## Science-driven Development of Supercomputer Fugaku



Hirofumi Tomita Flagship 2020 Project, Team Leader RIKEN Center for Computational Science

In this article, I would like to take a historical look at activities on the application side during supercomputer Fugaku's climb to the top of supercomputer rankings as seen from a vantage point in early July 2020. I would like to describe, in particular, the change in mindset that took place on the application side along the path toward Fugaku's top ranking in four benchmarks in 2020 that all began with the Application Working Group and Computer Architecture/Compiler/ System Software Working Group in 2011. Somewhere along this path, the application side joined forces with the architecture/system side based on the keyword "co-design." During this journey, there were times when our opinions clashed and when efforts to solve problems came to a halt, but there were also times when we took bold steps in a unified manner to surmount difficult obstacles. I will leave the description of specific technical debates to other articles, but here, I would like to describe the flow of Fugaku development from the application side based heavily on a "science-driven" approach along with some of my personal opinions. Actually, we are only halfway along this path. We look forward to the many scientific achievements and solutions to social problems that Fugaku is expected to bring about.

### 1. Introduction

On June 22, 2020, the supercomputer Fugaku became the first supercomputer in history to simultaneously rank No. 1 in four major supercomputer benchmarks (HPL, HPCG, HPL-AI, Graph500). Each of these benchmark tests represents a different aspect of computational techniques used in application software, so this achievement can be said to reflect the great potential of the Fugaku computer. Nevertheless, it was never the ultimate goal of the Fugaku Project to take the top position in any benchmark. Actual applications are composite programs that demand high performance in a number of benchmarks such as the ones above.

At a press briefing the next day, RIKEN Project Leader Yutaka Ishikawa commented, "We are almost there, like a climber near the summit of Mt. Fuji, but more work is needed to reach the top and complete the system." This is because improving the stability of the entire Fugaku system is essential and achieving a true victory will only come about once applications begin to produce real results. In this article, I would like to take a look back at the flow along the path toward Fugaku's top ranking as seen from the application side from a vantage point in early July 2020. During this journey, the application side worked together with the architecture/system side (hereafter, the system side), and there were times when our opinions differed and when efforts to solve problems came to a halt. However, there were also times when we took bold steps in a unified manner to overcome difficult barriers. In this article, I will try to be as objective as possible in relating this history, and I apologize if some of my subjective views should somehow seep in. Of course, there are different ways of thinking and points of view, and I welcome any comments or criticisms of what I present here.

### 2. Eve of Fugaku Development Project

The Fugaku Development Project adopted a "science-driven" approach from the start based on the idea of "getting the most out of a supercomputer

to maximize results." The application side and system side reached a consensus on this idea before actual development. Here, collaborative design (co-design) in which the application side contributes proactively to design work was taken to be an essential means of putting this idea into practice. In other words, both the system side and application side would have to work together diligently in architecture and system design.

Nine years ago in 2011, in conjunction with promotion of the High Performance Computing Infrastructure (HPCI) plan, the "Working Group for the Investigation of HPC Technology R&D of the Future" was established under the "Council for HPCI Plan and Promotion," an advisory commission to the Director General of the Research Promotion Bureau of the Ministry of Education, Culture, Sports, Science and Technology (MEXT). This WG, in turn, set up an Application Working Group and Computer Architecture/Compiler/System Software Working Group, which worked closely together to compile the "HPCI Technology Roadmap White Paper" and "Computational Science Roadmap White Paper" released in March 2012 [1]. On the application side, this was achieved through the participation of more than 100 Japanese computational-science researchers who compiled outlooks for a variety of fields in the space of just a half a year. Through this work, one could sense the expectations and air of excitement surrounding HPC at that time.

Next, to further organize the discussions held in these WGs and to explore the possibility of specific collaborative work, the Feasibility Study on Future HPC Infrastructures (Application Working Group) (abbreviated as "Application FS") was launched in July 2012 as MEXT-commissioned research. This study aimed to identify social problems and scientific breakthroughs that computational science could help solve and achieve and drew up a new Computational Science Roadmap as a result [2]. The participants involved in this roadmap formulation were not limited to researchers in the field of computational science. They also consisted of researchers from universities, research institutions, and corporations working on the front line of various academic communities involved in experimental, observational, and theoretical research. These researchers held intense discussions on computers that would be needed from here on and on their performance based on a well-balanced system.

These discussions were not restricted to what types of issues could be tackled by the successor to the K computer (post-K). There was also general agreement that contributions from science and computational science beyond the post-K era should offer some insight at that point in time. It was also agreed that these activities should be continued in some form after the Application FS came to an end.

Based on the Computational Science Roadmap drawn up by the Application FS, a MEXT committee comprised of specialists from the academic and industrial worlds defined those issues to be taken up by Fugaku (Priority Issues and Exploratory Challenges) from the following three points of view [3]:

- (1) Is the issue highly significant from a social and national perspective?
- (2) Can world-leading achievements be expected by tackling the issue?
- (3) Can Fugaku performance be effectively used in tackling the issue?

# 3. Role of application side in Fugaku development

The Fugaku Development Project began in April 2014 as a collaborative venture between RIKEN, a leading research institution, and Fujitsu, a development company, based heavily on a science-driven flow as pointed out in the previous section. At the end of that year, organizations for collaborating with RIKEN in tackling the priority issues mentioned above were selected and nine target applications were established (Table 1). Here, "target application" refers to the application code itself and the target problem to be solved by that application. A target problem has the potential of becoming associated with top-level science in the Fugaku era and involves the solving of pressing social issues, so a high target performance of more than 100 times that of the K computer was proposed for several applications. In the project, it was decided to pursue system design with the aim of achieving the target performance of those target applications. Consequently, as shown in Table 1, target applications were set so that major computational methods in computational science could be exhaustively used throughout those applications.

The true meaning of "exascale," which was often said to be the goal of Fugaku, was not a peak-time

Application	Main computational methods	Scale of target problern	Num. of nodes per problem	Structured- grid computa- tion	Non- structured- grid computa- tion	Particle- method computa- tion	Integer arithmetic	Dense- matrix computa- tion (local)	Dense- matrix computa- tion (global)	Commu- nications, low latency	Commu- nications, high band- width	Adjacent communi- cations	Global communi- cations	Large- capacity I/O
GENESIS	Molecular dynamics method	92,224 atoms	1			~						~	~	
Genomon	Large-capacity data analysis	Read length: 150, 1.4 billion reads	96				~							~
GAMERA	Non-structured-/ structured-grid stencil composite finite element method	1 trillion degrees of freedom	Entire system		>					~	>		>	
NICAM +LETKF	Structured-grid stencil finite volume method + local ensemble transform Kalman filter	Global 3.5 km mesh, 1,024-member ensemble assimilation	131,072	>				~			>	~		>
NTChem	High-precision molecular orbital method (sparse + dense-matrix computation)	720 atoms, 19,680 atomic orbitals	20,732					~	~		~	~	~	
ADVENTURE	Non-structured-grid, finite element method	1.65 billion degrees of freedom	4,096		~			~	~	~		~	~	
RSDFT	Density functional method (dense-matrix computation)	110,592 atoms, 221,184 bands	10,368	~				~			>	~	~	
FFB	Non-structured- grid, finite element method	674.8 billion elements	Entire system		~							~	~	
LQCD	Structured-grid path integral Monte Carlo method	1924 grids	147,456	~						~	~	~	~	

#### Table 1 Nine target applications and co-design items.

machine performance of 1 EFLOPS (exaFLOPS). On the system side, it referred to a speed improvement of approximately 50 times, and on the application side, to a speed improvement of several times including algorithm upgrading, which meant, in total, an improvement in speed of more than 100 times that of the K computer. In this sense, the meaning of "exascale" at that time was the ability to perform computations in a manner equivalent to a 1 EFLOPS computer. It is therefore more correct to say that "100 times faster" refers to an achievement made possible by a strong tie-up between the application side and system side. We had intended to make this meaning clear in an announcement, but articles that mistook this for a 100times improvement in machine performance appeared, which was somewhat of a disappointment for the application side.

Details of this co-design work between the system and application sides can be found in other articles in this special feature, but here I will give a brief description of how they were carried out. The overall image of the project is shown in **Figure 1**. Here, the Application Review Meeting would be held once or twice a month to hold discussions with the system side on issues raised by the nine priority issues WGs centered about needs, ideas, problem reports, and performance measurements. The system side, in turn, would take these issues back with them and discuss proposals regarding revisions and feasibility within individual WGs under the System Review Meeting. Repeating these cycles would eventually determine specifications.

A major problem at this time was not co-design in a single field of a single computational method but rather co-design of many fields and many computational methods. It was not a matter of enhancing only one of the evaluation items listed in Table 1. Additionally, even if technically feasible, it would be necessary to make adjustments within a limited budget taking even running costs into account. At times, compromise would be unavoidable. For example, when



Figure 1 Overall project meeting system and role of Application Review Meeting..

performing computations on a massively parallel computer like Fugaku, large-scale collective communication can easily result in bottlenecks. The application side consequently proposed that a separate specialized network be designed, but as a result of discussions on the high cost of such a network, the limited number of applications that would benefit from it, etc., it was decided to put aside this request.

One way of thinking that came out of this co-design effort was a method of establishing several modes for holding down power consumption and controlling those modes according to the characteristics of the target application. This approach would enable many applications with different levels of quality to run with good efficiency.

Today, in the world of science too, the problems that need solving are changing at a dizzying pace.

Design modifications are even being made at an intermediate stage in application development to increase support for data science including the increasingly popular fields of machine learning, deep learning, and artificial intelligence (AI), to enable high-efficiency half-precision computations that are less accurate but faster, etc.

### 4. Conclusion

In this article, I have described the nearly ten-year development of the supercomputer Fugaku from the application side. This development project was a great test not only for the system side but also for the application side. Here, in addition to making effective use of the developed machine, the aim of co-design was to decide on specifications in a science-driven manner across many fields. This co-design effort was perhaps the first of its kind. I, myself, was deeply involved in Fugaku development from the Application Working Group, but I also had the opportunity to experience many other activities such as coordinating between different application fields and coordinating with the system side. I often reflect on my past approaches to problems thinking that perhaps I could have handled them in a better manner.

From here on, I wonder whether HPC will continue through the development of general-purpose machines such as Fugaku or will alternatively turn to dedicated machines specific to a certain computational method. In any case, I believe that the way in which many researchers in diverse application fields including myself participated proactively in this project in formulating computer specifications has set a strong precedent for future projects. In a broad sense, this is one key achievement of the Fugaku Development Project.

I also believe that another achievement of this project in a broad sense is the formation of a computational science community. The active involvement of the application side in computer design and the birth of activities running throughout the diverse aspects of computational science gave great meaning to this project. In actuality, the Computational Science Roadmap continues to be updated. The computational science community formed by the Application FS was reestablished in 2014 as the Committee for Further Advancement of Computational Science Using HPC Infrastructure (Computational Science Committee). This committee continues to train, maintain, and develop a cross-sectional community in computational science fields, recruit young researchers, create forums for exchanging information and presenting results, and secure computational resources for computational science. Among these activities, the Computational Science Forum, an open cross-disciplinary academic gathering, is being held regularly to implement revisions to the Computational Science Roadmap including new fields. At present, this activity is affiliated with the HPCI Consortium, and in 2017, further revisions were made as the "Computational Science Roadmap (2017)" [4].

As I described at the beginning of this article, we are still at the halfway point in Fugaku development as of early July 2020–reaching the true summit still lies sometime in the future. Nevertheless, the great potential of Fugaku has been proven. This era is being heavily affected by the COVID-19 pandemic and will likely undergo major changes as a result. Against this background, I, as someone who participated deeply in this project from the application side, sincerely hope that Fugaku comes to be used to the fullest to solve social issues and give birth to many groundbreaking scientific achievements. At the same time, from the viewpoint of a single climatologist, I am thrilled at the prospect of using Fugaku for my own actual research as a tool for developing my own scientific field.

All company and product names mentioned herein are trademarks or registered trademarks of their respective owners.

### **References and Notes**

- [1] Archive of Strategic Development of High Performance Computing (SDHPC) Workshop. Report on Future HPCI Technology Development. (in Japanese) https://drive.google.com/file/d/ 12EeAF1q75MW-pFk5hvaqsyHy820rX9Ti/view
- [2] The Feasibility Study on Future HPC Infrastructures (Application Working Group). Computational Science Roadmap. https://hpci-aplfs.r-ccs.riken.jp/document/roadmap/ roadmap\_e\_1405.pdf
- [3] RIKEN Center for Computational Science. Priority Issues. https://www.r-ccs.riken.jp/en/fugaku/outcome
- [4] The Committee for Further Advancement of Computational Science Using HPC Infrastructure. Computational Science Roadmap 2017. (in Japanese) https://hpci-aplfs.r-ccs.riken.jp/document/ roadmap2017/roadmap\_170713.pdf

