

Green500 BoF

Fujitsu A64FX prototype

a.k.a. micro-Fugaku!

Toshiyuki Shimizu

Nov. 20th, 2019

FUJITSU LIMITED

- TOP500 / Green500 achievements
- How Fujitsu worked
- Green500 measurement conditions and results
- Supercomputer standard and future
- Summary

Green500, Nov. 2019

A64FX prototype –
Fujitsu A64FX 48C 2GHz
ranked **#1** on the list

768x general purpose A64FX
CPU w/o accelerators

- 1.9995 PFLOPS @ HPL, 84.75%
- 16.876 GF/W
- Power quality level 2

The GREEN 500

HOME GREEN500 LISTS RESOURCES ABOUT MEDIA KIT

Home / Lists / November 2019

NOVEMBER 2019

- The most energy-efficient system and No. 1 on the Green500 is a new Fujitsu A64FX prototype installed at Fujitsu, Japan. It achieved 16.9 GFlops/Watt power-efficiency during its 2.0 Pflop/s Linpack performance run. It is listed on position 160 in the TOP500.
- In second position is the NA-1 system, a PEZY Computing / Exascaler Inc. system which is currently being readied at PEZY Computing, Japan for a future installation at NA Simulation in Japan. It achieve 16.3 GFlops/Watt power efficiency. It is on position 421 in the TOP500.
- The No.3 on the Green500 is AiMOS, a new IBM Power systems at the Rensselaer Polytechnic Institute Center for Computational Innovations (CCI), New York, USA. It achieved 15.8 GFlops/Watt and is listed at position 25 in the TOP500.

Green500 List for November 2019

Listed below are the November 2019 The Green500's energy-efficient supercomputers ranked from 1 to 10.

Note: Shaded entries in the table below mean the power data is derived and not measured.

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	159	A64FX prototype - Fujitsu A64FX, Fujitsu A64FX 48C 2GHz, Tofu interconnect D, Fujitsu Fujitsu Numazu Plant Japan	36,864	1,999.5	118	16.876
2	420	NA-1 - ZettaScaler-2.2, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-5C2 700Mhz, PEZY Computing / Exascaler Inc. PEZY Computing K.K. Japan	1,271,040	1,303.2	80	16.256
3	24	AiMOS - IBM Power System AC922, IBM POWER9 20C 3.45GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Volta GV100, IBM Rensselaer Polytechnic Institute Center for Computational Innovations (CCI) United States	130,000	8,045.0	510	15.771
4	373	Satori - IBM Power System AC922, IBM POWER9 20C 2.4GHz, Infiniband EDR, NVIDIA Tesla V100 SXM2, IBM MIT/MGH/PC Holyoke, MA United States	23,040	1,464.0	94	15.574
5	1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	10,096	14.719



FUJITSU

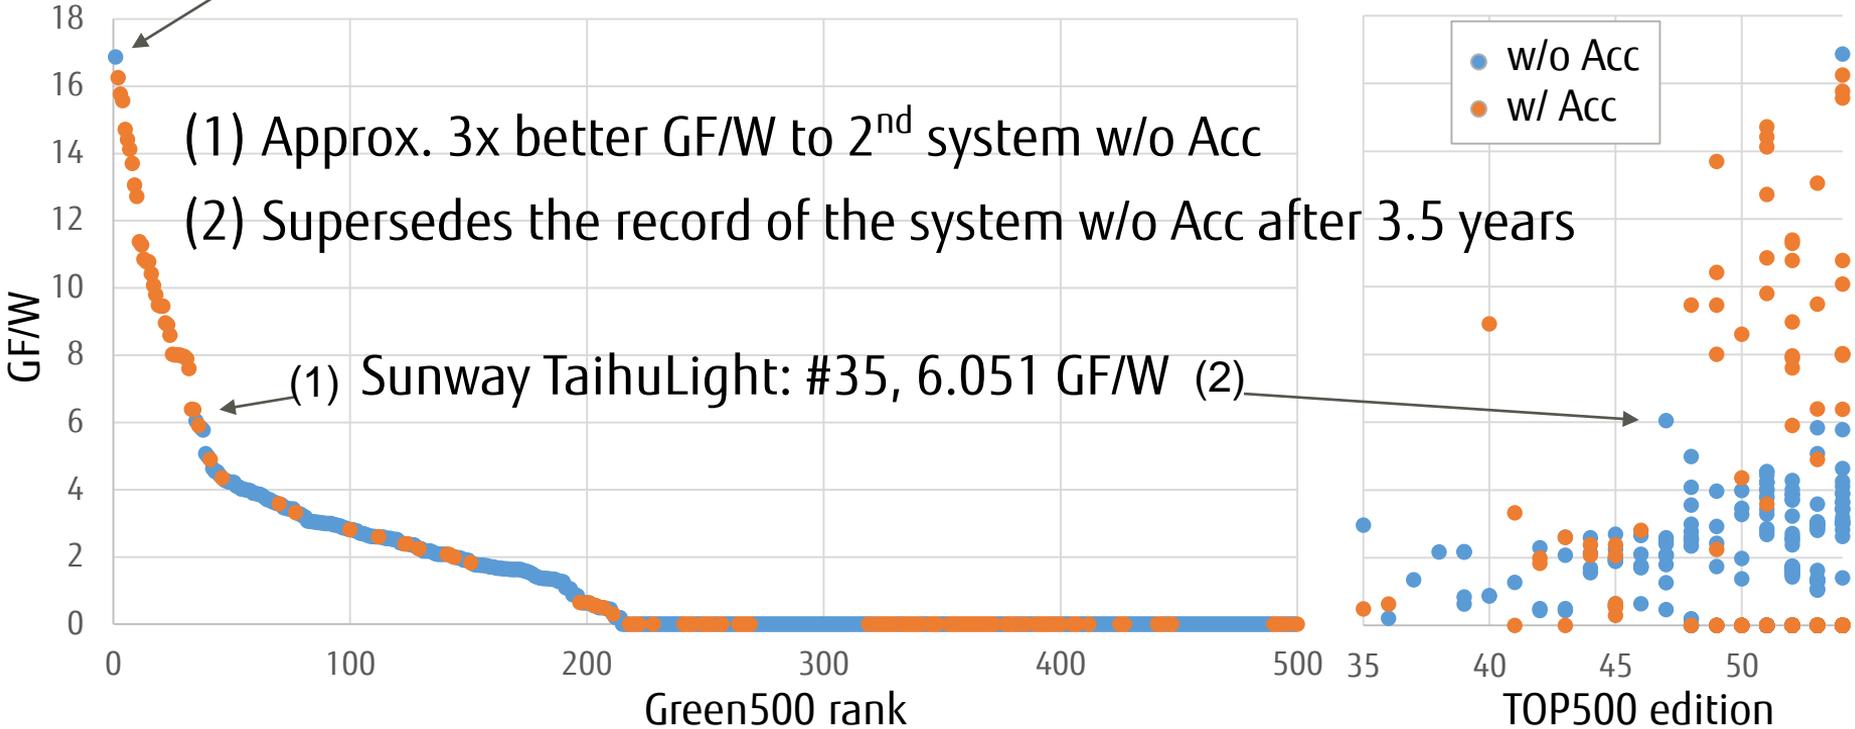
<https://www.top500.org/green500/lists/2019/11/>

SC19 Green500 ranking and 1st appeared TOP500 edition

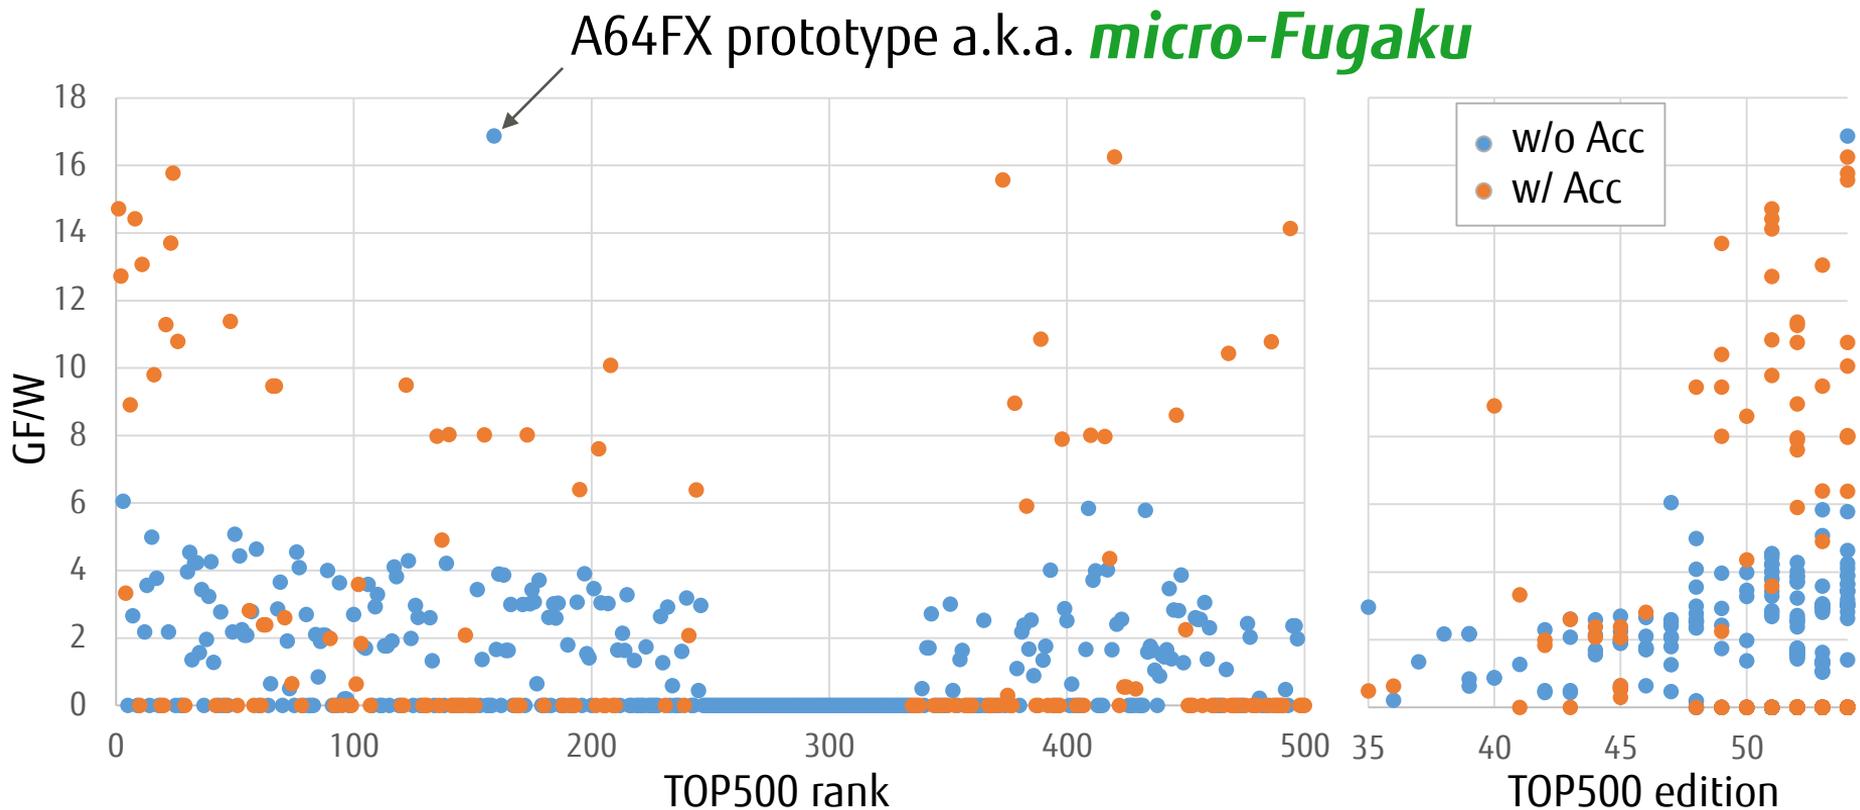
A64FX prototype: #1, 16.876 GF/W

(1) Approx. 3x better GF/W to 2nd system w/o Acc

(2) Supersedes the record of the system w/o Acc after 3.5 years



SC19 TOP500 ranking and GF/W



- TOP500 / Green500 achievements
- How Fujitsu worked
- Green500 measurement conditions and results
- Supercomputer standard and future
- Summary

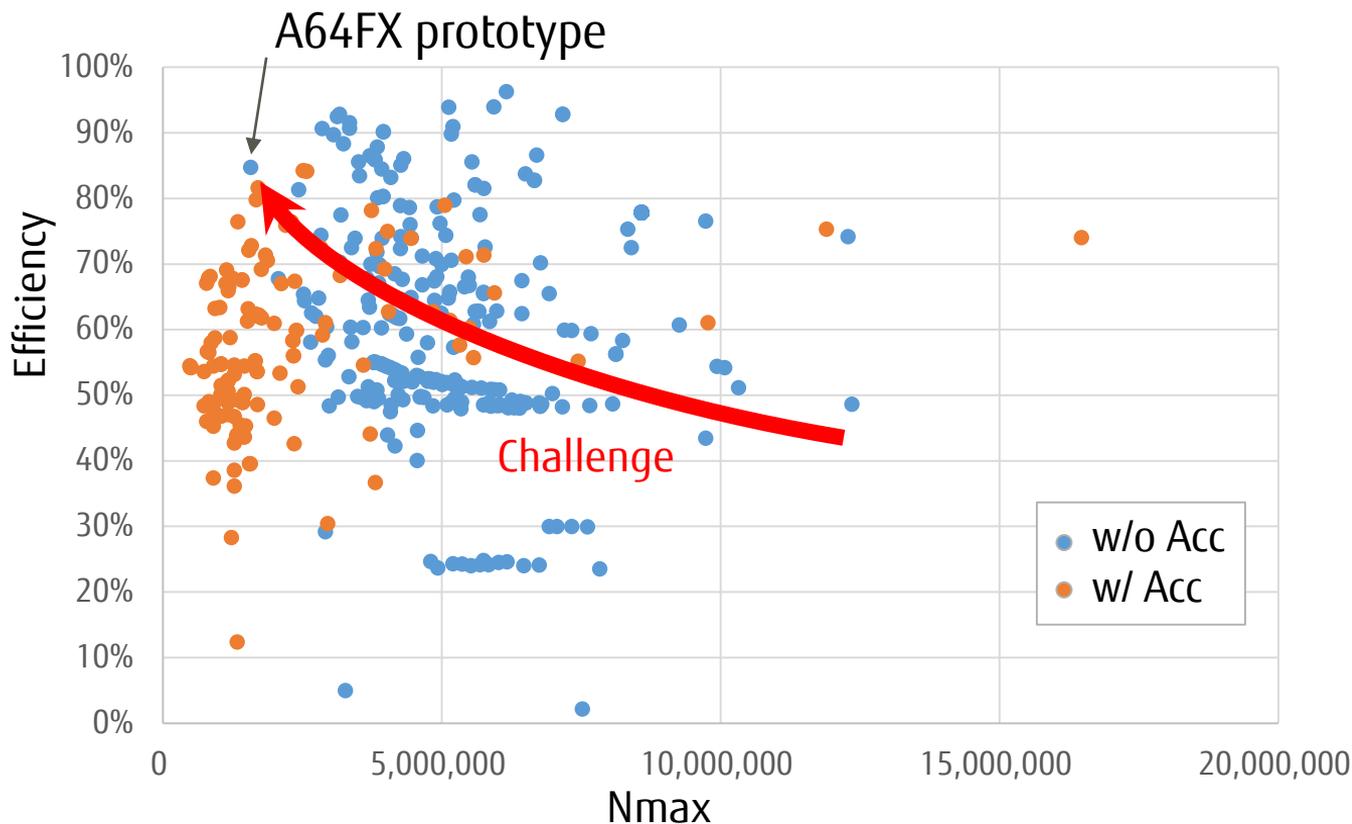
- Key for GF/W is {energy efficient HW} x {parallel/exec efficiency}
- A64FX is designed for energy efficient
 - Fujitsu's proven CPU microarchitecture & 7nm FinFET
 - SoC design: Tofu interconnect D integrated
 - CoWoS: 4x HBM2 for main memory integrated
- Superior parallel/exec efficiency
 - Math. libraries are tuned for application efficiency
 - Comm. libs are also tuned utilizing long experience of Tofu @ K computer
 - Performance tuning is efficiently done utilizing rich performance analyzer/monitor

SC19 TOP500 calculation efficiency

■ Superior efficiency
84.75% with small
Nmax

■ Results of:

- Optimized communication and math. libs
- Optimization of overlapped communication



- TOP500 / Green500 achievements
- How Fujitsu worked
- Green500 measurement conditions and results
- Supercomputer standard and future
- Summary

Power quality level

- Level 1 is difficult to report to be compared with other systems
- Level 2 is the choice of current technology and economy
- Level 3 is the best to evaluate greenness precisely but cost much

SC19 Green500 list: {power source} x {power quality level},
XXX, #n~ : # of systems is **XXX** and the top rank in the cell is **n**

Power quality level	Submitted	Optimized	Derived	Total
Level 1	18 #2~	2 #8, 18	-	20
Level 2	20 #1~	1 #6	-	21
Level 3	10 #3~	-	-	10
Blank & Derived	138 #7~	-	25 #136~	163
Total	186	3	25	214

Power quality level discussion

1. Increase # of systems w/ power consumption regardless of conditions
2. Classify quality level as much as possible
3. Increase Level 2 & 3 share

Target?
+ Yield insight into pwr *if* data shared
- Need complicated effort & cost

SC19 Green500 list: {power source} x {power quality level},
XXX, #n~ : # of systems is **XXX** and the top rank in the cell is **n**

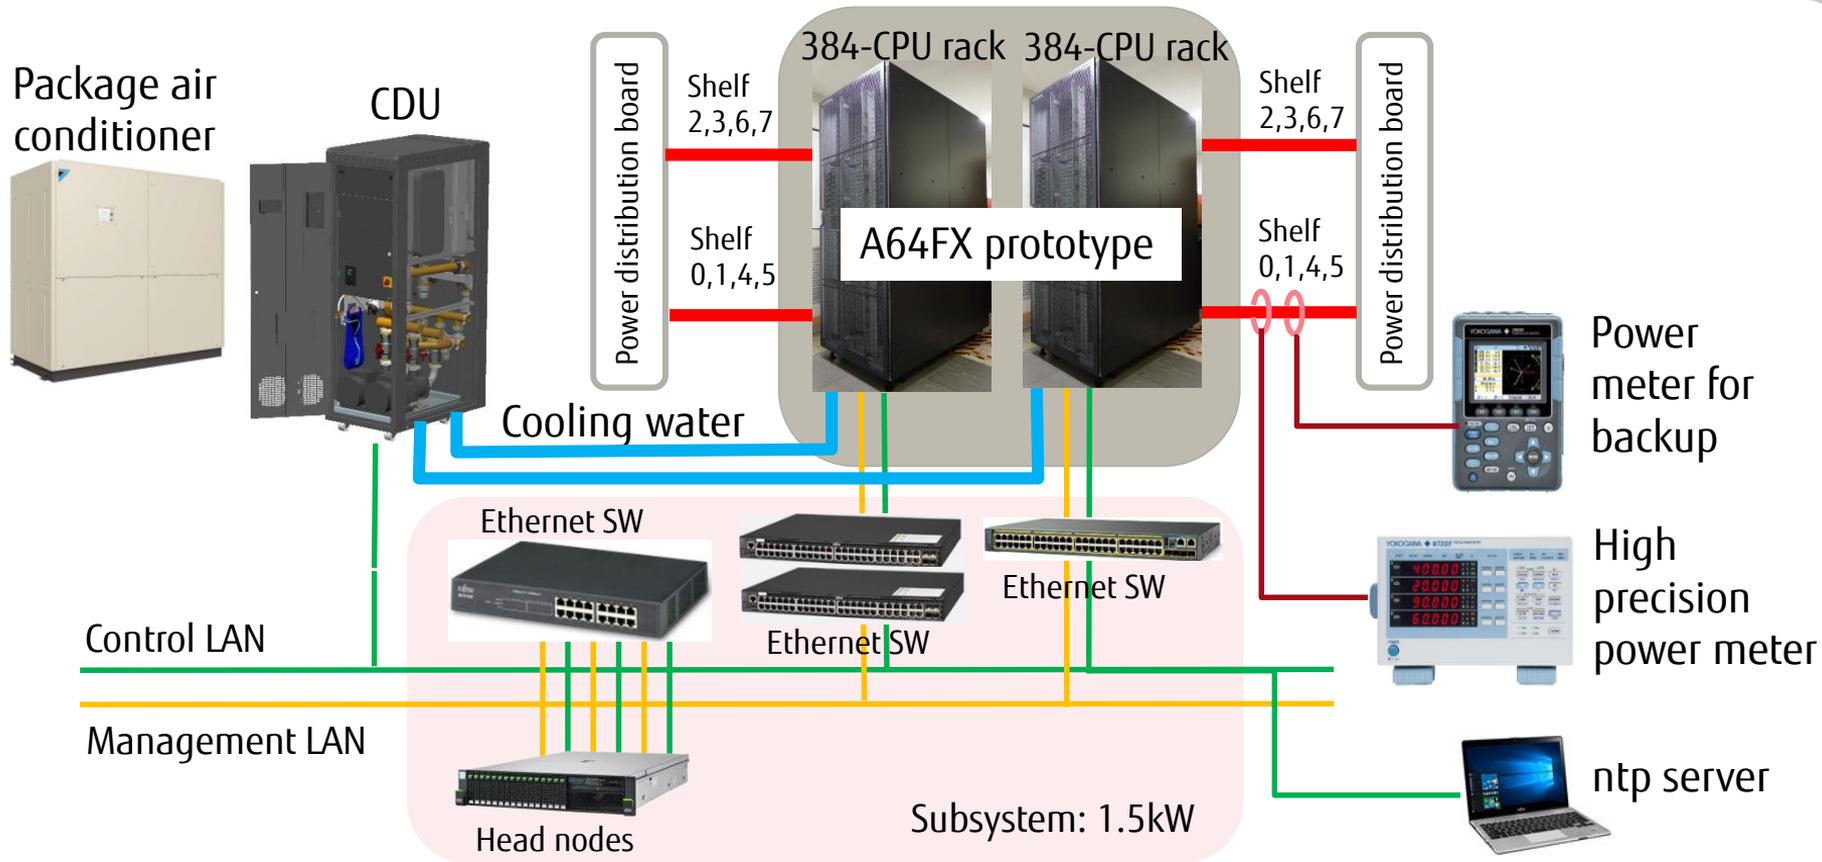
Power quality level	Submitted	Optimized	Derived	Total
Level 1	18 #2~	2 #8, 18	-	20
Level 2	20 #1~	1 #6	-	21
Level 3	10 #3~	-	-	10
Blank & Derived	138 #7~	-	25 #136~	163
Total	186	3	25	214

3

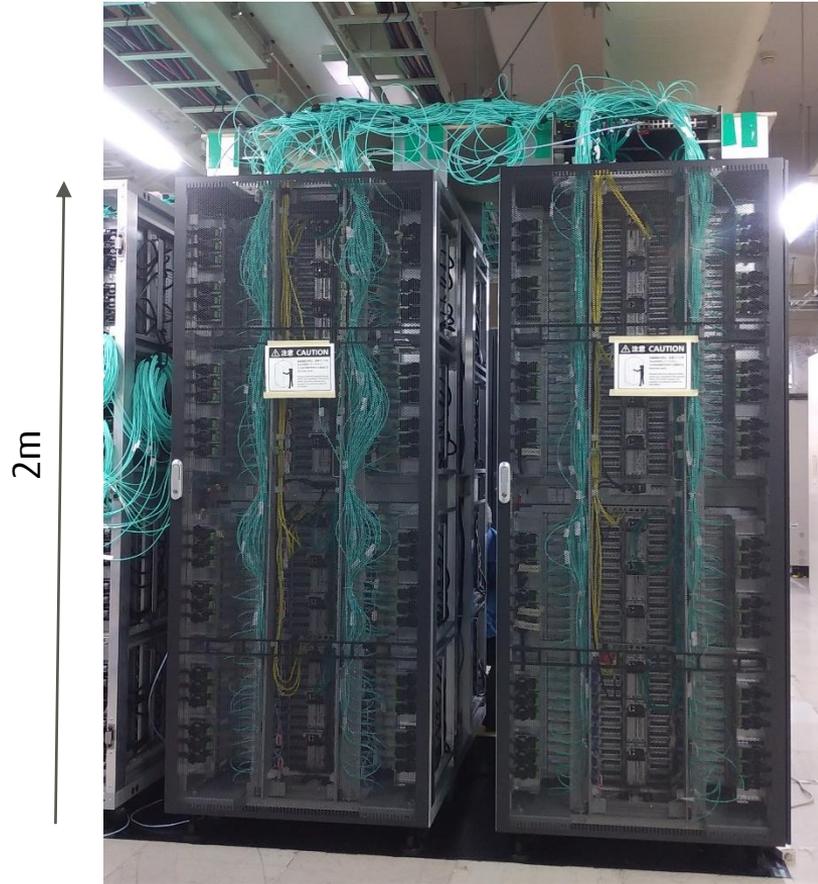
2

1

Fujitsu Numazu Plant: A64FX prototype system



A64FX prototype @ Fujitsu Numazu Plant

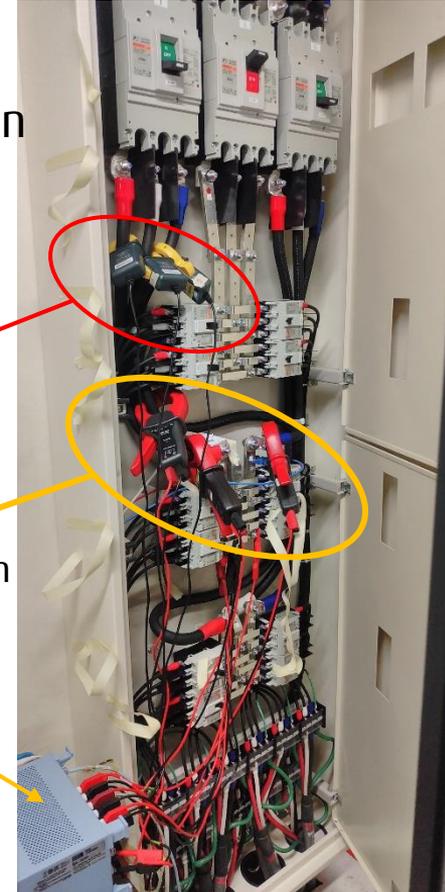


Power distribution board

Clamps for backup

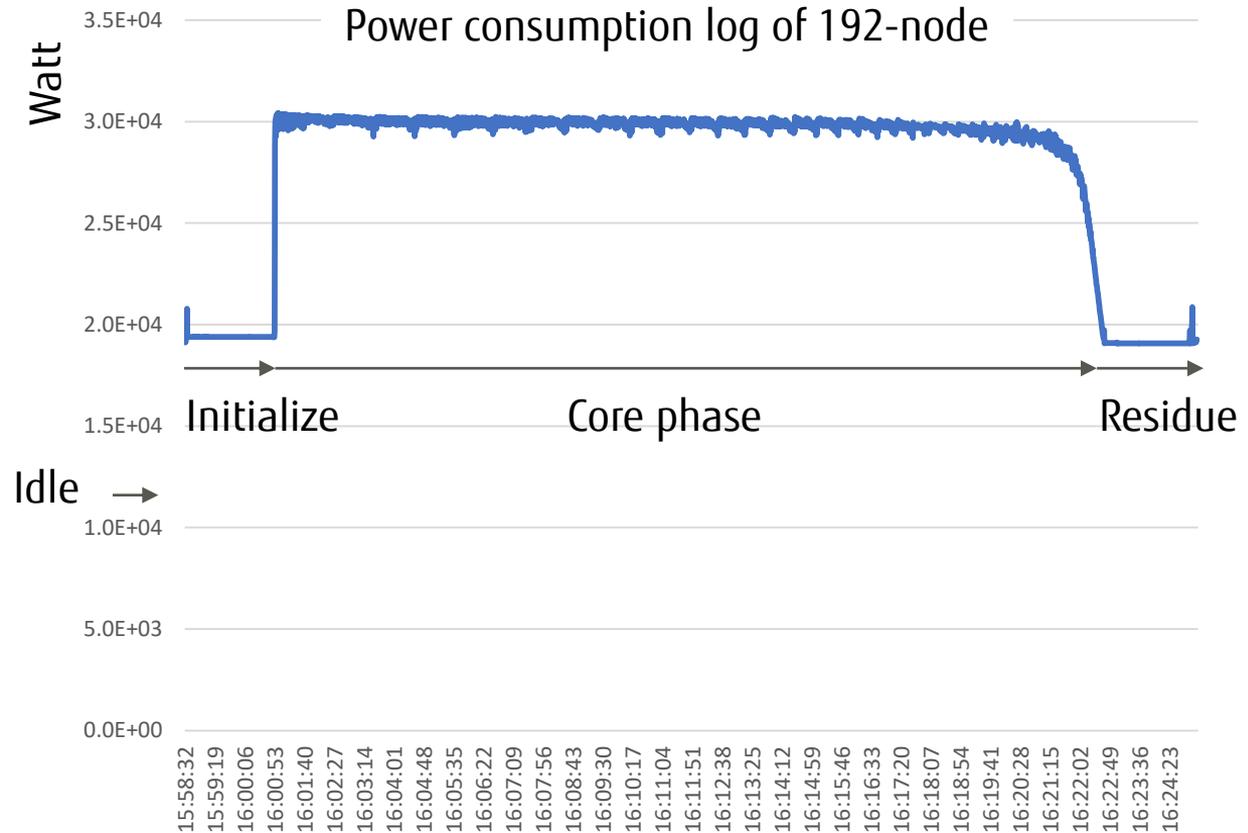
Clamps for high precision power meter

High precision power meter



Power consumption of 1/4 A64FX prototype system

- Steady and high efficient computation are observed even from the point of power consumption
- Ave. of core phase: 118.48 kW / system
- Idle power: 46.92 kW / system



- Fugaku「富岳」
- PRIMEHPC FX1000 / FX700
- Cray CS500

Supercomputer Fugaku named after Mt. Fuji



- Highest mountain (performance)
- Very broad gradual slopes (applications and users)

Fugaku 「富岳」



Focus	Approach
 Application performance	Co-design w/ application developers and Fujitsu's high memory bandwidth utilizing HBM2
 Power efficiency	Leading-edge Si-technology, Fujitsu's proven performance logic design, and power-control
 Usability	Arm®v8-A ISA with Scalable Vector Extension Linux

New PRIMEHPC Lineup

FUJITSU

PRIMEHPC FX1000

Supercomputer optimized for large scale computing

High Scalability

High Density

Superior power efficiency

A64FX processor
384 nodes/Rack
Tofu-D Interconnect
Water Cooling
Fujitsu Software Stack
for Supercomputing

PRIMEHPC FX700

Supercomputer based on
standard technologies

Ease to use

Installation

A64FX Processor
8 nodes/2U Rackmount
InfiniBand
Air Cooling
Utilize ISV and Open Source Software Stack

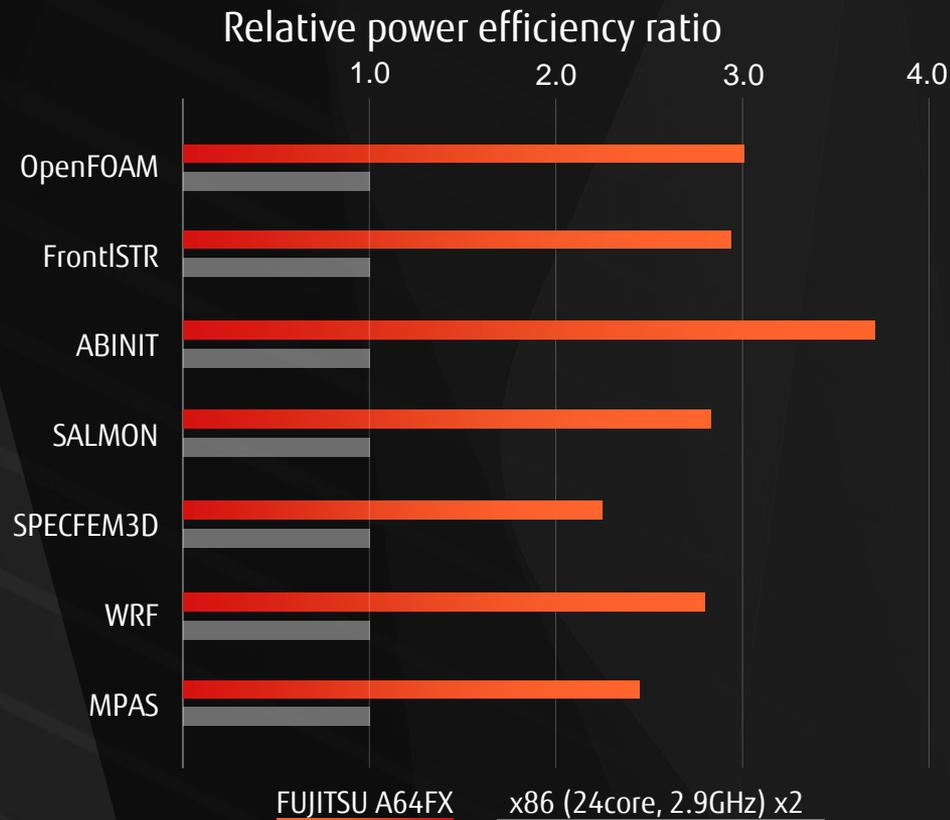


A64FX Superior Power Efficiency for Real Apps



Performance / Energy consumption on an A64FX, 2.2GHz of PRIMEHPC FX1000

- Up to 3.7x more efficient over the latest x86 processor (24core, 2.9GHz) x2
- High efficiency is achieved by energy-conscious design and implementation



INTRODUCING THE CRAY CS500 - FUJITSU A64FX ARM SERVER

- Next generation Arm® solution enabled through Cray Fujitsu Technology Agreement
- Builds on Cray and Fujitsu strong history with vector processing and supercomputing
- Supported in Cray CS500 infrastructure including Cray Programming Environment
- Leadership performance for many memory intensive HPC applications
- Provides customers with more choice and flexibility
- GA in mid'2020



Courtesy of Cray

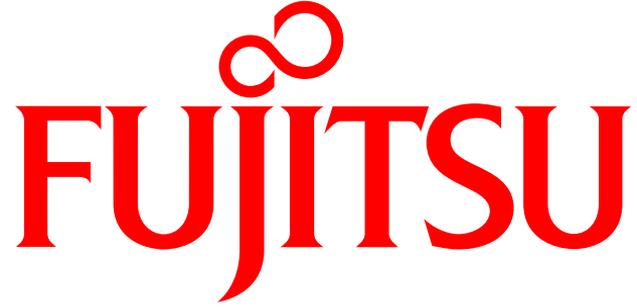


- TOP500 / Green500 achievements
- How Fujitsu worked
- Green500 measurement conditions and results
- Supercomputer standard and future
- Summary

- Energy saving is primary social issue for the future, even now
 - Greenness is not only for the environment but also for good economy
 - Power consumption or cost of electricity is a major part of TCO of supercomputer
- Future issues
 - Be more conscious for the supercomputer power consumption
 - Green500 is the good benchmark and will be more important
 - TOP500 better to ask to submit power consumption mandatory
 - Not only for competition but to raise the bottom and recognition, only 43% systems report the power now
 - Level 3
 - Cost much for the ranking event
 - To invest, consider to utilize existing functions or utilize power meters after the event

- A64FX prototype @ SC19
 - #1 @ Green500: Highest energy efficiency 16.876 GF/W, power quality level 2
 - #159 @ TOP500: High calc. efficiency 84.75% with small Nmax 1,576,960
- General purpose A64FX CPU will be widely used by many customers from Fujitsu directly and through many partners
 - Supercomputer Fugaku
 - Cray CS500
 - Fujitsu PRIMEHPC FX1000 / FX700





shaping tomorrow with you