An introduction to Fugaku

Fumiyoshi Shoji(shoji@riken.jp)
Operations and Computer Technologies Div.
R-CCS, RIKEN

December 15, 2021
NSCC Fugaku Call Briefing Session

RIKEN

- **One of the Japan's largest comprehensive research institution**
- **Founded in 1917**
- **about 3,000 researchers in various scientific domain from over the world**
- **RIKEN consists of many research centers**

- Center for Advanced Intelligence Project
- Center for Advanced Photonics
- Center for Biosystems Dynamics Research
- Center for Brain Science
- Center for Emergent Matter Science
- The Center for Emergent Matter Science
- Interdisciplinary Theoretical and Mathematical Sciences Program
- Center for Integrative Medical Sciences
- Nishina Center for Accelerator-Based Science
- Center for Quantum Computing
- Center for Sustainable Resource Science
- **Center for computational science (R-CCS)**
- BioResource Research Center
- SPring-8 Center

RIKEN Center for Computational Science (R-CCS)

- **Established on April 1, 2018 (former name: RIKEN AICS(2010-2017)).**
- **Missions**
  - Manage the operations and enhancement of the Fugaku.
  - Promote collaborative projects with a focus on the disciplines of computational and computer sciences.
  - Plot and develop Japan's strategy for computational science, including defining the path to exa-scale computing. -> Flagship2020 project (development of Fugaku)

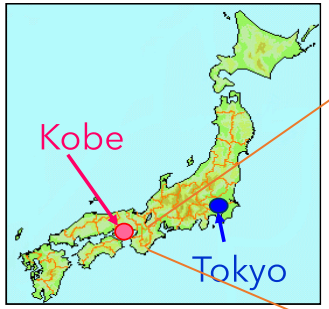

423km (263miles) west of Tokyo

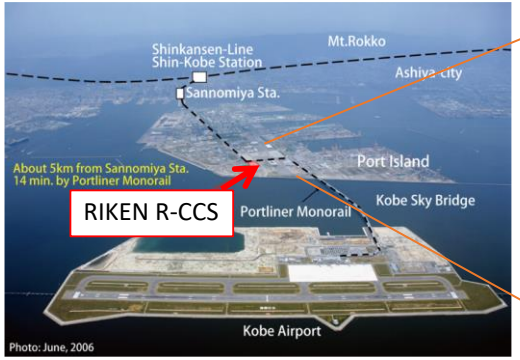

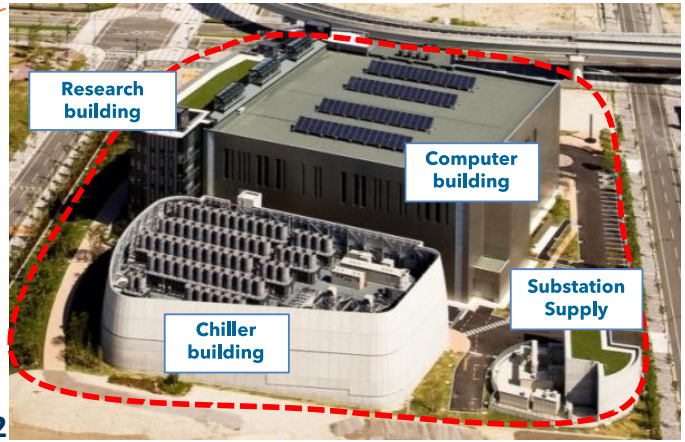

~20,000m²

A brief history of AICS and R-CCS

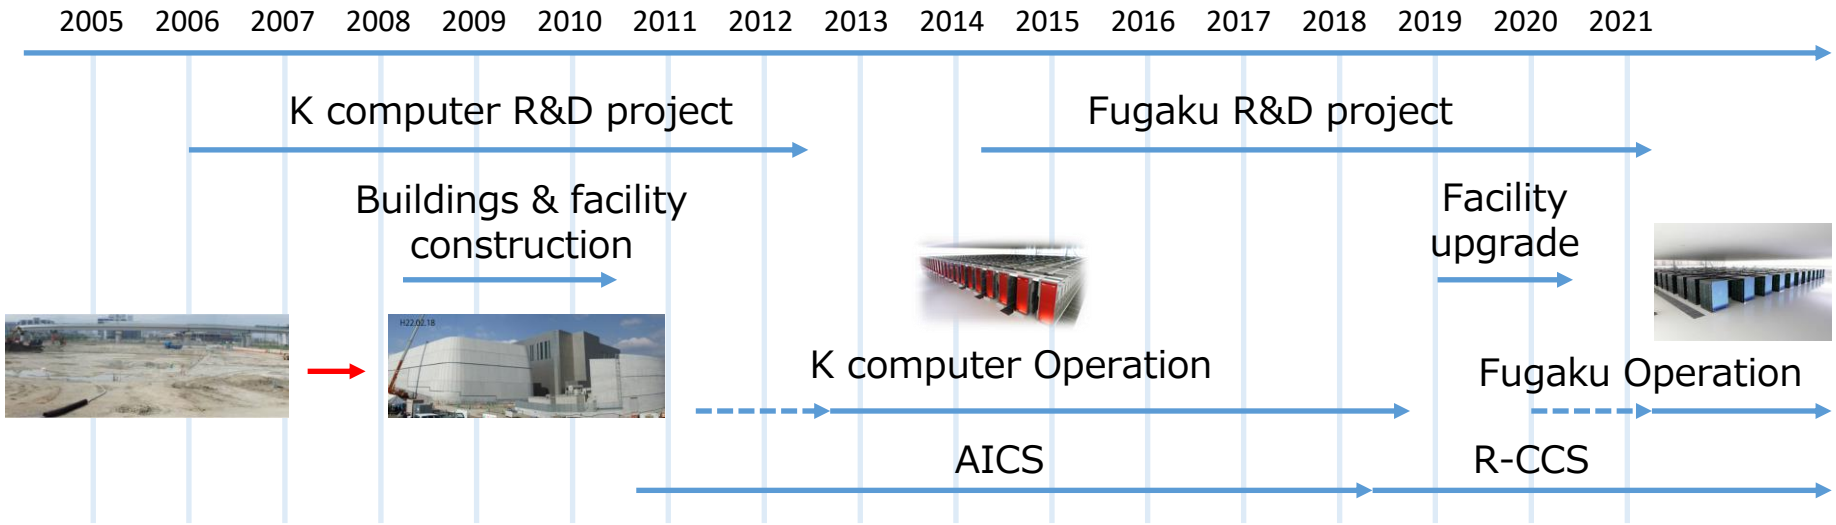

- RIKEN selected as the developer of next-generation super computer (later K computer) at October 2005
- Act on the Promotion of Public Utilization of the Specific Advanced Large Research Facilities modified to inclu de “Specified High-speed Computer Facilities” at April 2006
- Kobe selected as the construction site from 5 candidates by various aspects (location, available area, power s upply capacity, cooperation of local government, etc.) at 2008
- K computer was the first ”all Japan” HPC project which involved various field of science and engineering in Ja pan and AICS was the first national flagship HPC center in Japan.

Fugaku in SC21

Tue 16, 2021	Fugaku earns “quadruple crown” for 4th consecutive term Fugaku extends its reign as champion of supercomputers (RIKEN website)	
Thu 18, 2021	Fugaku took the first place for MLPerf HPC Benchmark Supercomputer Fugaku earns top rank on key machine learning benchmark (RIKEN website)	
Fri 19, 2021	R-CCS researchers received ACM Gordon Bell Special Prize for HPC-Based COVID-19 Research for novel aerosolized droplet simulation	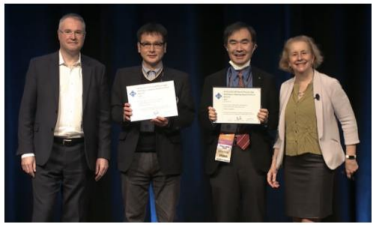

https://www.r-ccs.riken.jp/en/outreach/events/20211114-19/

Fugaku – successor of K computer -

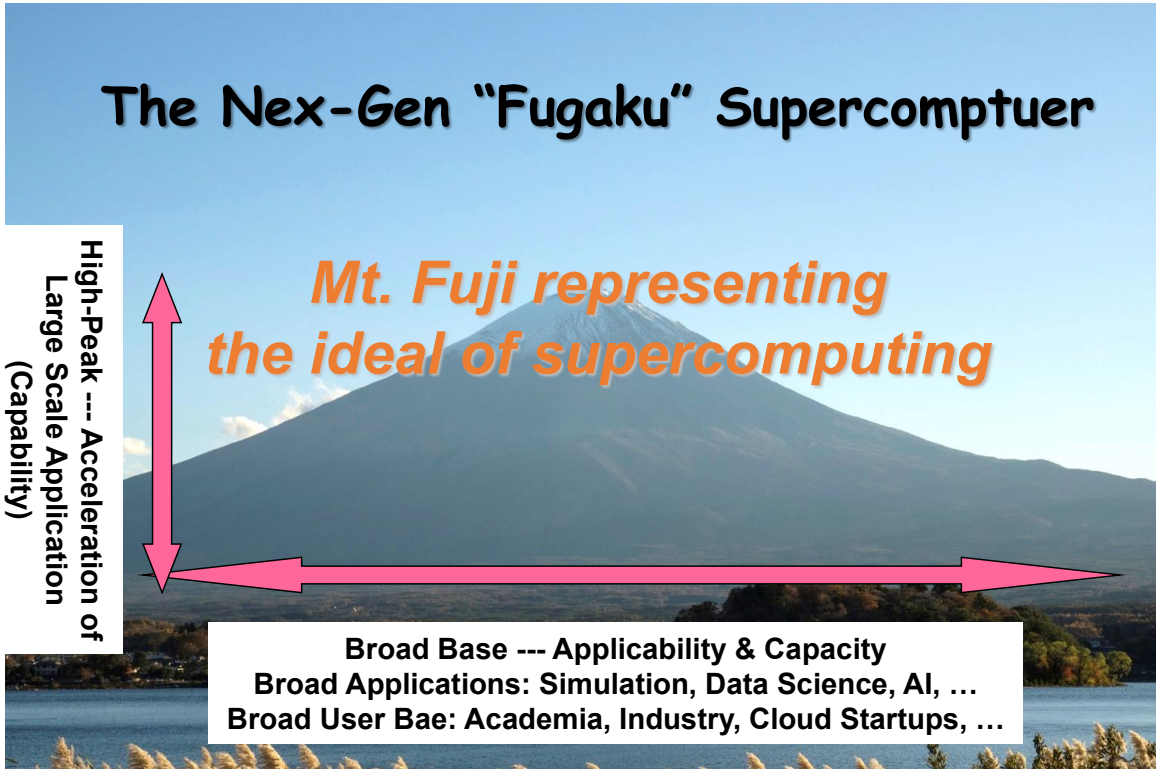

ふ が く
富 岳

Presentation by Satoshi Matsuoka @EEHPC SOP Workshop 2019

https://sites.google.com/view/eehpcsop2019/

Spec (Compute node)

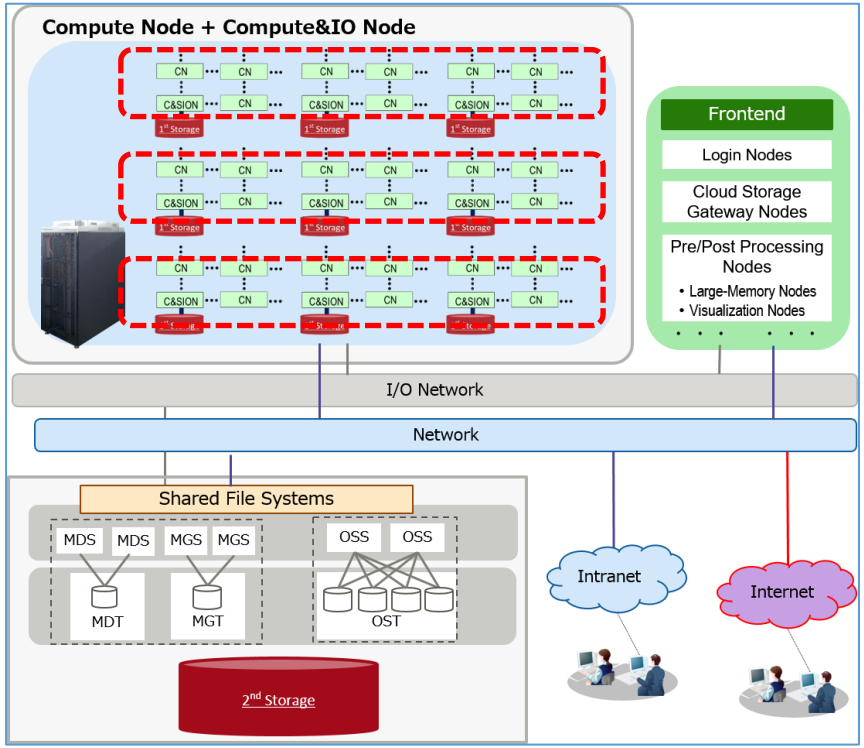

https://www.r-ccs.riken.jp/en/fugaku/about/

CPU A64FX

Architecture Information: Download from https://github.com/fujitsu/A64FX

	Description	
Architecture	Armv8.2-A SVE (512 bit SIMD)	
Core	48 cores for compute and 2/4 for OS activities	
	Normal: 2.0 GHz	DP: 3.072 TF, SP: 6.144 TF, HP: 12.288 TF
	Boost: 2.2 GHz	DP: 3.3792 TF, SP: 6.7584 TF, HP: 13.5168 TF
Cache L1	64 KiB, 4 way, 230+ GB/s (load), 115+ GB/s (store)	
Cache L2	CMG(NUMA): 8 MiB, 16 way Node: 3.6+ TB/s Core: 115+ GB/s (load), 57+ GB/s (store)	
Memory	HBM2 32 GiB, 1024 GB/s	
Interconnect	TofuD (28 Gbps x 2 lane x 10port)	
I/O	PCIe Gen3 x 16 lane	
Technology	7nm FinFET	

- Performance
  - Stream triad: 830+ GB/s
  - Dgemm: 2.5+ TF (90+% efficiency)
    - ref. Toshio Yoshida, "Fujitsu High Performance CPU for the Post-K Computer," IEEE Hot Chips: A Symposium on High Performance Chips, San Jose, August 21, 2018.

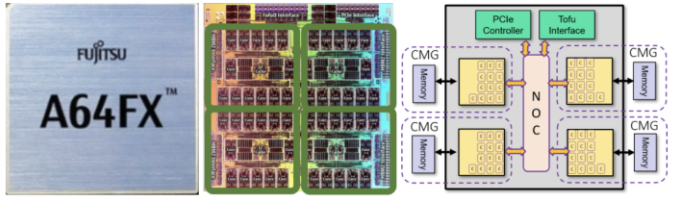

Spec (storage)

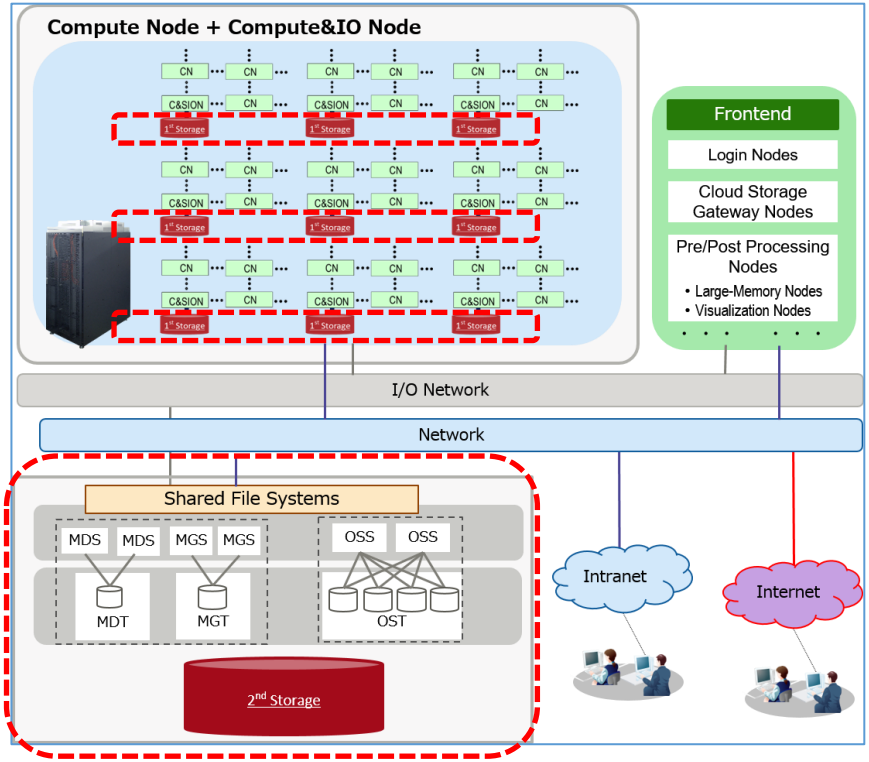

https://www.r-ccs.riken.jp/en/fugaku/about/

- 1st layer
  - SSD(1.6TB)/16node
    - cache of 2nd layer storage
    - temporary area
- 2nd layer
  - Luster based parallel file system (150PB)
- 3rd layer
  - offload to commercial cloud storage service
    - gateway servers for data transfer are available
    - fast connection between Oracle Cloud Infrastructure (OCI) has already been established via SINET5

login nodes and servers for pre-post processing

Fugaku users can use them with no-charge

Name	arch	# of core /node	Memory capacity	# of server	GPU
Login node	Intel	16x2	96/192GB	8	N/A
	Arm(TX2)	28x2	96GB	4	N/A
Login node with large memory	Intel	16x2	6,144GB	2	N/A
Visualization server	Intel	12x2	192GB	8	V100x2/node
Cloud storage gateway	Intel	16x2	256GB	2	N/A
Virtual Pre-Post server	Intel	26x2	384GB	30	N/A

Software stack

Programming environment

Compiler	Fortran 2008 and Fortran 2018 C11 with GNU and Clang extensions C++14 and C++17 with GNU and Clang extensions OpenMP 4.5 and OpenMP 5.0 Java
Parallel programing	XcalableMP [Details of XcalableMP (PDF 535 KB)] FDPS [Details of FDPS (PDF 260 KB)]
Script language	Python / Numpy / Scipy, Ruby
Numerical library	BLAS, LAPACK, ScaLAPACK SSL II (Fujitsu) EigenExa, Batched BLAS, 2.5D-PDGEMM

System software

Open-source management tool	Spack [Details of Spack (PDF 355 KB)]
Container, virtual machine	Singularity, KVM
OS	Red Hat Enterprise Linux 8 McKernel [Details of McKernel (PDF 641 KB)]
MPI	Fujitsu MPI (Based on OpenMPI), MPICH-Tofu (Based on MPICH) [Details of MPICH-Tofu (PDF 404 KB)]
File IO	LLIO DTF (Data Transfer Framework) [Details of DTF (PDF 220 KB)]

https://www.r-ccs.riken.jp/en/fugaku/about/

Spack: software package manager

- **Developed at LLNL and adopted as official software package manager of Exascale Computing Project (ECP)**
- **Enables easy software deployment for any architecture and scale.**
- **5,000+ packages have been registered on the latest version, e.g., abinit, lammps, openfoam, namd, blast, blat, …**
- **Simple CLI**
  - “spack find –x” provides list of installed packages
  - “spack load [app]” set up environment for the [app]
  - “spack list” provides list of compile-ready packages
  - “spack install [app]” build and install [app]
  - etc.

Open Source Software

Tuned by RIST
- OpenFOAM
- LAMMPS
- GROMACS
- Quantum ESPRESSO

Tuned by RIKEN and Fujitsu
- Chainer
- PyTorch
- Tensor Flow

https://www.hpci-office.jp/pages/hardware_software

OSS apps Performance

OSS Application Performance on FX1000/FX700

FUJITSU

Computation Speed of FX1000 and FX700 compared with x86 server

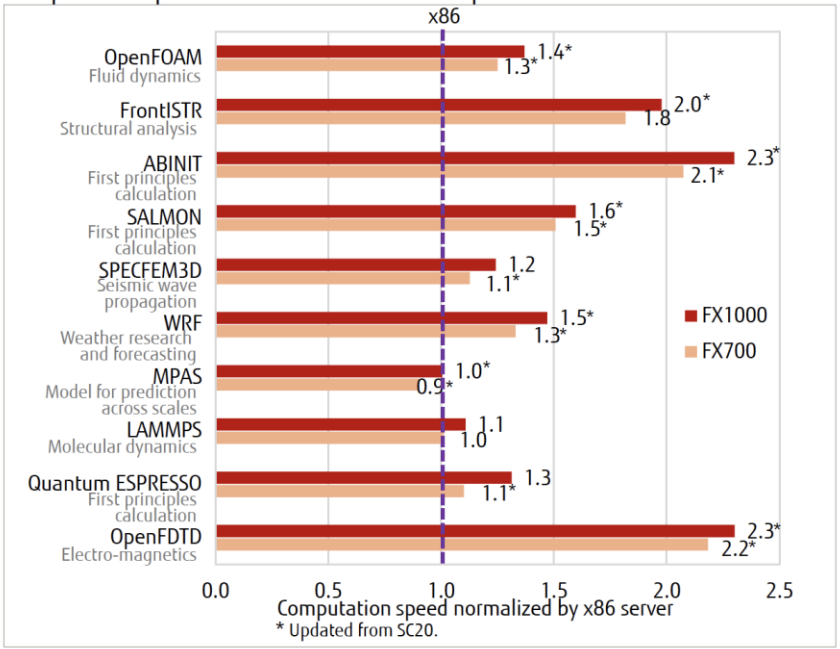

- Computation speed of FX1000 and FX700 is faster than x86 server up to 2.3 times with these OSS.
- Performance is improved by
  - Enhanced microarchitecture (SVE)
  - High bandwidth memory(HBM)
- Several software are improved by code tuning and enhancement of compiler and libraries.

FX1000

FX700

Benchmark Platform

FX1000: A64FX (2.2GHz x 48cores)
FX700: A64FX (2.0GHz x 48cores)
x86 server: Xeon 8268 (2.9GHz x 24cores) x2

12

Fujitsu in ISC21

https://www.fujitsu.com/downloads/SUPER/topics/isc21/applications-for-primehpc.pdf

Commercial software

Commercial Applications List

FUJITSU

- Availability of commercial applications for FX1000.
  Several applications were already verified on Fugaku.

Categories	Applications	Vendors	Availability	Notes
Engineering (Fluid dynamics)	CONVERGE	Convergent Science (East Asia distributor : IDAJ Co., LTD.)	Available in 3Q '21	Solver components only Verified on Fugaku
	Cradle CFD | scFLOW	Software Cradle Co., Ltd.	Beta ver. in 3Q '21	Verified on Fugaku
	Fluent	Ansys, Inc.	Alpha ver. in July '21	Solver components only
	HELYX	ENGYS Ltd.	Plan to be available *	
	Simcenter STAR-CCM+	Siemens Digital Industries Software, Inc.	Plan to be available *	
Engineering (Structural analysis)	ESI Virtual Performance Solution (VPS)	ESI Group	Available in 3Q '21	Explicit features only Verified on Fugaku
	LS-DYNA	Ansys, Inc.	**Available**	Verified on Fugaku
Engineering (Electromagnetics)	JMAG	JSOL Corporation	Plan to be available *	Verified on Fugaku
	Poynting	Fujitsu Limited	**Available**	Verified on Fugaku
Chemistry	Amber	University of California, San Francisco	**Available**	Collaboration with Australian National University
	Gaussian16	Gaussian, Inc.	**Available**	
Others	NAG Fortran Compiler	Numerical Algorithms Group Ltd	Available for FX700	Will be verified for FX1000 on customer request

*Release date will be announced later. **All application names used in this slide are trademarks or registered trademarks of their respective venders.

5

Fujitsu in ISC21
https://www.fujitsu.com/downloads/SUPER/topics/isc21/applications-for-primehpc.pdf

Deployment of commercial software

- **ISV apps which will be available till April are:**
  - Cradle CFD | scFLOW
  - CONVERGE
  - LS-DYNA
  - Poynting
  - Gaussian16

We plan to enrich the line-up of ISV apps.

Events for tuning technique promotion for A64FX

- Objective
  - To promote tuning technique for A64FX as quickly and widely as possible.
- Policy
  - Experts who have experience of tuning of large scale simulation on Fugaku introduce and share their technique and experience.
  - Materials and talks are in English and everything is recorded and opened.
  - Online based meeting and anyone (even non-Fugaku user) can join the meeting.
  - Organized by RIST and R-CCS
- Pointers for past meetings
  - https://www.hpci-office.jp/pages/e_meetings_A64FX
  - Slack workspace: https://join.slack.com/t/meetingforapp-16k4753/shared_invite/zt-ms93bjpy-BOGm1bvDsTTSJd5bsWL18g

Meeting	Date	Topics
The first meeting for application code tuning on A64FX computer systems	December 9, 2020	Performance tuning of Graph500 benchmark on Supercomputer Fugaku A64FX Tuning - SCALE on Fugaku -
The second meeting for application code tuning on A64FX computer systems	December 23, 2020	LQCD tuning on A64FX Optimization of GENESIS on Fugaku
The 3rd meeting for application code tuning on A64FX computer systems	February 3, 2021	Development of Massively Parallel DMRG for Fugaku HPL-AI bechmark on Fugaku
The 4th meeting for application code tuning on A64FX computer systems	March 17, 2021	Development of EigenExa from K to Fugaku, and beyond Fugaku Development of a deep neural network library for A64FX
The 5th meeting for application code tuning on A64FX computer systems	April 27, 2021	Performance tuning on LAMMPS for A64FX system CPU and Thread Parallelization Tuning of FFVHC-ACE on Fugaku
The 6th meeting for application code tuning on A64FX computer systems	June 30, 2021	Overview of Software Environment on Fugaku VELOC: Very Low Overhead Checkpointing System Hands-on session on VELOC

Portal web site of Fugaku

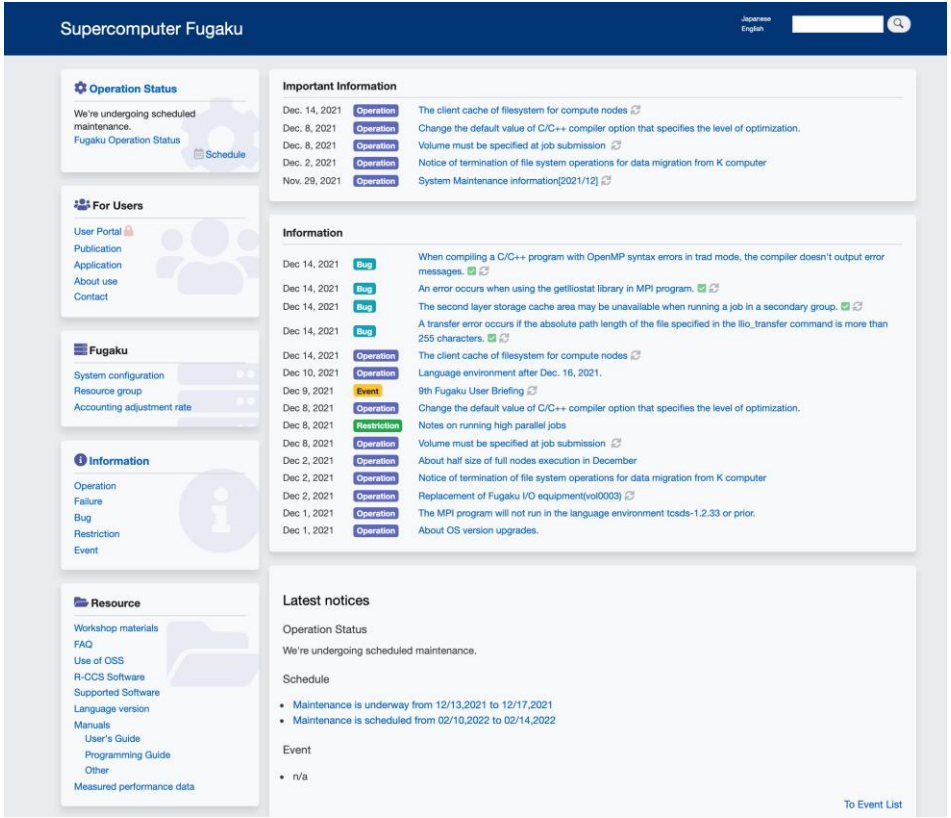

Fugaku users can access the web site via a client certificate authentication

Performance data of File I/O and MPI functions

- Benchmark results of File I/O and MPI functions are available in the portal site (Fugaku users only).

(Reference) LLIO throughput for large jobs

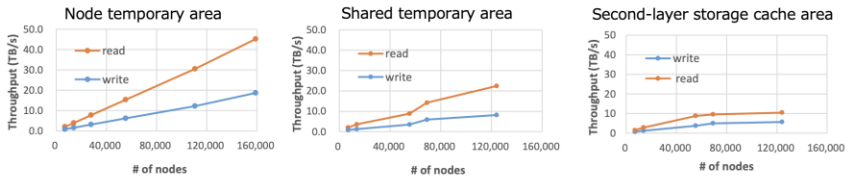

Write / SIO (MB/s)

# of ndoes	Node Temp	Shared Temp	2nd_cache
6,912	1,997	1,797	1,552
13,824	1,797	1,477	1,494
27,648	1,803		
55,296	1,785	1,009	1,142
69,120		1,363	1,210
124,000	1,764	1,053	770
158,963	1,876		

Read / SIO (MB/s)

# of ndoes	Node Temp	Shared Temp	2nd_cache
6,912	4,642	4,615	3,633
13,824	4,471	4,154	3,462
27,648	4,479		
55,296	4,425	2,559	2,666
69,120		3,294	2,323
124,000	4,405	2,891	1,418
158,963	4,548		

2.1 Allreduce

Communication time [sec]

Message size [bytes]	384 nodes 1536 procs	3072 nodes 12288 procs	27648 nodes 110592 procs
0	1.20E-07	1.30E-07	1.40E-07
4	4.55E-06	6.11E-06	1.24E-05
8	4.56E-06	6.29E-06	1.21E-05
16	3.11E-05	4.23E-05	6.62E-05
32	3.58E-05	5.10E-05	7.81E-05
64	3.56E-05	5.25E-05	7.80E-05
128	3.63E-05	5.42E-05	7.89E-05
256	3.91E-05	5.50E-05	8.54E-05
512	4.74E-05	6.46E-05	9.89E-05
1024	5.18E-05	7.26E-05	1.27E-04
2048	6.74E-05	9.42E-05	1.55E-04
4096	8.95E-05	1.36E-04	1.77E-04
8192	1.37E-04	2.20E-04	2.01E-04
16384	2.30E-04	3.94E-04	2.67E-04
32768	4.24E-04	7.40E-04	3.57E-04
65536	8.84E-04	4.22E-04	6.39E-04
131072	3.37E-04	4.81E-04	9.07E-04
262144	4.41E-04	6.04E-04	1.02E-03
524288	5.83E-04	7.37E-04	1.48E-03
1048576	8.27E-04	9.26E-04	1.67E-03
2097152	1.23E-03	1.31E-03	2.03E-03
4194304	2.02E-03	1.89E-03	2.74E-03

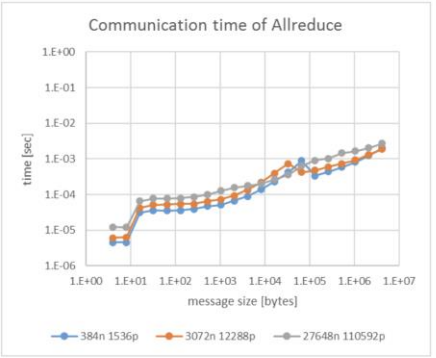

Summary

- **Fugaku is a general purpose supercomputer**
  - consisting of Arm-based CPU only (no accelerators)
  - developed by co-design of computer and computational scientists

- **Fugaku provides not only high-end computing capability but also high usability**
  - Software eco-system for Arm based architecture
  - Many optimized OSS/ISV software have already been available and Spack helps to use them easily.
  - Portal site provides basic and advanced information to use Fugaku efficiently

We look forward to serving Fugaku to all of you

Thank you very much for your attention