Applications for PRIMEHPC FX1000/FX700

November 15, 2021

Fujitsu Limited

FUJITSU

Contents

- Overview
- Commercial Applications
  - Commercial Applications
  - Commercial Applications List
  - Ansys Fluent
  - Ansys LS-DYNA
  - CONVERGE
  - Cradle CFD | scFLOW
  - JMAG
  - Poynting
  - Applications for Automotive Industry
- Open-Source Software (OSS) Applications
  - OSS Performance on FX1000 and FX700
  - OSS Power Efficiency on FX1000 and FX700
  - FrontISTR
  - OpenFOAM
  - WRF
  - LAMMPS
  - QUANTUM ESPRESSO
  - Information and Tools for OSS
- Open-Source Oil & Gas Applications evaluated with Arm ecosystem collaborators
  - A64FX advantage in Oil & Gas applications
  - Devito
  - GIRIH

Overview

Feature of PRIMEHPC FX1000 and FX700

○ **PRIMEHPC FX1000** and **FX700** are featured by **ARM** Processor **A64FX** which is used in super-computer **Fugaku**.

© RIKEN

A64FX processer

PRIMEHPC FX1000
Super-computer for very large-scale system

FUJITSU Supercomputer PRIMEHPC

✓ Scalable Vector Extension (SVE) 512bit SIMD is implemented.

✓ An extension of Armv8.2-A instruction set architecture is applied.

✓ High memory throughput is realized by HBM2.

✓ It is designed to obtain high power efficiency.

PRIMEHPC FX1000
Super-computer for very large-scale system

✓ 384nodes/rack
✓ Liquid cooling system
✓ Tofu Interconnect D

PRIMEHPC FX700
Introduction model based on standard technology

✓ 8nodes/2U chassis
✓ Air cooling system
✓ InfiniBand

Performance of FX1000 and FX700

○ In the five main categories, *FX1000* and *FX700* have better performance and power efficiency than x86 server.

Computation speed normalized by x86 server[*] **Power Efficiency normalized by x86 server**[**]

Geoscience
- Devito: FX700 3.34, FX1000 3.57 | Power: FX700 6.89, FX1000 9.21
- WRF: FX700 1.35, FX1000 1.44 | Power: FX700 3.23, FX1000 4.40

Material and chemistry
- LAMMPS: FX700 1.00, FX1000 1.11 | Power: FX700 2.83, FX1000 3.78
- QUANTUM ESPRESSO: FX700 1.08, FX1000 1.27 | Power: FX700 2.96, FX1000 4.07

Electromagnetics
- JMAG: FX1000 1.22[***] | Power: FX700 3.18, FX1000 3.81
- Poynting: FX700 2.82, FX1000 2.44 | Power: FX700 6.73, FX1000 7.46

Fluid dynamics
- OpenFOAM: FX700 1.24, FX1000 1.33 | Power: FX700 3.49, FX1000 4.45
- Cradle CFD | scFLOW: FX1000 3.18[***] | Power: FX700 2.07, FX1000 2.65

Structural analysis
- FrontISTR: FX700 1.90, FX1000 1.98 | Power: FX700 4.12, FX1000 5.06

Legend: ■ FX700 ■ FX1000

Benchmark Platform
FX1000: *A64FX* (2.2GHz x 48cores)
FX700: *A64FX* (2.0GHz x 48cores)
x86 server[***]: *Xeon 8268* (2.9GHz x 24cores) x2

[*] Some computational model is not the same as the one used to measure power efficiency.
[**] Measured at one node
[***] Benchmark platform is different. Please refer to individual applications page for details.

5

Commercial Applications

Commercial Applications（as of Nov. 2021）

FUJITSU

Amber
(by University of California, San Francisco.) ✓

Gaussian16
(by Gaussian, Inc.) ✓

Chemistry

COLMINA CAE Magnetic Simulator
(by Fujitsu Limited) ✓

JMAG®
Simulation Technology for Electromechanical Design
(by JSOL Corporation) ✓ F

Poynting
(by Fujitsu Limited) ✓ F

Electromagnetics

Ansys LS-DYNA
Ansys LS-DYNA ✓ F

ESI Virtual Performance Solution (VPS)
(by ESI Group) ✓ F

Structural analysis

PRIMEHPC FX1000

PRIMEHPC FX700

Fugaku
©RIKEN

Fluid dynamics

Ansys Fluent
Ansys FLUENT ✓

COLMINA CAE Particle-based Casting Simulator
(by Fujitsu Limited) ✓

CONVERGE CFD SOFTWARE
(by Convergent Science) ✓ F

Cradle | scFLOW
(by **HEXAGON** | **Cradle**) ✓ F

HELYX®
(by ENGYS Ltd.) ○

Simcenter STAR-CCM+ ○
(by Siemens Digital Industries Software)

Others

NAG Fortran Compiler ✓
(by **nag®**)

✓ = Already Available

○ = Plan to be Available

F = Verified on *Fugaku*

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

7

© 2021 Fujitsu Limited

Commercial Applications List（as of Nov. 2021）

FUJITSU

○ Availability of commercial applications for *FX1000* and *FX700*. Six applications were already verified on *Fugaku*.

| Categories	Applications	Vendors	Availability	Notes
Engineering (Fluid dynamics)	Ansys Fluent	ANSYS, Inc.	**Available** (alpha version)	Solver components only
	COLMINA CAE Particle-based Casting Simulator	Fujitsu Limited	**Available**	
	CONVERGE	Convergent Science (East Asia distributor : IDAJ Co., LTD.)	**Available**	Solver components only Verified on *Fugaku*
	Cradle CFD | scFLOW	Software Cradle Co., Ltd.	**Available**	Verified on *Fugaku*
	HELYX	ENGYS Ltd.	Plan to be available*	
	Simcenter STAR-CCM+	Siemens Digital Industries Software	Plan to be available*	
Engineering (Structural analysis)	Ansys LS-DYNA	ANSYS, Inc.	**Available**	Verified on *Fugaku*
	ESI Virtual Performance Solution(VPS)	ESI Group	**Available**	Explicit features only Verified on *Fugaku*
Engineering (Electromagnetics)	COLMINA CAE Magnetic Simulator	Fujitsu Limited	**Available**	
	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**	

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

Ansys Fluent

○ About **Ansys Fluent**

✓ **Commercial CFD software** capable of creating advanced physics models and analyzing a variety of fluids phenomena
✓ Contains the best-in class physics models and can accurately and efficiently solve large, complex models
✓ More information : https://www.ansys.com/products/fluids/ansys-fluent
✓ Available on **FX1000** and **FX700**

○ Verification and Performance evaluation

✓ **Ansys Fluent** alpha version (based on Ansys 2021 R2) for **FX1000**/**FX700**
 - The alpha version is built for generic Arm, and it includes scripts specific to Fujitsu PRIMEHPC
 - Available for **Ansys Fluent** users as requested
✓ Tested and evaluated on **FX1000**/**FX700**
 - The "Ansys Fluent Benchmarks" as the test cases
 - Up to 1,536 cores with f1_racecar_140m on **FX1000**

○ Challenges on **Fugaku***

✓ Started performance evaluation on **Fugaku** from Nov 2021 for 6 months
✓ Planning to evaluate the performance and scalability using the largest model (280 million cells) of "Ansys Fluent Benchmarks" with more than 1,000 nodes

 * This work is using computational resources of the supercomputer **Fugaku** provided by RIKEN through the HPCI System Research Project (Project ID: hp210283).

◆ Shows good scalability with the large model on **FX1000**

f1_racecar_140m

Ideal
Actual

Solver Rating (y-axis: 0, 200, 400, 600, 800, 1000, 1200)
# of cores (x-axis: 0, 400, 800, 1200, 1600)

Scalability by **Ansys Fluent**'s "Solver Rating"

Computational model:
External flow over a Formula-1 race car (140 million cells)

◆ Power efficiency on **FX1000** and **FX700** is relatively better than **Xeon Platinum 8268** (2.9GHz)

sedan_4m

Normalized power efficiency (y-axis: 0, 0.5, 1, 1.5, 2, 2.5)

Higher is better

1.98 — FX1000 (2.2GHz)
1.45 — FX700 (2.0GHz)
1.00 — Xeon (2.9GHz)

Power Efficiency of **Ansys Fluent** on 1 node

Computational model:
External flow over passenger sedan (4 million cells)

Ansys LS-DYNA

Ansys LS-DYNA

FUJITSU

○ About *Ansys LS-DYNA*

✓ *Ansys LS-DYNA* is the industry-leading multi-physics simulation software used for complex real-world applications.
✓ *Ansys LS-DYNA* is the first-ever commercial CAE structural simulation software available on *FX1000, FX700* and *Fugaku*.
✓ Single precision *Ansys LS-DYNA* has been released in 2021, and double precision *Ansys LS-DYNA* will be release in 2022.
  https://www.ansys.com/news-center/press-releases/10-15-2021-ansys-enables-more-sustainable-product-development-with-fujitsu

---

○ Verification and Performance evaluation

✓ Verified *Ansys LS-DYNA* key features.

✓ Verified parallel scalability up to 3,000 cores

✓ *Ansys LS-DYNA* on *FX1000* and *FX700* shows significant power efficiency.

Normalized power efficiency

2.39 — FX1000 (2.2GHz)
1.94 — FX700* (2.0GHz)
1.00 — Xeon (2.9GHz)

↑ **Higher is better**

**_Ansys LS-DYNA_ power efficiency on one node
(Car to car model, 2.4 million elements)**

*File system of FX700 is FEFS

○ Challenge on *Fugaku***

Study of a high-performance and simulating complex real world problems:
✓ Car crash models of 30 millions elements was evaluated using 30000 cores on *Fugaku*
✓ JAMA has been conducting 12,000 cases of car crash simulation utilizing *Fugaku* with supporting of Fujitsu. JAMA has been collecting FEM analysis results with skeleton layout and strength balance as parameters and using these as teacher data of machine learning. By clarifying the correlation between the accuracy of machine learning predictions and the number of FEM analysis results accumulated, JAMA will revolutionize the conventional development process.

Fujitsu will work together with ANSYS, Inc. for higher parallel performance on *Fugaku.*


Car crash simulation image

**This work used computational resources of the supercomputer *Fugaku* provided by RIKEN through the HPCI System Research Project (Project ID: hp200228, hp210089).

CONVERGE

About *CONVERGE*

- Commercial CFD software featuring **truly autonomous meshing, state-of-the-art physical models, a robust chemistry solver, and the ability to easily accommodate complex moving geometries**
- Used in various industries(engine cylinders, fuel injections, exhaust-gas aftertreatment devices, pumps, compressors, gas turbines, etc.)
- More information : https://convergecfd.com / https://www.idaj.co.jp/product/converge (in Japanese)
- Available on *FX1000, FX700* and *Fugaku*

Verification and Performance evaluation

- Verified all features of *CONVERGE* on *FX1000* and *FX700*

Application	Free spray
Turbulence model	LES
Spray model	Lagrange 2 phases
Sub models	Drag Break-up Collision Evaporation

Free spray analysis with using LES
(Max. 10 million cells, 384 parallel on *FX1000*)

- Confirmed good scalability over 300 cores on *FX1000*

Scalability of steady port flow analysis

CONVERGE results on *Fugaku** : ICE combustion(LES, detailed chem)

- In-cylinder combustion simulation of Internal Combustion Engine(ICE) aims to evaluate heat release rate, emissions, knocking, etc., and eventually engine performance.
- Obtained winkled flame structure by **LES calculation with full-domain fine mesh** considering detail chemical reactions **in 2 hours**
- Expected to contribute not only to improve energy efficiency of conventional gasoline engines but also to **accelerate development speed of e-fuel engines** which is recently gathering attention, by **accurate knocking prediction in a practical calculation time**
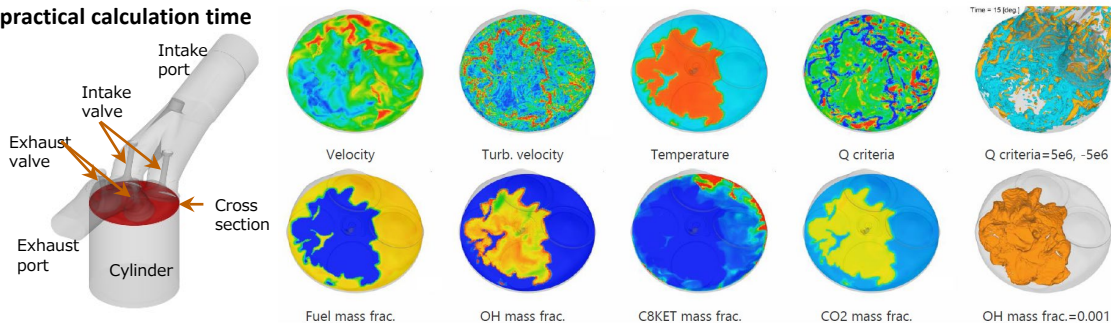
Internal Combustion Engine model

Calculation results of ICE model

Application	# of meshes	Mesh size	Turbulence model	Combustion model	# of parallel	Calculation time (combustion only)
ICE combustion	7 million (max.)	0.5mm (whole domain)	LES	Detailed chemistry	4,096	2 hours

11

© 2021 Fujitsu Limited

*This work used computational resources of the supercomputer *Fugaku* provided by RIKEN through the HPCI System Research Project (Project ID: hp200209).

Cradle CFD | scFLOW

HEXAGON | Cradle | Cr FL Cradle | scFLOW

FUJITSU

○ About *Cradle CFD | scFLOW*

✓ **Commercial CFD** software with **multiphysics** and **general purpose** capabilities.
✓ Used in **various industries** (aerospace, automotive, construction, electronics, heavy manufacturing, medical and pharmaceutical, etc.)
✓ More information : https://www.mscsoftware.com/product/scflow
✓ Available on *FX1000, FX700* and *Fugaku*.

○ Verification and performance evaluation

✓ Improved performance of the linear solver of system equation.

Details
• Model :
  Aerodynamics of flow around car
  (108 million cells, RANS)
• 384procs x1thread
• (*Xeon 8260 Platinum*) 8nodes,
  (*FX1000*) 48nodes
• 100steps

3.18
FX1000
(2.2GHz)

Higher is better

1.00
Xeon
(2.4GHz)

Normalized performance

Performance comparison in
Cradle CFD | scFLOW solver

✓ Power efficiency on *FX1000* and *FX700* is better than *Xeon Platinum 8268.*

Details
• Model : Aerodynamics of
  transonic-flow around aircraft
  (2 million cells, RANS)
• 48procs x1thread
• (*Xeon 8260 Platinum*) 1node,
  (*FX1000*) 1node,
  (*FX700*) 1node
• 2,000steps

2.65
FX1000
(2.2GHz)

2.07
FX700
(2.0GHz)

Higher is better

1.00
Xeon
(2.9GHz)

Normalized power efficiency

Power efficiency of *Cradle CFD | scFLOW*

○ Challenges on *Fugaku**

✓ **Achieved solving large-scale LES containing billions of element faces over the limitation of 32bit integer**

◆ Large-scale LES is important for the evaluation of stability and safety of aircrafts operations by a numerical simulation, especially for off-design conditions.
◆ Modified *Cradle CFD | scFLOW* code. It makes performance stably at over 48,000 hybrid parallel.
◆ **Max. 216,000 hybrid parallel / 4,500 nodes** was evaluated with 787 million cells model (containing 2.4 billion element faces).

Mach number : 0.847
Attack angle : 2.94 deg
# of Cells : **787 million**

Ideal

Parallel scale
Number of cores

Parallel scale (LES, 100steps)

# of Cells : 237 million

Surface pressure coefficient of aircraft
(LES, 12,000procs.x4threads/1,000nodes,
50,000steps)

✓ **Succeeded to create hundreds of millions of unstructured cells by parallelized mesher**

◆ Applied MPI parallelization into mesher of *Cradle CFD | scFLOW* in order to eliminates limitation of memory from meshing.
◆ Confirmed that the mesher succeeded to create max. 511 million of unstructured cells (model : National Stadium)

Wind speed :
2.9m/s at 41m above ground
# of Cells : 103 million

*This work used computational resources of the supercomputer *Fugaku* provided by RIKEN through the HPCI System Research Project (Project ID: hp200209 and hp200302).

Wind analysis around the National Stadium
(2,400procs. x 4threads/200nodes, 2,000steps)

JMAG

○ About **JMAG**

- ✓ Commercial software for electric device design and development. Developed & provided by JSOL Corporation.
- ✓ Used in various industries (automotive, electric appliances, digital equipment, electric power equipment, factory automation, etc.)
- ✓ More information : https://www.jmag-international.com/
- ✓ Plan to be available on **FX1000, FX700** and **Fugaku**. Release date will be announced later.

○ Verification and Performance evaluation

- ✓ Max 20 million elements model & max 8,192 hybrid parallel was evaluated.
- ✓ Verified all major features of **JMAG HPC solver**.

Preliminary results of **JMAG** on **FX1000** and **FX700**

- ✓ **JMAG** performance on **FX1000** is x1.22 faster than **Xeon Platinum 8280** 2.7GHz
- ✓ **JMAG** power efficiencies on **FX1000** and **FX700** are better than **Xeon Platinum 8268** 2.9GHz

Higher is better

1.22 FX1000 (2.2GHz) / 1.00 Xeon (2.7GHz)

Higher is better

3.81 FX1000 (2.2GHz) / 3.18 FX700 (2.0GHz) / 1.00 Xeon (2.9GHz)
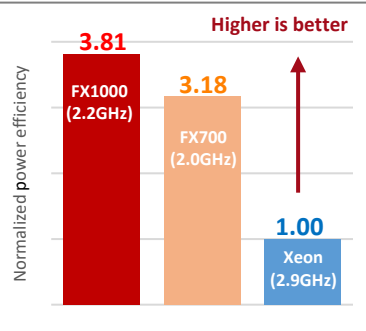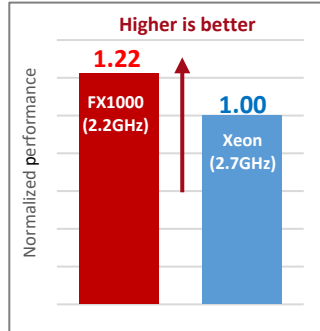
Performance of **JMAG**
(IPM motor with sinusoidal wave input, 10 million elements, 50 steps, 128 hybrid parallel)

1-node power efficiency of **JMAG**
(IPM motor with sinusoidal wave input, 1 million elements, 10 steps, 48 hybrid parallel)

○ Challenges on **Fugaku***

- ✓ In addition to the evaluation of **JMAG** for IPM motor published in ISC2021, we have evaluated the several magnetic field analysis functions of **JMAG** for various large-scale models on **Fugaku**.
- ✓ These results demonstrate that **JMAG** has capability to calculate in a practical time on **Fugaku** for most industrial needs.

New results

IPM motor with PWM input, 15 million elements

✓ We have checked scalability of **JMAG** on **Fugaku** by multiple models. Two models on the right that are often calculated in hundreds of steps are evaluated excluding the initialization time.


7-fold speedup

Analysis time per step

Conventional solver, MPP 512 / New solver in Fugaku, MPP 8192

Induction heating for crankshaft, 17 million elements, 1step

Scalability of total time

Alternator, 15 million elements, 10steps

Scalability of time excluding initialize

Linear motor, 5 million elements, 10steps

Scalability of time excluding initialize

Poynting

○ About **Poynting**

✓ **Poynting** is an electromagnetic wave analysis software based on FDTD method* developed by Fujitsu.
✓ **Poynting** is an extremely advantageous electromagnetic wave analysis software on **FX1000** and **FX700**, because the FDTD* method provides massive parallel efficiency while requires memory bandwidth.
✓ More information : https://www.fujitsu.com/jp/solutions/business-technology/tc/sol/poynting/ (in Japanese)
✓ Available on **FX1000**, **FX700** and **Fugaku**

○ Verification and Performance evaluation

✓ Verified major features of Poynting solver
✓ Verified parallel scalability up to 115,200 cores
✓ Fujitsu Poynting on **FX1000** and **FX700** shows significant power efficiency.

- Normalized performance
- Normalized power efficiency

○ Challenge on **Fugaku******

✓ Verified extreme large-scale electromagnetic wave simulation with Poynting on **Fugaku**
✓ Verified 0.5 trillion cells model of 5G "Vehicle to Vehicle" and "Vehicle to Infrastructure" communications around a crossroad in 2 hours
✓ Verified that it is possible to consider the influence of fine(~5cm) structures in a space of the size of a crossroad(~100m), which was not possible with approximation methods such as ray tracing simulation

Number of Cells: **0.5 Trillion cells**
Calculation time:
2 hours on **115,200 cores of Fugaku**

Results of Simulation:
Electric field distributions around moving vehicles across a crossroad.

*FDTD method : finite-difference time-domain method   **File system of FX700 is FEFS   ***Intel **Xeon Platinum 8268** x2sockets(2.9GHz x 24core)
****This work used computational resources of the supercomputer **Fugaku** provided by RIKEN through the HPCI System Research Project (Project ID: ra010012).

Applications for Automotive Industry

○ **Japan Automobile Manufacturers Association** (**JAMA**) has started evaluating leading-edge computer aided engineering for automobile on supercomputer *Fugaku* with **Fujitsu**'s supports since FY2020

- ✓ In FY2020, car crash models of 4 – 30 millions elements were evaluated, using *Ansys LS-DYNA* and *VPS* on *Fugaku*[*].

- ✓ In FY2021, to develop technologies improving the accuracy of crash performance prediction, JAMA has been conducting 12,000 cases of crash simulation utilizing *Fugaku*[*]. Also, fluid dynamics using *Simcenter STAR-CCM+* will be evaluated on *Fugaku*.

*This work used computational resources of the supercomputer *Fugaku* provided by RIKEN through the HPCI System Research Project (Project ID: hp200228, hp210089).

Ansys LS-DYNA
- ✓ *Ansys LS-DYNA* is a multiphysics simulation software capable of solving complex real world problems.

ESI Virtual Performance Solution (VPS)
- ✓ *VPS* is the world first car crash simulation software and later expanded to the other structural simulations by leveraging its Single Core Modeling concept.

Simcenter STAR-CCM+
- ✓ *Simcenter STAR-CCM+* is a multiphysics computational fluid dynamics (CFD) software for the simulation of products operating under real-world conditions.

Open-Source Software (OSS) Applications

OSS Performance on FX1000 and FX700

FUJITSU

○ Computation speed of *FX1000* and *FX700* compared with x86 server

Chart: Computation speed normalized by x86 server

Software	FX1000	FX700
OpenFOAM (Fluid dynamics)	1.3	1.2
FrontISTR (Structural analysis)	2.0	1.9
OpenFDTD (Electro-magnetics)	2.3	2.2
ABINIT (First principles calculation)	2.3	2.2
SALMON (First principles calculation)	1.6	1.5
QUANTUM ESPRESSO (First principles calculation)	1.3	1.1
LAMMPS (Molecular dynamics)	1.1	1.0
WRF (Weather research and forecasting)	1.4	1.3
MPAS (Model for prediction across scales)	1.0	0.9
SPECFEM3D (Seismic wave propagation)	1.2	1.2

x86 = 1.0 (reference line)

○ 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.

FX700

FX1000

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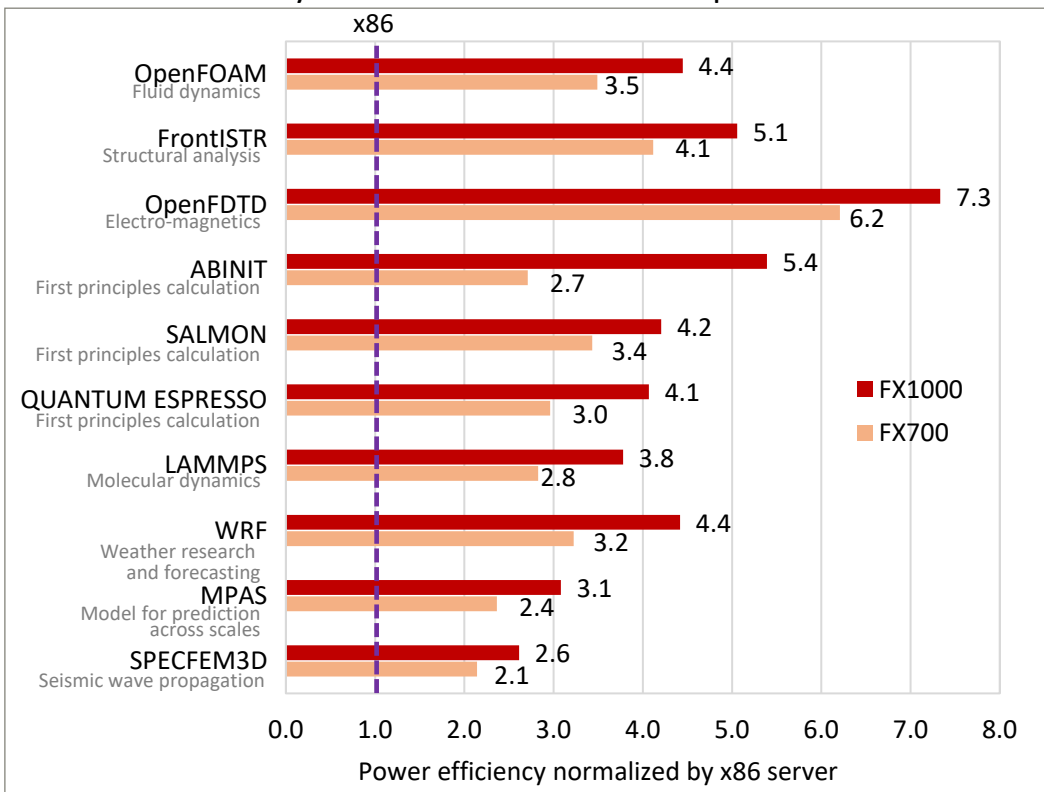*FX1000*: *A64FX* (2.2GHz x 48cores)
FX700: *A64FX* (2.0GHz x 48cores)
x86 server: *Xeon 8268* (2.9GHz x 24cores) x2

A64FX

OSS Power Efficiency on FX1000 and FX700

○ Power efficiency of *FX1000* and *FX700* compared with x86 server

Power efficiency normalized by x86 server

Benchmark	FX1000	FX700
OpenFOAM — Fluid dynamics	4.4	3.5
FrontISTR — Structural analysis	5.1	4.1
OpenFDTD — Electro-magnetics	7.3	6.2
ABINIT — First principles calculation	5.4	2.7
SALMON — First principles calculation	4.2	3.4
QUANTUM ESPRESSO — First principles calculation	4.1	3.0
LAMMPS — Molecular dynamics	3.8	2.8
WRF — Weather research and forecasting	4.4	3.2
MPAS — Model for prediction across scales	3.1	2.4
SPECFEM3D — Seismic wave propagation	2.6	2.1

○ Power efficiency of *FX1000* and *FX700* is at least twice greater than that of x86 server. *FX1000* and *FX700* have significant advantage of power efficiency in comparison with x86 server.

○ The energy-saving technology developed for *Fugaku* which won the top of Green500 is applied to *FX1000*. This technology provides exceptional performance/power rates in large-scale systems. *FX700* contains the air-cooling system in each chassis and power efficiency is a bit less than *FX1000*.

FX700

FX1000

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

FrontISTR

About *FrontISTR*

✓ *FrontISTR* is an open-source large-scale parallel FEM program for structural analysis. It is implemented to run various nonlinear structural analyses including large deformation, contact, and various nonlinear materials on a parallel computer.

✓ *FrontISTR* is developed by *FrontISTR Commons* and distributed at https://www.frontistr.com/ (in Japanese). A tuned code for *A64FX* is available at github.com (https://github.com/fujitsu/FrontISTR)

Performance evaluation

We verified scalability of *FrontISTR* on *Fugaku* up to 1600nodes using a 76 million elements short-fiber composite rubber analysis model.

930sec@1600nodes
1298sec @800nodes
2447sec @400nodes
4368sec @200nodes

Strong scalability of *FrontISTR* (considering geometrical and material nonlinearity, 13 timesteps, CG with SSOR solver, 4p12t/node)

Challenge on *Fugaku**

Mitsuboshi Belting Ltd. studies on short fiber composite rubber with Fujitsu's support

✓ A large-scale analysis is required to evaluate mechanical properties of short fiber composite rubber used is used in friction transmission belts. Only confirming the stress concentration area is possible in current analysis scale.

✓ We confirmed that large-scale analysis can be done in a practical time through a benchmark using 76 million elements model with 200 to 1,600 nodes at *Fugaku*.

✓ It is planned to run multiple cases in the future to clarify the relationship between mechanical properties and fiber state such as fiber length, swell and orientation variation.

Current: 2mm x 1mm x 1mm area, single case

Target: 5mm x 5mm x 1mm area, multiple cases

Embedded short fiber

Scale up with *Fugaku*

3 mil. elems

$\times N$

Pull

76 mil. elems

3 mil. elems

*This works was conducted as part of the RIKEN-FOCUS-VINAS collaborative research aimed at the verification of industrial applications for the *Fugaku* Cloud Platform.

OpenFOAM

- About **OpenFOAM**
  - ✓ **Open-source Field Operation And Manipulation (OpenFOAM)** is Open-source CFD software which has various functionality of fluid dynamics simulation, including thermal conduction and turbulence analysis (https://www.openfoam.com/).

- Performance evaluation
  - ✓ Computation speed is better than x86 in most cases. It is much enhanced by code tunning. The optimization of loop structure to enhance vectorization (more efficient use of SIMD) in code tunning has significant effect with Computation speed for **A64FX**.
  - ✓ Two benchmark models were calculated. One is aero-dynamics of motorBike (LES) and another is thermal fluid of multiRegionHeater (RANS). Two matrix solvers PCG and GAMG are applied to computation of motorBike model.

multiRegionHeater

motorBike

OpenFOAM: v1812
Benchmark Platform
  FX1000 A64FX (2.2GHz x 48cores)
  x86: **Xeon 8268** (2.9GHz x 24cores) x2
**Input data**: multiRegionHeater and motorBike

Computation speed normalized by x86

		Value
multiRegion Heater	FX1000 A64FX TUNED	1.21
	FX1000 A64FX ASIS	0.89
	x86	1.00
motorBike GAMG	FX1000 A64FX TUNED	1.26
	FX1000 A64FX ASIS	0.92
	x86	1.00
motorBike PCG	FX1000 A64FX TUNED	1.33
	FX1000 A64FX ASIS	0.99
	x86	1.00

Two benchmarks were calculated. Asis and tuned code are used in **FX1000**. High Computation speed is obtained compared with x86. (multiRegionHeater: 15M cells, motorBike:10M cells)

Weak scaling of FX1000 A64FX

- moterBike GAMG TUNED
- moterBike PCG TUNED

80M cells
40M cells
20M cells
10M cells

Weak scaling is shown. The size of input data is changed with the number of cores. **FX1000** has good weak scalability and computation scale increases with the number of nodes.

WRF

- **About *WRF***
  - ✓ ***The Weather Research and Forecasting* (*WRF*) *Model*** is a next-generation mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting applications. It features two dynamical cores, a data assimilation system, and a software architecture supporting parallel computation and system extensibility. (https://www.mmm.ucar.edu/weather-research-and-forecasting-model)

- **Performance evaluation**
  - ✓ The performance of ***WRF*** on ***A64FX*** systems shows better performance than other platforms because of high memory bandwidth.
  - ✓ The benchmark model "CONUS 2.5km" was applied to the estimation. Two node are used to calculate this benchmark model.

WRF v3.9.1.1, CONUS 2.5km   2 Nodes

Platform	Value
FX1000 A64FX	1.93
FX700 A64FX	1.72
EPYC	1.66
x86	1.00

Computation speed normalized by x86
Computation speed is measured by "Seconds / timestep"

***WRF*: version 3.9.1.1 (Tuned code)**
Benchmark platform
   FX1000 A64FX (2.2GHz x 48cores)
   FX700   A64FX (2.0GHz x 48cores)
   x86: ***Xeon P8260M*** (2.4GHz x 24cores) x2
   EPYC : ***7742*** (2.25GHz x 64cores) x2

Input data
  CONUS 2.5km
   - 1501x1201x35
   - 3 hours forecasting

(https://www2.mmm.ucar.edu/wrf/bench/benchdata_v3911.html)

LAMMPS

- ○ About *LAMMPS*
  - ✓ *Large-scale Atomic / Molecular Massively Parallel Simulator* (*LAMMPS*) is a classical molecular dynamics simulation code with a focus on materials modeling. It was designed to run efficiently on parallel computers. (https://www.lammps.org/)

- ○ Performance evaluation
  - ✓ The performance on *FX1000* was achieved 1.12x faster than x86. SIMD and efficient data access enhance performance of *FX1000* and *FX700*.
  - ✓ The benchmark model CHUTE was applied to this estimation. Two sections "PAIR" and "MODIFY" which account for 90% of loop time with CHUTE were tuned to obtain fine performance of SIMD and high memory bandwidth.

LAMMPS: version 7Aug19 (Tuned code)
Benchmark platform
   FX1000 A64FX (2.2GHz x 48cores)
   FX700 A64FX (2.0GHz x 48cores)
   x86: *Xeon P8260M* (2.4GHz x 24cores) x2
Input data: "CHUTE" Granular chute flow
   12,800,000 atoms (/node) for 100 timesteps

Computation speed is measured by "atom-timesteps/s" and is normalized by the value of x86-1node case.

Ideal scalability is shown on *FX1000* and *FX700*

22

QUANTUM ESPRESSO

○ About **QUANTUM ESPRESSO**

   ✓ **QUANTUM ESPRESSO** is an integrated suite of Open-Source computer codes for electronic-structure calculations and materials modeling at the nanoscale. It is based on density-functional theory, plane waves, and pseudopotentials (https://www.quantum-espresso.org/).

○ Performance evaluation

   ✓ **A64FX** performance measurement for **QUANTUM ESPRESSO 6.4.1**

   ✓ First Principles Calculation

   Molecular Geometry

   High
   Low

   Calculation of Electron Density

   ✓ **A64FX** is 1.5x faster over the x86 (2.4GHz x 24cores) x2 in QEF benchmark

Computation speed normalized by x86

PSIWAT: QEF/benchmarks/PSIWAT/psiwat.in, 8 nodes
GRIR443: QEF/benchmarks/GRIR443/grir443.in, 8 nodes
Benchmark platform
   FX1000 A64FX (2.2GHz x 48cores)
   x86: **Xeon P8260M** (2.4GHz x 24cores) x2

   ✓ **A64FX** is 1.3x faster over the x86 processor (2.9GHz x 24cores) x2 in structural optimization simulation

   • QE[*1] :1.4x faster by high memory b/w and using SIMD instructions
   • FFT[*2] :1.3x faster by long data-length and using SIMD instructions
   • BLAS[*3]: Equivalent performance by using cache and SIMD instructions

Computation time of libraries normalized by x86

*Input-data: Structural Optimization Simulation (SCF calculation),
   Based on QEF/benchmarks/small-benchmarks/test_1.in, 224 atoms
*1: ELPA used for deriving eigenvalue (https://github.com/fujitsu/elpa)
*2: FFTW3 Library tuned by FUJITSU (https://github.com/fujitsu/fftw3)
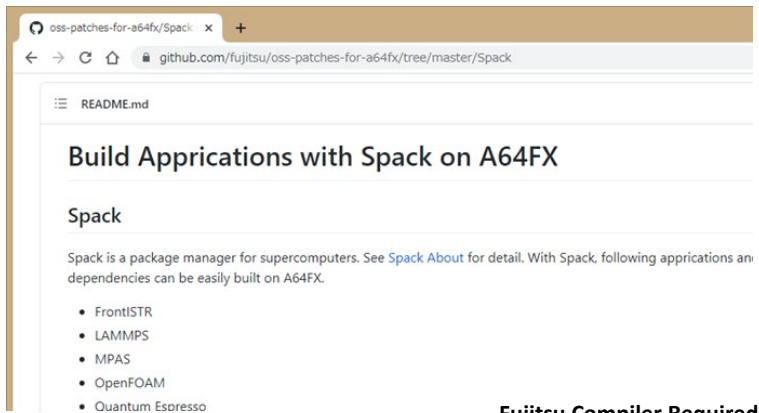*3: FUJITSU SSL Ⅱ Library
   Benchmark platform
   FX1000 A64FX (2.2GHz x 48cores)
   x86: **Xeon 8268** (2.9GHz x 24cores) x2

Information and Tools for OSS

Installation guides and build scripts using Spack	Performance tuning information and patches
Fujitsu Compiler Required	Fujitsu Compiler Required
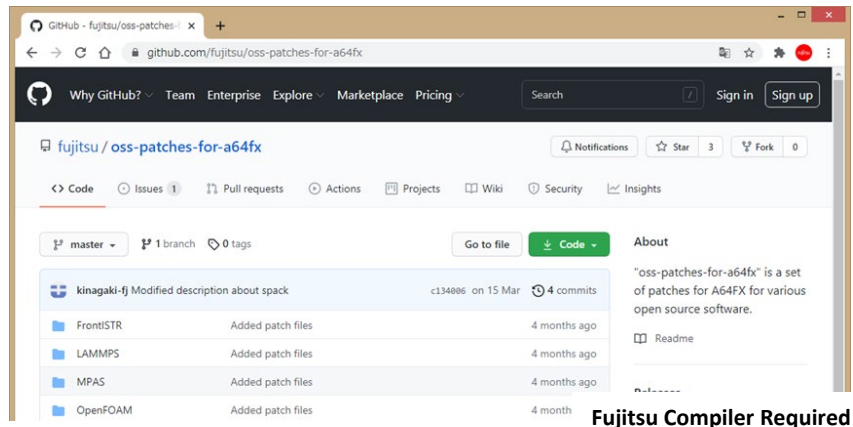
○ Target OSS

OpenFOAM	FrontISTR	ABINIT	SALMON
SPECFEM3D	MPAS	LAMMPS	QUANTUM ESPRESSO

○ Target OSS

OpenFOAM	FrontISTR	SPECFEM3D	QUANTUM ESPRESSO
MPAS	LAMMPS		

○ Information site   (Provided by Spack*)
https://github.com/fujitsu/oss-patches-for-a64fx/tree/master/Spack

○ Information site   (Fujitsu official Github)
https://github.com/fujitsu/oss-patches-for-a64fx

*Spack ： An OSS package manager developed at Lawrence Livermore National Laboratory. For various computing environments and applications, a build script called a recipe is prepared, and the tool can build the application including the dependent library with one command. https://spack.readthedocs.io/en/latest/

Open-Source Oil & Gas Applications evaluated with Arm ecosystem collaborators

A64FX advantage in Oil & Gas applications

✓ *Seismic simulation* = important technology for Oil & Gas vendors

✓ Requires a huge amount of 3D data calculation to identify the geological structure

✓ Can benefit from *High Bandwidth Memory* and *Scalable Vector Extension* of *A64FX*

➤ *Seismic Simulation* image

➤ Example of *Seismic Simulation* Performance on *A64FX*

survey ship
source of shock waves (air gun)
hydrophones
sea bed
sedimentary rock layers
unconformity
path of reflected waves
oil
gas
oil
water
porous reservoir rock

Devito evito
https://www.devitoproject.org/

GIRIH
https://github.com/ecrc/girih

Devito

FD-Gpts/s

Faster

- A64FX (48cores)
- EPYC (48cores)
- Cascade Lake (48cores)

Acoustic SO4

GIRIH 8th order in space & 2nd order in time, w/o temporal blocking

Gstencil/s

Faster

- A64FX(48cores)
- Epyc(128cores)
- Cascade Lake(40cores)

Benchmark 0

Devito

○ About **Devito**

✓ Domain-specific Language (DSL) and code generation framework for the design of highly optimized finite difference kernels for use in inversion methods, developed by **Imperial college London** https://www.devitoproject.org/

○ Performance Data

✓ **A64FX** shows remarkable performance over other CPUs especially acoustic model of 4th space order (SO4)

✓ **HBM** (**High Bandwidth Memory**) and **SVE** (**Scalable Vector Extension**) contribute this advantage.

✓ Reference: https://www.youtube.com/watch?v=--2LrXL6Y2g&list=PLcsG4X8Zn_UAdbYQODr5PQLCcdCo_sovO&index=36&t=61s

• **A64FX**: 2.2GHz 48cores 32GB HBM2
  12 threads x 4 MPI process to fit A64FX Core Memory Group

• **Xeon 9275CL**: 2.4GHz 24cores x 2 sockets

• **EPYC 7R32**: 3.3GHz 48cores

GIRIH

FUJITSU

○ About **GIRIH**

✓ A high-performance stencil framework using multicore wave front diamond tiling, developed by **Extreme Computing Research Center** at **KAUST** *

✓ Spatial and temporal blocking techniques are implemented

○ Performance Data

✓ **A64FX** shows remarkable performance and scalability with spatial blocking, compared with spatial/temporal blocking performance on **Cascade Lake**.

✓ **HBM** (**High Bandwidth Memory**) and **SVE** (**Scalable Vector Extension**) contribute this advantage.

- 8th order, domain size 1024x1024x512, Single precision (benchmark 0)

- **A64FX**: 2.2GHz 48cores 32GB **HBM2** (Tuned code) 1/3/6/9/12 threads x 4 MPI process to fit **A64FX CMG** (**CMG**=**Core Memory Group)**

- **Cascade Lake** performance are quoted from **github** https://github.com/ecrc/girih and provided by **KAUST**.

* **KAUST**: **King Abdullah University of Science and Technology**

GIRIH performance — A64FX(48cores), Cascade Lake(40cores); GStencil/s vs Spatial Blocking / Temporal Blocking

A64FX scalability — Spacial Blocking; Gstencil/s vs #cores

Thank you