

Supercomputer Fugaku and Fujitsu commercial supercomputers

Toshiyuki Shimizu

2020.11.17




FUJITSU LIMITED

Update at SC20 in appendix

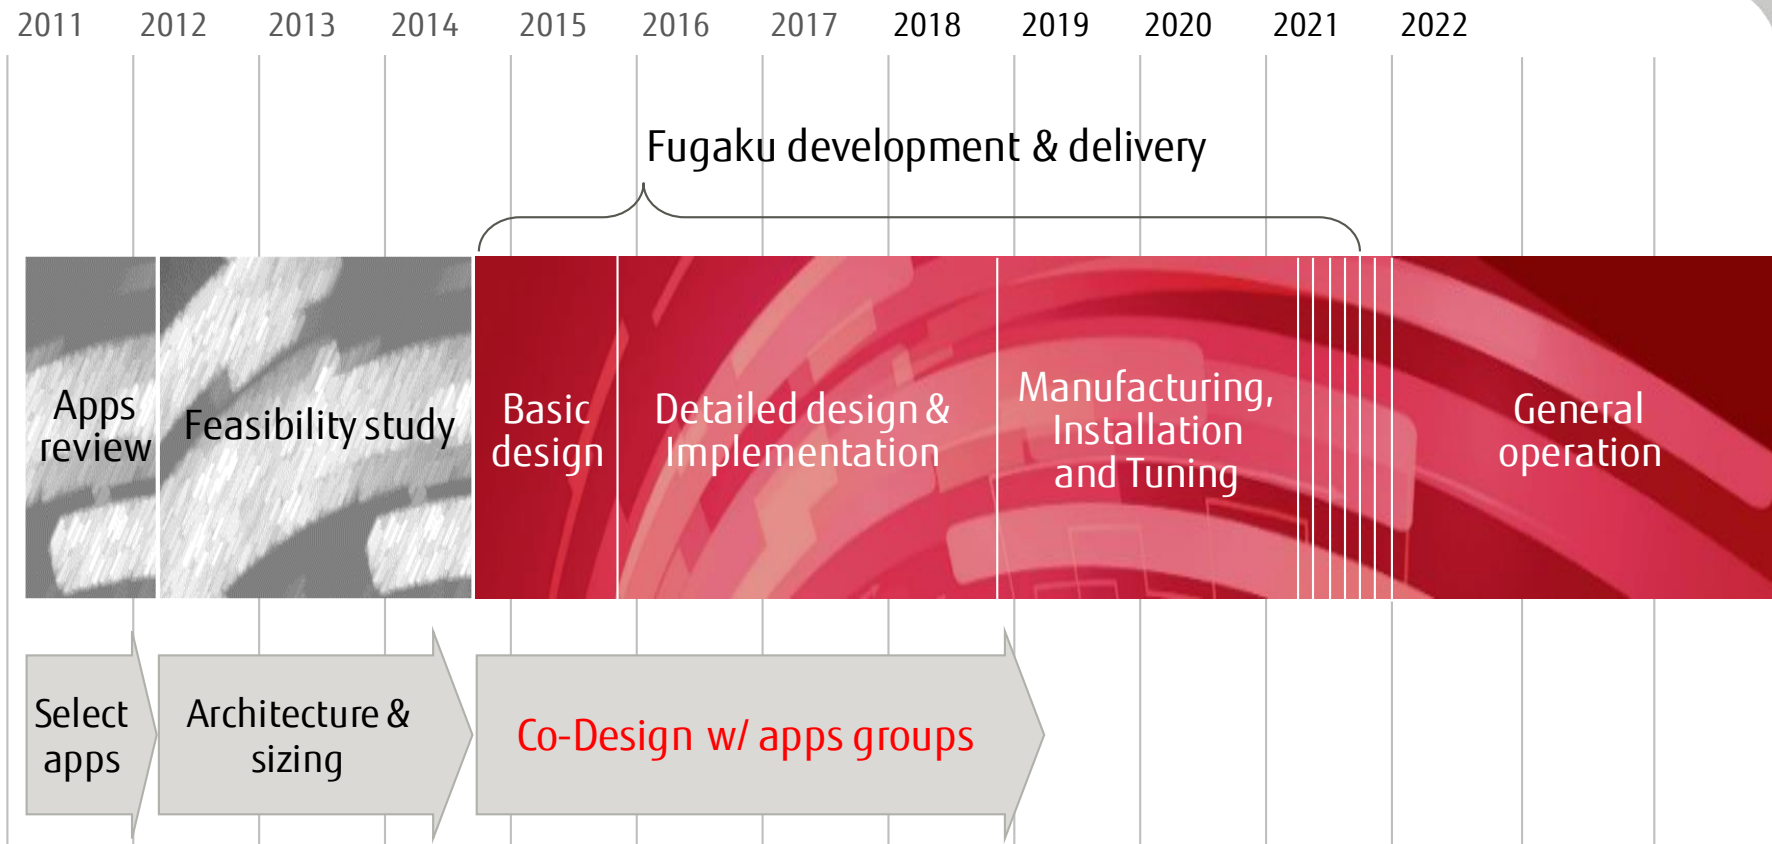
- Supercomputer Fugaku project
 - Approach & design results
 - Benchmark results & analysis
- Fugaku and FUJITSU Supercomputer PRIMEHPC FX1000/FX700
 - Specification and software stack
 - OSS and ISV applications
- Approach and status for AI
- Summary

Design targets and approaches for Fugaku



Target	Approach
 Application performance	Co-design w/ application developers and Fujitsu-designed CPU core w/ high memory bandwidth utilizing HBM2
 Power efficiency	Leading-edge Si-technology, Fujitsu's proven low power & high performance logic design, and "Power Knobs"
 Usability	Armv8-A ISA with Scalable Vector Extension ("SVE"), and Arm standard Linux

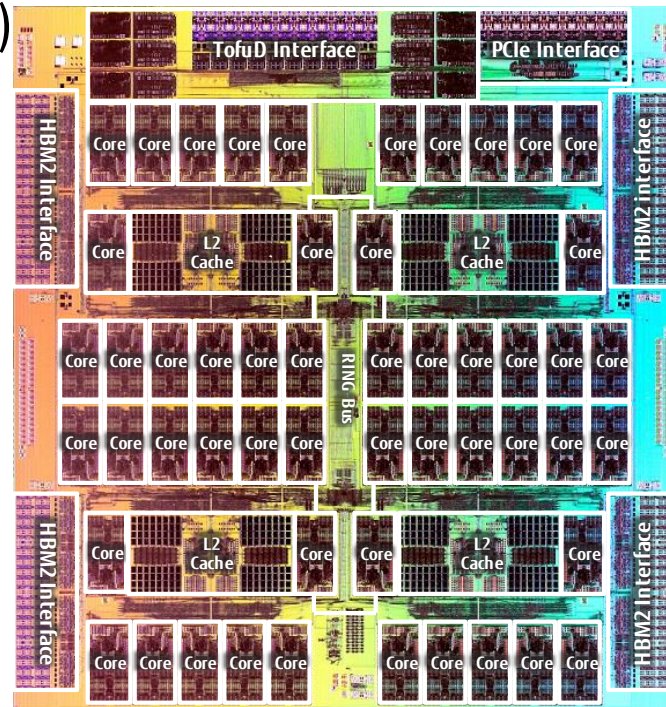
Fugaku project schedule



A64FX CPU for supercomputers

- All-in-one 7nm SoC w/ low power consumption
 - Armv8.2-A, 512-bit SVE (Scalable Vector Extension)
 - Four HBM2, 32 GiB per package
 - Tofu Interconnect D integrated
 - HW inter-core barrier & sector cache
 - 48 compute cores & 4 assistant cores for OS daemon & MPI offload

CPU core frequency	1.8	2.0	2.2	GHz
Peak DP perf (FP64)	2.7	3.0	3.3	TFLOPS
Peak SP perf (FP32)	5.5	6.1	6.7	TFLOPS
Peak HP perf (FP16)	11	12	13	TFLOPS
Memory peak bandwidth	1024			GB/s



Fugaku ranked at 1st place in all categories, June 22, 2020



2.8x of 2nd's



4.6x of 2nd's




2.6x of 2nd's



3.0x of 2nd's

TOP500 #1 history from K computer to Fugaku

- Performance improvement from K computer is about 40x while power consumption is only 2.2x => 12.66MW vs 28.33MW
- Good scalability and HPL execution efficiency of Tofu Interconnect D, MPI, job management software, and reliable hardware





Scalability & HPL execution efficiency

System	# of nodes	HPL eff	Interconnect
Fugaku (富岳)	152,064	80.87%	TofuD
Summit	4,356	74.01%	Infiniband
SunwayTaihuLight	40,960	74.15%	Custom
Tianhe-2	16,000	61.68%	Custom(Fat tree)
Titan	18,688	64.88%	Gemini
Sequoia	98,304	81.09%	Custom(5D torus)
K computer (京)	88,128	93.17%	Tofu

TOP500 #1 history from K computer to Fugaku

- Flagship machine power consumption (GF/W) is important due to societal demand on limiting power consumptions of facilities




TOP500 #1 history and Fugaku's choice

- Fugaku improved CPU core performance, avoiding external accelerators for apps execution performance

Way for perf.	Accelerator core	CPU core
GF/W improvement	Easier	Not easy
Apps development	Not easy	Easier
Apps domains	Narrow	Wider

Fugaku's choice




TOP500 #1 history and Fugaku's choice

■ Good results in all four benchmarks

■ **TOP500 #1**, **HPCG #1**, HPL-AI #1, Graph500 #1

Way for perf.	Accelerator core	CPU core
GF/W improvement	Easier	Not easy
Apps development	Not easy	Easier
Apps domains	Narrow	Wider

Fugaku's choice




HPCG number of Tianhe-2A is not published

HPCG results of TOP500 @ ISC20 + Fugaku 216-rack



- Fugaku's efficiency is very high and the same in both system sizes ~3%
- Optimized by splitting Symmetric Gauss-Seidel loops[**]
- Neighbor comm. using Tofu's 6-dir simultaneous comm.
- Tofu HW reduction for MPI_allreduce is effective for scalability



[*] Fractions of Fugaku are calculated by the peak performance of each sizes measured

[**] http://www.hpcg-benchmark.org/downloads/sc16/HPCG_on_the_K_Computer.pdf


Fugaku and A64FX greenness on TOP500, June 22, 2020

■ Power efficiency in GF/W, w/ ACC and w/o ACC

Way for perf.	Accelerator core	CPU core
GF/W improvement	Easier	Not easy
Apps development	Not easy	Easier
Apps domains	Narrow	Wider

Fugaku's choice

A64FX demonstrating power efficiency comparable to latest accelerators, and 3x superiority cf. traditional CPUs



- Supercomputer Fugaku project
 - Approach & design results
 - Benchmark results & analysis
- Fugaku and Fujitsu supercomputer FX1000/FX700
 - Specification and software stack
 - OSS and ISV applications
- Approach and status for AI
- Summary

Fugaku and Fujitsu supercomputers

FUJITSU



© RIKEN

Model	FX700	FX1000	Fugaku「富岳」
Concept	Cooperation with standard technologies	Application performance, energy efficiency, and scalability	
CPU	A64FX x8 / chassis	A64FX x384 / rack	
Max CPU clock freq.	1.8 GHz / 2.0 GHz	2.2 GHz	
Interconnect	InfiniBand EDR	Tofu Interconnect D	
Cooling	Air	Water	
Dimension	2U rack mountable	Custom: 800 mm x 1,400 mm x 2,000 mm	
Mgmnt software	BCM, Warewulf, PBS, SLURM	Fujitsu Technical Computing Suite	

Powerful software stack for A64FX systems

■ RIKEN & Fujitsu software for scalability, Open env. for easier deployment

Applications

Fujitsu Technical Computing Suite / RIKEN-developed Open Source System Software

Management software

System management
for high availability &
power saving
operation

Job management for
higher system
utilization & power
efficiency

File system

FEFS
Lustre-based
distributed file system

LLIO*
NVM-based file I/O
Accelerator for Fugaku

Programming environment

XcalableMP, FDPS

Open MPI, MPICH(PiP), DTF

OpenMP, COARRAY, Math.libs.

C, C++, Fortran, AI frameworks

Debugging, Tuning tools, Containers

Open environment

Management software

BCM, Warewulf
PBS, SLURM

Programming env.

Fujitsu Compiler Package,
GCC, OpneMPI,
Arm Performance Lib.

Multi-Kernel System: RHEL8 & light-weight kernel (IHK/McKernel)

RHEL8 / CentOS



Fugaku

© RIKEN



FX1000


* LLIO for Fugaku only



FX700

OSS application performance results on FX1000

FX1000 (48 cores) vs x86 sever (48 cores) *



■ Xeon and A64FX w/ the same number of cores*

■ Up to 1.9 times faster

■ Up to 4.8 times lower energy

* A64FX on FX1000: 48 cores × 1 CPU (2.2 GHz)
Xeon Platinum 8268: 24 cores × 2 CPU (2.9 GHz)

■ Scalable performance obtained by


■ Enhanced microarchitecture for HPC

■ Energy-saving design & implementation

■ Performance update of OpenFOAM & FrontISTR from SC19 are apps tuning and compiler enhancement for optimization in SIMD operations

CAE application performance results on FX700

FX700 (48 cores) vs x86 sever (48 cores) *



■ Xeon and A64FX w/ the same number of cores*

■ Up to 2.8 times faster

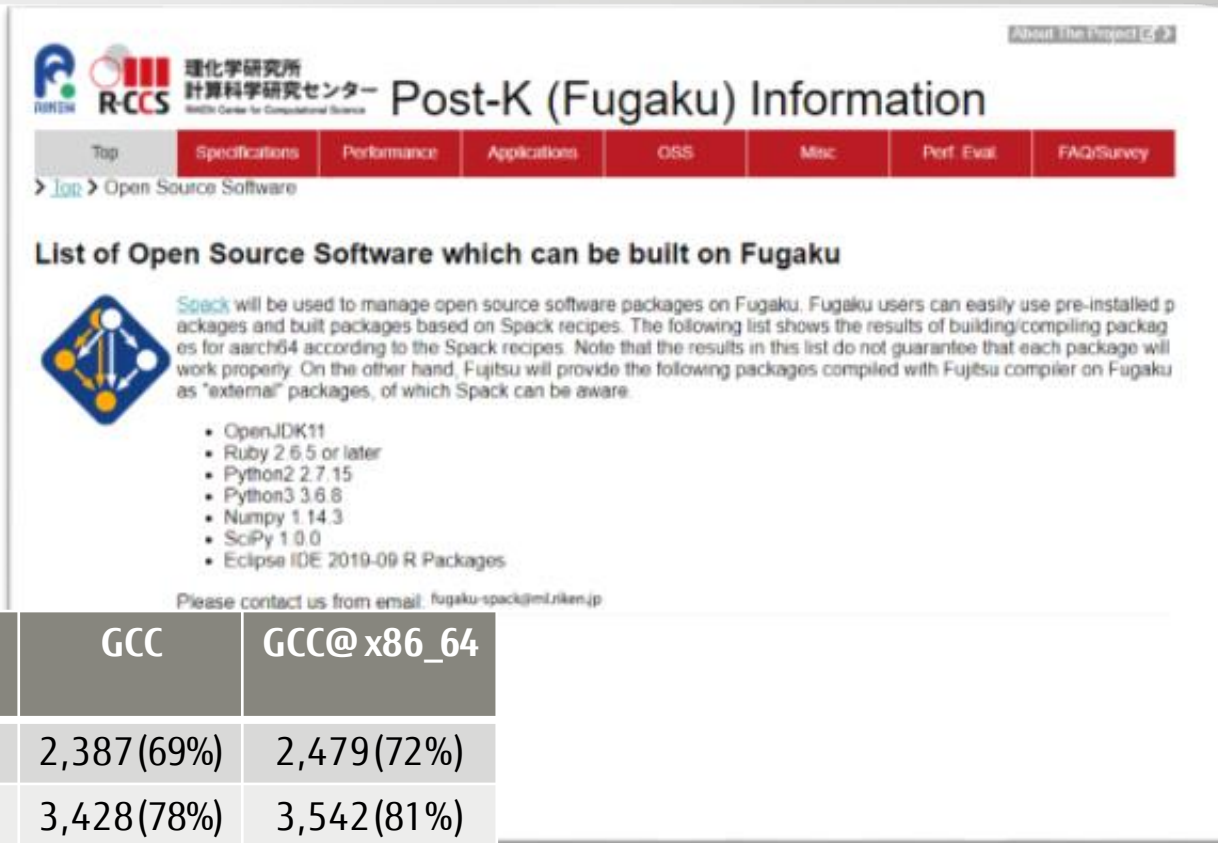
■ Up to 6.7 times lower energy

* A64FX on FX700: 48 cores × 1 CPU (2.0 GHz)
Xeon Platinum 8268: 24 cores × 2 CPU (2.9 GHz)

■ Power efficiency is also good even on the air cooled FX700 compare to the x86 server

Line-up of OSS on AArch64 with Spack

- RIKEN and Fujitsu are maintaining OSS packages to be built / compiled for *aarch64*
- Over 3000 OSS are built successfully for A64FX and close to x86_64



The screenshot shows the 'Post-K (Fugaku) Information' page from the RIKEN Center for Computational Science. It features a navigation bar with links like 'Top', 'Specifications', 'Performance', 'Applications', 'OSS', 'Misc', 'Perf. Eval.', and 'FAQ/Survey'. The 'OSS' link is highlighted. Below the navigation bar, there is a section titled 'List of Open Source Software which can be built on Fugaku'. This section includes a Spack logo and a paragraph explaining that Spack will be used to manage open source software packages on Fugaku. A list of supported packages is provided, including OpenJDK11, Ruby 2.6.5 or later, Python2 2.7.15, Python3 3.6.8, Numpy 1.14.3, SciPy 1.0.0, and Eclipse IDE 2019-09 R Packages. At the bottom, there is a contact information link: 'Please contact us from email: fugaku-spack@ml.riken.jp'.

Time	Reg. apps	Fujitsu compiler	GCC	GCC@ x86_64
SC19	3,451	2,072 (60%)	2,387 (69%)	2,479 (72%)
SC20	4,335	3,194 (73%)	3,428 (78%)	3,542 (81%)

<https://postk-web.r-ccs.riken.jp/oss/public/> captured on Oct. 8th, 2020

Commercial applications

- Fujitsu works with vendors to make commercial apps available for FX1000, also for FX700 and Fugaku with binary compatibility

Available

In a research & development phase (as of October 2020)

Engineering (Structural analysis, Fluid dynamics and Electronics)

Available in
Q4 '20

• LS-DYNA
(by Ansys, Inc.)

ADVENTURECluster
(by Allied Engineering Co.)

Altair Radioss™
(by Altair Engineering, Inc.)

Ansys Fluent
(by Ansys, Inc.)

Available in
Oct. '20

• Poynting
(by Fujitsu Limited)

 **CONVERGE**
CFD SOFTWARE
(by Convergent Science)

HELYX
(by ENGYS Ltd. & VINAS Co., Ltd.)

JMAG
Simulation Technology for Electromechanical Design
(by JSOL Corporation)

Chemistry*

Available for
FX1000.
Installed on
first customer
in Jul. '20

• Amber
• Gaussian16
(by Gaussian, Inc.)

Marc
(by MSC Software Ltd.)

scFLOW
(by Software Cradle Co., Ltd.)

Simcenter STAR-CCM+
(by Siemens Industry Software Inc.)

VASP

VPS (PAM-CRASH)
(by ESI Group)

*Collaboration with Australian National University

**All application names used in this slide are trademarks or registered trademarks of their respective vendors.

- Supercomputer Fugaku project
 - Approach & design results
 - Benchmark results & analysis
- Fugaku and Fujitsu supercomputer FX1000/FX700
 - Specification and software stack
 - OSS and ISV applications
- Approach and status for AI
- Summary

Deep Learning environment for A64FX systems

■ Deep Learning software stack is now available

- Fugaku: Pre-build environment is available (DNN library + {TensorFlow, Pytorch, Chainer})
- FX1000/700: DNN library published as OSS (<https://github.com/fujitsu/oneDNN>)

DL framework



DNN lib. (OneDNN_aarch64)

Fugaku, FX1000/700 hardware



© RIKEN



■ Benefit of supporting Deep Learning on A64FX systems

- The world's #1 performance of Fugaku can be utilized for Deep Learning
- Huge simulation results generated by Fugaku can be utilized to AI training efficiently


A64FX preliminary results for Deep Learning

■ Setup


- Using the same number of CPU cores
 - FX1000 single node (A64FX 2.2 GHz) vs. Xeon Platinum 8268 (24 core, 2.9GHz) x2
- ResNet50 (image classification)
- OpenNMT (natural lang. processing)

■ Results

- Performance:
 - Almost the same performance as Xeon
- Energy efficiency:
 - Up to 2.8x more efficient over Xeon



Training using fp32, PyTorch v1.5.0, OneDNN_aarch64, batch size 75 x 4proc.



Training using fp32, PyTorch v1.6.0, OneDNN_aarch64, batch size 3850 x 2proc.


Fugaku scalability of Deep Learning, ongoing project

- ResNet50 on multi-node training
 - Good performance scalability is observed
- Beyond 10K node evaluation is ongoing
 - Evaluation in other neural network models
 - Leveraging Data and Model parallelism

Fugaku

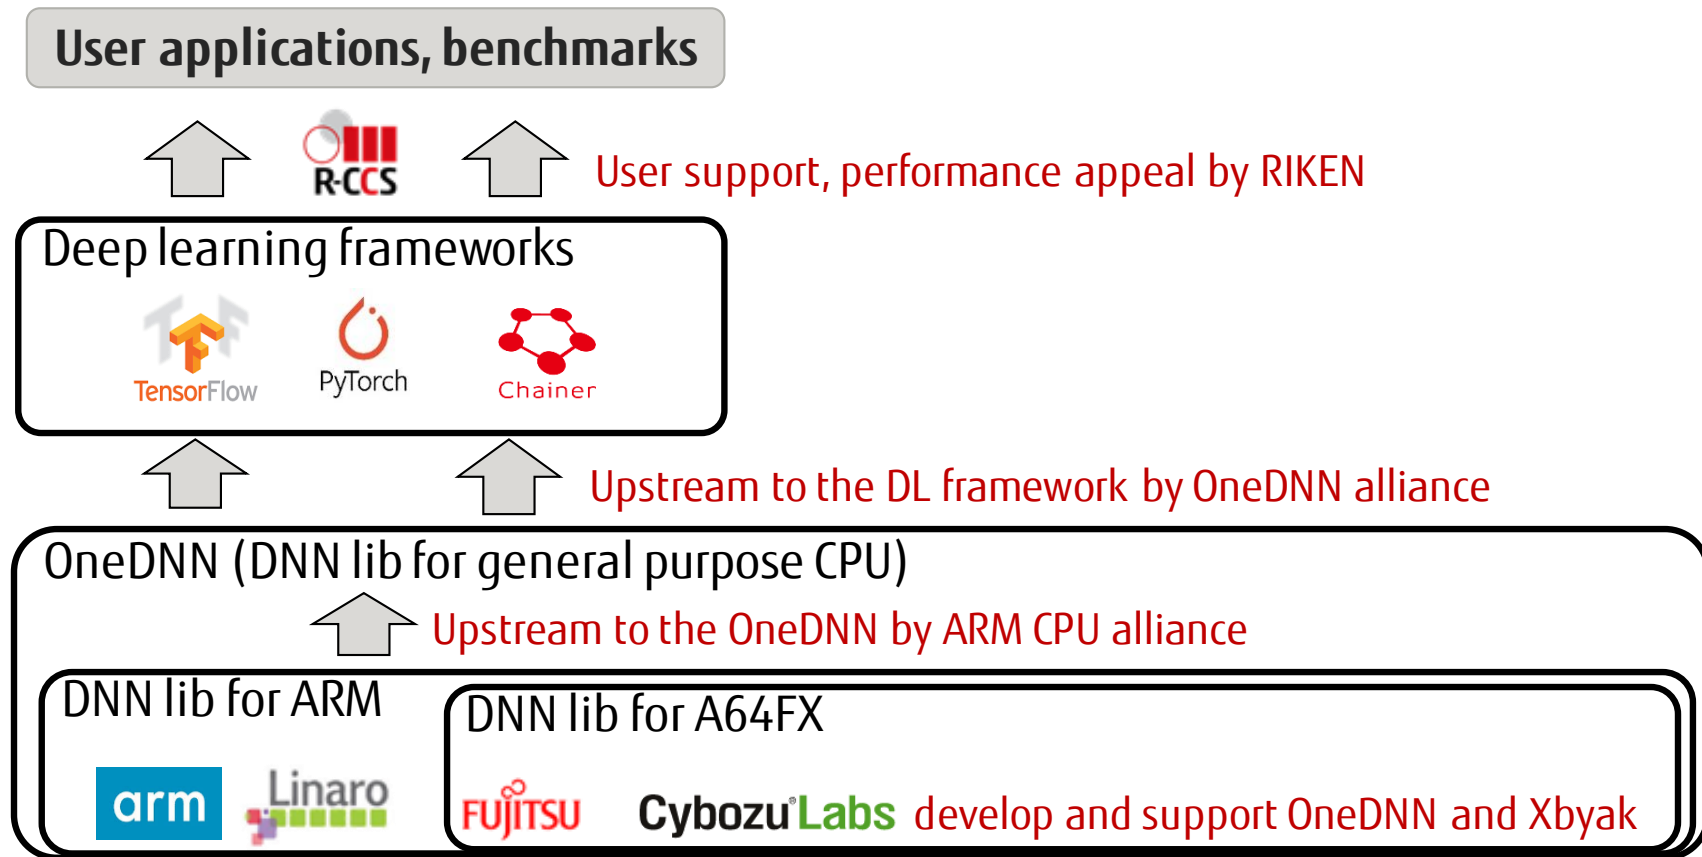


© RIKEN



Training using fp32, PyTorch v1.5.0, OneDNN_aarch64, Dummy data, weak scaling

Collaboration in building software stacks for AI



Summary

- Fugaku with its been co-designed approach, was ranked #1 in 4 major supercomputer rankings at ISC20, and runs apps at high performance w/ optimal power consumption

Visit Fujitsu virtual booth@SC20 for the latest info

Fugaku



© RIKEN

- Fugaku, FX1000, and FX700 equipped w/ Fujitsu designed A64FX CPU are utilized for many apps & AI research



PRIMEHPC
FX1000



PRIMEHPC FX700



Appendix:

Update at SC20

2020.11.17

FUJITSU LIMITED

Results of SC20 Fugaku Rankings at a Glance


■ Fugaku ranked #1 by large margin in ALL performance benchmarks

Benchmark	Unit	#1	Score@June	@Nov	#2	Score	#1/#2
TOP500	PFLOPS	Fugaku	415.53	442.01	Summit	148.60	3.0
HPCG	PFLOPS	Fugaku	13.37	16.00	Summit	2.93	5.5
HPL-AI	EFLOPS	Fugaku	1.42	2.00	Summit	0.55	3.6
Graph500	TTEPS	Fugaku	70.98	102.95	TaihuLight	23.76	4.3

Note: all the benchmarks on Fugaku were conducted on the full machine @ Nov. 2020

TOP500 #1 History from K computer to Fugaku

- Performance improvement from K computer is about 42x while power consumption is only 2.4x => 12.66MW vs 29.90MW
- Good scalability and HPL execution efficiency of TofuD interconnect, MPI, job management software, and reliable hardware




Scalability & HPL execution efficiency

System	# of nodes	HPL eff	Interconnect
Fugaku (富岳)	158,976	82.28%	TofuD
Summit	4,356	74.01%	Infiniband
SunwayTaihuLight	40,960	74.15%	Custom
Tianhe-2	16,000	61.68%	Custom(Fat tree)
Titan	18,688	64.88%	Gemini
Sequoia	98,304	81.09%	Custom(5D torus)
K computer (京)	88,128	93.17%	Tofu

TOP500 #1 History from K computer to Fugaku

- Flagship machine power consumption (GF/W) is important due to societal demand on limiting power consumptions of facilities

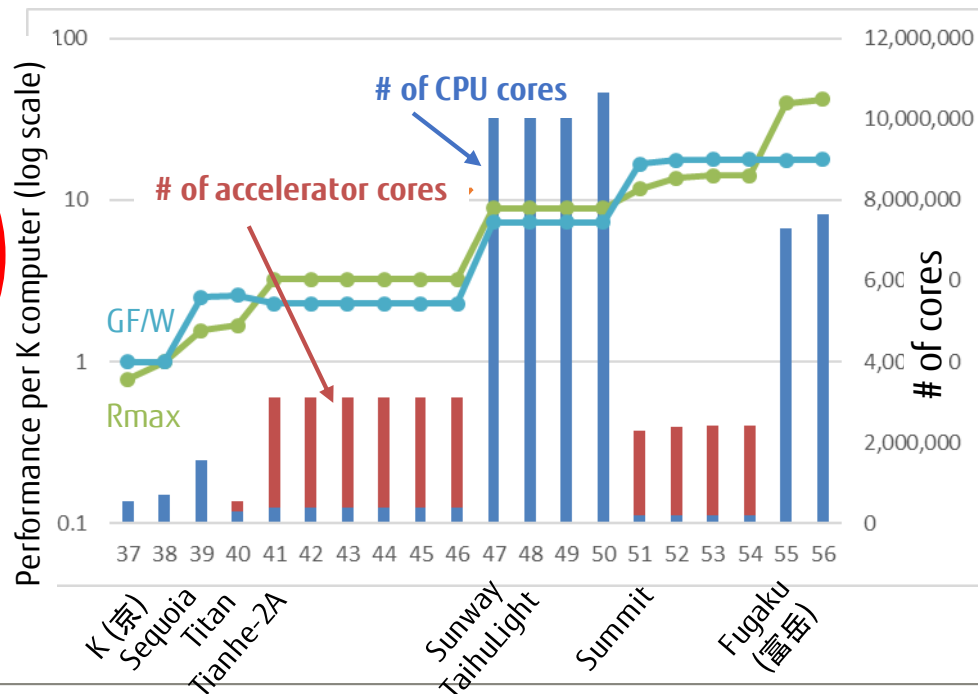


TOP500 #1 History and Fugaku's Choice

- Fugaku improved CPU core performance, avoiding external accelerators for apps execution performance

Way for perf.	Accelerator core	CPU core
GF/W improvement	Easier	Not easy
Apps development	Not easy	Easier
Apps domains	Narrow	Wider

Fugaku's choice



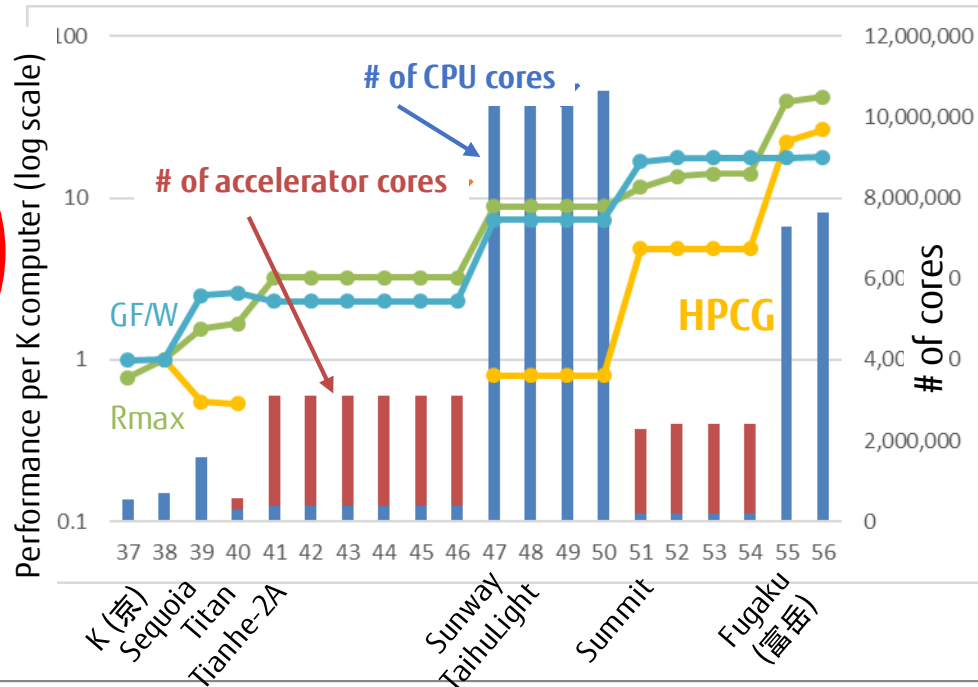
TOP500 #1 History and Fugaku's Choice

■ Good results in other benchmarks

■ **HPCG #1**, HPL-AI #1, Graph500 #1

Way for perf.	Accelerator core	CPU core
GF/W improvement	Easier	Not easy
Apps development	Not easy	Easier
Apps domains	Narrow	Wider

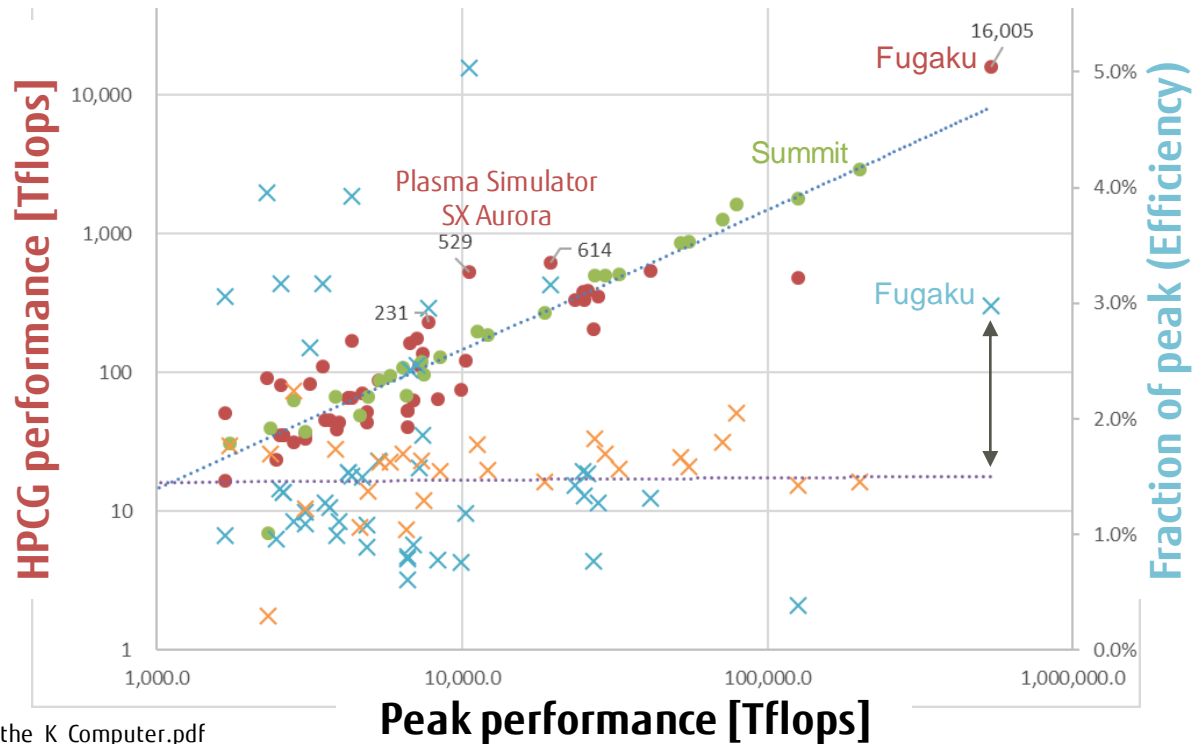
Fugaku's choice



HPCG number of Tianhe-2A is not published

HPCG Results of TOP500 @ SC20

- Fugaku's efficiency is very high and the same in both system sizes ~3%
- Optimized by splitting Symmetric Gauss-Seidel loops[*]
- Neighbor comm. using Tofu's 6-dir simultaneous comm.
- Tofu HW reduction for MPI_allreduce is effective for scalability




[*] http://www.hpcg-benchmark.org/downloads/sc16/HPCG_on_the_K_Computer.pdf


Fugaku and A64FX Greenness on TOP500

■ Power efficiency in GF/W, w/ ACC and w/o ACC

Way for perf.	Accelerator core	CPU core
GF/W improvement	Easier	Not easy
Apps development	Not easy	Easier
Apps domains	Narrow	Wider

Fugaku's choice





shaping tomorrow with you