

Supercomputer Fugaku takes first place in the Graph 500 for the two consecutive terms

Further performance improvements were achieved
by utilizing the entire system of Supercomputer Fugaku

A collaborative research group* consisting of RIKEN, Kyushu University, Fixstars Corporation, and Fujitsu Limited has utilized the entire system of Supercomputer Fugaku (Fugaku) at RIKEN and won first place for two consecutive terms following June 2020 (ISC20) in the Graph 500, an international performance ranking of supercomputers for large-scale graph analysis. The rankings were announced on Nov. 12 (Nov. 13 in Japan) to coincide with the 2020 International Conference for High Performance Computing, Networking, Storage and Analysis (SC20), an international conference on high-performance computing (HPC).

The performance of large-scale graph analysis is an important indicator in the analysis of big data that requires large-scale and complex data processing. Fugaku has more than tripled the performance of the K computer, which won first place for nine consecutive terms since June 2015.





※A collaborative research group

Center for Computational Science, RIKEN

Architecture Development Team

Team Leader Mitsuhsa Sato

Senior Scientist Yuetsu Kodama

Researcher Masahiro Nakao

Institute of Mathematics for Industry, Kyushu University

Professor Katsuki Fujisawa

Fixstars Corporation

Executive Engineer Koji Ueno

The extremely large-scale graphs have recently emerged in various application fields, such as transportation, social networks, cyber-security, and bioinformatics. The number of vertices in the graphs has grown from billions to trillions and edges from hundreds of billions to tens of trillions. In recent years, high-performance graph analysis demands have been remarkably increased through converting the real-world data into graph data. Therefore, the large-scale graph analysis has attracted significant attention as a new application of the next-generation supercomputer.

The Graph500 benchmark was designed to measure a computer system's performance for applications that require irregular memory and network access patterns, such as large-scale graph analysis. Following its announcement in June 2010, the Graph500 list was released in November 2010, since it has been updated semiannually. The Graph500 benchmark measures the supercomputer's performance performing a breadth-first search (BFS) in terms of traversed edges per second (TEPS). Thus, the elapsed time of a BFS and the total number of traversed edges determine the performance of the benchmark;

Graph500 Website

<https://graph500.org>

The collaborative research group uses 158,976 nodes, the entire system of Fugaku, and successfully performed a BFS for a huge-scale graph consisting of about 2.2 trillion vertices and 35.2 trillion edges in a harmonic mean of 0.34 seconds. The Graph500 score was 102,955 GTEPS (GigaTEPS), well above the ISC2020 (as of June 2020) score of 70,980 GTEPS. Similarly, the measurement result of the K computer was 31,302 GTEPS (as of June 2019), which is more than three times the performance improvement of the K computer.

As of November 2020, the second place in the ranking of Graph500 is Sunway TaihuLight in China. The measurement result is 23,756 GTEPS. In other words, the performance difference is about four times or more from second place.



The collaborative research group has advanced the development of the software package that performs high-speed processing of large-scale graphs as follows;

- 1: Algorithms to reduce redundant graph searches,
- 2: Optimization of communication performance on massively parallel computers connected by thousands to tens of thousands of high-speed networks,
- 3: Optimization of memory access on multicore processors.

The collaborative research group finally succeeded in coping with large-scale and complicated real data expected in the future and developing graph search software with the world's highest performance.

The fact that Fugaku was ranked first in the Graph 500 demonstrates that it can perform not only regular calculations commonly used in scientific and engineering calculations but also graph analysis involving many irregular calculations. It shows the versatility of Fugaku, which allows it to be used in a wide range of applications.

It also shows the high level of technical capabilities of the collaborative research group that has developed software to maximize hardware performance. In addition to further optimizing communication performance, the collaborative research team will also reduce redundant searches and equalize memory usage at each node.

The collaborative research group has released source codes that have been achieved high-performance on Fugaku and the K computer so far from the following GitHub repositories.

<https://github.com/suzumura/graph500/> (for the K computer)
<https://github.com/mnakao/graph500/> (for Fugaku)