FUJITSU

shaping tomorrow with you

FUJITSU Supercomputer PRIMEHPC FX100

The K computer and the evolution of PRIMEHPC

FUJITSU

Fujitsu has been developing supercomputers over 30 years, and will continue its development to deliver the best application performance.

Post-
K computer

K computer

SPARC64 VIIIfx: 8 cores / 128 GF
11 PF, 2010~

PRIMEHPC FX10

SPARC64 IXfx: 16 cores / 236.5 GF
23 PF, 2012~

PRIMEHPC FX100

SPARC64 XIfx: 32 cores / over 1 TF
Over 100 PF, 2015~

PRIMEHPC FX100 design concept

FUJITSU

Designed for massively parallel supercomputer system

- High performance for a wide range of real applications

Enhance and inherit K computer features

- Many-core CPU-based architecture for application productivity
- Enhanced VISIMPACT (hardware barrier synchronization, sector cache, etc.)

Introduce new technologies to Exascale computing

- HPC-ACE2 : Wide SIMD enhancements
- Assistant cores : Dedicated cores for non-calculation operation
- HMC : Leading-edge memory technology

SPARC64™ XIfx

FUJITSU

Over 1 TF high performance processor

- 32 compute cores
- 2 assistant cores: Offloading non-calculation operations
  ⟶ Daemons, IOs, non-blocking MPI functions, etc.

HPC-ACE2: ISA enhancements

- Two 256-bit wide SIMD units per core
- Addressing mode (stride load/store, indirect load/store)
- Cross lane operation (compress, permutation)

HMC support

- 480GB/s/node of theoretical memory throughput

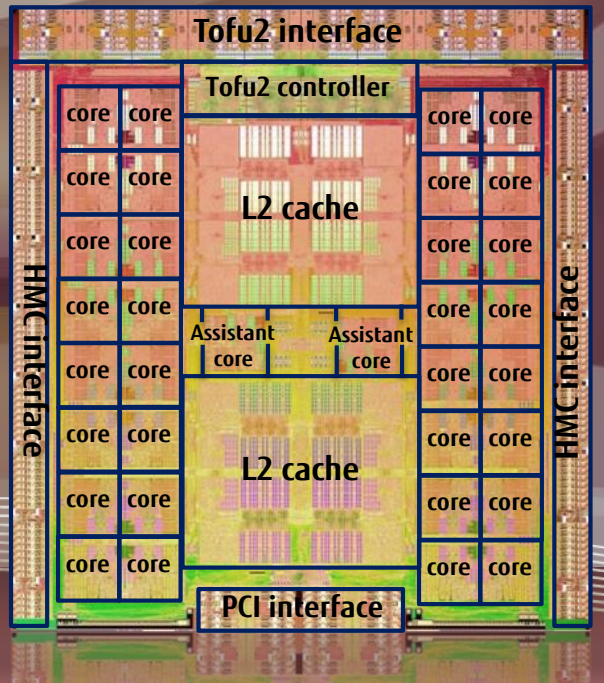

Tofu interconnect 2

FUJITSU

Enhanced Tofu interconnect

- Highly scalable, 6-dimensional mesh/torus topology
- Increased link bandwidth by 2.5 times to 12.5GB/s

CPU-integrated interconnect controller

- Reduced communication latency
- Improved packaging density and energy efficiency

Optical cable connection between chassis

- Enable flexible installation

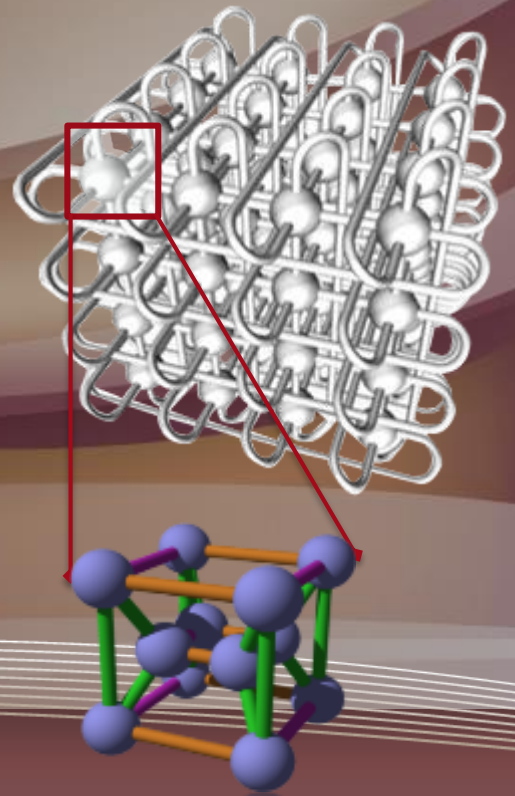

Technical Computing Suite

FUJITSU

Enhanced software stack developed by FUJITSU

Applications

Technical Computing Suite

Management software	High Performance File System FEFS	Programming Environment
System management	Lustre-based distributed file system (enhanced for FX100)	MPI, OpenMP, XPFortran
Job management		Compilers(C,C++,Fortan) Mathematical libraries
		Debugging and tuning tools

Linux-based OS enhanced for FX100

PRIMEHPC FX100

Performance highlights

FUJITSU

Node performance is 2.9 to 4.7 times higher than previous generation.

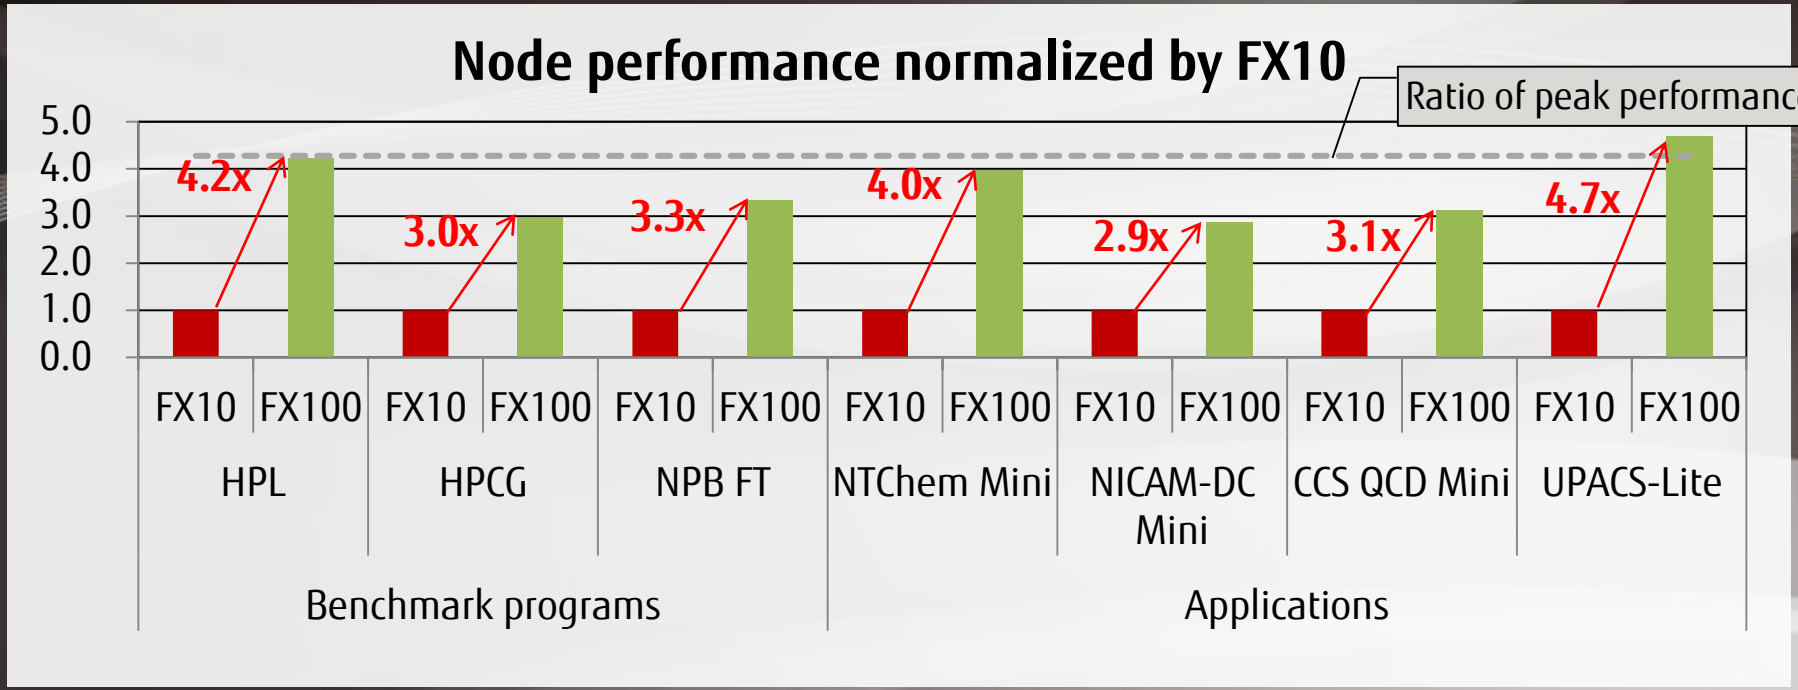

Results of major benchmarks - LINPACK

FUJITSU

- The first four FX100 systems have been delivered to customers, with >90% efficiency.

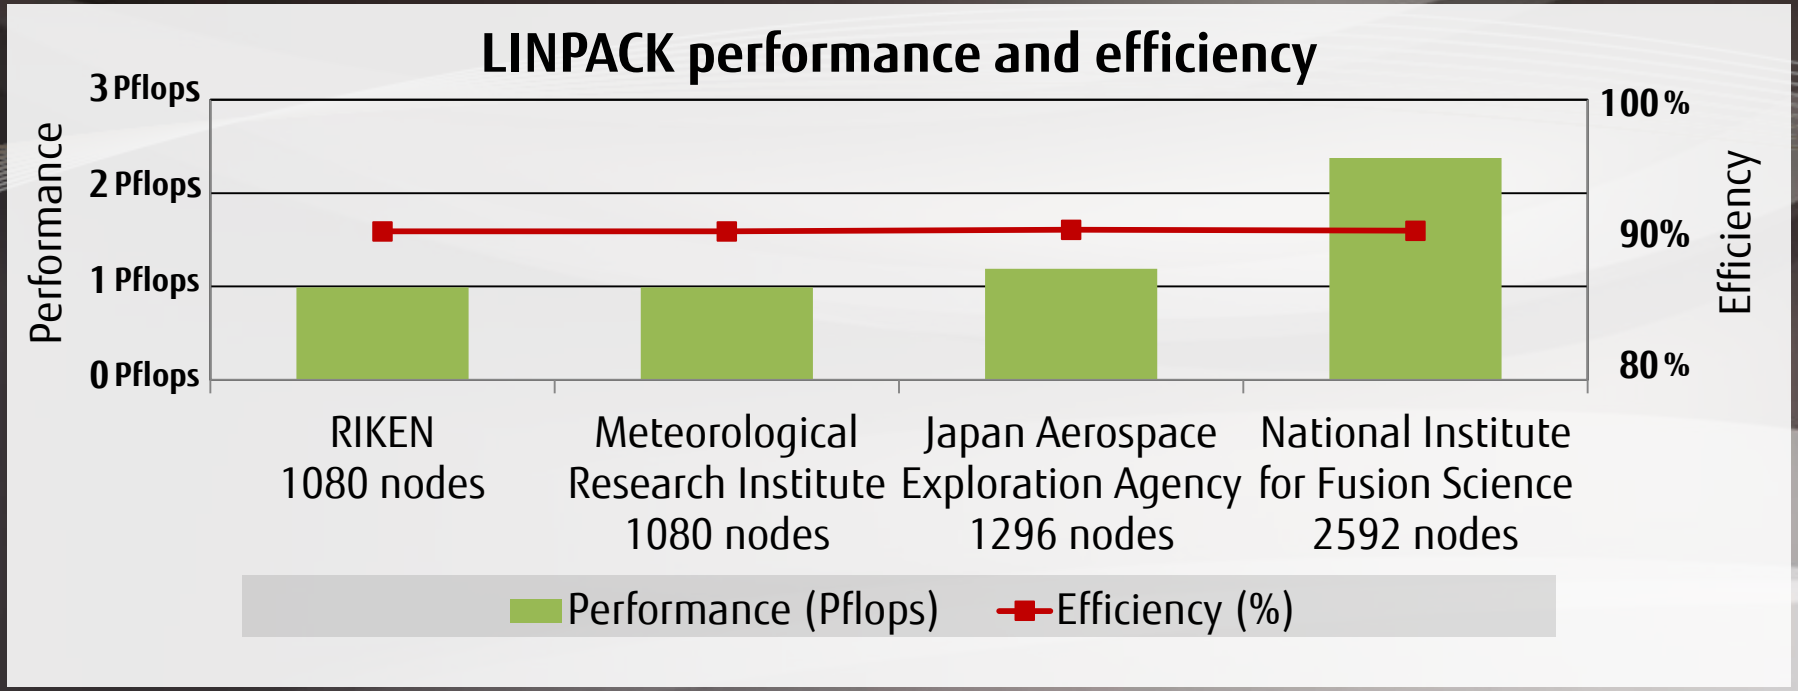

Results of major benchmarks - HPCG

FUJITSU

- Solves large sparse linear systems, using Conjugate Gradient method.
- Enhanced memory bandwidth enables higher performance, keeping ease-of-use for programming.

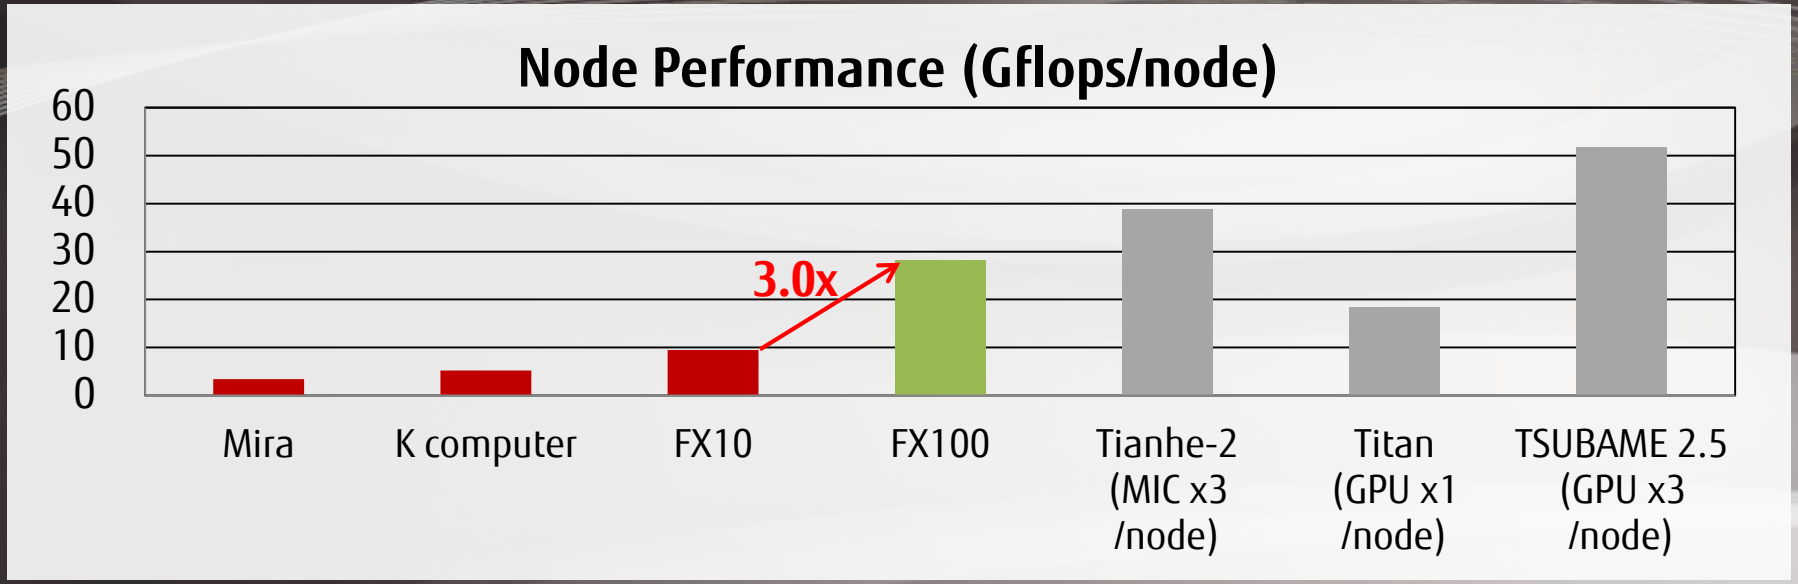

Results of major benchmarks – NPB FT

FUJITSU

- Solves a 3D partial differential equation using FFT.
- Improved by higher cache/memory throughput, as well as increased CPU cores and SIMD width.

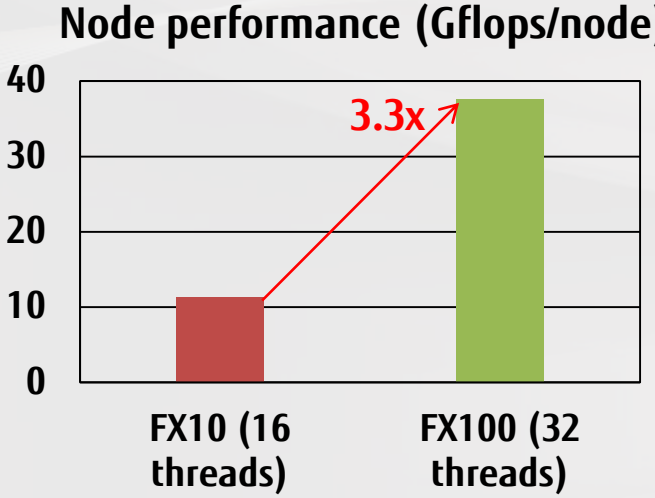

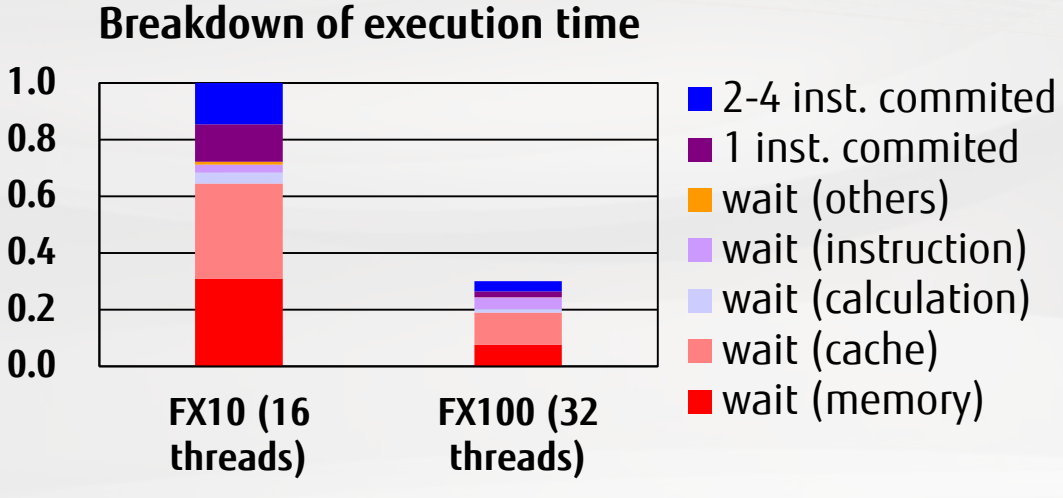

Application performance - NTChem-MINI†

FUJITSU

- A mini-app. of the ab-initio quantum chemistry for the molecular electronic structure calculation
- 4.0x faster on FX100, close to the ratio of peak performance (4.4x)

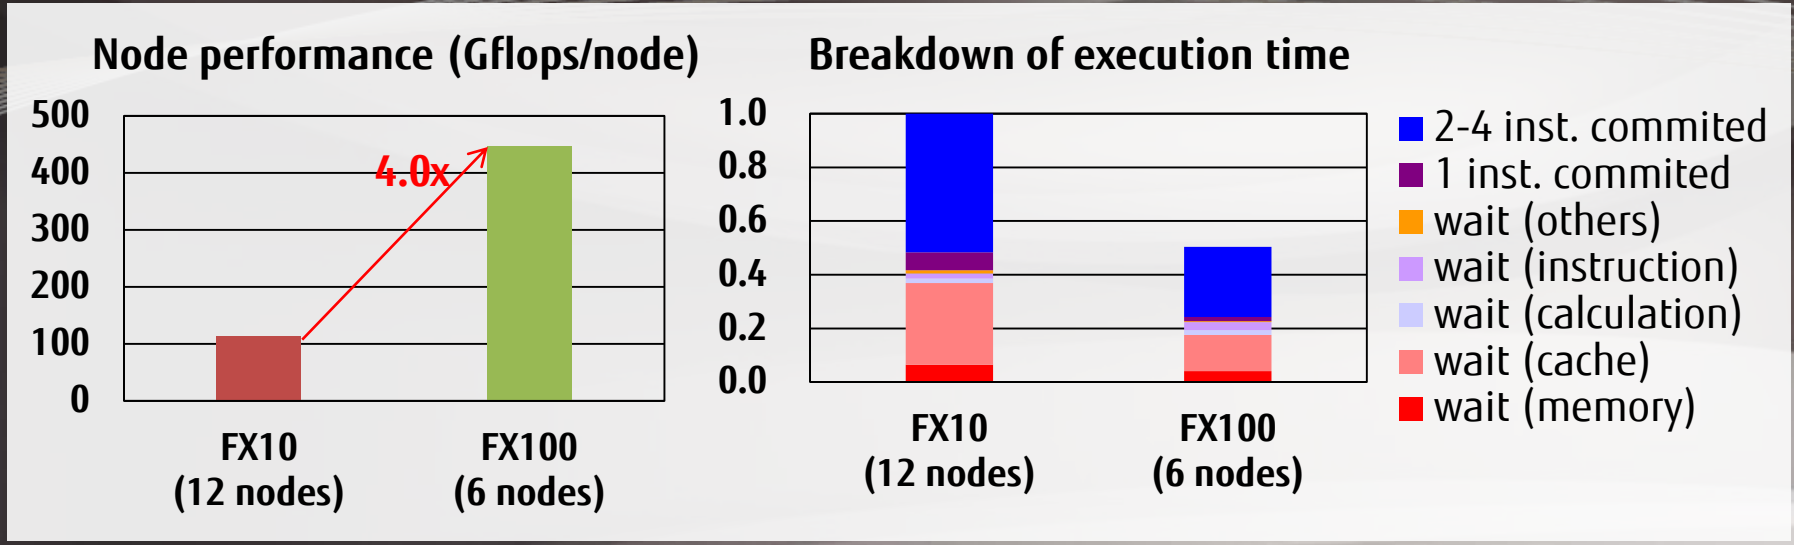

† https://github.com/fiber-miniapp/ntchem-mini

Application performance - NICAM-DC-MINI†

FUJITSU

- A mini-app. derived from NICAM-DC, dynamic phase of NICAM (Nonhydrostatic ICosahedral Atmospheric Model)
- 3.1x higher memory throughput in “Vertical Implicit” calculation speeds up 1.5x with half the number of FX10 nodes.

Node performance (Gflops/node)

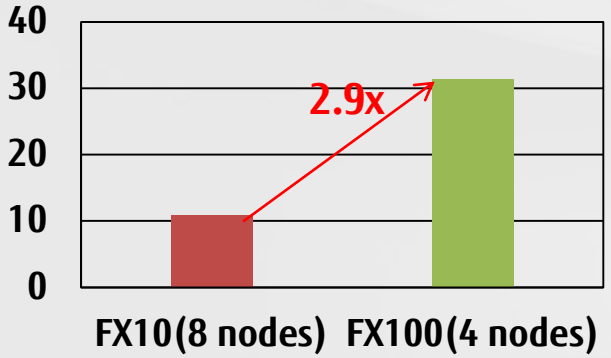

Breakdown of execution time ("vi_path2" routine)

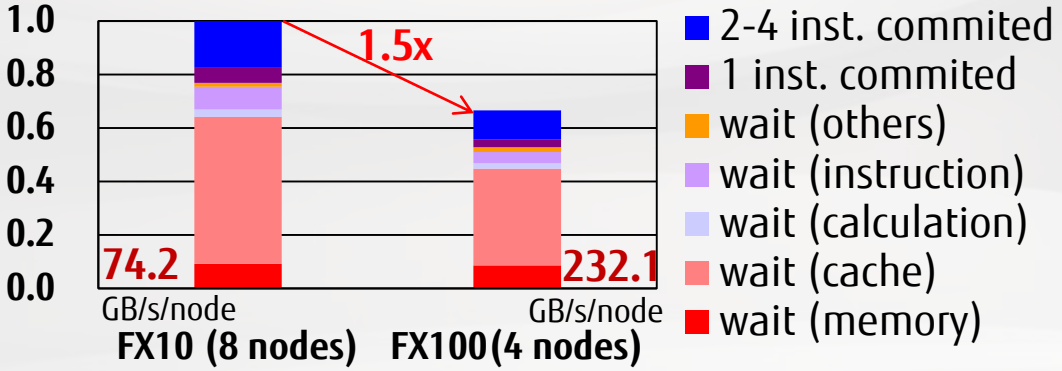

† https://github.com/fiber-miniapp/nicam-dc-mini

Application performance – CCS QCD-MINI†

FUJITSU

- A mini-app. code for a lattice QCD problem (32^4 size)
- Enhanced memory bandwidth boosts the performance and Sector Cache mechanism promotes data reuse on L2$.

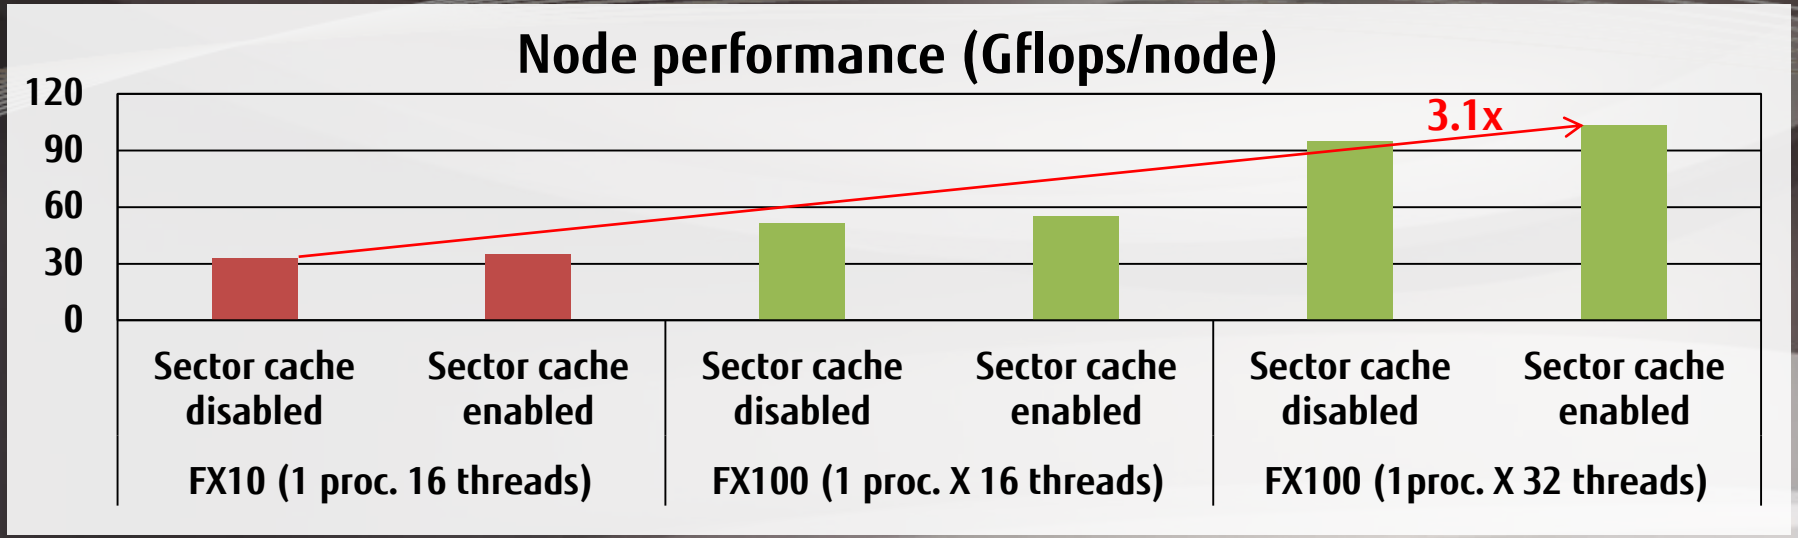

† https://github.com/fiber-miniapp/ccs-qcd

Application performance – UPACS-Lite

FUJITSU

- Calculation of compressible fluid dynamics, 160^3 grid points per node
- Enhanced memory bandwidth and other CPU improvements boost the performance

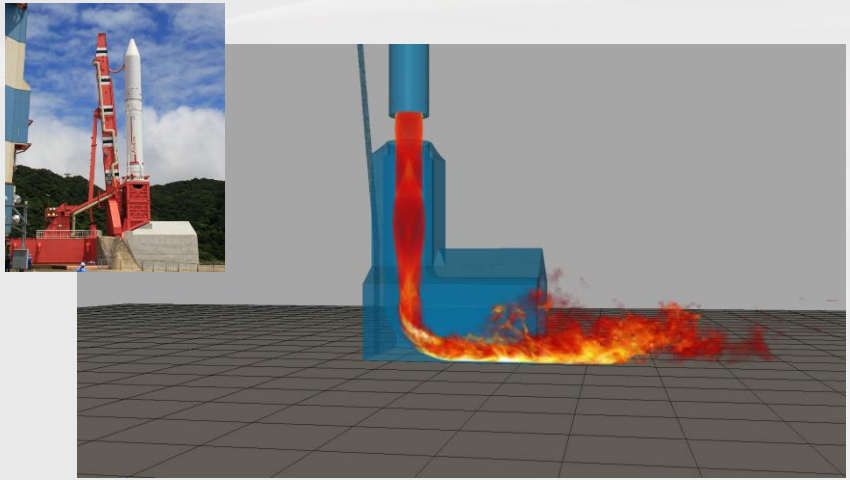

Example: simulation for acoustic design of a launch pad

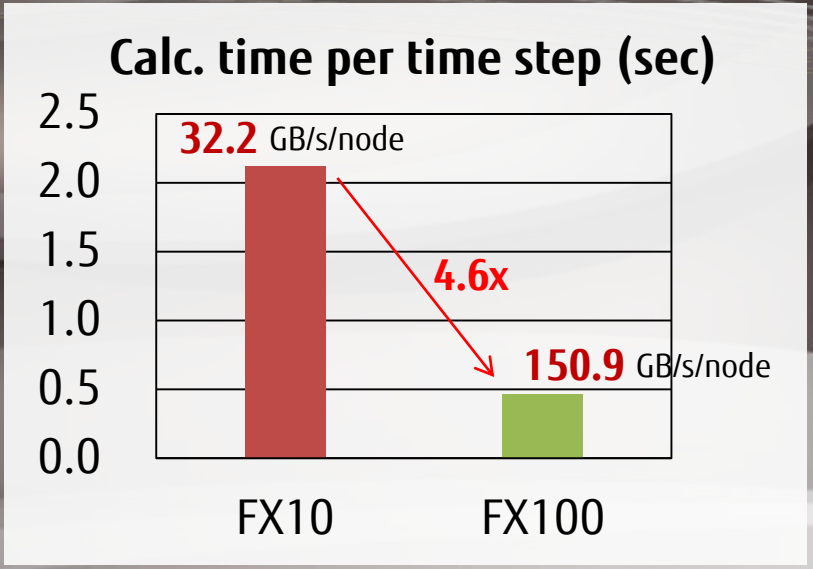

(Courtesy of Japan Aerospace Exploration Agency)

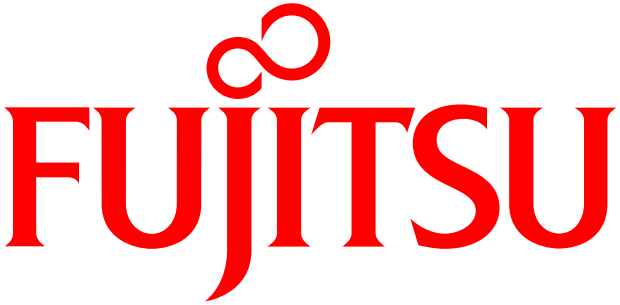

shaping tomorrow with you