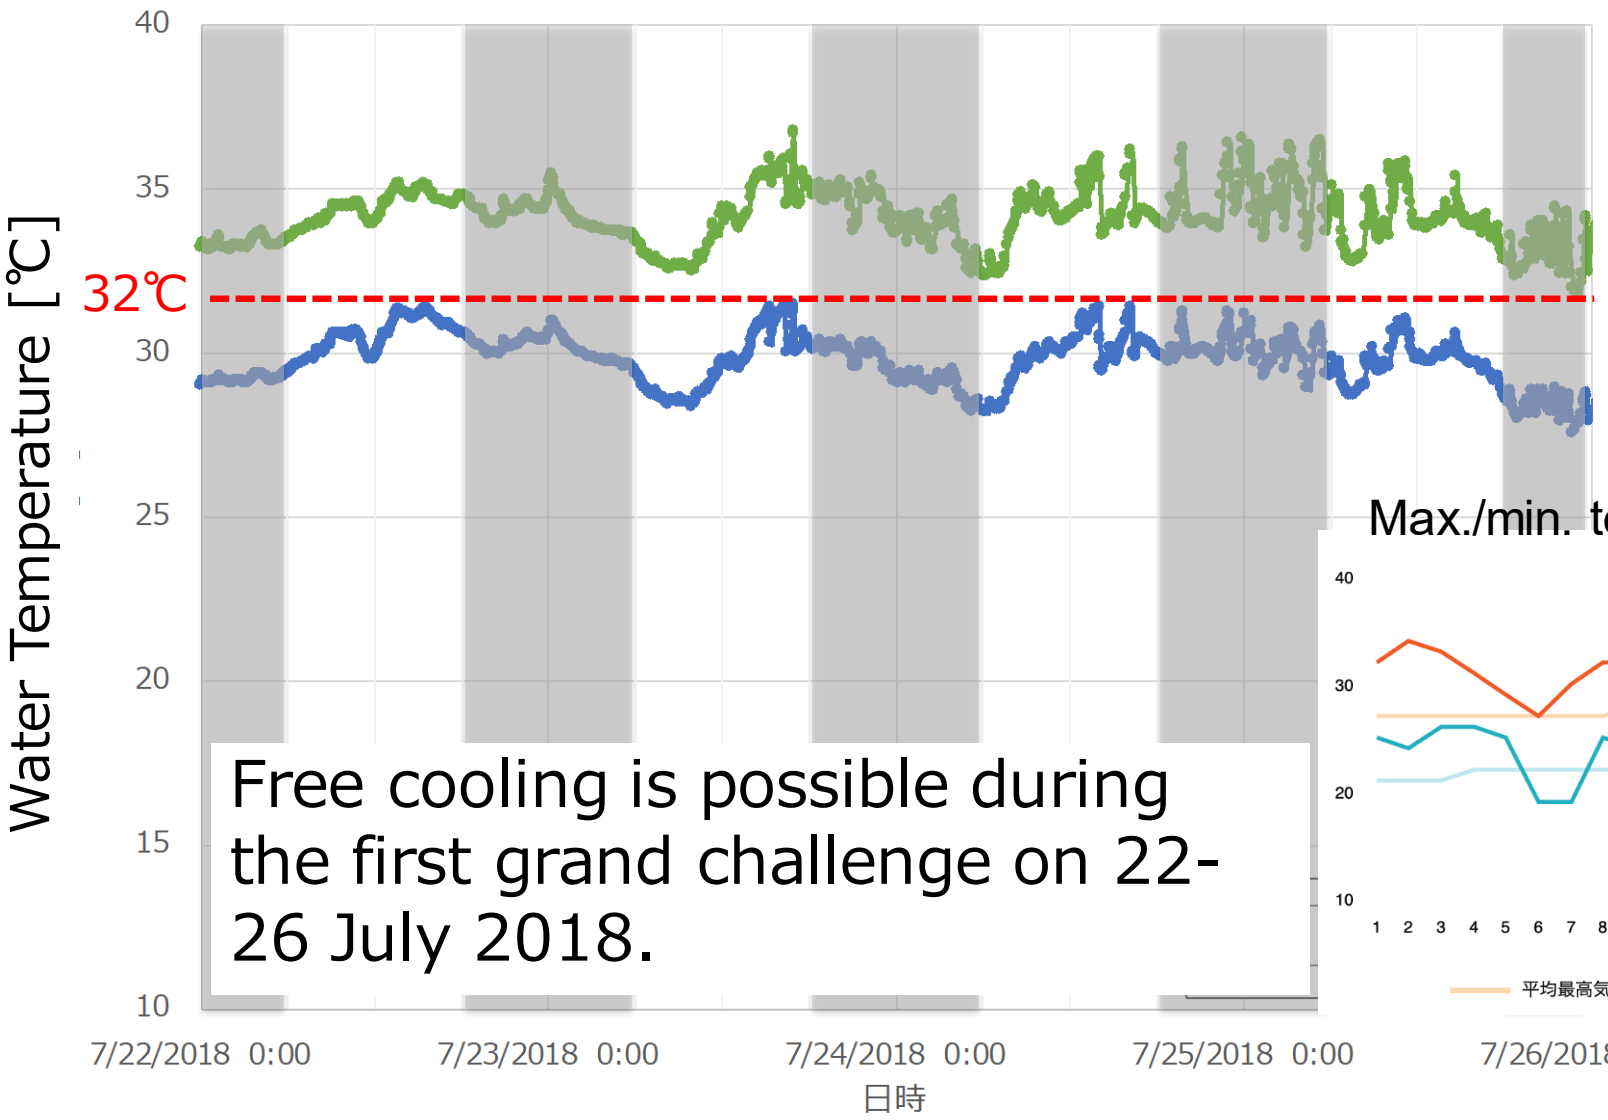

Cooling Water (From ABCI): <40℃

Cooling Water (To ABCI): <32℃

Free cooling is possible during the first grand challenge on 22-26 July 2018.

Max./min. temperature in Kashiwa in July 2018

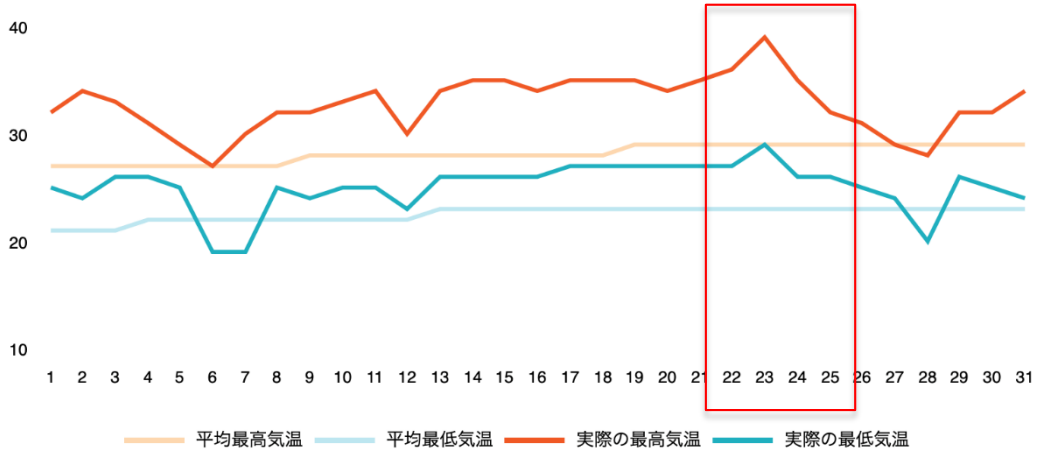

	AAIC（AIST AI Cloud）	TSUBAME3.0	ABCI (AI Bridging Cloud)	Summit	TPU 3.0 Pod
Institution	AIST	Tokyo Tech	**AIST**	ORNL	Google
Deployment	2017	2017	**2018**	2018/2019	2018
#Nodes	50	540	**1088**	4608	NA
Throughput AI Processor	NVIDIA Tesla P100	NVIDIA Tesla P100	**NVIDIA Tesla V100**	NVIDIA Tesla V100	TPU 3.0
#GPUs	400	2160	**4352**	27648	NA
Peak FP64	2.2 PF	12.2 PF	**37.2 PF**	200 PF	NA
Peak DL	8.6 PF	47.2 PF	**550 PF**	3.3 EF	100 PF / Pod
TOP500*	287	19	**5 (#1 Japan)**	1	NA
GREEN500*	7	6	**8**	5	NA
Nodes / Rack	6	36	**34**	16	NA
GPUs / Rack	48	144	**136**	96	NA
kW / Rack	22 kW	64.8 kW	**67.33 kW**	45-55 kW (est.)	NA
AI-FLOPS / rack	0.9 PF	3.1 PF	**17 PF**	12 PF	12.5 PF (100 PF /8 racks)

(As of June 2018)

TOP500 List (June 2018)

#	System	Architecture	Country	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Summit, ORNL	IBM, POWER9 + GPU(GV100)	USA	122,300.0	187,659.3	8,806
2	TaihuLight, NSCW	Sunway, SW26010	China	93,014.6	125,435.9	15,371
3	Sierra, LLNL	IBM POWER9 + GPU(GV100)	USA	71,610.0	119,193.6	-
4	Tianhe-2, NSCG	NUDT, CPU + Matrix-2000	China	61,444.5	100,678.7	18,482
5	ABCI, AIST	Fujitsu, CPU + GPU(V100)	Japan	19,880.0	32,576.6	1,649
6	Piz Daint, CSCS	Cray, CPU + GPU(P100)	Switzerland	19,590.0	25,326.3	2,272
7	Titan, ORNL	Cray, CPU + GPU(K20x)	USA	17,590.0	27,112.5	8,209
8	Sequoia, LLNL	IBM, BlueGene/Q	USA	17,173.2	20,132.7	7,890
9	Trinity, LANL/SNL	Cray, MIC(KNL)	USA	14,137.3	43,902.6	3,844
10	Cori, NERSC-LBNL	Cray, MIC(KNL)	USA	14,014.7	27,880.7	3,939

Comparing ABCI to Classical IDC
AI IDC CAPX/OPEX accelerartion by > x100

Perf > x400~600

Power Eff > x200~300

Traditional Xeon IDC

~10KW/rack PUE 1.5~2
15~20 1U Xeon Servers
2 Tera AI-FLOPS(SFP) / server
30~40 Tera AI-FLOP / rack
Low cooling efficiency

ABCI "Open Source" IDC

~70KW/rack PUE 1.0x
~500 Tera AI-FLOPS(HFP) / server
~17 Peta AI-FLOPs / rack
Inexpensive, very high cooling efficiency (PUE~1.1)

Fusion Recurrent Neural Network (FRNN)

Alexey Svyatkovskiy, Julian Kates-Harbeck, Bill Tang

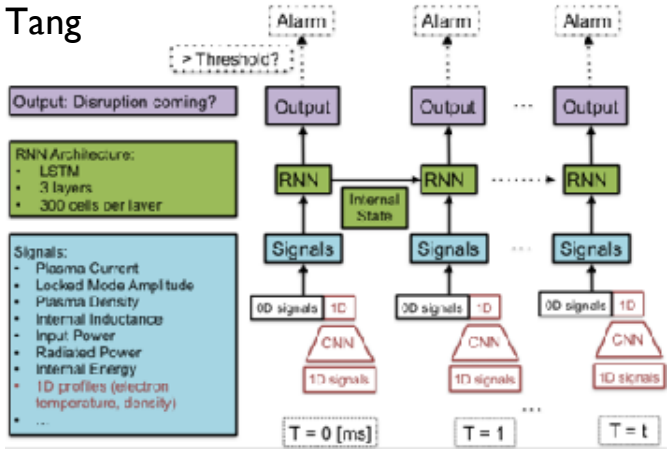

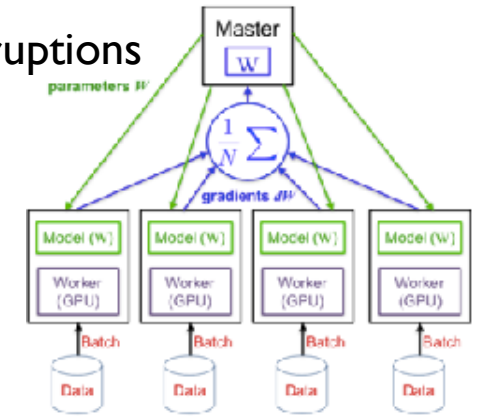

Objective

Use machine learning to predict disruptions in ITER class fusion reactors

Methods

3 layer LSTM in Tensorflow
300 cells per layer

Contributions

Prediction 30ms before disruption is 96% which is much better than SVM and RF

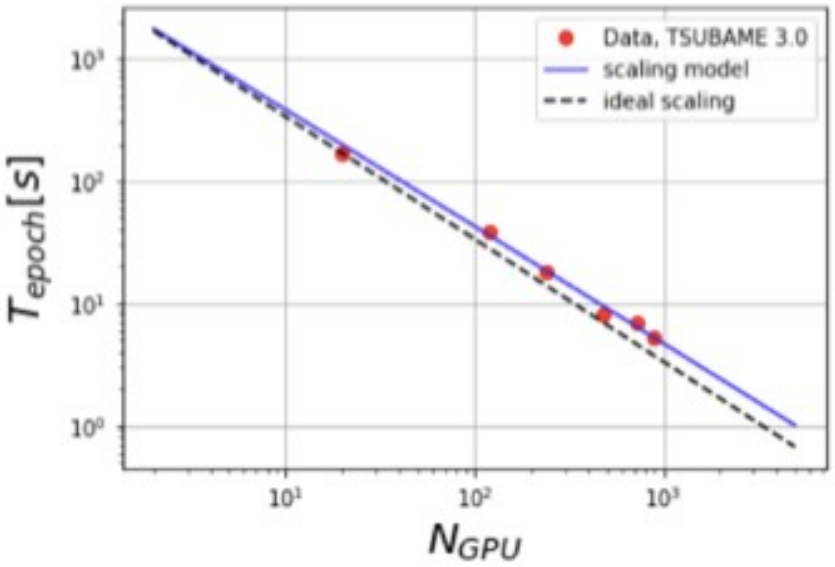

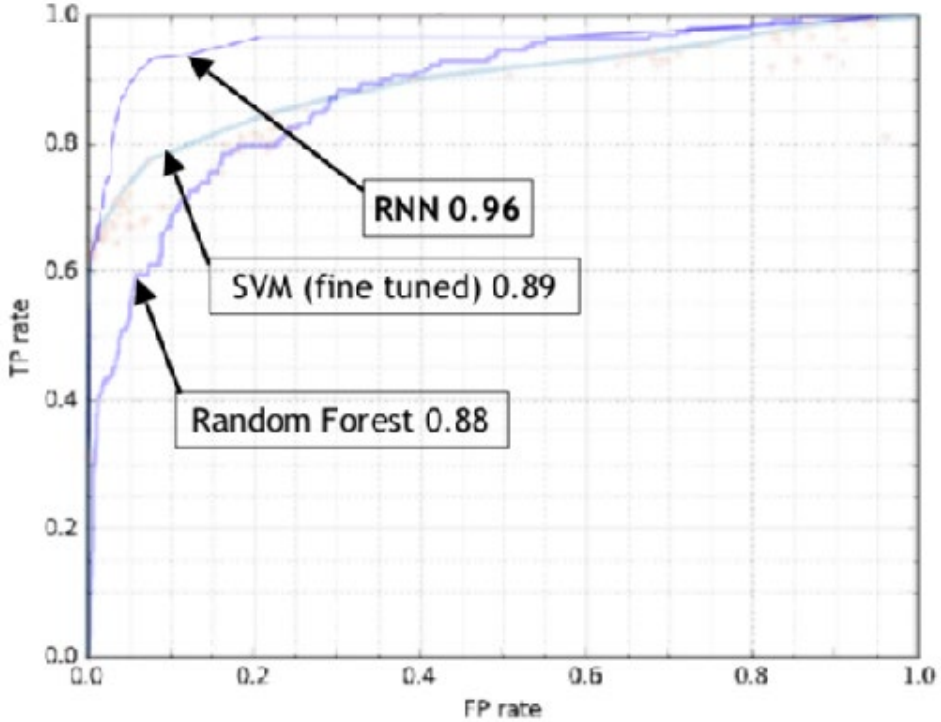

ABCI-PRP: Grand Challenge Project

- Gerald Pao, Salk Institute; George Sugihara, SIO
- Creation of neuromorphic deep learning architectures by large-scale dynamic modeling of transparent fish brains on ABCI

ABCI is ideally suited to perform the convergent cross mapping (CCM) to interrogate the relationships between the ~120,000 neurons of the larval zebrafish brain during exposure to various stimuli

Check Wassapon's Poster (PP25)!

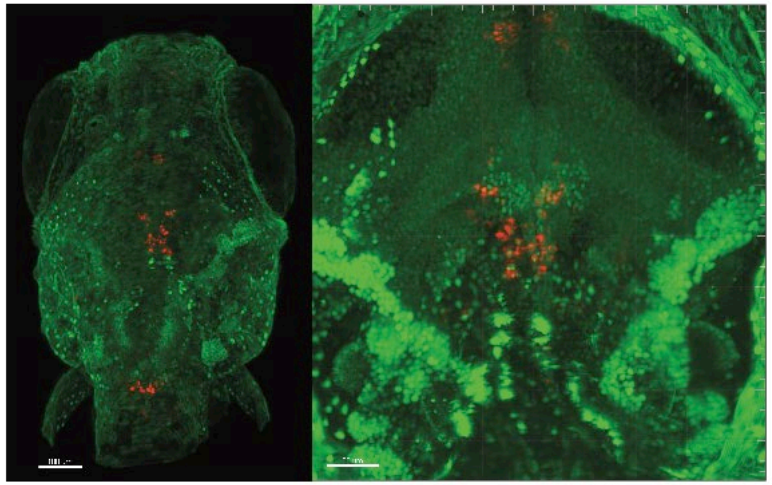

From Pao, 2018. Green are neurons, red are active neurons responding to hypoxic environment.

Training ImageNet in Minutes

Rio Yokota, Kazuki Osawa, Yohei Tsuji, Yuichiro Ueno, Hiroki Naganuma, Shun Iwase, Kaku Linsho
Tokyo Institute of Technology
+ Akira Naruse (NVIDIA)

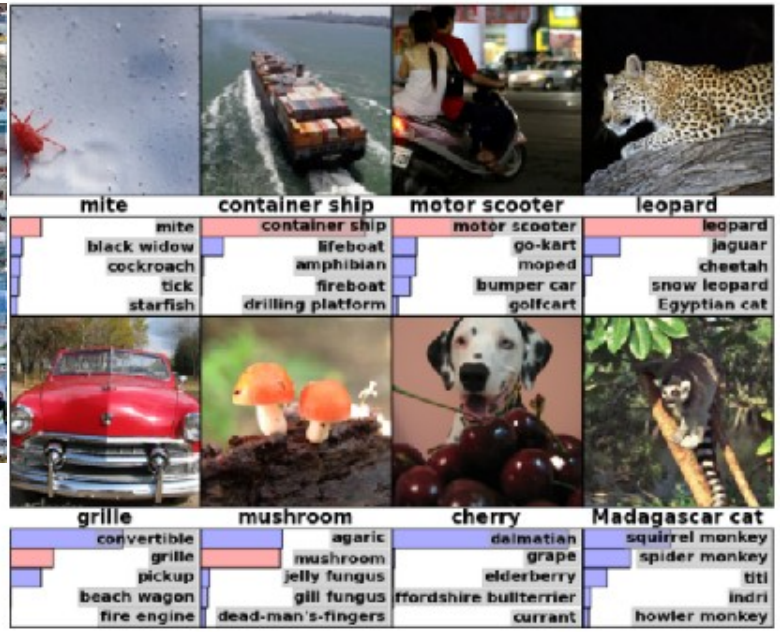

TSUBAME3.0

ABCI

	GPU/KNL	Time
Facebook	512	30 min
Preferred Networks	1024	15 min
UC Berkeley	2048	14 min
Tencent	2048	6.6 min
Tokyo Tech + AIST	4096	? min

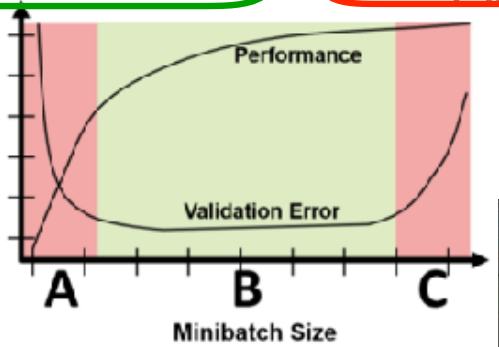

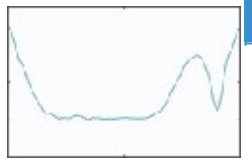

Source Ben-nun & Hoefler https://arxiv.org/pdf/1802.09941.pdf

ImageNet/ResNet-50 Training in XXX Seconds (world record)

Yoshiki Tanaka, Hisahiro Suganuma, Hiroaki Mikami, Pongsakorn U-chupala, Yuichi Kageyama Sony Corporation

Successfully trained ImageNet/ResNet-50 in **XXX seconds** with Neural Network Libraries developed by Sony without significant accuracy loss on ABCI

Neural Network Libraries

Issues of Large-scale Distributed DNN training with a massive GPU cluster

Issue 1

Communication overhead of gradient synchronization among GPUs

Approach

2D-Torus topology comprises of multiple rings in horizontal and vertical orientations

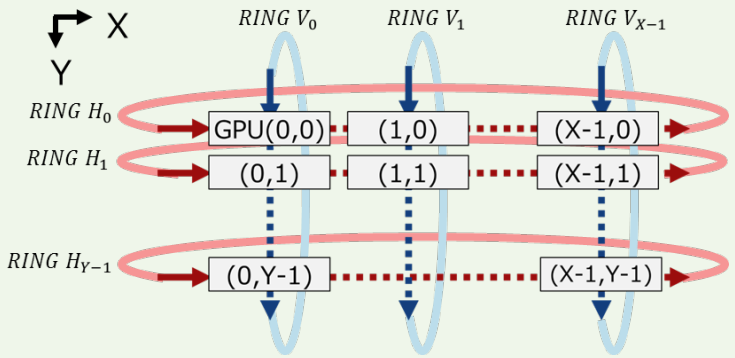

Issue 2

Accuracy loss with large mini-batch training

Approach

Increase batch size in response to the progress on DNN training

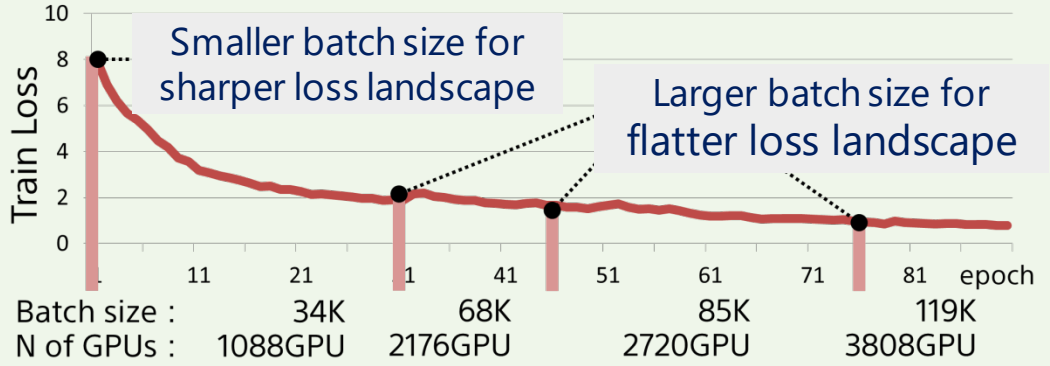

Apr 1 2018 Became Director of Riken-CCS:
Science, of Computing, by Computing, and for Computing

Riken Center for Computational Science（R-CCS）
World Leading HPC Research, active collaborations w/Universities, national labs, & Industry

Sci. of Computing

Foundational research on computing in high performance for K, Post-K, and beyond towards the "Post-Moore" era, including future high performance architectures, new computing and programming models, system software, large scale systems modeling, big data analytics, and scalable artificial intelligence / machine learning

Sci. by Computing

Breakthrough Science & Technology using high performance computing capabilities of K, Post-K and beyond to address the issues of high public concern, in areas such as life sciences, climate & environment, disaster prediction & prevention, advanced manufacturing, applications of machine learning for Society 5.0.

High Resolution, High Fidelity Analysis & Simulation

Mutual Synergy

Novel Future High Performance Computing Architectures & Algorithms

Sci. for Computing

New Materials & Electronic Devices e.g., Photonics, Neuromorphics, Quantum, Reconfigurable

The K computer & the facility

- **The K computer:**
  - Computing Capability 10.51 PetaFlops (HPL score)
  - Power Consumption 12-15MW
  - Water/Air cooling ratio 70 : 30
  - Operation start 28 September, 2012
- **Facility:**
  - Gas Turbine Power Generator 5MW x 2
  - Chiller(absorption) 1700USRt(5.98MW) x 2
  - Chiller(Centrifugal) 1400USRt(4.93MW) x 2, 700USRt(2.46MW) x 2
  - Air handlers, voltage transformers, etc.

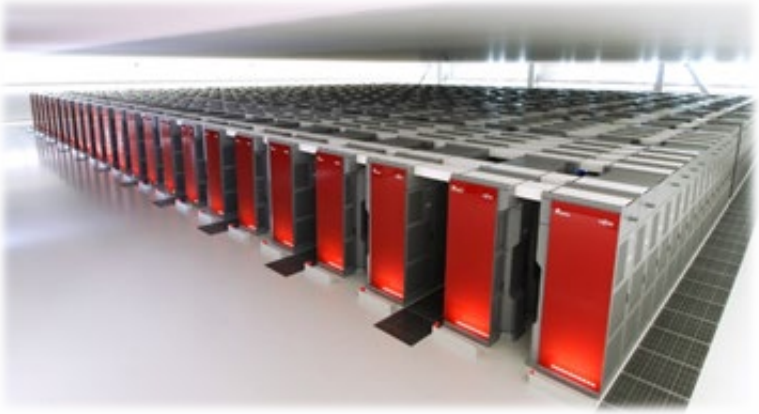

Electric power source balance
Provider: Generator = 70 : 30

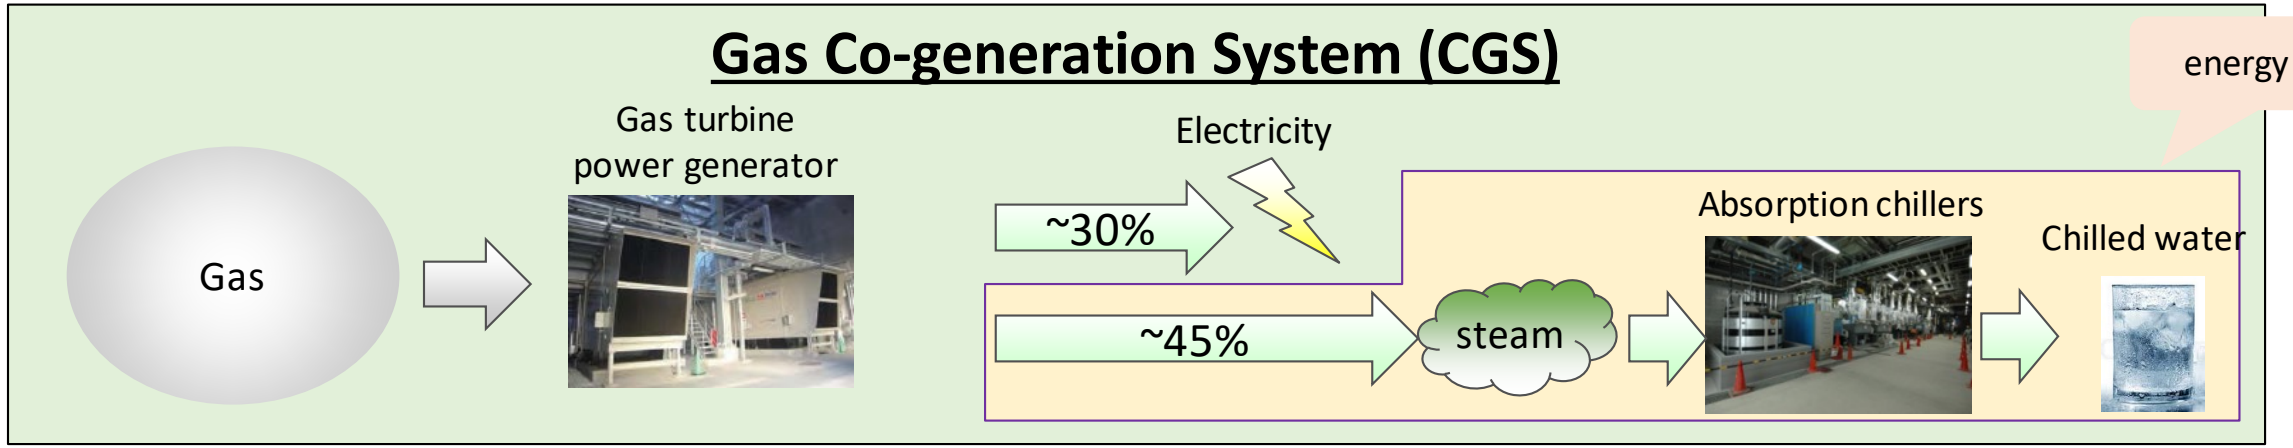

Efforts for better contract

- **Electric power**
  - A private contract with a regional power supply company until FY2016.
  - A competitor appeared for FY2017 supply contract renewal, then we introduced competing bidding.
  - ---> The challenger beat the champion then the electric power cost decreased by 10% and no minimal consumption constraint is included.
  - ---> No quality degradation observed so far.

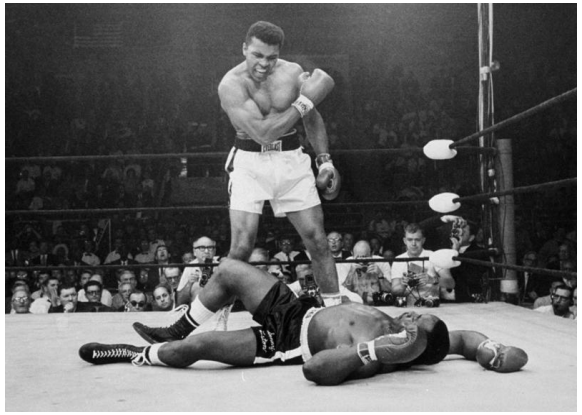

- **Gas**
  - A contract based on a floating gas rate until Jun.2015.
  - A fixed rate service has started since Aug.2015.
    - In our case, a fixed rate, which is determined 6 month ago, is applied for 1 year.
    - ---> Consequently, the fixed rate level was higher than the floating for 2 years, but the trend was reversed recently.
    - ---> If the fixed rate contract needs more cost, it is beneficial for us because no shortage/residue is guaranteed.

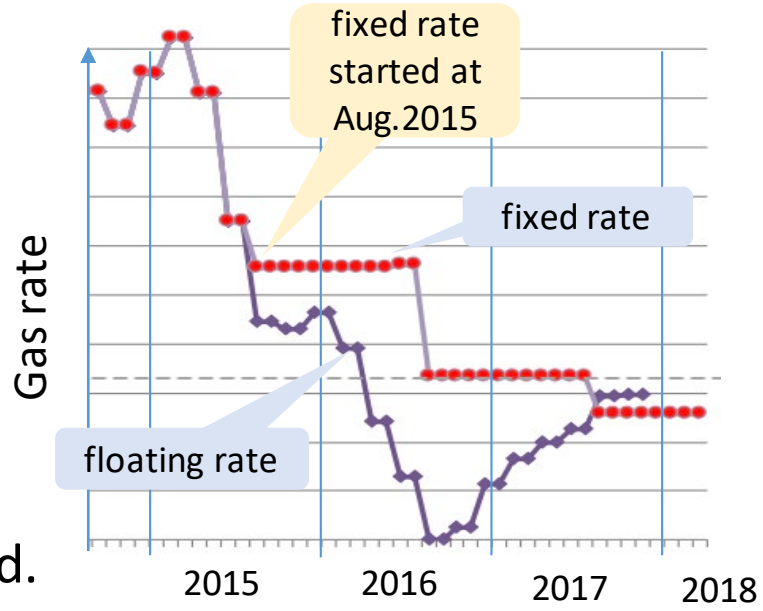

Lessons learned for CGS operation

1. Generation efficiency of gas turbine power generator increases when it is driven at higher output level.
   - ex. when generator is driven at 50% output level, efficiency decreases by 30%.
2. When power generator is driven at higher level, richer steam is also generated then power consumption of chillers can be saved.
   - ex. when generator is driven at 50% output level, power consumption of chillers increases by 30%.

Generation level vs efficiency

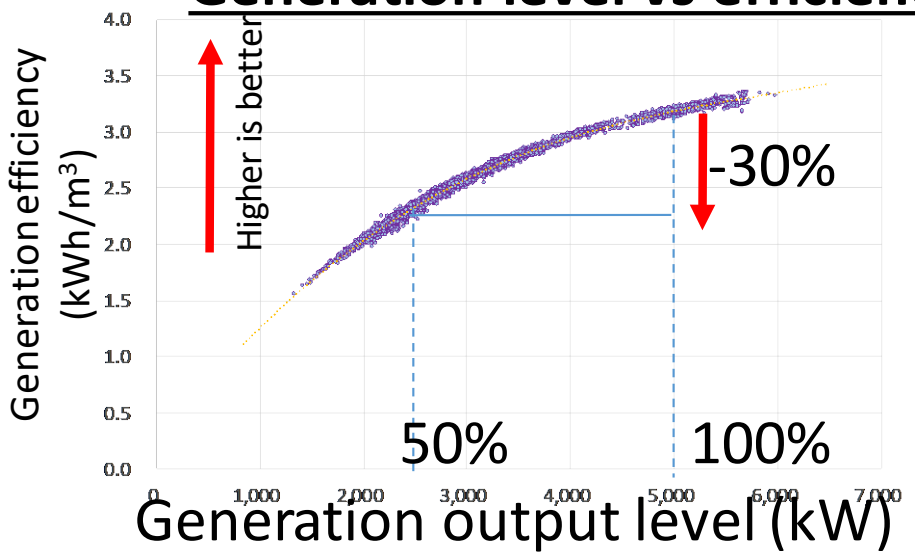

Generation level vs power consumption of c

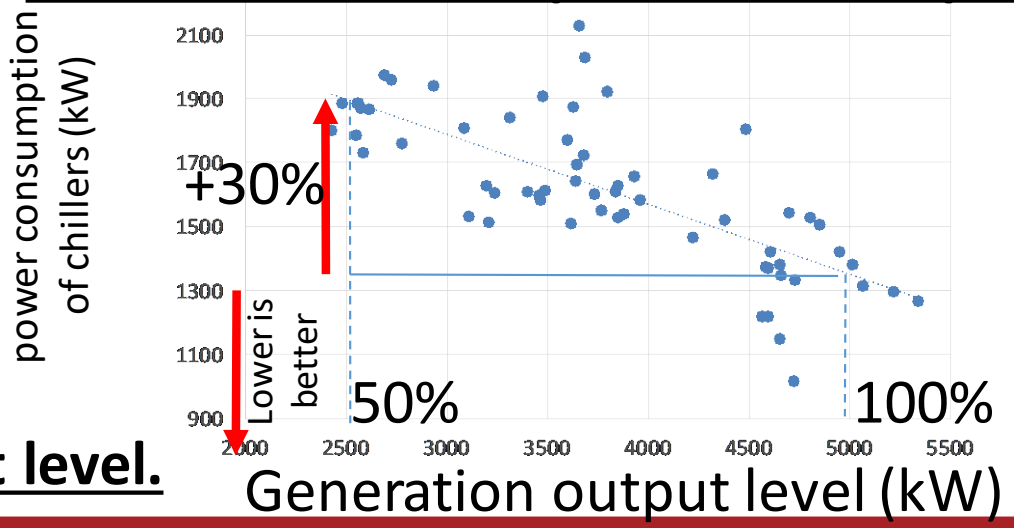

Power generator should be driven at higher output level.

CGS operation improvement

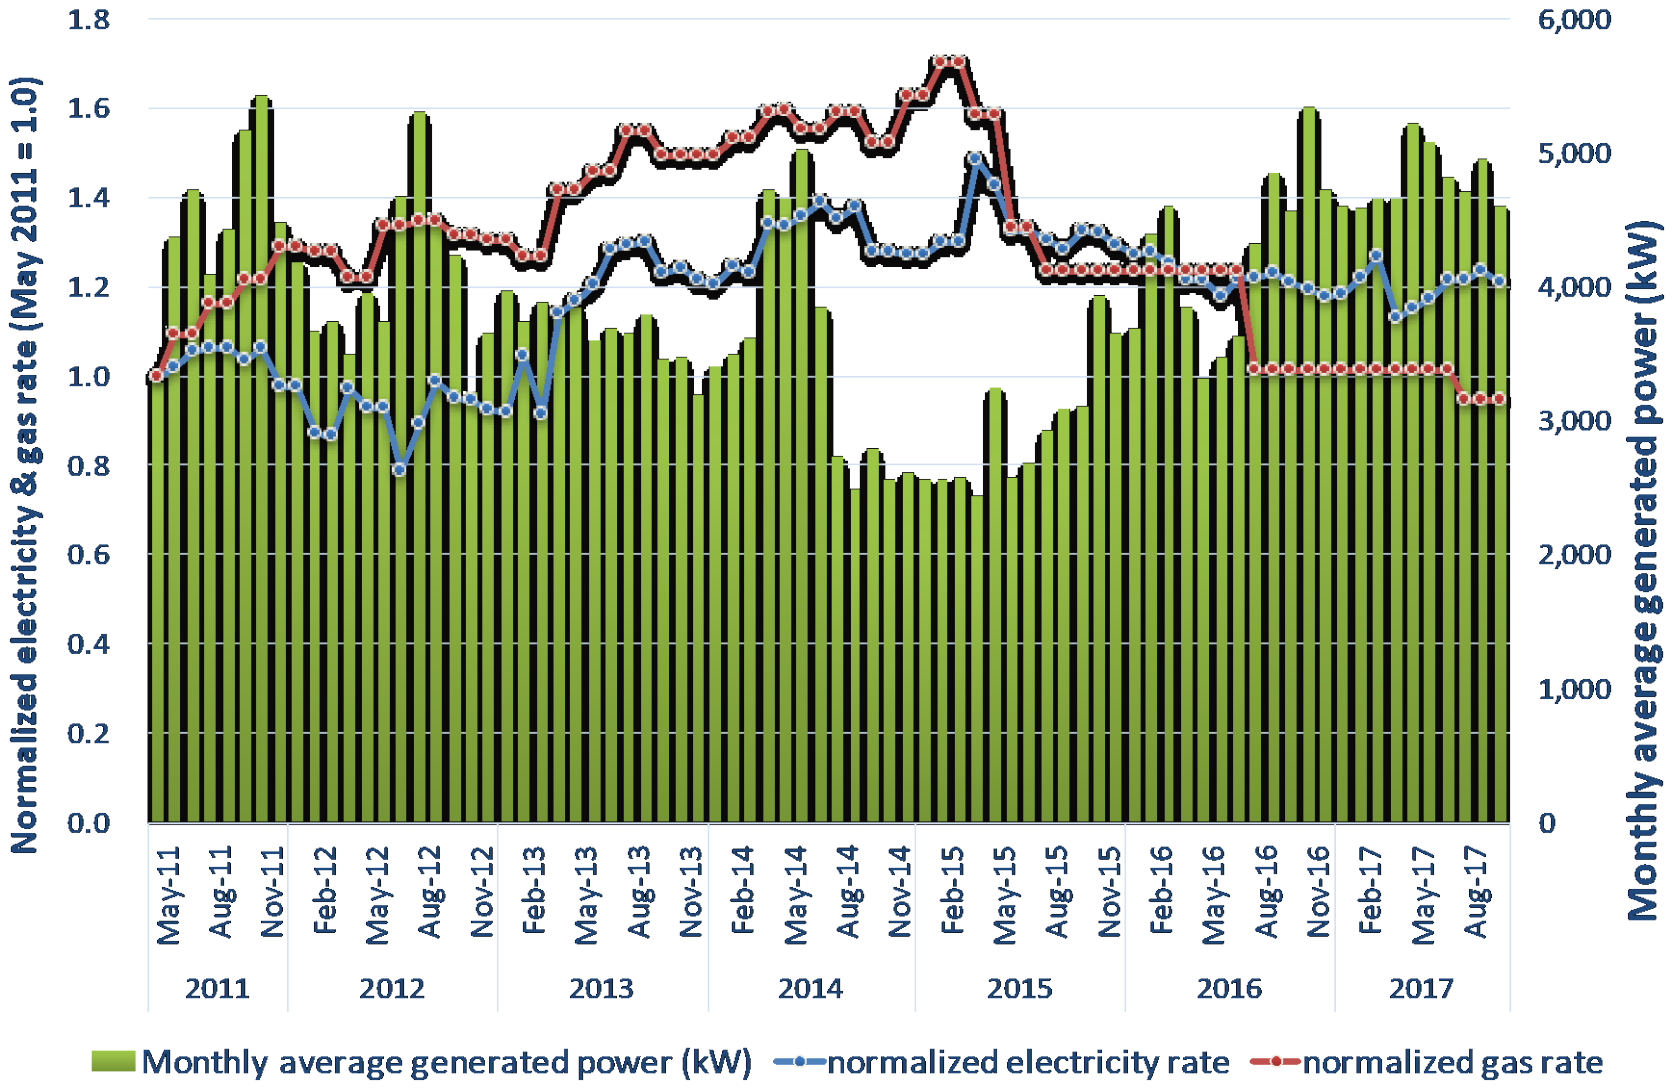

- In actual operation, optimal operation could not be achieved for some periods due to higher gas rate or constraint from contract (ex. minimal power consumption)
- New power supply contract including no minimal power consumption and gas supply contract with fixed rate also help us to improve energy efficiency.

Power consumption and PUE

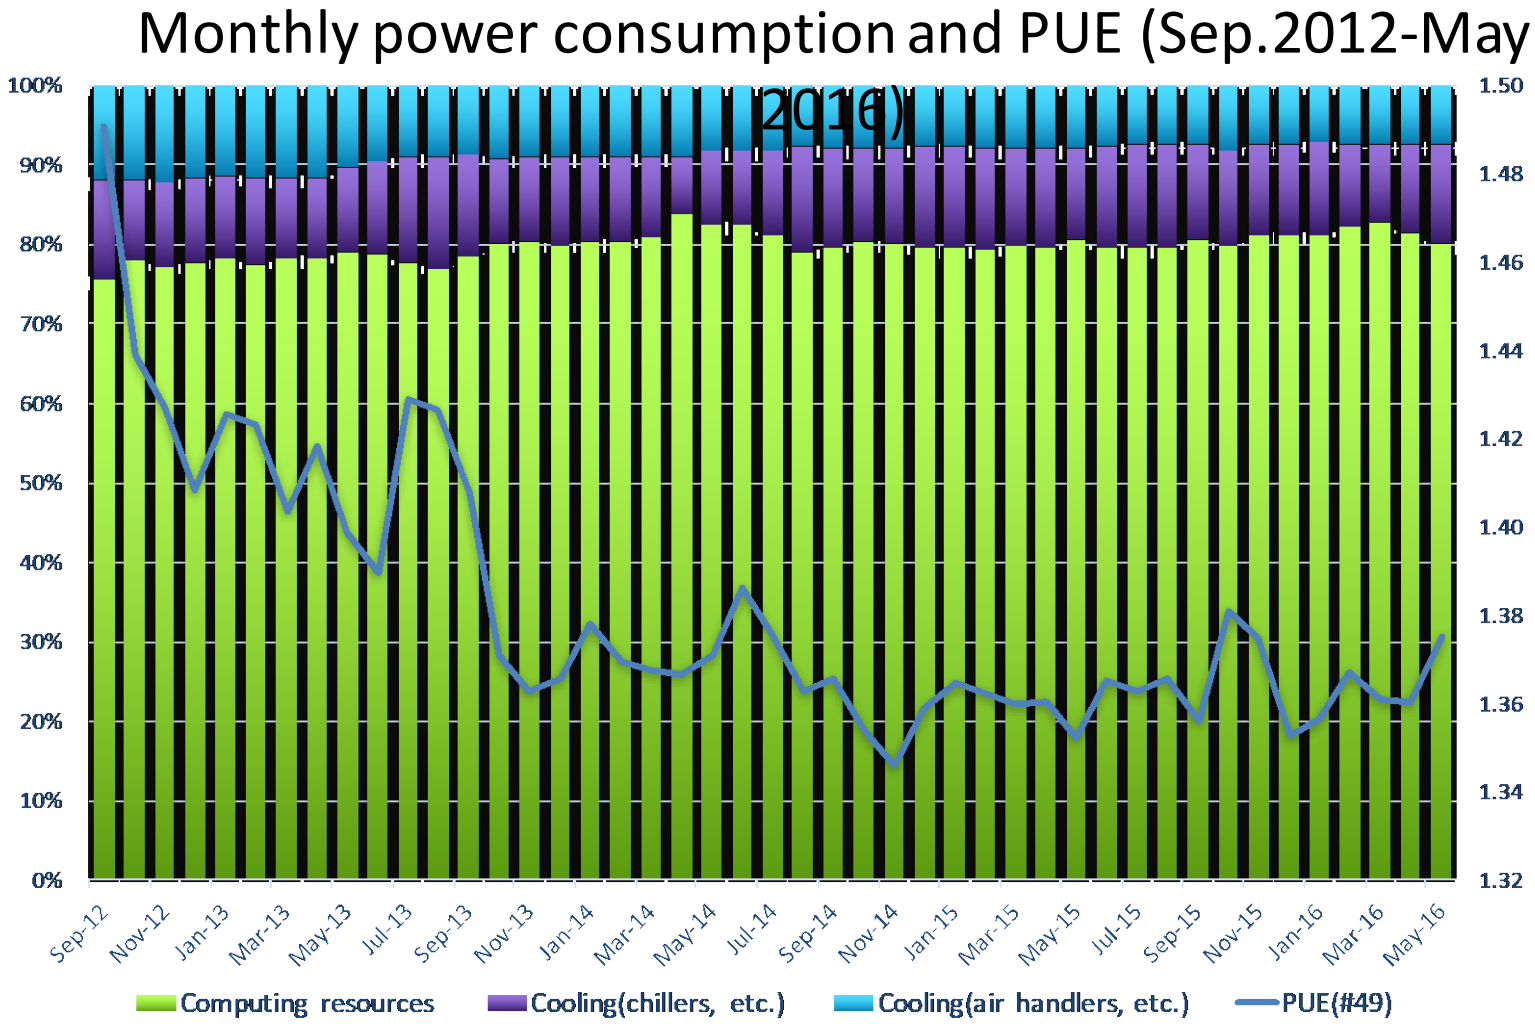

Reduction of power consumption of air handlers contributed to decrease PUE.

Japan Flagship 2020 "Post K" Supercomputer

- ✓CPU
  - Many core, Xeon-Class ARM v8 cores + 512 bit SVE (scalable vector extensions)
  - Multi-hundred petaflops peak total
  - Power Knob feature
- ✓Memory
  - ✓3-D stacked DRAM, Terabyte/s BW /chip
- ✓Interconnect
  - TOFU-D CPU-integrated 6-D torus network
- I/O acceleration with massive SDs
- 30MW+ Power, liquid cooled
- Riken co-design with **Fujitsu**
- **? Million cores in system**

Prime Minister Abe visiting K Computer 2013

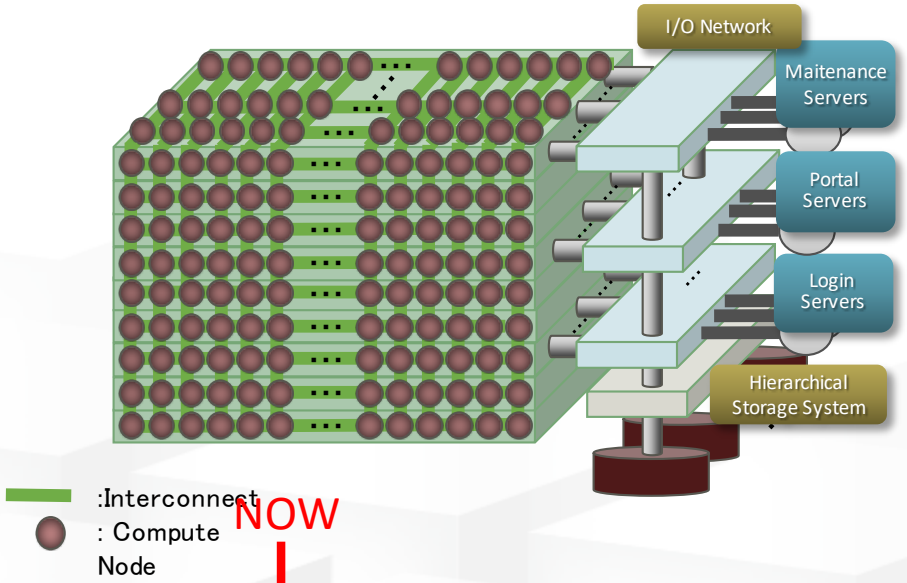

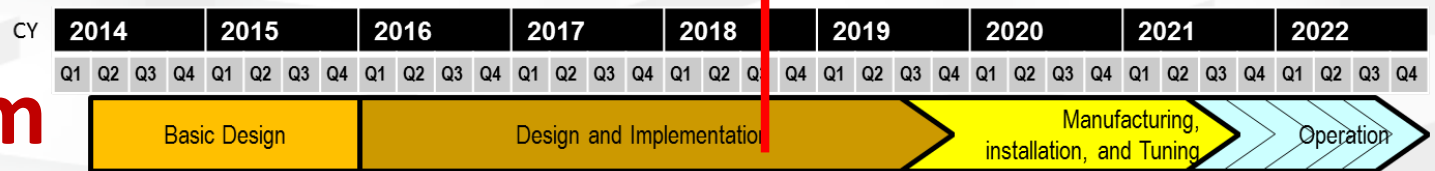

Post-K: The Game Changer

1. **Heritage of the K-Computer, HP in simulation via extensive Co-Design**

- High performance: up to x100 performance of K in real applications
- Multitudes of Scientific Breakthroughs via Post-K application programs
- Simultaneous high performance and ease-of-programming

2. New Technology Innovations of Post-K

- **High Performance, esp. via high memory BW**
  Performance boost by "factors" c.f. mainstream CPUs in many HPC & Society5.0 apps
- **Very Green e.g. extreme power efficiency**
  Ultra Power efficient design & various power control knobs
- **Arm Global Ecosystem & SVE contribution**
  Top CPU in ARM Ecosystem of 21 billion chips/year, SVE co-design and world's first implementation by Fujitsu
- **High Perf. on Society5.0 apps incl. AI**
  Architectural features for high perf on Society 5.0 apps based on Big Data, AI/ML, CAE/EDA, Blockchain security, etc.

Global leadership not just in the machine & apps, but as cutting edge IT

ARM: Massive ecosystem from embedded to HPC

Technology not just limited to Post-K, but into societal IT infrastructures e.g. Clouds

From K to Post K

	K computer	Post K
Voltage	3-phase AC 200V	->
Peak Power	15MW	30-40MW
Idol Power	~10MW	>10MW
Fluctuation width	~5MW	10~18MW
Rack weight	~1t	~2t
Cooling ratio(water:air)	65:35	90:10
Water supply	44l/rack/min	117l/rack/min

Power consumption fluctuation on K

Typical power consumption history (4/14-15, 2016)

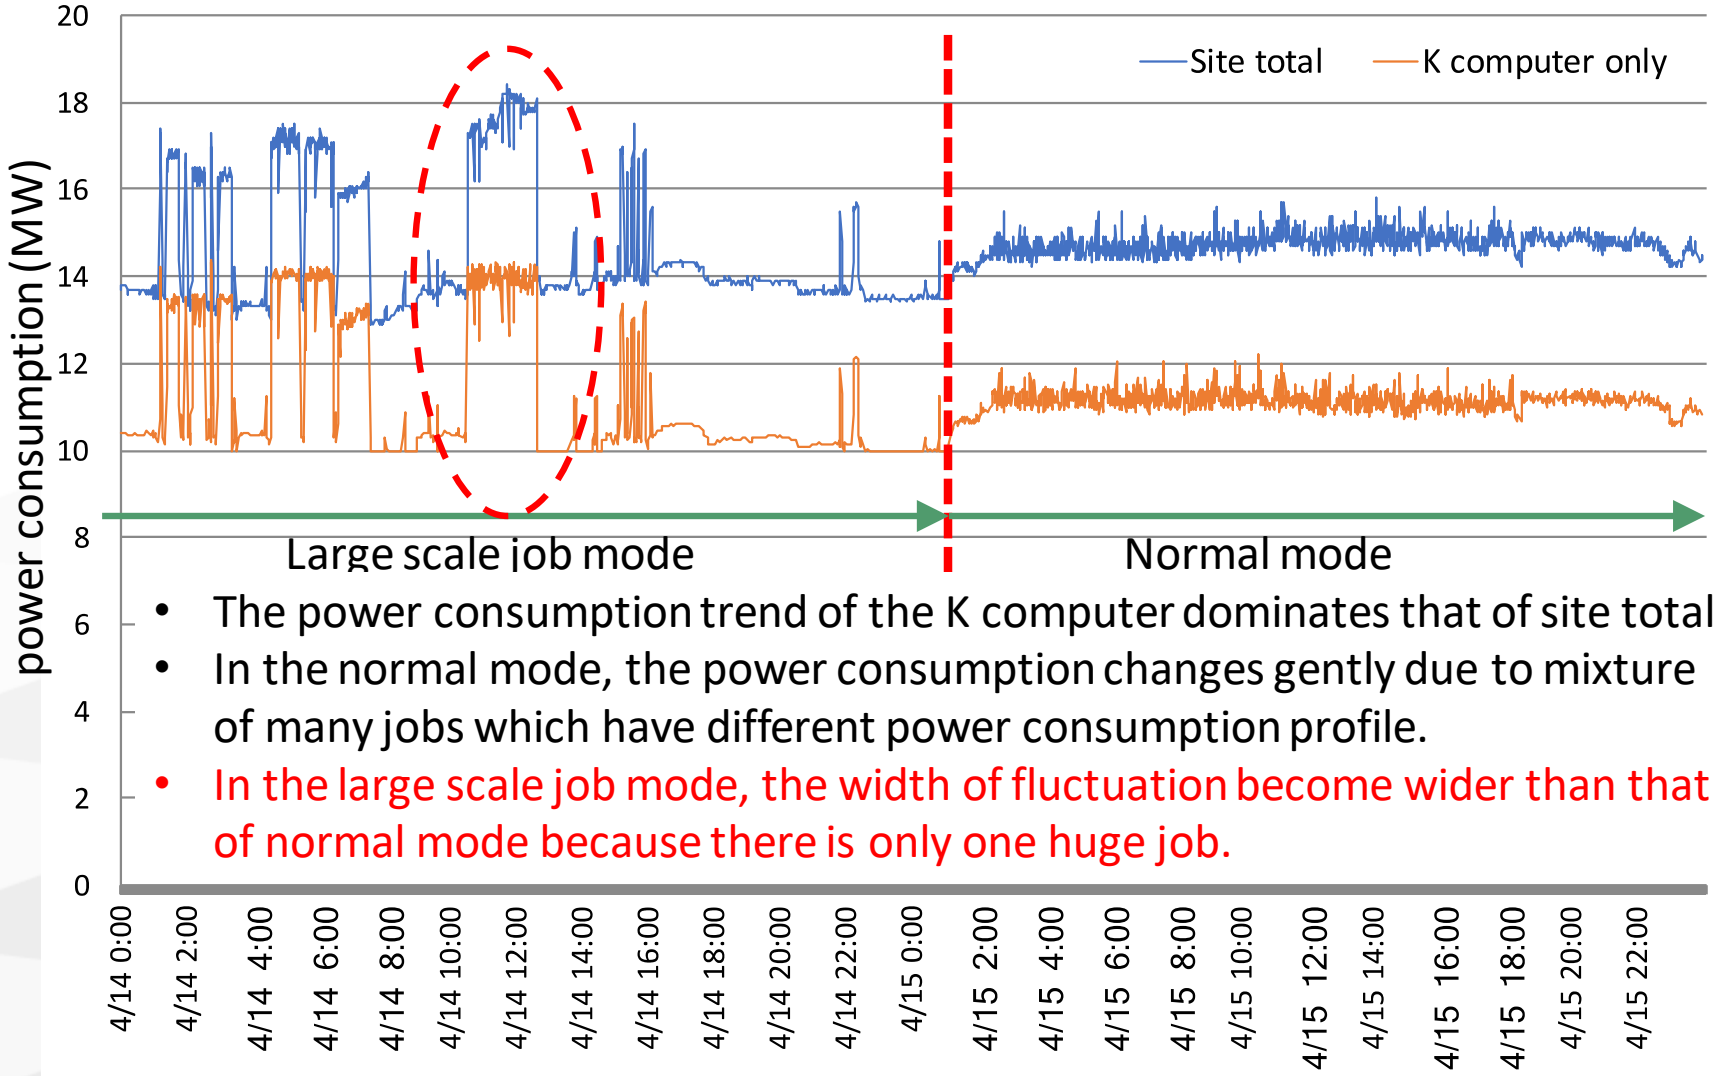

- The power consumption trend of the K computer dominates that of site total.
- In the normal mode, the power consumption changes gently due to mixture of many jobs which have different power consumption profile.
- In the large scale job mode, the width of fluctuation become wider than that of normal mode because there is only one huge job.

Cooling in the idle and peak (in future)

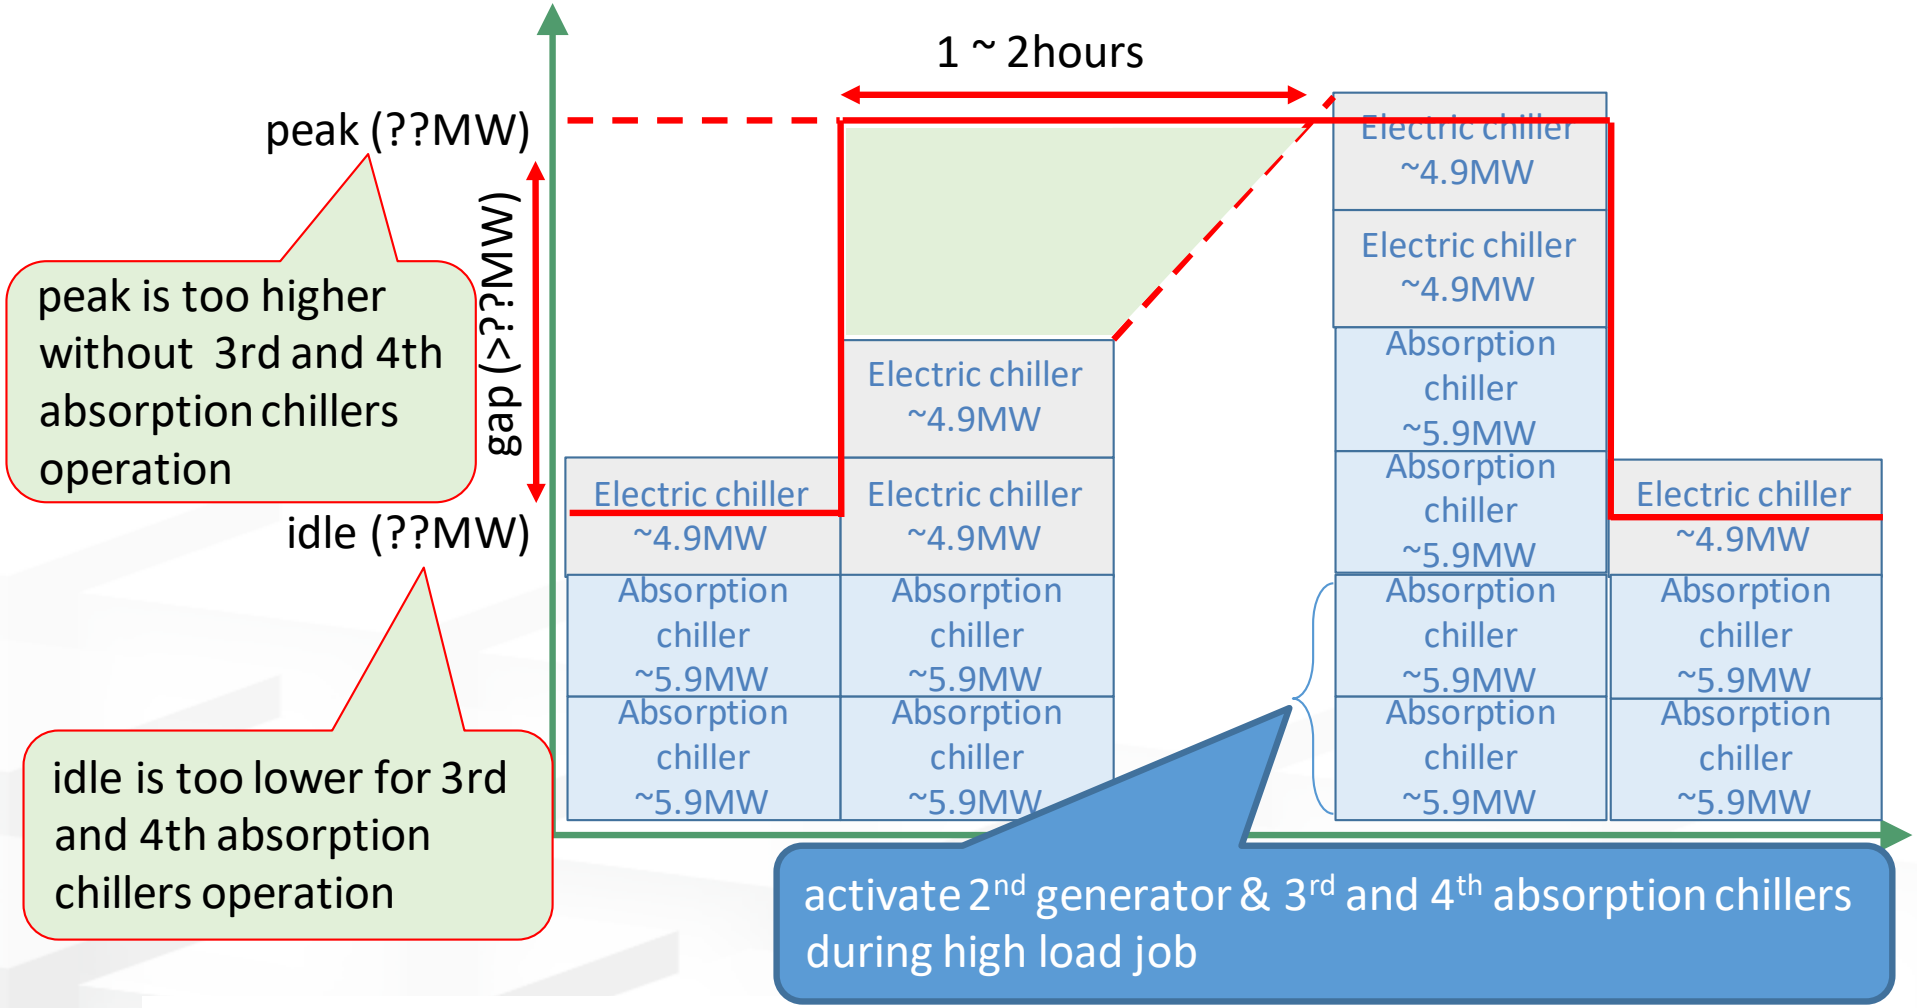

- If the gap between the idle and peak become huge, cooling supply might not in time due to the activation latency of absorption chillers.
- How can we bridge the gap?

Thermal storage tank (long duration peak)

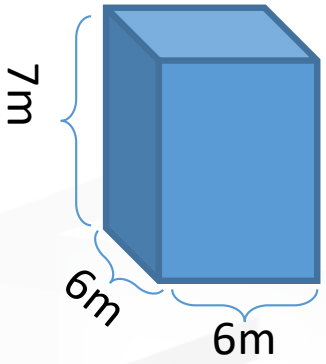

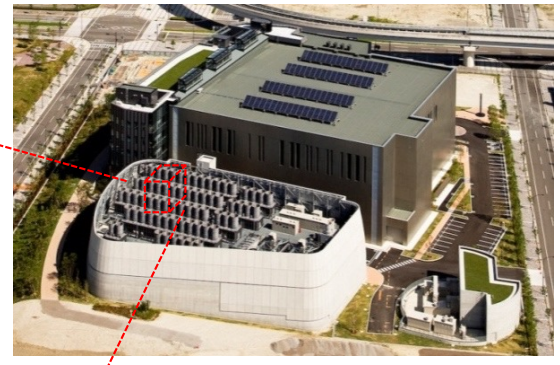

- Specs
  - Water capacity : 210 tons
  - Charge/Discharge time:
    - charge : 200 min.
    - discharge : 40 min. (2.9MW) — 80 min. (1.4MW)
  - Water temp. : 8~18C
  - Cost : 2 million USD (approx.)

Supercapacitor (short duration peak)

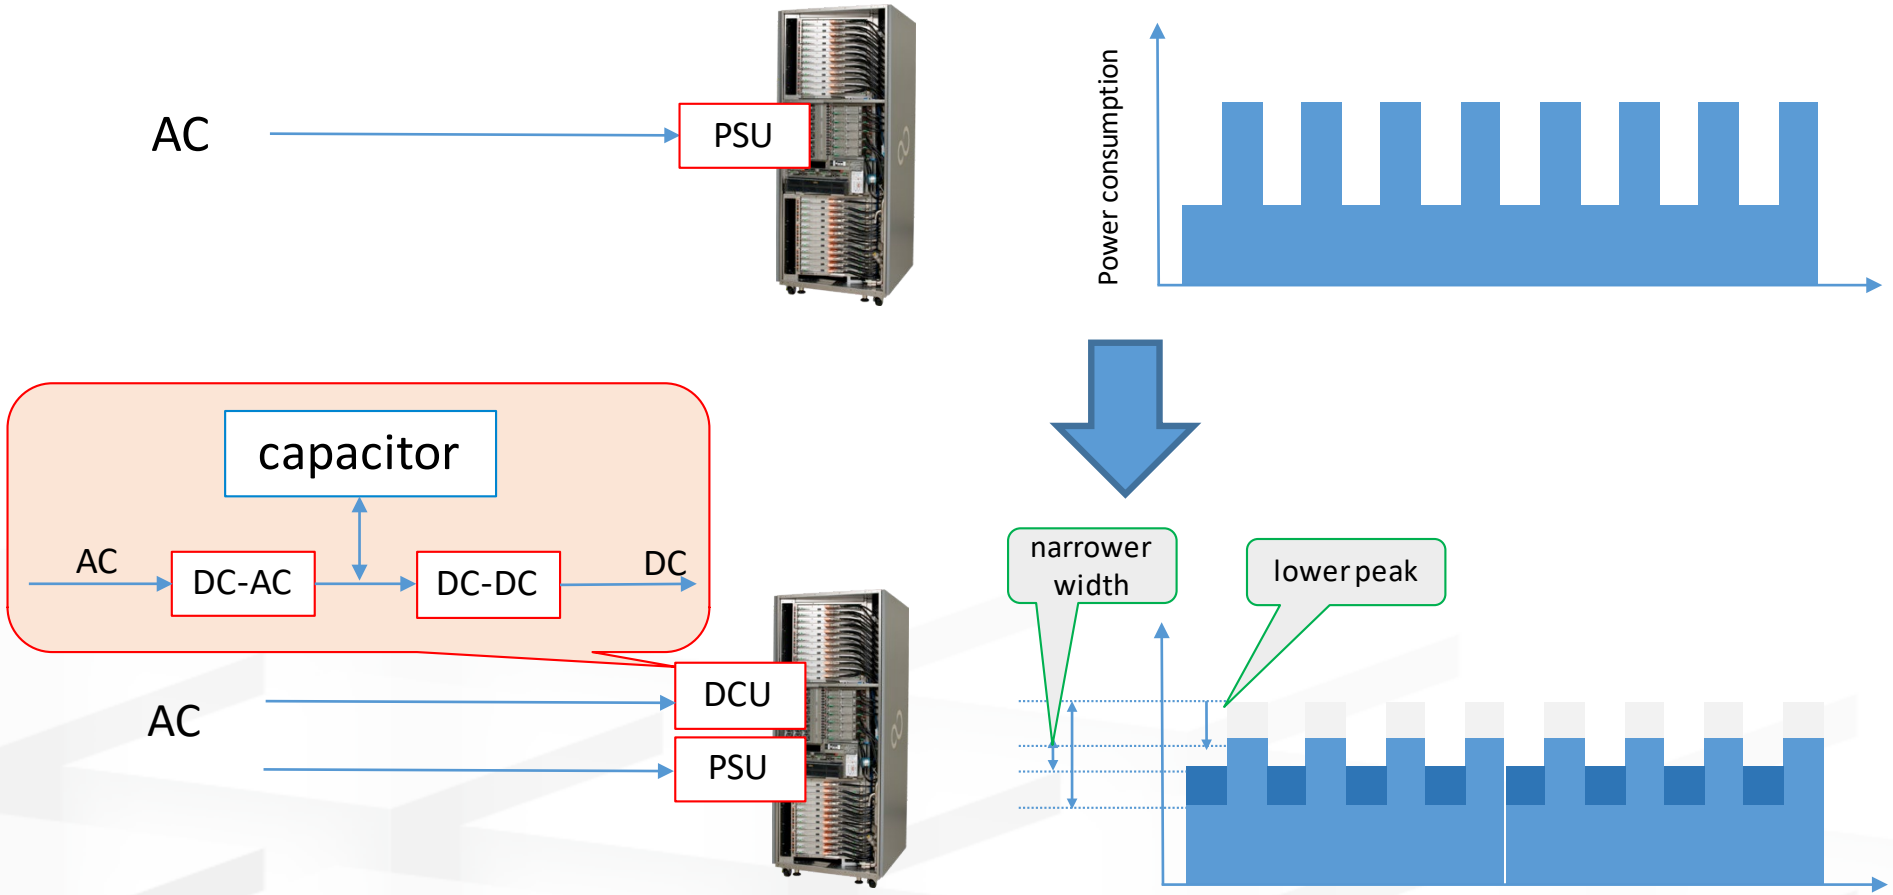

- Capacitor is expected to work to suppress the sharp fluctuation.
  - height of peaks will be lower
  - width of fluctuation will be narrower

Post K Fujitsu A64fx Processor is…

- **a Many-Core ARM CPU…**
  - 48 compute + 2 or 4 assistant (OS) cores
  - Brand new core design
  - Near Xeon-Class Integer performance core
  - ARM V8.2 --- 64bit ARM ecosystem
  - Tofu-D + PCIe 3 external connection

- **…but also a GPU-like power efficient processor**
  - SVE 512 bit vector x 2 (ARM & Fujitsu)
    - Integer (1, 2, 4, 8B) + Float (16, 32, 64B)
  - Shallow cache + localization (sector cache)
  - HBM2 3D memory – Massive Mem BW (Bytes/DPF ~0.4), 840GB/s
    - Streaming access, strided, scatter/gather etc.
  - Intra-chip barrier synch. and other memory enhancing features

- **GPU-like High performance in HPC, AI/Big Data, Auto Driving…**

ARM 64fx
~3TF DFP / ~1TB/s
~160W TDP incl mem, HCA

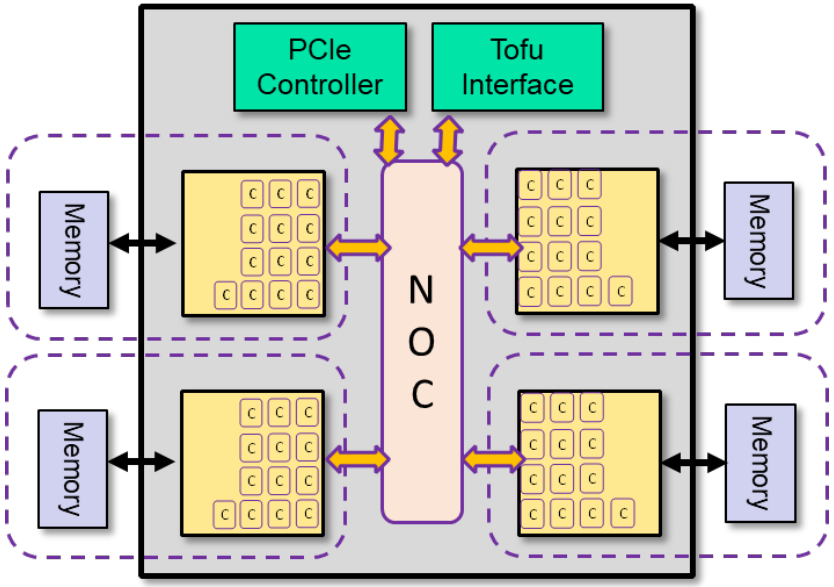

Skylake Xeon
6ch DDR4 2666
1.7TF DFP/21.3GB/s x 6 / chip
250~300W TDP incl mem, HCA

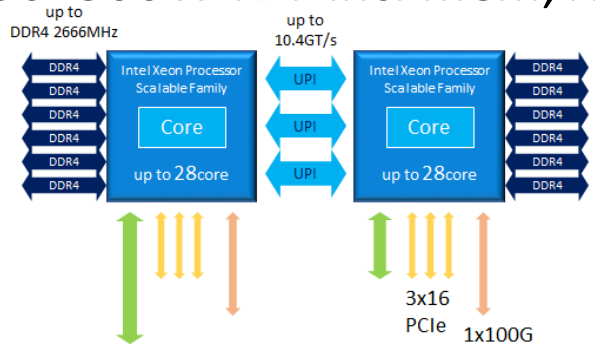

Post-K A64fx A0 (ES) performance

	Performance / CPU					Machine Performance (HPC)		
	Peak TF (DFP)	Peak Mem. BW	Stream Triad	Theor etical B/F	DGEMM Efficiency	Linpack Efficiency	GF/W	Network BW Per Chip
Post-K A64fx (A0 Eng. Sample)	2.764/ 3.072	1024GB/s	840GB/s	0.37/ 0.33	94 %	87.7 %	>15	TOFU-D 40.8GB/s (6.8x 6)
Intel KNL	3.0464	600GB/s	490GB/s	0.20	66%	54.4 %	4.9	12.5 GB/s
Intel Skylake	1.6128	127.8GB/s	97 GB/s	0.08	80 %	66.7 %	4.5	6.2GB/s
NVIDIA V100 (DGX-2)	7.8	900 GB/s	855GB/s	0.12		76 %	15.113	160GB/s 6.2GB/s

Performance

- A64FX boosts performance up by microarchitectural enhancements, 512-bit wide SIMD, HBM2 and process technology
  - > 2.5x faster in HPC/AI benchmarks than SPARC64 XIfx (Fujitsu's previous HPC CPU)
    - The results are based on the Fujitsu compiler optimized for our microarchitecture and SVE

A64FX Benchmark Kernel Performance (Preliminary results)

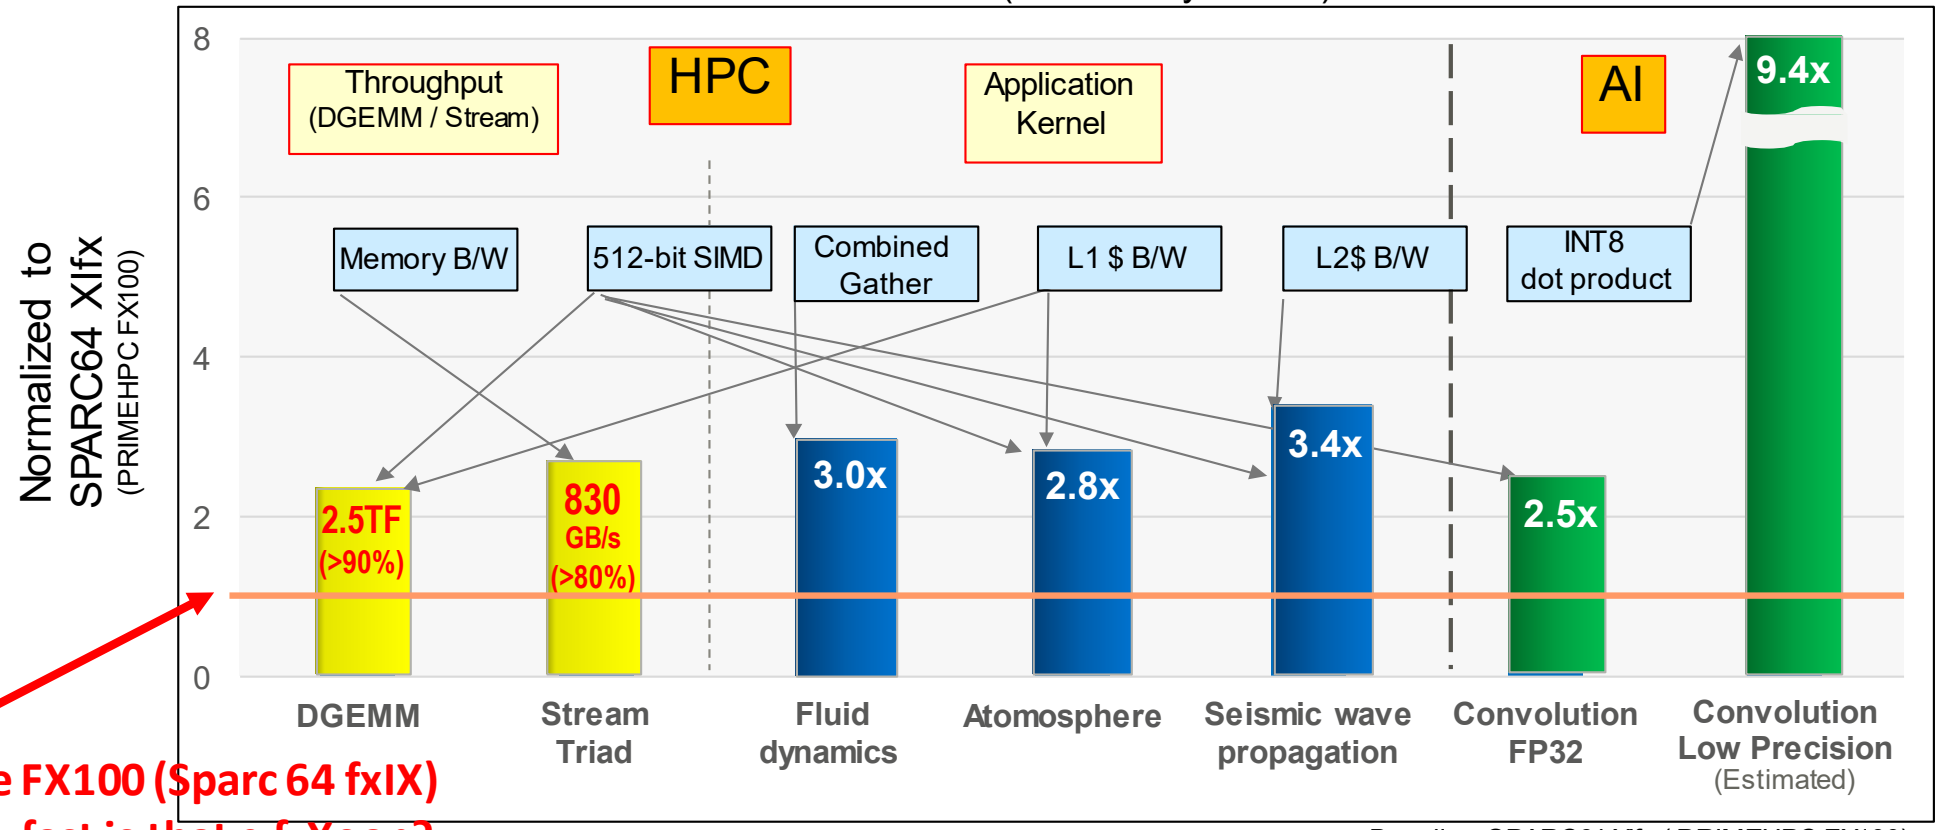

Baseline FX100 (Sparc 64 fxIX)

But how fast is that c.f. Xeon?

All Rights Reserved. Copyright © FUJITSU LIMITED 2018

NAS Parallel Benchmark of FX100

FUJITSU

[Slide by Ikuo Miyoshi, Fujitsu, SSKen2015]

■ OpenMP版を用いてノードあたり演算性能を評価

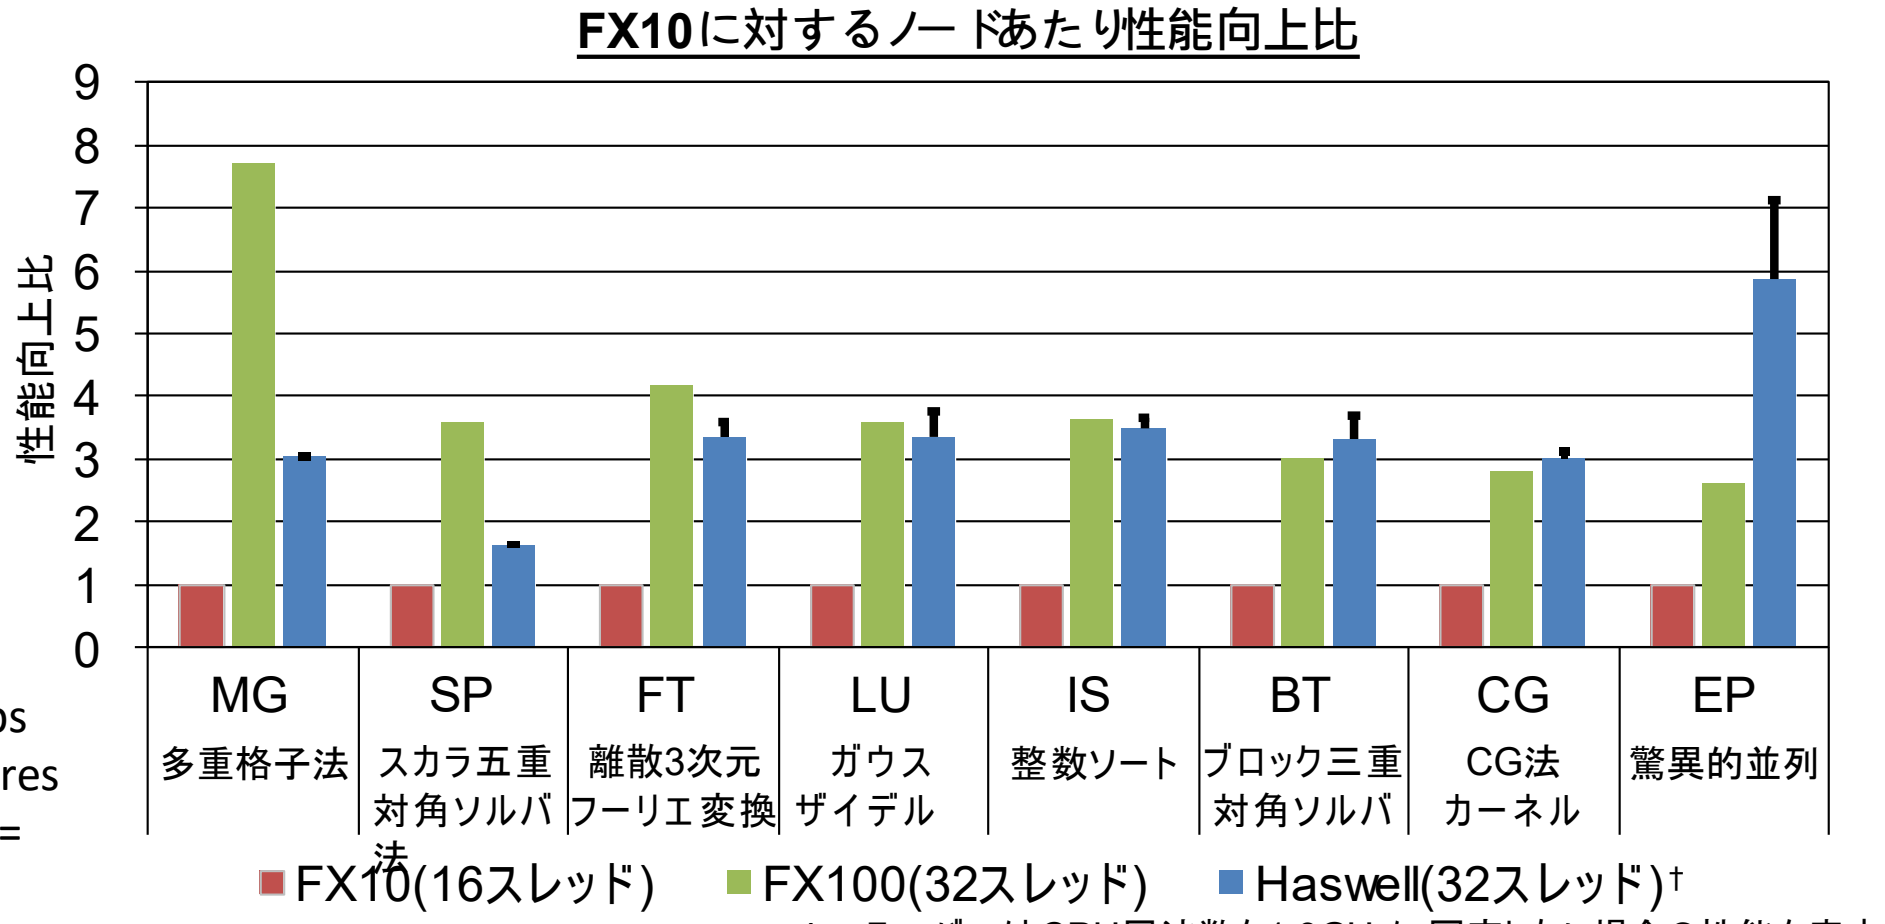

†エラーバーはCPU周波数を1.9GHzに固定しない場合の性能を表す

Note: Haswell:
1 node = 2 chips
32 threads&cores
FX100: 1 node =
1 chip
32 threads&cores

FX100は、FX10比平均3.9倍、Haswell比平均1.3倍のノードあたり演算性能

使用コード: NAS Parallel Benchmarks Ver. 3.3.1 OpenMP版 クラスC

Copyright 2015 FUJITSU LIMITED

Fiber (Post-K) MiniApp on FX100

FUJITSU

[Slide by Ikuo Miyoshi, Fujitsu, SSKen2015]

■ 評価条件

アプリ名	問題サイズ	評価区間	スレッド数×プロセス数 (左からFX10、FX100、Haswell)		
CCS QCD	32×32×32×32	BiCGStab	16t×2p	32t×1p	16t×2p
NICAM-DC-MINI	gl05rl00z80pe10	Dynamics	3t×10p		
FFB-MINI	1,048,576要素	MAIN_LOOP	1t×32p	8t×4p	1t×32p
FFVC-MINI	256×256×256	Total	4t×8p		16t×2p
NTChem-MINI	taxol	RIMP2_Driver	1t×32p	16t×2p	2t×16p

■ 測定結果

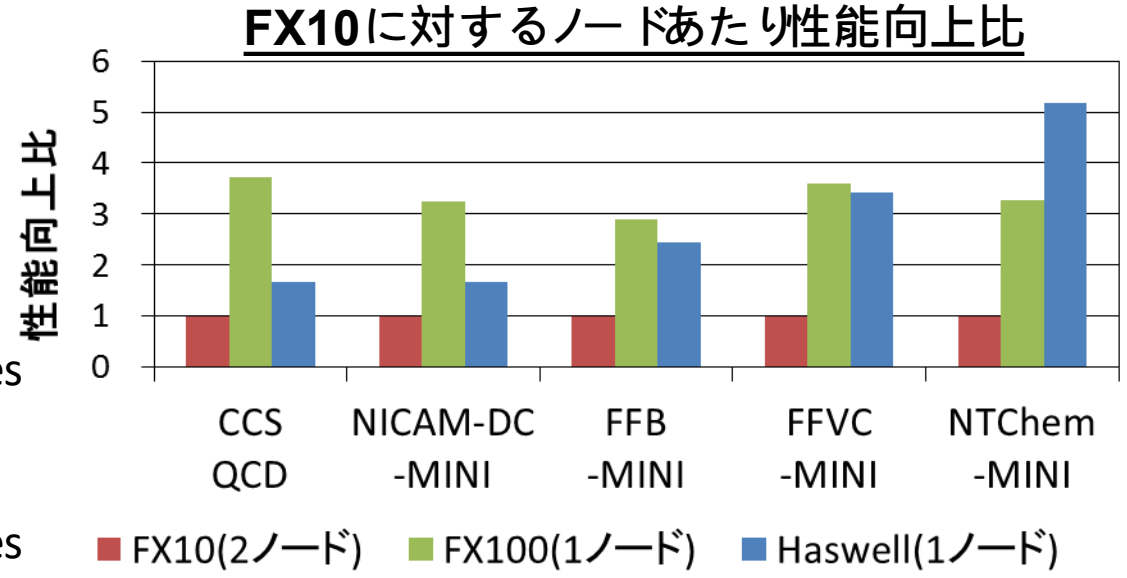

Note: Haswell:
1 node = 2 chips
32 threads&cores
FX100: 1 node =
1 chip
32 threads&cores

FX100は、FX10比平均3.3倍、Haswell比平均1.4倍の ノードあたり性能

注) FFVCには開発版コンパイラ、NTChemには開発版数学ライブラリを使用。QCDではセクタキャッシュ利用、FFBではループリローリングのコード変更を実施

Copyright 2015 FUJITSU LIMITED

Power Management (Cont.)

- "Power knob" for power optimization
  - A64FX provides power management function called "Power Knob"
    - Applications can change hardware configurations for power optimization

→ Power knobs and Energy monitor/analyzer will help users to optimize power consumption of their applications

<A64FX Power Knob Diagram>

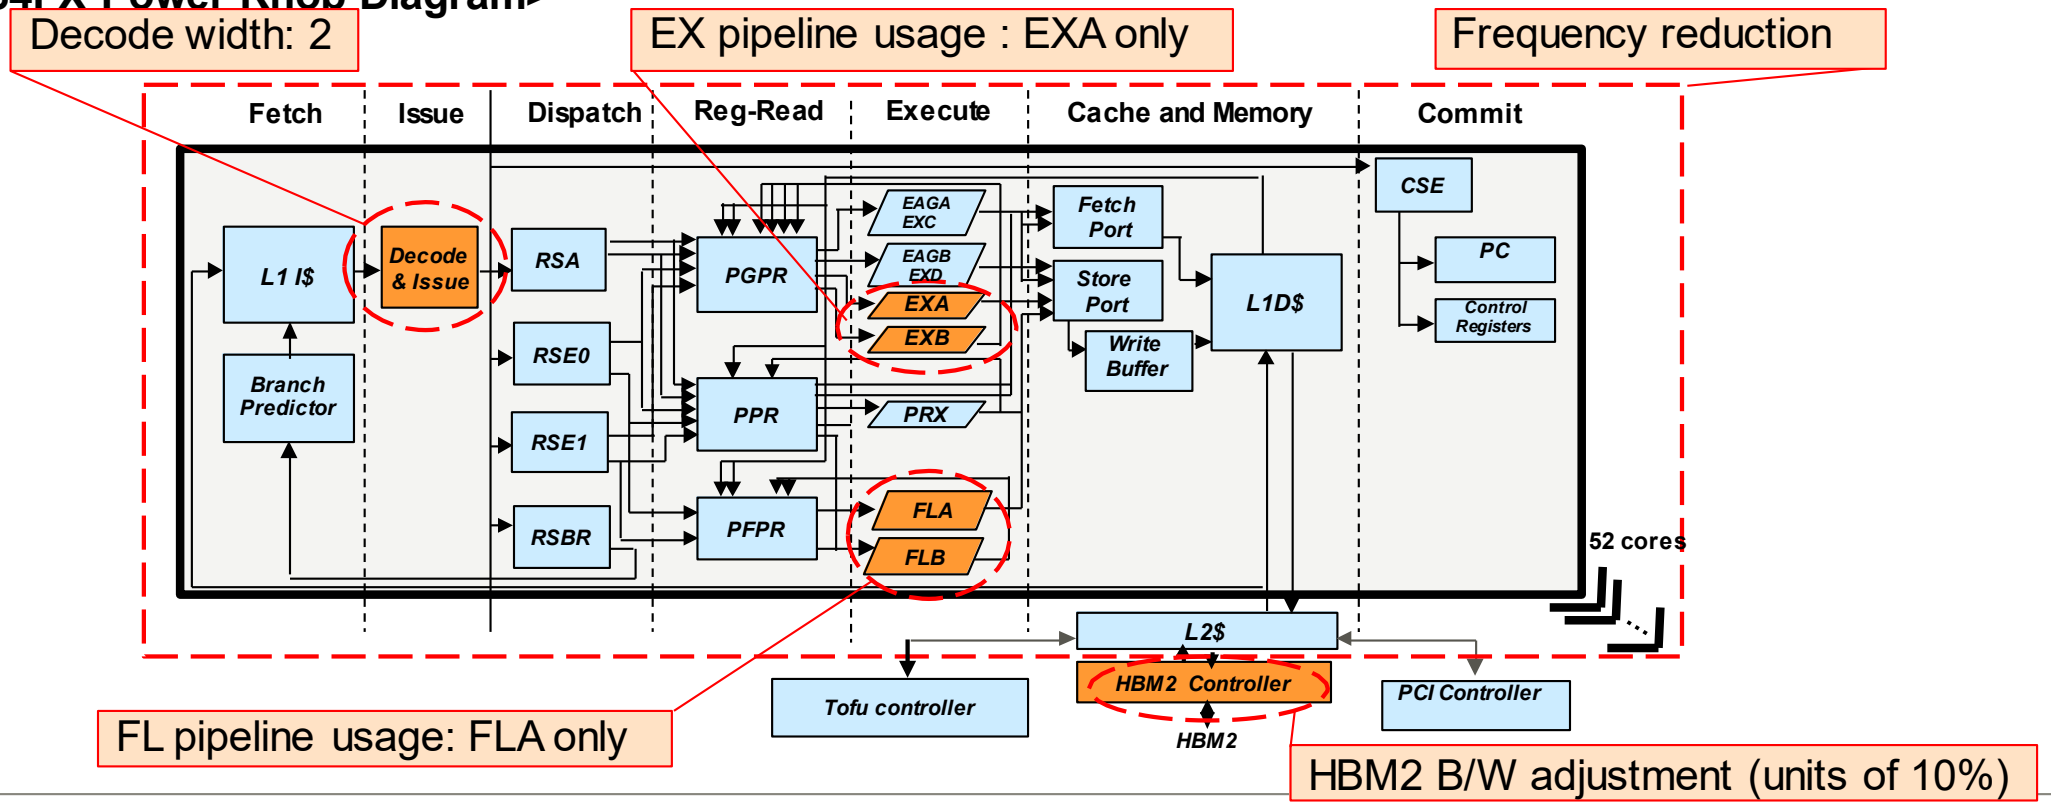

All Rights Reserved. Copyright © FUJITSU LIMITED 2018

Post-K Chassis, PCB (w/DLC), and A64fx CPU Package

FUJITSU

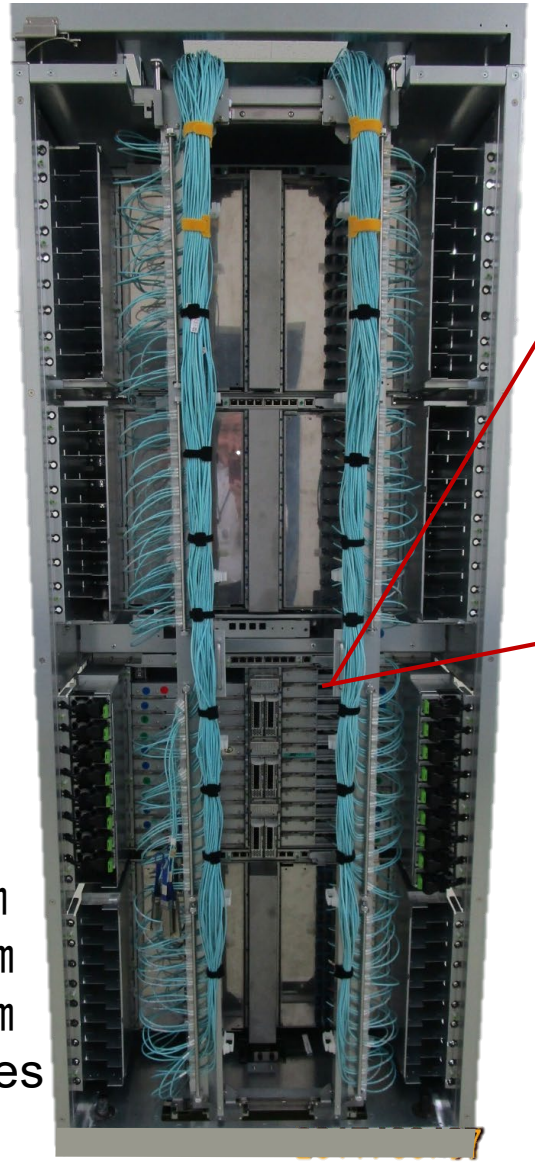

W 800mm
D1400mm
H2000mm
384 nodes

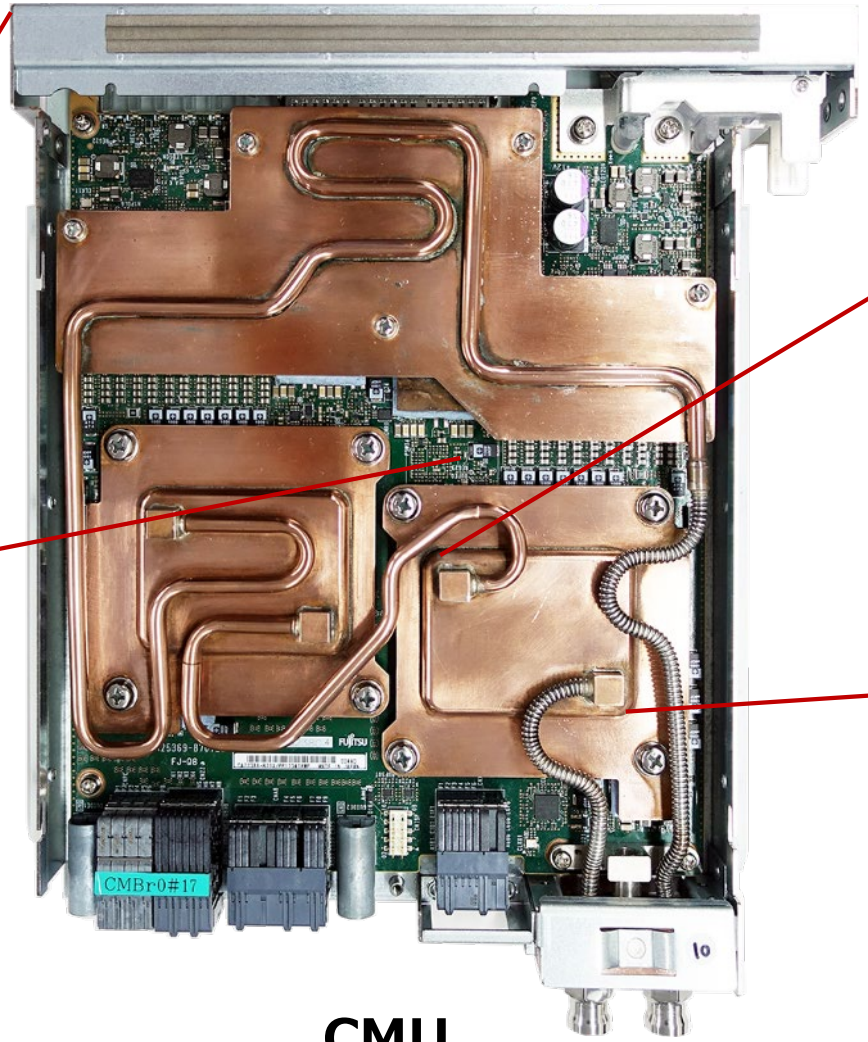

CMU

CPU Package

A0 Chip Booted in June
Undergoing Tests
B0 underway

Copyright 2018 FUJITSU LIMITED

Higher-density Node Configuration

FUJITSU

- The CPU is smaller and the off-chip channels are halved
  - The number of 3D-stacked memories was halved from 8 to 4
  - Each Tofu link was reduced from 4 lanes to 2 lanes
- More resources are integrated into the CPU
  - The number of CPU Memory Groups (NUMA nodes) doubled from 2 to 4
  - The number of Tofu Network Interfaces increased from 4 to 6

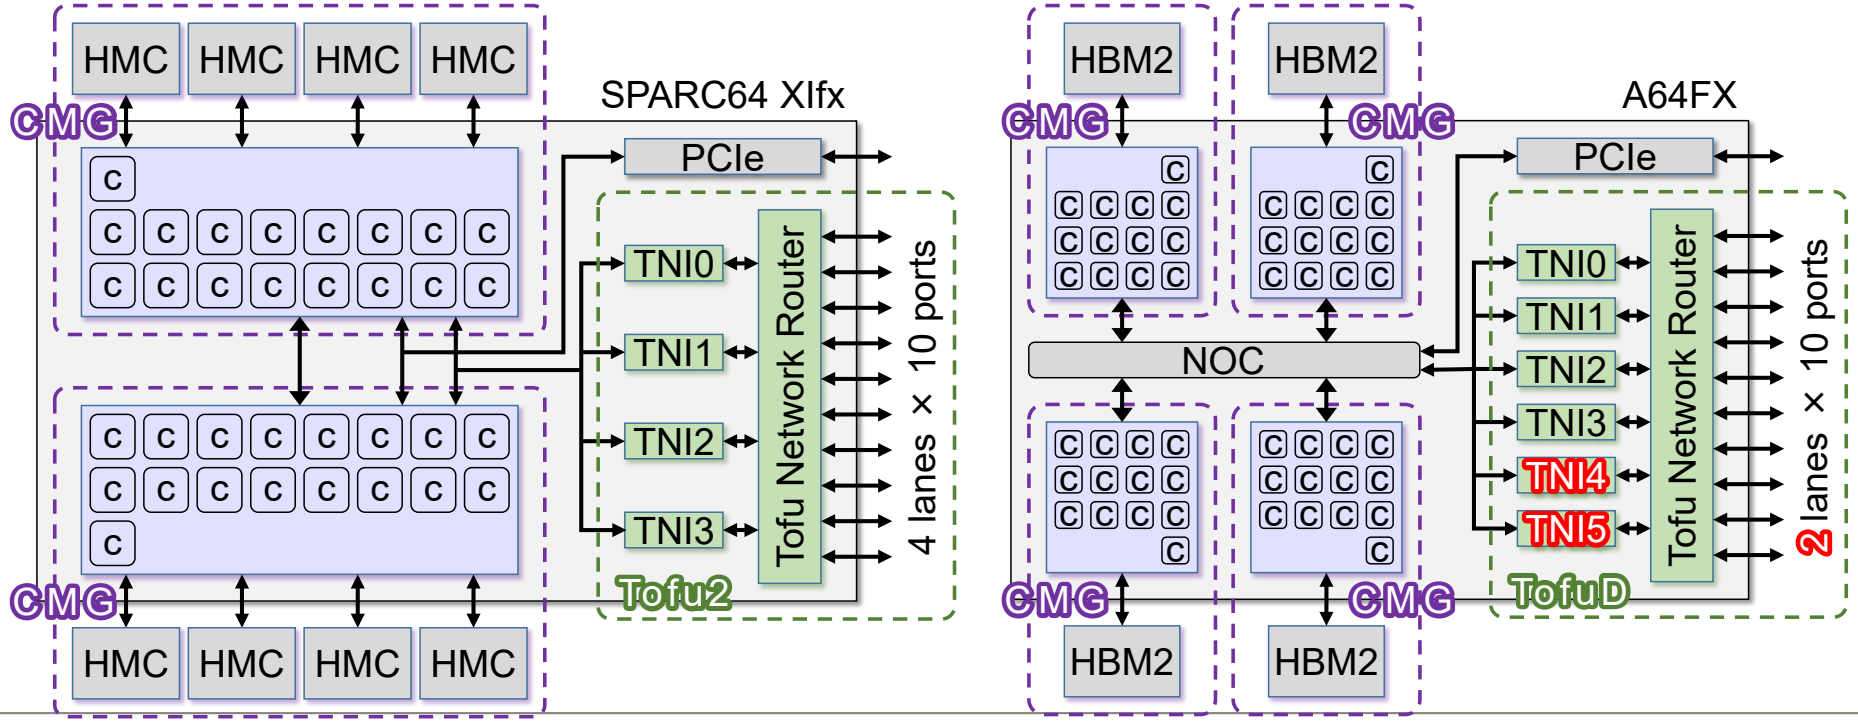

Copyright 2018 FUJITSU LIMITED

Injection Rates per Node

FUJITSU

- Simultaneous Put transfers to multiple nearest-neighbor nodes
  - Tofu1 and Tofu2 used 4 TNIs, and TofuD used 6 TNIs

	Injection rate	Efficiency
Tofu1 (K)	15.0 GB/s	77 %
Tofu1 (FX10)	17.6 GB/s	88 %
Tofu2	45.8 GB/s	92 %
TofuD	38.1 GB/s	93 %

- The injection rate of TofuD was approximately 83% that of Tofu2
- The efficiencies of Tofu1 were lower than 90%
  - Because of a bottleneck in the bus that connects CPU and ICC
- The efficiencies of Tofu2 and TofuD exceeded 90 %
  - Integration into the processor chip removed the bottleneck

Copyright 2018 FUJITSU LIMITED

An Overview of Post-K Hardware

- **Compute Node, Compute + I/O Node connected by 6D mesh/torus Interconnect**
- **3-level hierarchical storage system**
  - 1st Layer
    - Cache for global file system
    - Temporary file systems
      - Local file system for compute node
      - Shared file system for a job
  - 2nd Layer
    - Lustre-based global file system
  - 3rd Layer
    - Storage for archive
  - **>100,000 nodes**
  - **x100PF, >100PB/s mem BW**
  - **Approaching 10 million CPU cores**

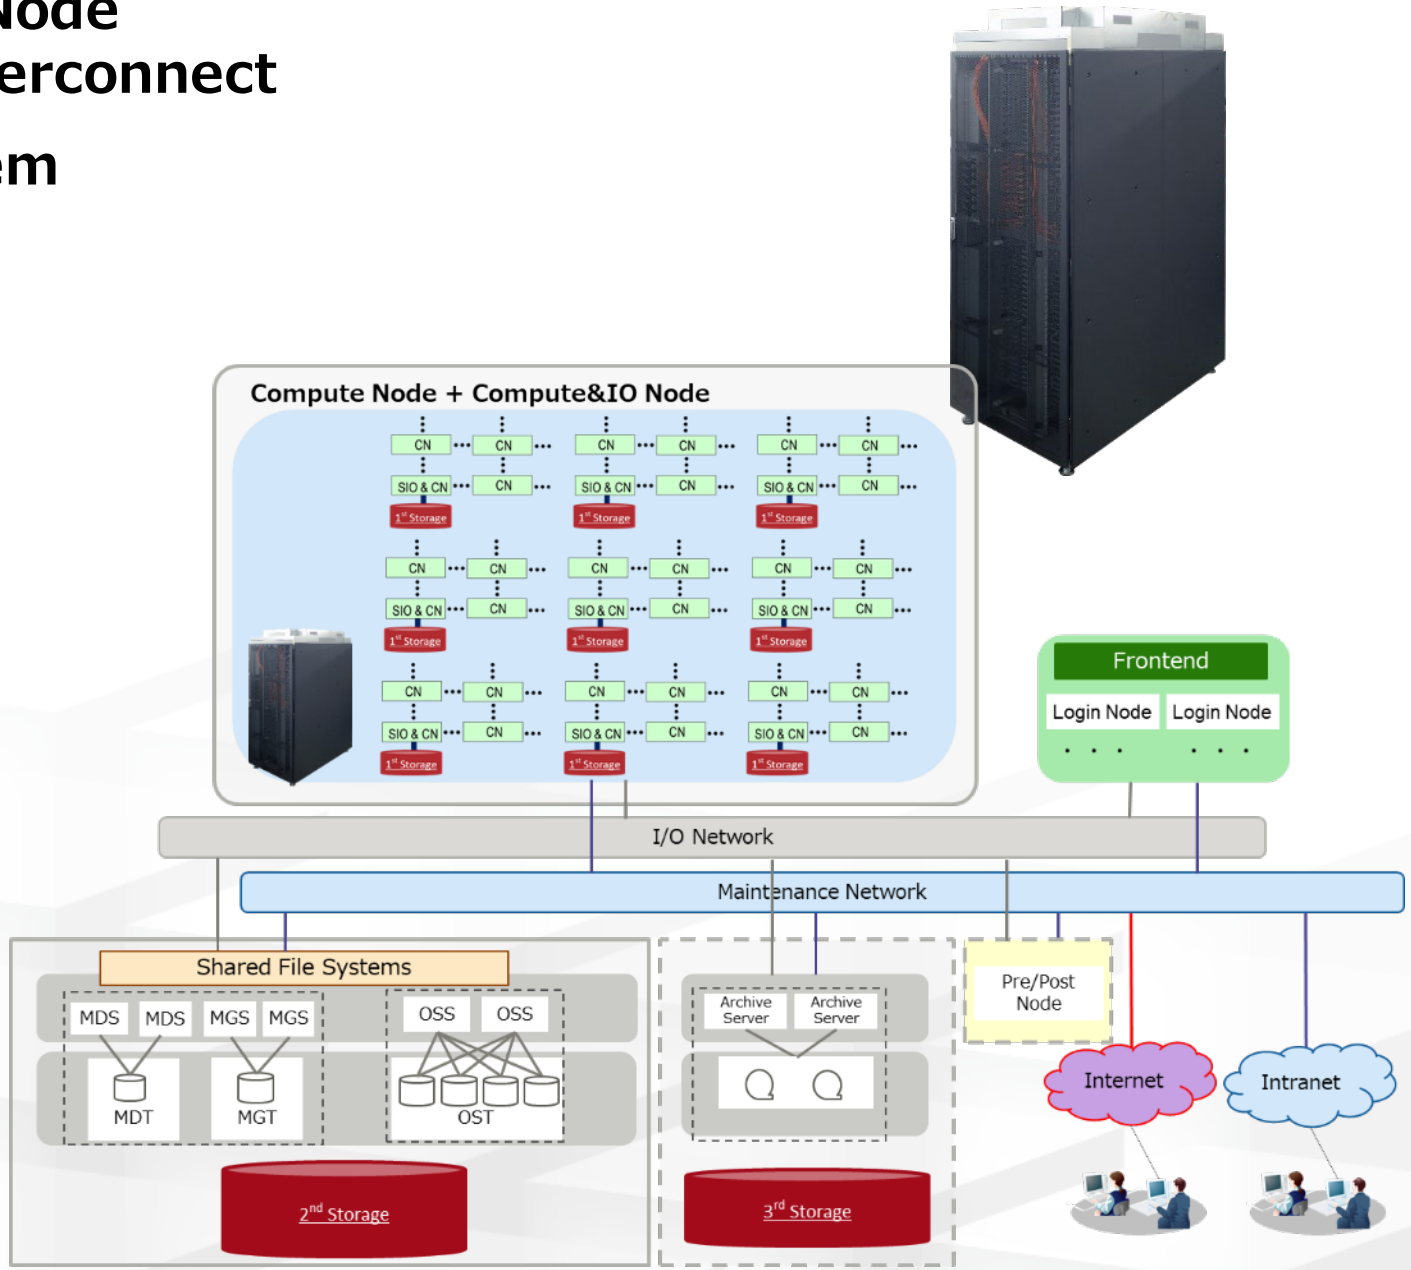

Massive Scale Deep Learning on Post-K

Post-K Processor

- ◆ High perf FP16&Int8
- ◆ **High mem BW for convolution**
- ◆ **Built-in scalable Tofu network**

High Performance DNN Convolution

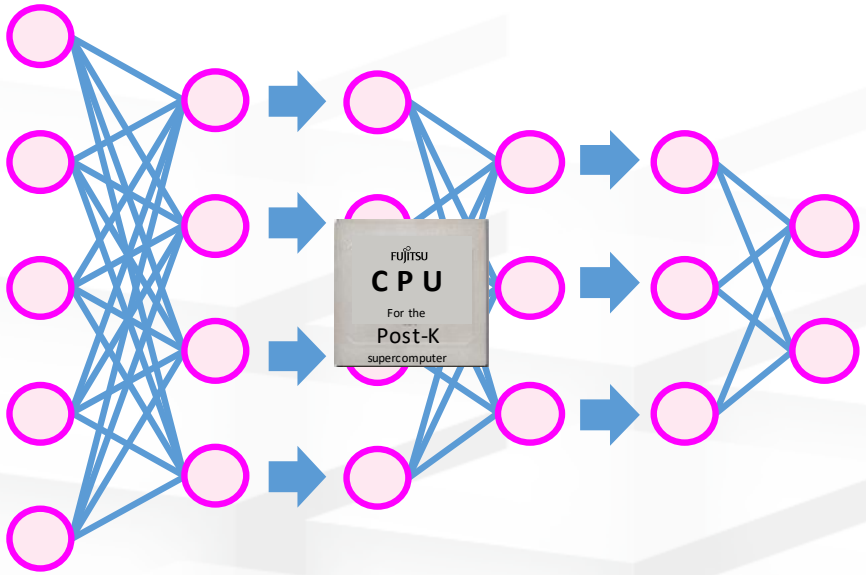

Low Precision ALU + High Memory Bandwidth + Advanced Combining of Convolution Algorithms (FFT+Winograd+GEMM)

Unprecedened DL scalability

High Performance and Ultra-Scalable Network for massive scaling model & data parallelism

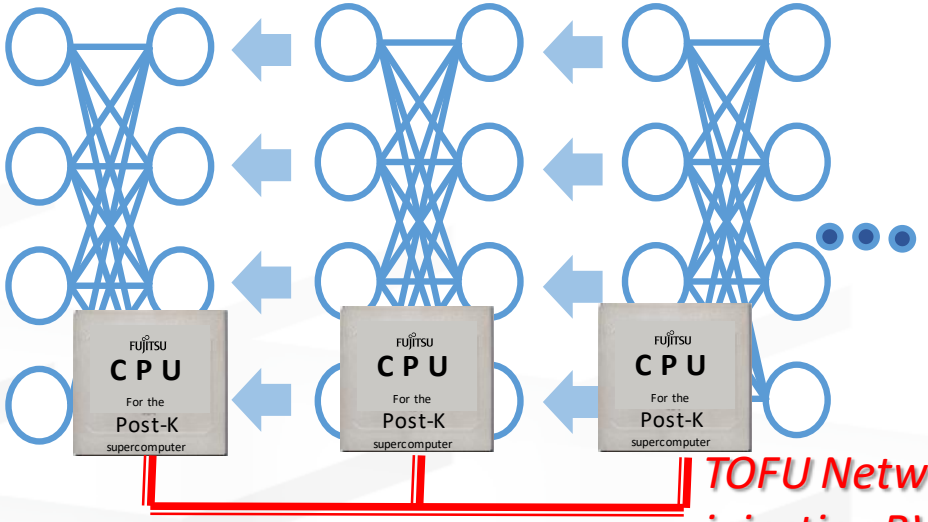

TOFU Network w/high injection BW for fast reduction

Unprecedented Scalability of Data/

Selecting the Optimal Convolution Kernel

- **NEW! Micro Batching: Tokyo Tech. and ETH [Oyama, Tan, Hoefler & Matsuoka, IEEE Cluster 2019]**
  - Use the "micro-batch" technique to select the best convolution kernel
    - Direct, GEMM, FFT, Winograd
    - Optimize both speed and memory size
  - On high-end GPUs, in many cases Winograd or FFT chosen over GEMM
    - They are faster but use more memory
  - Currently implemented as cuDNN wrapper, applicable to all frameworks
  - For Post-K, (1) Winograd/FFT are selected more often, and (2) performance will be similar to GPUs in such cases

Evaluation: WR using Dynamic Programming

- μ-cuDNN achieved **2.33x** speedup on forward convolution of AlexNet conv2

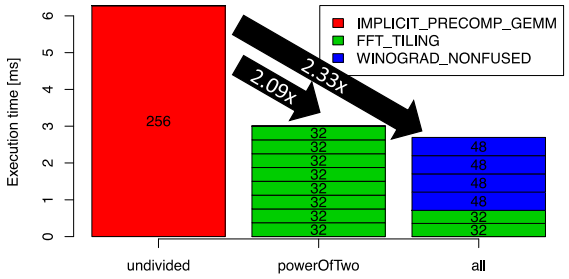

`cudnnConvolutionForward` of AlexNet conv2 on NVIDIA Tesla P100-SXM2
Workspace size of 64 MiB, mini-batch size of 256
Numbers on each rectangles represent micro-batch sizes

Evaluation: WD using Integer LP

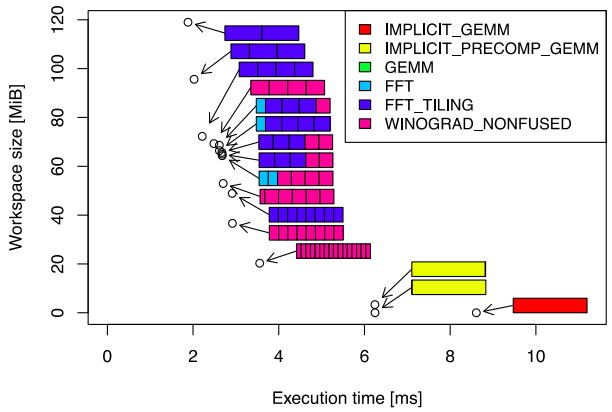

A desirable configuration set of AlexNet conv2 (Forward)
Mini-batch size of 256, P100-SXM2
Each bar represents proportion of micro-batch sizes and algorithms

Large Scale simulation and AI coming together
[Ichimura et. al. Univ. of Tokyo, IEEE/ACM SC17 Best Poster]

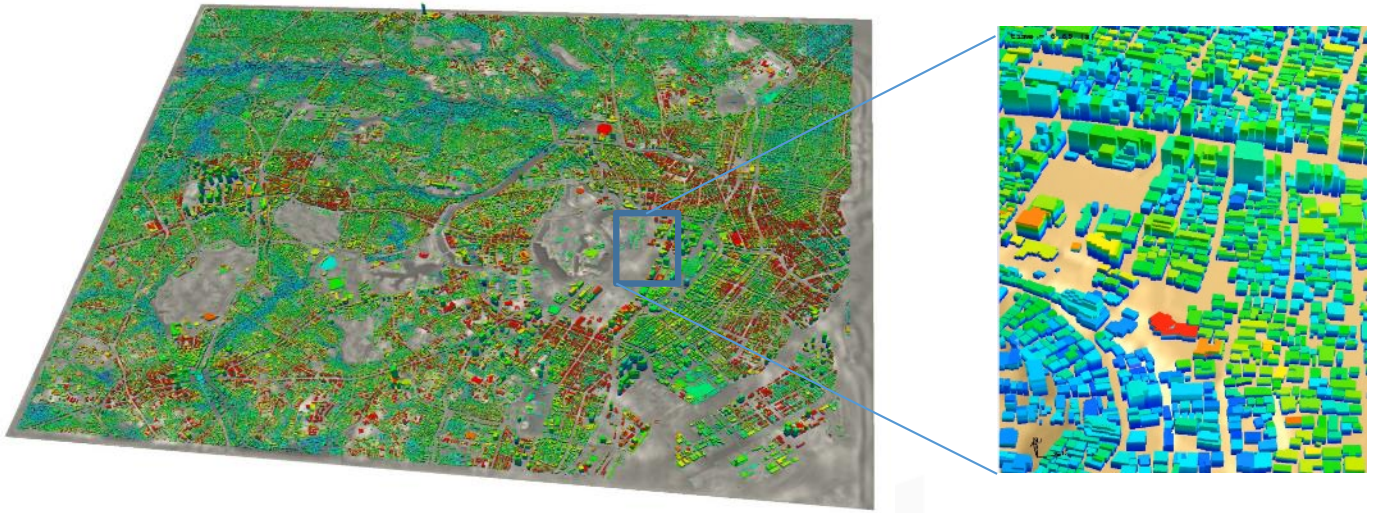

130 billion freedom earthquake of entire Tokyo on K-Computer (ACM Gordon Bell Prize Finalist, SC16,17 Best Poster)

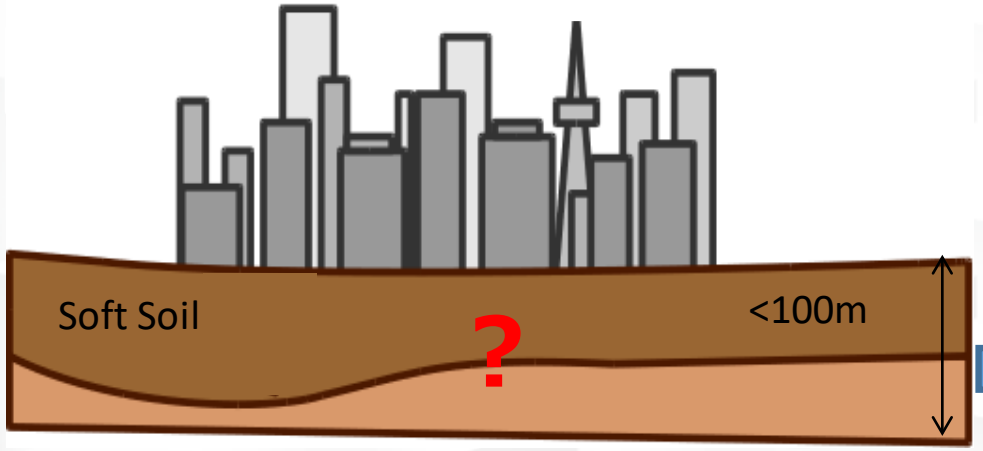

Earthquake

Too Many Instances

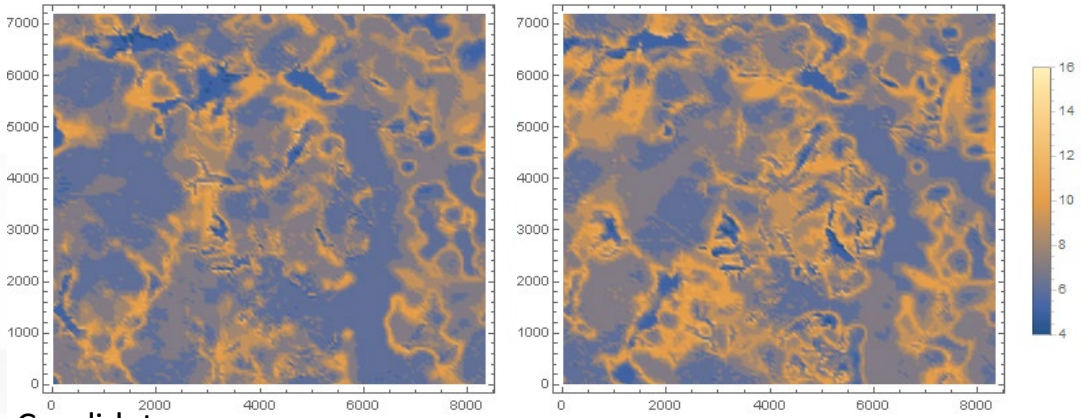

Candidate Underground Structure 1

Candidate Underground Structure 2

AI Trained by Simulation to generate candidate soft soil

Cutting Edge Research AI Infrastructures in Japan Accelerating BD/AI with HPC (w/accompanying BYTES)

(and my effort to design & build them)

Tokyo Tech

AI-ExaFlop era

In Production

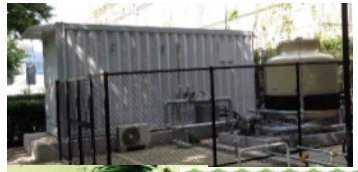

x5.8

Oct. 2015
TSUBAME-KFC/DL
(Tokyo Tech./NEC)
1.4 AI-PF(Petaflops)

In Production

Mar. 2017 x5.8
AIST AI Cloud
(AIST-AIRC/NEC)
8.2 AI-PF

In Production

Aug. 2017 x11.7
TSUBAME3.0 (Tokyo Tech./HPE)
47.2 AI-PF (65.8 AI-PF w/Tsubame2.5)

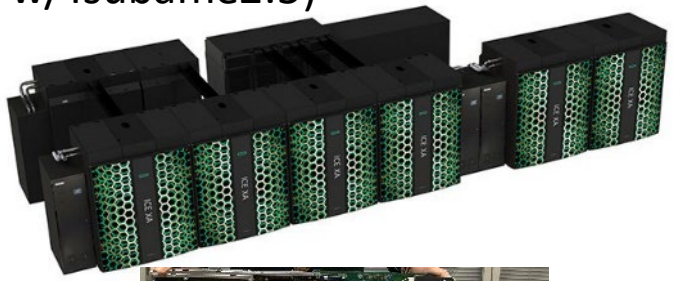

In Preparation

X4~6?

Aug 2018
ABCI (AIST-AIRC)
550 AI-Petaflops

2020 Post-K Multi AI-Exaflops

order of magnitude over ABCI

still under plans

~x400 in 3 years

Built/funded

Mar. 2017
AI Supercomputer
Riken AIP/Fujitsu
4.1 AI-PF

R&D Investments into world leading AI/BD HW & SW & Algorithms and their co-design for cutting edge Infrastructure absolutely necessary (just as is with Japan Post-K and US ECP in HPC)

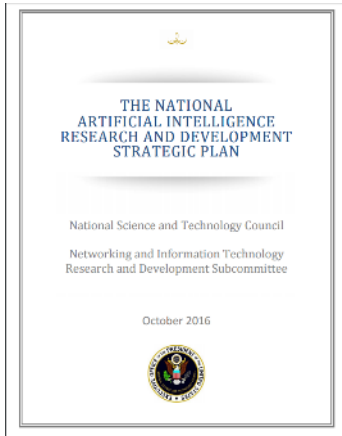
THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN

National Science and Technology Council

Networking and Information Technology Research and Development Subcommittee

October 2016