Performance Evaluation of Supercomputer Fugaku using Breadth-First Search Benchmark in Graph500

Masahiro Nakao†, Koji Ueno‡, Katsuki Fujisawa*, Yuetsu Kodama†, Mitsuhisa Sato†
(†RIKEN Center for Computational Science, ‡Fixstars Corporation, * Kyusyu University)

The supercomputer Fugaku

- The supercomputer Fugaku is under installation at RIKEN Center for Computational Science as the successor system of the K computer
- A64FX CPU on Fugaku has four Core Memory Groups (CMGs) consisting of 12+1 cores and HBM2. They are connected by a network on chip (NOC).
- A64FX has six Tofu Network Interfaces (TNIs)

富岳

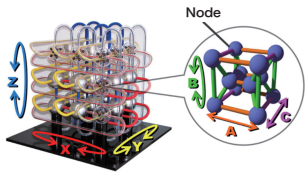

Each node is connected with Tofu-D, 6d mesh/torus network

CPU	A64FX, 48+2/4 cores, 2.0/2.2GHz, 3,072/3,379GFlops(DP)
Memory	HBM2 32 GiB, 1,024GB/s
Interconnect	Tofu-D, 6-dimensional mesh/torus 28.05Gbps × 2 lane ×10 ports

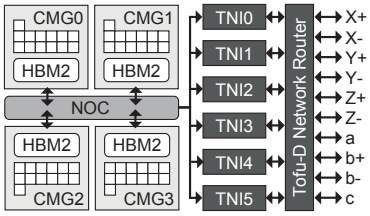

Tuning on Fugaku

Note that Fugaku is currently an evaluation environment before sharing, so the results may not reflect the performance after sharing.

Rendezvous v.s. Eager

The MPI implementation in Fugaku can set the MPI_send() protocol to Rendezvous or Eager. By default, all MPI_send() in our BFS are executed with Rendezvous. Thus, we also executed them with Eager.

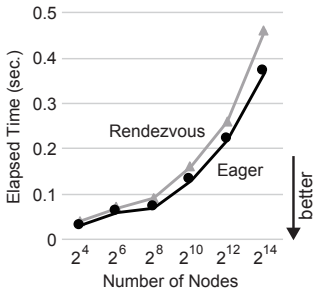

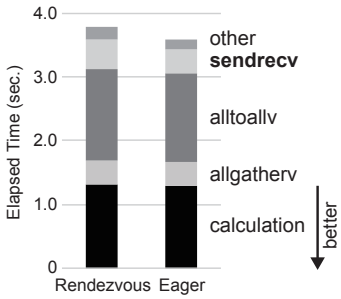

Eager protocol has higher performance

Using Boost mode and Eco mode

A64FX operating at 2.0 and 2.2 GHz is said to be in **Normal mode** and **Boost mode**, respectively. Furthermore, to reduce energy consumption, A64FX can be set to **Eco mode**, in which the floating-point arithmetic pipelines are reduced from two to one.

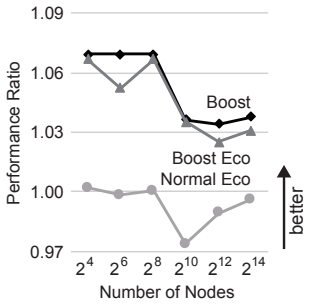

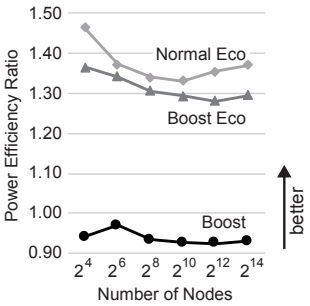

We evaluate the performance and power efficiency (performance/W). Since the Boost/Eco modes are set orthogonally, evaluation is performed for four combinations, namely **Normal**, **Boost**, **Normal Eco**, and **Boost Eco**. A value higher than 1.00 indicates performance higher than that of Normal. The result shows that **Boost Eco is suitable for BFS.**

Evaluation

Six-dimensional process mapping

Fugaku is not currently fully available, so we evaluated BFS using 92,160 nodes. The size of each axis of Tofu-D was (X,Y,Z,a,b,c) = (20,16,24,2,3,2). BFS operated on a two-dimensional process grid (R×C). We set (R,C) = (XY,Zabc) = (320,288), which yields the highest communication efficiency. In the process mapping for C-axis, adjacent communication is suitable, so we made all nodes adjacent nodes. Following figure shows an example of assigning abc-axis(2x3x2) to C axis.

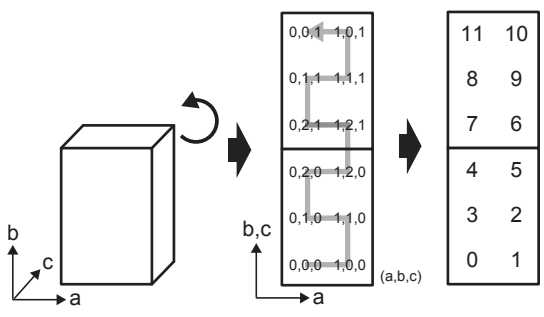

Result

- Set the number of vertices to 2^{40} and Boost Eco using 92,160 nodes
- **The performance and power efficiency were 70,980 GTEPS (Traversed Edges Per Second) and 8.55 MTEPS/W**
- Performance was 2.27 times that of the K computer (82,944 nodes), and power efficiency was 1.93 times that of Mira in ANL
- The performance ratios when executing 16,384 (2^{38} vertices) and 92,160 nodes are as follows