

A64FX: The Fastest Arm-based CPU for HPC

Fugaku is the Successor to the K Computer

FUJITSU

Supercomputer Fugaku
Ranked #1 Green500 2019
Top500, HPCG, Graph500, HPL-AI, 2020

K computer

Ranked Top500 No.1 in 2011



© RIKEN



PRIMEHPC FX1000

© RIKEN



PRIMEHPC FX700

- Fugaku is the successor to K computer that has been in operation for 7 years
- Powered by A64FX, an Arm-based processor designed for HPC
- Commercial models, PRIMEHPC FX1000/FX700, are now available

Fugaku Project



- Joint Development by RIKEN and Fujitsu
- Flagship supercomputer for Japan
- Arm64-based fast and Green Supercomputer
- Solving the Social issues in the World*



Application
Performance



Usability



Energy
Efficiency

Fugaku Ultra-Dense Hardware

FUJITSU

Single rack (384 CPUs) = 1 Peta flops
 10^{15} floating point operation per second

K computer

100 rack



Fugaku
1 rack

Energy Efficient HPC



- Supercomputers consume vast amount electric power
 - K computer 13MW = 15K US households
 - K computer X 100 1.3GW = 1.5M US households
- Green500, a global ranking for the energy efficiency
 - Fugaku prototype took No.1 in the Green500 at the SC19,

TOP500							
Rank	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	


- Fugaku has 40x the performance using only 3x the power of K computer!

High Performance and Power Efficiency



The power efficiency of major OSS HPC applications

- Measured on A64FX 2.2GHz normalized to IA CPU 2.9 GHz x2
- High efficiency is achieved by energy-conscious design and implementation



Fugaku ranked at 1st place in all categories

FUJITSU




2.8x of 2nd's



4.6x of 2nd's



2.6x of 2nd's



3.0x of 2nd's

FUJITSU

The First **arm** based HPC processor

A64FX

Architecture
and
Features



Road to Fujitsu's ARM-based Supercomputer



- Strong collaboration with Arm
 - Plus Fujitsu's Supercomputer DNA
- Added features for HPC applications
 - Optimized toward new supercomputer standard
 - Best performance/power
 - High-performance CPU core for HPC usage
 - High memory bandwidth and B/F ratio
 - High scalability of Tofu interconnect
 - Targeting to not only current HPC but also AI applications
- Expanding eco-system with HPC technology and applications
 - Roadmap to general purpose Enterprise
 - Working with Arm/Linux community and OSS



Application
Performance



Usability



Energy
Efficiency


A64FX Overview

■ Fujitsu's Arm Processor

- Lead partner with Arm for SVE architecture
- High speed
- Low Power Design

■ Goal

- CPU with GPU-class performance
- High performance, high efficiency
 - DP performance >3 TFLOPS, > 90%@DGEMM
 - Memory BW 1024GB/s, > 80%@STREAM Triad
- Power efficiency
 - 2-5x performance/watt over IA
 - 16.9GFLOPS/W (HPL)



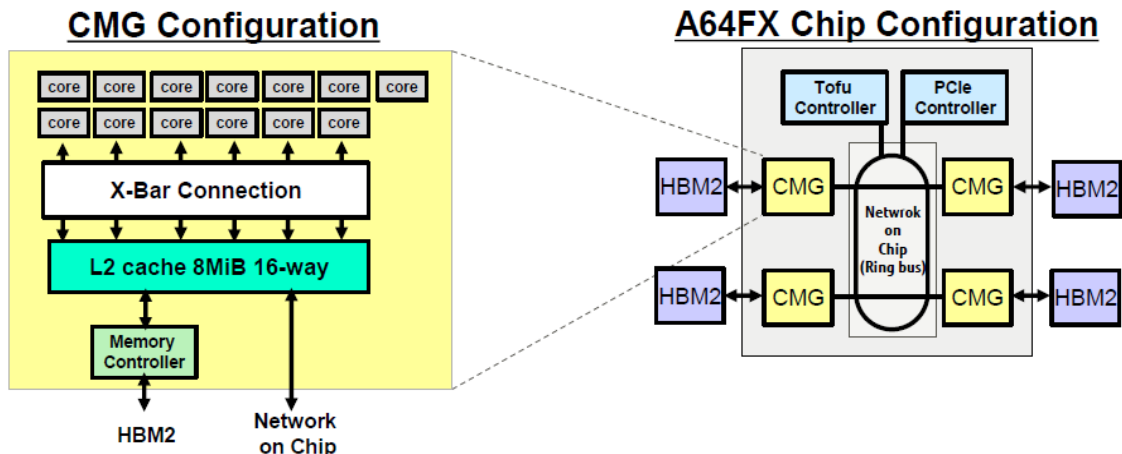
Collaboration with Arm

- A64FX is the first processor of the ARMv8-A SVE architecture
- Collaboration with ARM to develop and optimize SVE for a wide range of applications
 - 512 bit SIMD
 - Added FP16 and INT16/8 dot product for AI applications
 - Many other enhancements
 - Gather/Scatter
 - Four operand FMA
 - Predicated operations
 - Math acceleration
 - Compression
 - HW barrier and sector cache




Many Core Architecture

- A64FX consists of four CMGs (Core Memory Group)
 - A CMG consists of 13 cores, an L2 cache and a memory controller
 - One out of 13 cores is an assistant core which handles daemon, I/O, etc.
 - Four CMGs keep cache coherency by ccNUMA with on-chip directory
 - X-bar connection in a CMG maximizes high efficiency for throughput of the L2 cache
 - Process binding in a CMG allows linear scalability up to 48 cores
- On-chip-network with a wide ring bus secures I/O performance



HPCG results of TOP500 @ ISC20

- Fugaku's efficiency is very high at ~3%




Fugaku and A64FX Greenness on TOP500

■ Power efficiency in GF/W, **w/ ACCEL** and **w/o ACCEL**

■ Top500 - June 22, 2020

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



Full Software Stack for the A64FX CPU



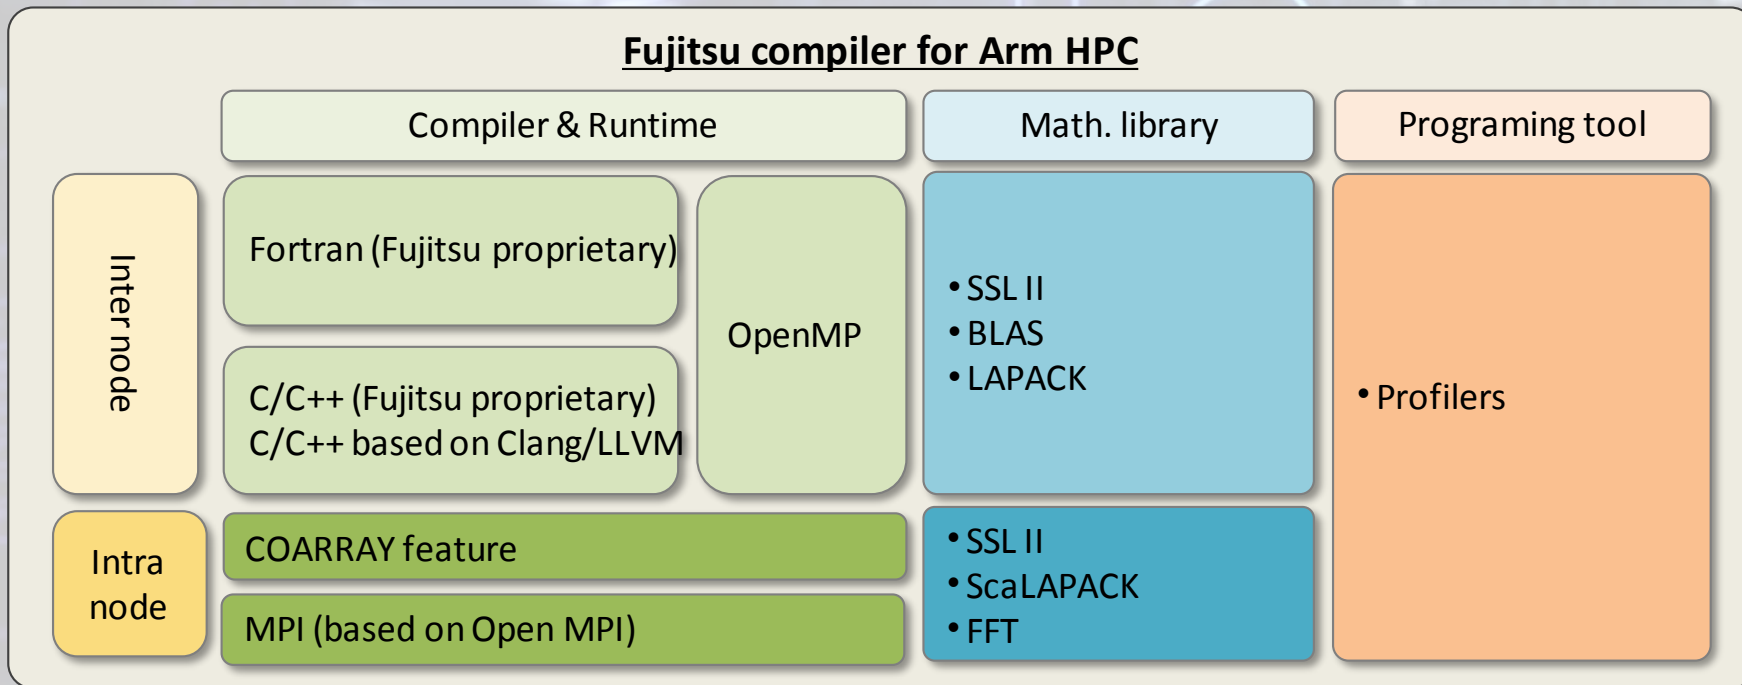
- From Kernel, middleware, compiler to HPC application
- OSS stack provides high usability
- Keeping binary compatibility with the other Arm64 v8.2+SVE systems



Fujitsu's Programming Environment



Fujitsu compiler for Arm HPC

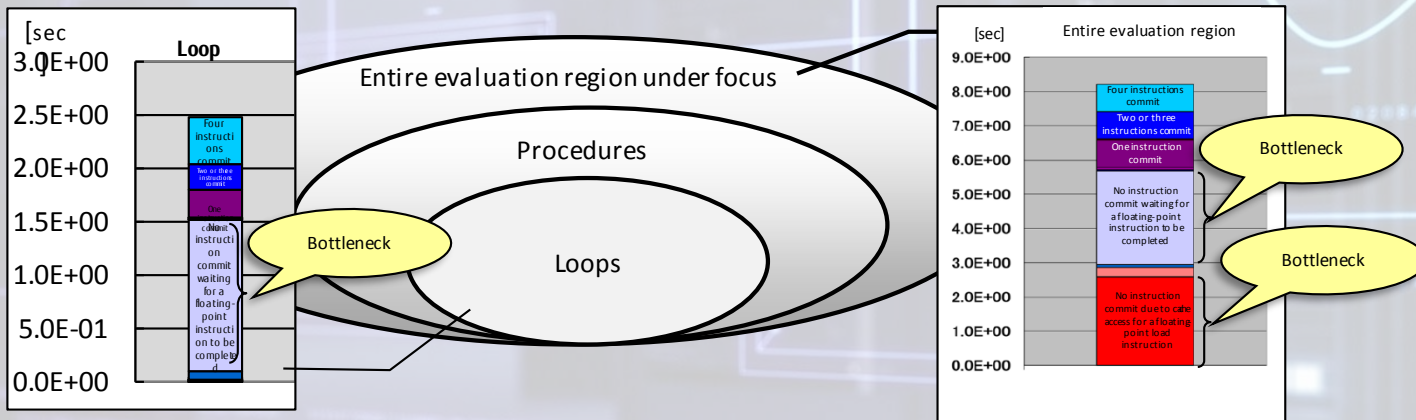


- Supports major HPC programming environment software and latest international standard
- Supports Arm v8 SVE for optimized application performance

Bottleneck Diagnostics with Fujitsu Profiler



- CPU Performance Analysis information produced by FUJITSU Profiler enables you to
 - observe summarized visualization of activities in the CPU
 - determine bottleneck in the CPU
 - obtain data in the level of loop, procedure, or entire region
- Major part of CPU Performance Analysis information is the cycle accounting graph such as the following



- What is a reason of the bottleneck?
- To what extent can the bottleneck be improved?

CPU Performance Analysis information produced by FUJITSU Profiler is the answer to these questions.

A64FX-based PRIMEHPC Family



PRIMEHPC FX1000



PRIMEHPC FX700

A64FX-based PRIMEHPC Platforms



FX1000



High density custom cabinet,
384 nodes/cabinet, Scale to ~400k nodes
Water cooled,
Tofu interconnect D,
Fujitsu software stack

FX700





Standard 19' Rack mount,
8 nodes(=CPU)/2U,
Air cooled,
InfiniBand EDR/HDR100
Utilize ISV and Open Source

PRIMEHPC FX700 Components




Bird's eye View (Front Side)

Bird's eye View (Rear Side)




Blade

CPU(Node)




Bird's eye View of Blade

System Layout of PRIMEHPC FX1000



Layer	No de #	description
CPU	1	CPU and Main Memory CPU includes execution cores, assistant cores, L2\$, memory controller, interconnect controller (ICC) and PCIe interface
CMU	2	CMU (CPU Memory Unit) consists 2x CPU
BoB	16	BoB (Bunch of Blades) consists 8x CMU
Shelf	48	Shelf consists 3x BoB
Main Rack	384	Main Rack consists 8x shelves



shaping tomorrow with you