FUJITSU

shaping tomorrow with you

Fujitsu PRIMEHPC FX100 Features and Performance

PRIMEHPC FX100 Design

FUJITSU

Highly scalable, high performance supercomputer system

- Introducing cutting edge technologies

High application productivity

- Strong hardware/software for application development and tuning

Highly reliable, available and operable system

- Implementing/facilitating Fujitsu's long HPC experience

PRIMEHPC FX100 Features

FUJITSU

Improved CPU & interconnect from the K computer & FX10

- General purpose CPU architecture for application productivity
- Hardware barrier synchronization, sector cache, etc.

Introduced technologies for next-generation computing

- HPC-ACE2, assistant cores, HMC, Tofu interconnect 2

	FX100	FX10	K computer
Double precision flops / CPU	**Over 1TF** (x8)	236.5 GF (x2)	128 GF (1)
Single precision flops / CPU	**Over 2 TF** (x16)	236.5 GF (x2)	128 GF (1)
# of cores	**32** (x4)	16 (x2)	8 (1)
SIMD width	**256 bit** (x2)	128 bit (x1)	128 bit (1)
Byte per flop	**0.4 ~ 0.5**		
Interconnect	**Tofu 6D mesh/torus**		

Software Stack: Technical Computing Suite

FUJITSU

- Applications can exploit maximum performance with Fujitsu software stack

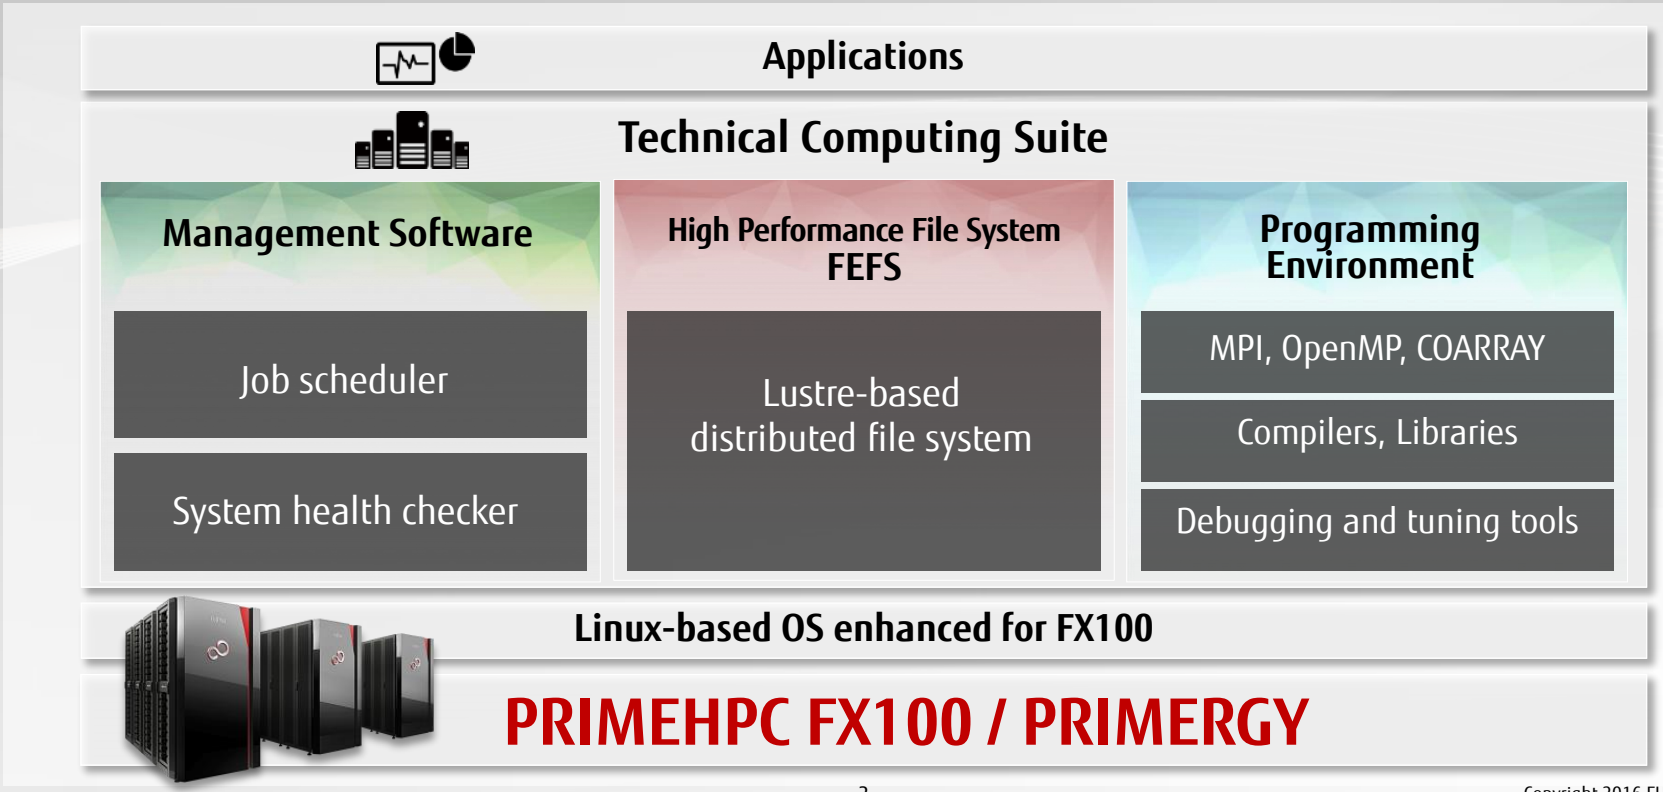

PRIMEHPC FX100
Updated Performance Results

Balanced Enhancement of FLOPS and Memory

FUJITSU

- Over 1 TFLOPS and 480 GB/s memory bandwidth per chip
- PRIMEHPC series shows high performance for both HPL and HPCG

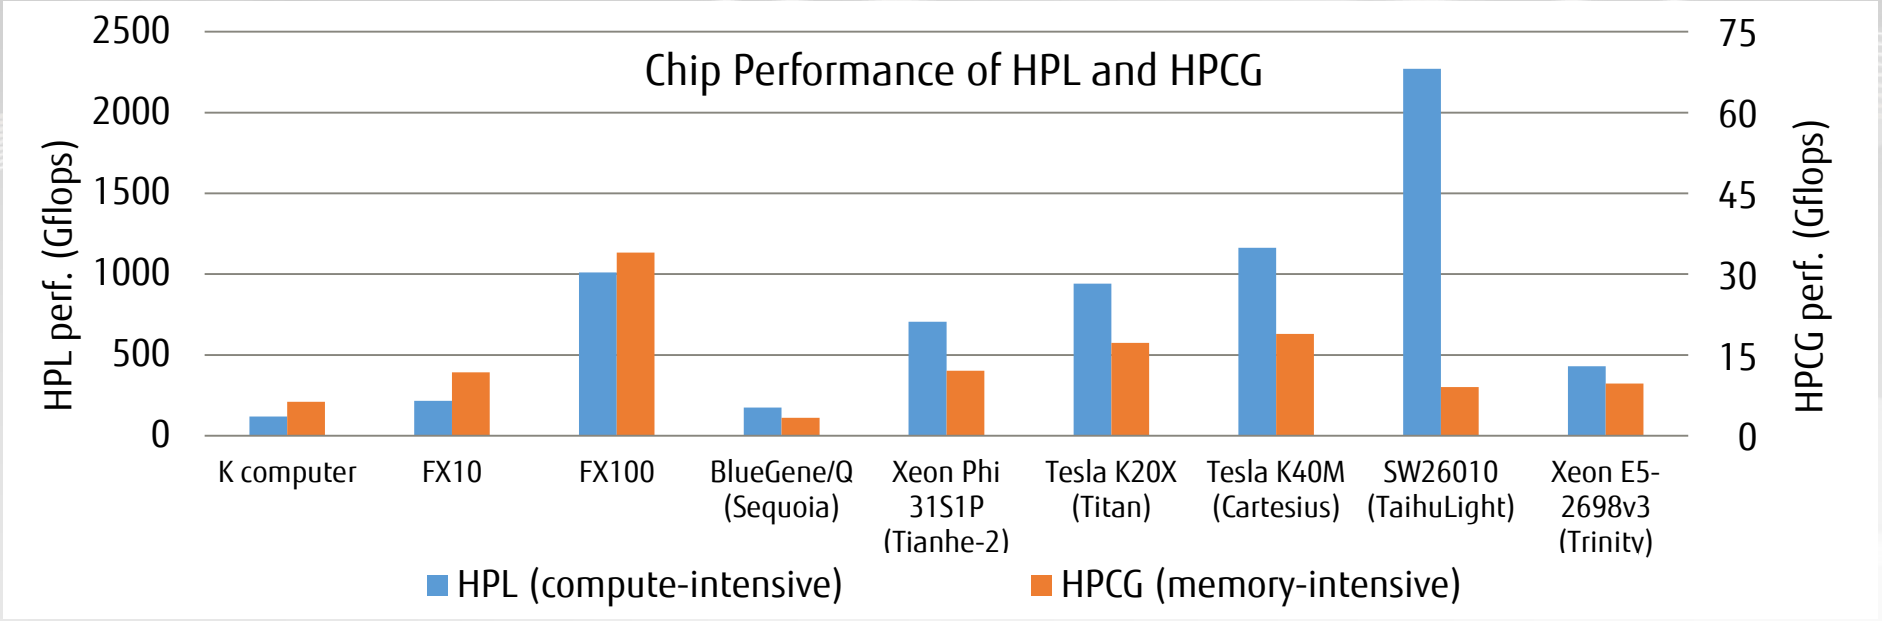

Comparison Using Highly-tuned App. Kernels

FUJITSU

- Performance of application kernels in fusion plasma codes[1]
  - GYSELA: semi-Lagrangian scheme, compute-intensive
  - GT5D: finite-difference scheme, memory-intensive

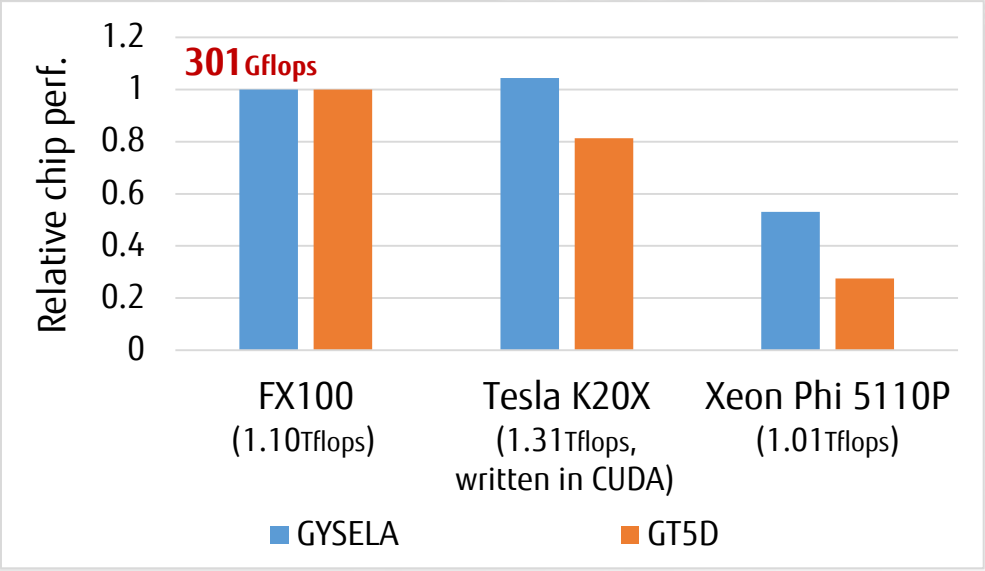

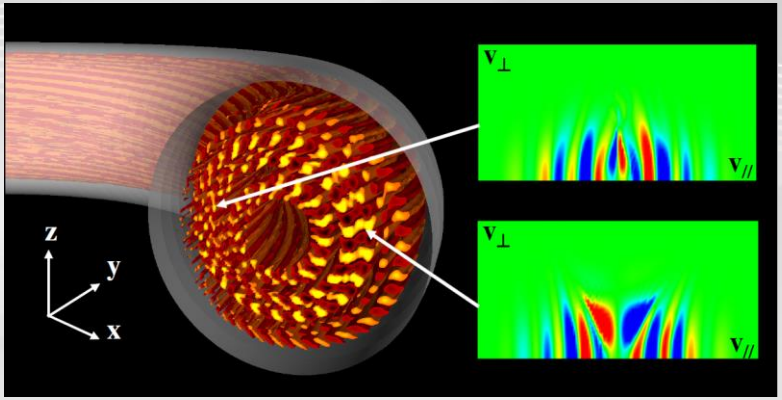

Fusion plasma turbulence in 5D phase space
(By courtesy of JAEA)

- FX100 realizes both high efficiency and ease-of-coding

[1] Asahi et al., submitted to IEEE TPDS

Node Performance and Node Scalability

FUJITSU

- Performance of magnetohydrodynamic simulations[1]
  - Modified leap frog method, memory-intensive

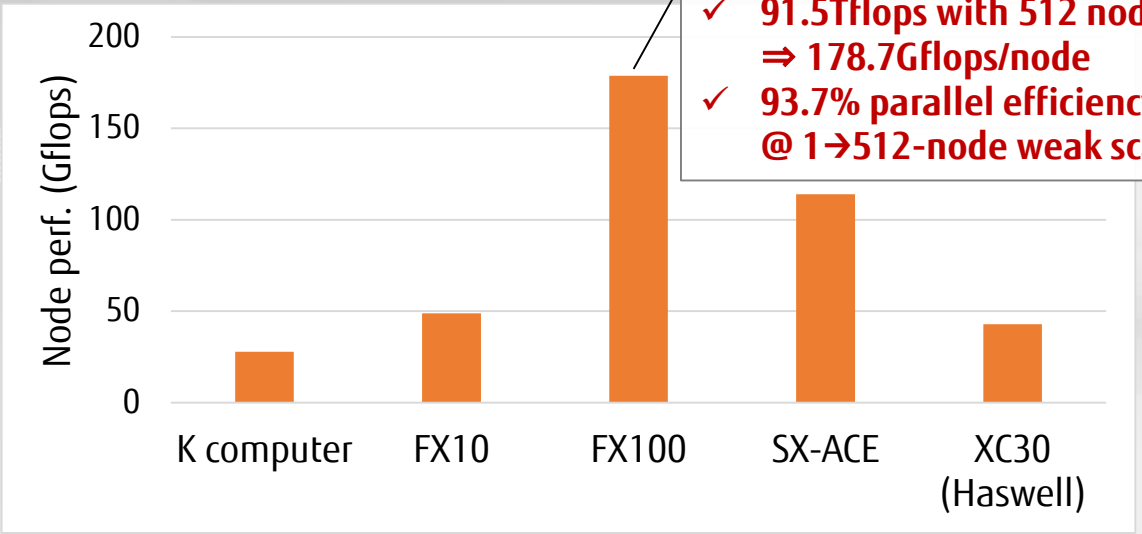

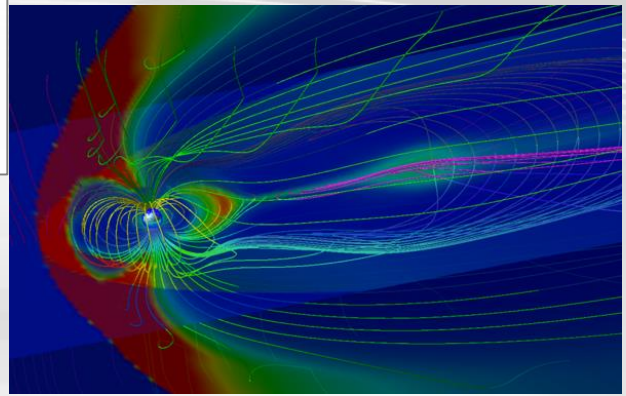

Simulation of Terrestrial magnetosphere
(By courtesy of Dr. Fukazawa @ Kyoto Univ.)

[1] Fukazawa et al., HPDC'16 poster 07

- FX100 shows higher performance per node and good node scalability

Production Runs on FX100 in JAXA†

FUJITSU

- Simulations for acoustic design of clustered supersonic jets[1]
  - LS-FLOW-HO: flux reconstruction method, memory-intensive

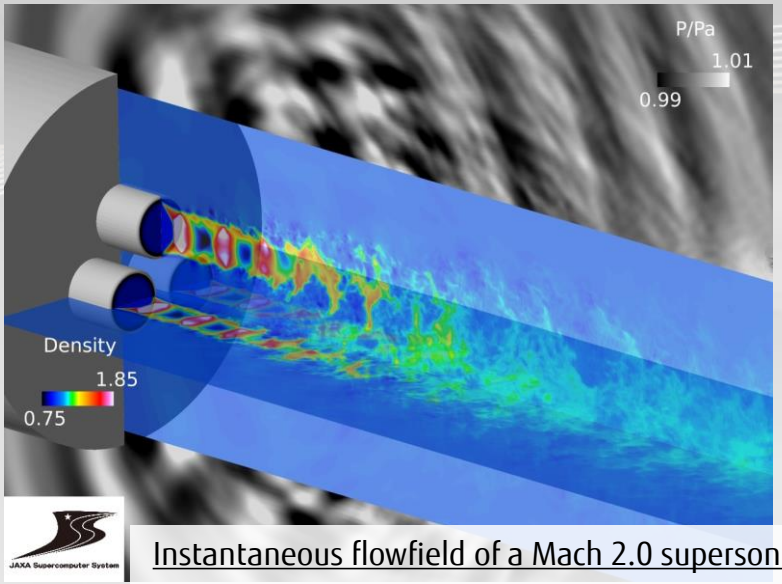

Instantaneous flowfield of a Mach 2.0 supersonic jet

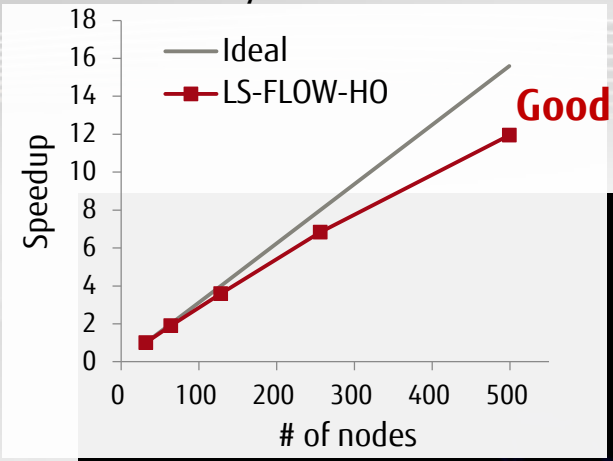

Good strong scalability is reported

Time averaged Mach number

(By courtesy of Dr. Haga @ JAXA)

[1] Haga et al., 48th Fluid Dynamics Conference / 34th Aerospace Numerical Simulation Symposium, 2C07

†: Japan Aerospace Exploration Agency

Real High System Utilization Rate

- High system utilization on RIKEN's HOKUSAI GreatWave (1,080 nodes/34,560 cores)[1]
- Semi-rectangular node allocation achieves **94% utilization**

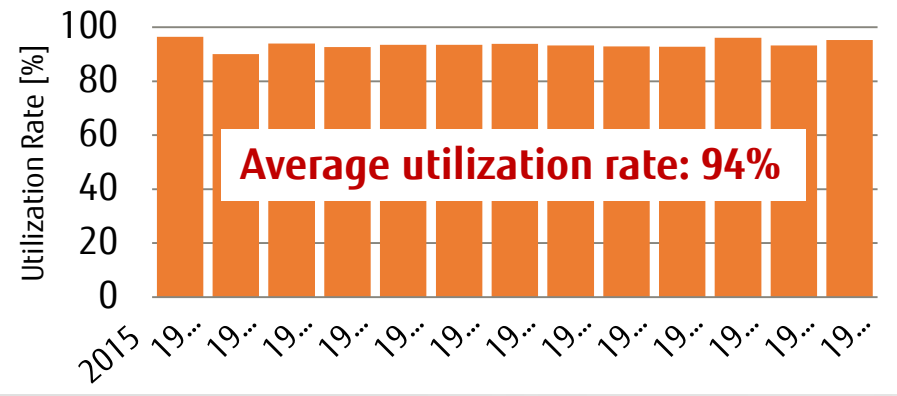

a. Rectangular node allocation

b. Semi-rectangular node allocation

Node allocation on 6D mesh/torus network

a. Rectangular node allocation (conventional)

- Best for MPI performance
- Fragmentation decreases utilization rate

b. Semi-rectangular node allocation (introduced in FX100)

- Good for MPI performance
- Very high utilization rate

[1] RIKEN symposium 2016 "Supercomputer HOKUSAI and Shoubu"

Toward Next Generation HPC

FUJITSU

- FX100 achieves high performance on various applications through continuously improving HPC technologies
- This technological evolution is advancing toward Post-K development

Post-K

K computer

PRIMEHPC FX10

PRIMEHPC FX100

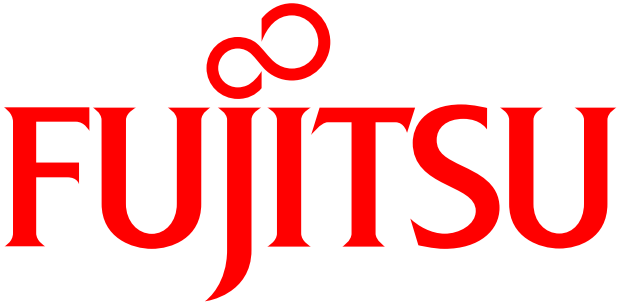

shaping tomorrow with you