Final Version of July 13, 2023

Report of the 2023 RIKEN R-CCS Advisory Council

Meeting of the Advisory Council of the RIKEN Center for Computational Science (R-CCS), May 28 – May 31, 2023, Kobe, Japan

Tasks of the R-CCS AC

The R-CCS AC was tasked with reviewing the activities of R-CCS from 2018 to 2022 following the ToR for review as laid out by the RIKEN President, which is supplemented by the ToR laid out by the R-CCS Director. Both ToR are included in the White Paper for the 2023 R-CCS AC.

Members of the R-CCS AC

Prof. Dr. Dr. Thomas Lippert (Chair)	Director of the Institute for Advanced Simulation Head of Jülich Supercomputing Centre
Prof. Dr. Jack Dongarra	Emeritus Professor Department of Electrical Engineering and Computer Science University of Tennessee
Prof. Dr. William D. Gropp (Acting Chair)	Director National Center for Supercomputing Applications
Prof. Dr. Thomas C. Schulthess	Director Swiss National Supercomputing Center
Dr. Fred Streitz	Chief Computational Scientist, Lawrence Livermore National Laboratory Senior Advisor, Center for Forecasting and Outbreak Analytics (CFA) in the Centers for Disease Control and Prevention (CDC)

Prof. Dr. Katherine Yelick	Vice Chancellor for Research and Distinguished Professor Electrical Engineering and Computer Sciences UC Berkeley
Prof. Dr. Catherine Lambert	Director Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique
新野宏 Dr. Hiroshi Niino	東京大学大気海洋研究所名誉教授 Professor Emeritus Atmosphere and Ocean Research Institute the University of Tokyo
関口智嗣 Dr. Satoshi Sekiguchi	・産総研フェロー AIST Fellow, National Institute of Advanced Industrial Science and Technology (AIST) ・株式会社 AIST Solutions 専務取締役 Senior Vice President (COO), AIST Solutions Co.
下條真司 Dr. Shinji Shimojo (Co-Chair)	青森大学　Aomori　University ソフトウェア情報学部 Faculty of Software and Information Technology 教授
常行真司 Dr. Shinji Tsuneyuki	東京大学大学院理学系研究科・理学部 教授 Professor, Department of Physics, The University of Tokyo

Final Version of July 13, 2023

水口賢司 Dr. Kenji Mizuguchi	大阪大学蛋白質研究所　教授 国立研究開発法人医薬基盤・健康・栄養研究所 AI 健康・医薬研究センターセンター長 Professor of Computational Biology, Institute for Protein Research, Osaka University Director, Artificial Intelligence Center for Health and Biomedical Research, National Institutes of Biomedical Innovation, Health and Nutrition

Information provided to the R-CCS AC included:

Activity Report
White Paper
RIKEN's Vision on the 2030 Horizon
Trip Concept
Agenda (2023 AC Program Chart)
AC Member List
Fugaku Tour Info
Presentations as PDFs for all 3 days

The material was shared in a box folder to which all AC members had access. Some additional materials were shared by email with all AC members.

1 Introduction

The reporting period 2019 to 2023 of the meeting of the Advisory Council (AC) of the RIKEN Center for Computational Science (R-CCS) is characterized by the successful deployment and operation of the Fugaku computer, which continues to occupy leading positions in benchmark rankings, especially on the HPCG and Graph500 lists.

The evaluation of the R-CCS covers not only the deployment and operation of Fugaku, but the activities of research groups both building on the success of Fugaku and preparing for the next system, Fugaku-next. As Fugaku-next will be well into its design and deployment before the next AC meeting, this is the last opportunity for the AC to make suggestions on the strategy and direction of the R-CCS research, operation and support activities.

The evaluation of R-CCS by the AC follows the main items of the ToR of the R-CCS director. The criteria of the two ToRs, that of the RIKEN President and that of the Director of the R-CCS, form the basis of the AC's assessment.

In section 2, the AC presents the assessment of the R-CCS in terms of its strategy, operation and infrastructure development. Section 3 deals with an in-depth evaluation of the 26 research and operation teams and units. Section 4 provides overall conclusions, and section 5 gives the AC's recommendations.

2 General Findings

This section provides general findings about R-CCS's strategy, operations, and infrastructure development. An important development since the 2019 AC review, as recommended by the 2019 AC, is the creation of six divisions. This section ends with a brief summary of findings for each of the divisions, along with some recommendations.

2.1 R-CCS Strategy

The strategy for R-CCS as presented by its director is built around a tight integration of computing with science and the needs of society. The vision is well captured by the phrase "Science of Computing, by Computing, and for Computing". The AC endorses this approach and the focus that this gives not only to delivering science results but to improving the state-of-the-art in computational science.

The director described how R-CCS is implementing this strategy. A major part of this has been the very successful deployment and operation of Fugaku, one of the most powerful computers in the world. The use of Fugaku, combined with expertise in the research teams at R-CCS, to respond to the COVID-19 pandemic is an excellent example of the value of such a world-leading resource to society. The director described a strong and integrated approach to delivering value to society from Fugaku. These include the development of the "Fugaku Society 5.0 initiative," a cloud strategy based on lowering barriers to making use of software on both Fugaku ("Cloud APIs on Fugaku") and on the cloud ("Virtual Fugaku"), and initiatives in important application areas, such as HPC- and AI-driven Drug Development Platforms.

As part of the implementation of the R-CCS strategy, the director has made changes to the management structure and to the research teams and units to better meet the goals of R-CCS. The AC endorses the changes, in particular the creation of new teams to address new application areas and to improve the state-of-the-art in computational science. The AC is pleased to see the strong growth in visiting researchers which strengthen the teams and provide more diverse perspectives.

Finally, as part of its strategy, R-CCS is pursuing a strong, application-oriented approach to planning Fugaku-next, including support for AI and data science, and integration with new computing approaches such as quantum computing.

2.2 Operation, Management, Development of Computing Infrastructure

The AC notes the exemplary execution of the operation, management, and development of the computing infrastructure.

The Operations and Computer Technologies Division leader, Dr. Shoji, presented details of the division's mission and activities. As power and cooling are a major component of the operating cost of Fugaku, this was presented in detail.

Fugaku was designed with energy efficiency in mind. When energy costs roughly doubled in late 2022 and early 2023, R-CCS shut down about ⅓ of the system for three months (roughly August-October 2022) to reduce operating costs. Taking advantage of power knobs in the Fugaku design, the operations division worked with users to reduce the power consumption without significantly impacting the science productivity. After the shutdown, these energy saving features have reduced the power consumption by about 10%. The AC was impressed with the effective collaboration with the user community.

The availability of Fugaku is high (slightly better than the K Computer) and the "job filling rate," a measure of how efficiently the available nodes are being used by applications, is significantly higher than for the K Computer. This is evidence of the high quality of the operation efforts for Fugaku.

The AC was pleased to see the development of a strategy to support applications that can use the commercial cloud. This cloud strategy is built around providing APIs for applications so that they can more easily move between systems. For example, "Cloud APIs on Fugaku" means supporting APIs such as the AWS S3 protocol on Fugaku, and "Fugaku-fying the Cloud", which provides Fugaku applications and software environment on AWS. This exploits similarities between the Fugaku ARM+SVE CPU and the AWS Graviton 3/3E, which is also ARM based.

One feature of the management of the system by RIKEN is that user support is divided between RIST and R-CCS (this has some parallels to how NSF supports users in the US, with NSF ACCESS (and formerly XSEDE) having some of the roles that RIST has). This separation sometimes caused problems for users. The AC applauds the deployment of an improved and more integrated platform, Zendesk, which provides a more effective user experience. The AC is pleased to see steps to improve coordination between R-CCS and RIST in the support of users, and recommends that it be broadened to include selection of projects to receive time on Fugaku.

The AC was pleased to see growing collaborations with HPC centers and users world-wide. The future plans described are reasonable, with an emphasis on continuing to develop the Cloud strategy and a focus on a better user experience. One challenge is in staffing, and the division is working to hire more staff.

2.3 Fugaku and Beyond

The AC finds that the Flagship2020 project is a resounding success. More than 3 years after deployment, Fugaku remains the fastest system for some leading benchmarks, such as HPCG. The AC notes that HPCG is representative of an important class of applications, and reflects the focus on designing and operating a system best suited to science applications. The thought and care in the design and operation of Fugaku shows in many of its features. The AC found two particularly noteworthy. One was the innovative and effective response to the power crisis, which also highlights the foresight in developing power controls in the architecture. The other was the addition of FP16 to the Fugaku CPU. This has enabled AI efforts, supporting rapid growth in AI research and applications. The overall success of Fugaku is reflected in the many awards and outstanding science results resulting from the use of Fugaku.

Unfortunately, the lifetime of leadership computer systems is relatively short, and planning for the next system must begin almost as soon as the current system is deployed. The AC was impressed with the planning for FugakuNEXT, and endorses the strong, application-oriented approach to planning FugakuNEXT The AC was impressed with the analysis of tradeoffs between technical difficulty and application performance impact as well as the power efficiency of different approaches. The AC notes that the likely architecture will have more complex computing and memory components than Fugaku, and will require careful co-design with software, algorithms, applications, and hardware. Sufficient time for development of the full software stack and of applications will need to be included in planning. In summary, the AC believes the approach being taken for the planning for the very difficult task of FugakuNEXT is strong.

2.4 Center Management

The AC is very pleased with the changes to R-CCS center management since the 2019 review. At the top level, the number of direct reports to the Director has been reduced to an appropriate number. This also provides more focus for and organized synergies in the work of the teams and units. The AC also supports the approach of leveraging visiting researchers and students to strengthen the team. Not only does this provide additional human resources, it is a value source of diverse perspectives, which are known to improve the quality of research.

The AC notes that, like many other HPC centers, R-CCS has trouble attracting and retaining staff, and also continues to struggle to improve the diversity of the staff. The AC approves of the steps taken to address these issues, particularly the “long view” in the support for education programs, summer schools, and internship programs. The AC encourages further steps in these directions, and encourages R-CCS to work with RIKEN to address challenges that are beyond the control of R-CCS.

2.5 Cross-cutting Observations

The AC noted a number of cross-cutting observations. Some of these are common issues faced by several teams while others are common themes and strengths.
Of the cross-cutting issues identified, there were three that stood out. First and foremost was the need for a national strategy for AI that would strengthen coordination and collaboration and the prioritization of research directions and investments in resources. Second was diversity in the research staff; while the AC applauds the efforts already undertaken, more is needed. Third was hiring issues, ranging from student interns though research staff. In this latter case, one element of these hiring issues is related to the salary and working conditions offered to computer science workers by other institutions, which are often significantly better than offered to researchers in other fields.

The AC observed many cross-cutting strengths:

- The world leadership in HPC and AI for science was evident.
- The AC applauds the use of interns and visitors to strengthen activities. Benefits of this approach include diversity, outreach, impact. This also helps to leverage both capabilities of R-CCS staff and the Fugaku system.
- The AC was pleased to see common themes, such as for AI/ML, Quantum computing and information, life sciences, and Fugaku-next, in the presentations by the teams and units.
- The AC noted the great progress since the 2019 review and appreciated the detailed responses to the recommendations of the 2019 review.
- Fugaku is a tremendous success. The design, operation, and utilization are world-leading. The AC was impressed with the innovative response to the power cost crisis, including working effectively with the user community to reduce power consumption.
- The research teams are strong, with excellent visitor and student programs. There are many good collaborations but there are more opportunities to pursue.

2.6 Highlights

The AC was impressed with the progress and accomplishments of R-CCS. In this section we briefly list a few highlights.

The deployment and operation of Fugaku serves as an example of how a world-class center should be managed and operated. The AC was impressed with the innovative approach to the power crisis, which took advantage of power “knobs” built into the design as well as a strong collaborative approach with the user community. It is notable that the Japan Meteorological Agency adopted this architecture for experimental operational weather prediction! Another highlight was the early application of Fugaku to the COVID-19 emergency. The AC also found interesting the connections between Fugaku and the commercial cloud, particularly the ability to move software between the systems.

Several software packages have been developed and are impacting the community. One example is the development and release of GENESIS packages, which are versatile and are delivering value to multiple teams.

Notable since the last review are the strong external collaborations and outreach to the HPC community.

R-CCS has addressed very well the management recommendations from the 2019 review. The new organization positions R-CCS for success. The AC applauds the difficult step of sunsetting some teams, and notes that action provides resources for emerging research directions. The streamlined organization of R-CCS is well-positioned to deliver on science.

2.7 Concerns

The AC noted a few areas of concern. Some of these are not fully under the control of R-CCS but do impact the success of R-CCS

1. Fugaku is a uniquely valuable resource in which a great deal has been invested. To get the most value from this investment, the cost of operation of Fugaku needs to be fully funded.
2. For maximum impact and effectiveness, there needs to be a national AI strategy. The AC recommends that R-CCS should lead the development of the HPC/AI part of that strategy, and be a part of the development of the overall national AI strategy.
3. Staffing levels are critically low in some groups.
4. Diversity remains an issue.
5. Architecturally Fugaku is designed for science. Promises to develop commercializable platforms may conflict with this design.
6. The growth of AI research is laudable, but science has other big data problems beyond or in addition to AI.

2.8 Brief Summary by Division

This section provides a brief summary by the six divisions of R-CCS, including findings and high-level recommendations. More details about each team, including specific findings and recommendations, are presented in Section 3.

2.8.1 Computer Science

This division includes the teams of Programming Environment, Advanced Processor Architectures, Parallel Numerical Technology, Next Generation High Performance Architecture, High Performance Big Data Systems, High Performance AI Systems, and Supercomputing Performance Research.

Findings:
- Excellent engagement and interactions with the international community regarding computer science-related matters.
- Particular strength in computer architecture critical to FugakuNEXT.
- The units have adapted to emerging problems, e.g., mixed precision algorithms, edge computing hardware, etc.
- Some activities have been appropriately deemphasized to make room for new research areas, such as Quantum/HPC hybrid.
- Successful recruitment of new talent.

Recommendations:
- Maintain a strong focus on addressing the needs of applications (current and future) throughout the development process.
- Continue to track technology trends at all levels (e.g., processor to cloud).

2.8.2 Computational Science

This division includes the teams of Field Theory, Discrete Event Simulation, Molecular Science, Quantum Physics, Biophysics, Climate Science, HPC Engineering Applications, Data Assimilation, Structural Biology, and Disaster Mitigation and Reduction.

Findings:
- World class research at the Center with strong application teams that collaborate well with the user community.
- These teams are delivering the lighthouses against which FugakuNEXT can be developed. (e.g., data assimilation that bridges "big data" and numerical models, incorporating ML, has been contributing to innovative predictions of various socially important problems).

Recommendations:
- Collaborations with the computer science teams are encouraged, particularly for the adoption of modern software engineering methodologies and modern programming models. This will help with the “overfitting problem” seen on Fugaku and give more design & implementation options for next generation systems.
- Consider integrating the science teams in data science and AI as well, in order to develop end-to-end solutions and pipelines.

2.8.3 HPC and AI-driven Drug Development Platform

This division includes the teams of Biomedical Computational Intelligence, Medicinal Chemistry Applied AI, Molecular Design Computational Intelligence, and AI-driven Drug Discovery Collaborative.

Findings:
- Exciting new division demonstrates the TRIP framework to address Society 5.0 challenges.

- A good overall vision is presented with promising new results in some areas including target discovery and designing anti-SARS-CoV2 molecules.
- Building an AI-driven drug discovery pipeline is an ambitious goal. Good collaborations are on-going with other divisions of RIKEN and external organizations (including pharmaceutical companies) but more will be needed with clear division of labor and overall coordination.

Recommendations:

- The AC notes that the Division Leader is currently also leading two of the teams. Given the large expanse of this project we recommend identifying as quickly as possible independent Team Leads.
- The AC recommends that more regular communication be established with other related elements across R-CCS (and RIKEN) to maximize opportunities for collaboration and contribution.

2.8.4 Quantum-HPC Hybrid Platform

This division includes the teams of Quantum-HPC Hybrid Software Environment, Quantum Computing Simulation, and Quantum-HPC Hybrid Platform Operations.

Findings:

- R-CCS is currently looking at tools for integrating Quantum Computing with HPC.
- R-CCS has strong computational science strength in quantum chemistry, condensed matter / materials, and field theory that should be involved in designing quantum (computing/science) experiments for RQC.

Recommendations:

- The low hanging fruit is in the collaboration with RQC with a focus on quantum physics.
- Include computational scientists in the analysis of experiments and support RQC, along with the computer scientists, in designing error correcting and resilient QC architectures.
- In the long term, R-CCS computer engineers, with their expertise in designing HPC architectures, should support better integration of digital computing architecture inside the cryogenic environment of future QC.

2.8.5 Operations and Computer Technologies

This division includes the teams of HPC Usability and Development, Facility Operations and Development, System Operations and Development, Software Development Technology, and Advanced Operation Technologies.

Findings:

- The successful transition from K-computer to Fugaku computer over the past four years is a remarkable achievement that deserves recognition.
- The AC acknowledges that these teams are actively addressing numerous practical challenges. Machine Learning of operations data will contribute to the automation of the system operation.

- Regrettably, the AC has yet to see sufficient progress in hiring - human resources is the most crucial agenda to sustain the operation of such a large facility.

Recommendations:
- To enhance their effectiveness and drive meaningful outcomes, the AC recommends that they prioritize clarifying key performance indicators (KPIs) that are of utmost significance to both the users and the teams involved.
- It is challenging to secure human resources, but the AC advises you to identify the obstacles and take appropriate measures.

2.8.6 Fugaku Society 5.0 Initiative

This is a new office with cross-cutting connections to the application of Fugaku to the challenges of society.

Findings:
- We welcome the creation of the Initiative for Fugaku to significantly contribute to the Society 5.0.
- Rapid simulation of huge cases of air drop of COVID-19 infection contributes to the society largely.
- Digital Transformation (DX) platform for drug discovery expects to show a large impact on society.
- Digital Twin Simulation with real data mining of mobile phones contributes to evacuation planning for Kobe city.
- Virtual Fugaku and Platforms contribute to easy access from a wider array of customers.

Recommendation:
- For realization of Platforms for Society 5.0, more input and collaboration with the R-CCS teams are required.

3 Summary of Evaluation of Research Groups

In this section, the AC provides a detailed evaluation of the individual research teams and operation and development units. The assessment follows the following criteria:

Evaluation criteria:
- The ToR of the President of RIKEN
- The ToR of the Director of the R-CCS
- Recommendations of the last AC from 2019 and their fulfillment
- Quality of work during 2018-2022
- Strategy for 2025-2031

Final Version of July 13, 2023

The terms of reference (ToR) of the RIKEN President and of the R-CCS Director provide more specifics about the evaluation criteria, and these are reproduced below. There are also specific criteria for research infrastructure centers such as R-CCS.

ToR of the RIKEN President:

1. Evaluate the responses to the 2019 AC recommendations.
2. Based on the results of the Center's self-analysis, evaluate operations and R&D activities for the 4th Mid- to Long-Term Plan period (FY 2018 - 2024).
3. Evaluate the policies of the 5th Mid- to Long-Term Plan period (FY 2025 - 2031) and recommend new directions for operations and R&D that should be implemented and promoted.

Criteria for Research Infrastructure Centers:

- Governance and management to maximize R&D achievements
- Research initiatives for operation, sharing, upgrading, and utilization of R&D infrastructure
- Creation of outstanding research results for advancement and utilization
- Creation of results that contribute to the development of science and technology, the economy and society through external sharing of research and development infrastructure, etc.
- Returning research achievements to society including public relations activities
- Industry-society tie-ups
- Nurturing and recruiting research talent: Brains Without Borders, Diversity

ToR of R-CCS Director:

- Overall contribution to the core research directives of R-CCS
    - Especially in relevance to our mission, science of/by/for computing
    - International competitiveness of the HPC research as well as the infrastructure
    - Contributions and impact to the society, e.g., “Society5.0 (satisfaction of societal sustainability goals through digital transformations”
    - Plans for next generation of computing --- large scale AI platform, FugakuNEXT, Quantum-hybrid testbed, etc.
- Contribution to the “Research DX”, or digital transformation of research
    - Digital transformation of scientific process itself through digital transformation, e.g., Riken TRIP project, and how HPC can serve as a platform for such transformations
    - R&D activities by R-CCS to contribute to Research DX
    - How organizational changes and upgrades of R-CCS would be contributing to the changes
- Internal and External Collaborations in the area of HPC as a global top-tier Supercomputing Center
    - How existing collaborations are contributing to R-CCS objectives and goals

    - How new collaborations are expected to provide additional contributions.
    - Whether procedural/administrative changes are contributing positively to such collaborations
- Outreaching to the HPC Community at large and to the industry
    - Outreaching activities at various levels
        - To Japanese and International research community
        - To industries, not just Fugaku as a machine but the research portfolio
        - To the society-at-large, governments, startups, people, etc.
    - Establishment of the “Society 5.0 office”
- Human resource development, Diversity and Career plan
    - Various educational activities
        - Domestic and international HPC schools
        - Public outreach to encourage the next generation HPC scientists
    - Diversity - Nationality, gender, school/professional career, age/experience
    - Career plan – hiring, attrition to other institutions, opportunity for internal career changes

3.1 Dr. Aoki, Field Theory Research Team

The team aims to answer fundamental questions in particle and nuclear physics by precisely verifying the Standard Model in particle theory and pioneering theories beyond the Standard Model through computational approaches based on quantum field theory. To this end, we are developing various algorithms for Lattice QCD and developing code to exploit the performance of Fugaku.

To date, they have achieved the world's first performance of over 100 PFLOPS (more than 38 times faster than K-Computer) using Fugaku. AC appreciates this milestone achievement. They have also implemented the most efficient domain-wall fermion method among chiral fermion formulations and started applying it to nuclear material calculations at finite temperatures.

In addition to these research activities, he contributes to the Lattice QCD research community through JLDG (Japan Lattice Data Grid) and ILDG (International Lattice Data Grid). He also represents the basic science field in the FS of FugakuNEXT and contributes to the future planning of HPC. These activities are in line with the 2019 AC recommendations.

Collaboration with different disciplines was also recommended in 2019, including the transfer of their simulation technology to materials science. The team has very recently started working with Dr. Yunoki’s team on tensor networks, a promising method common to QCD and materials science, and is embarking on cross-disciplinary collaboration. The AC recommends further strengthening such interdisciplinary collaboration, not only within R-CCS but also with external researchers.

3.2 Dr. Ito, Discrete Event Simulation Research Team

The aim of this team to challenge microscopic simulations for macroscopic natural and social phenomena with discrete-element and/or agent-based models using the Fugaku supercomputer has become increasingly significant in light of the ongoing pandemic and the emergence of smart cities. The integration of extensive real-world data, such as mobile data provided by mobile companies, has rendered these simulations highly applicable to policy-making.

The Advisory Committee (AC) recognizes and supports the group's efforts in applying High-Performance Computing (HPC) to the concept of Digital Twin, which effectively combines data and simulation. In response to the Covid-19 pandemic, the group has collaborated with Softbank to conduct agent-based simulations and big data mining. They have also collaborated with local governments, including Kobe City, to address the challenges posed by the pandemic and contribute to the realization of Society 5.0. The AC commends the group for embarking on research related to quantum computing simulations.

Furthermore, the AC recommends adopting a more convergent approach to Digital Twin within the Research Center for Computational Science (RCCS). Given the importance of technologies and platforms that enable the realization of Digital Twin in the social sciences through the fusion of data and HPC, a more convergent approach is deemed necessary.

3.3 Dr. Nakajima, Computational Molecular Science Research Team

This team aims to explore the frontiers of molecular science through the development of quantum chemistry theory, algorithms, and software and the use of Fugaku. On the theoretical side, they have proposed various new methods such as a coupled cluster method using explicitly-correlated wave functions, a high-precision transition state theory, a two-component relativistic electronic structure calculation method that is comparable to a four-component method, and a method to efficiently calculate electron correlations by dividing the range.

The team develops and provides a general-purpose, multi-functional, massively parallel electronic structure calculation package called NTChem. The code is highly tuned for Fugaku and achieves, for example, 70 times better performance than K-Computer for MP2 calculations. Globally, this field is an oligopoly of a few packages, but NTChem has outstanding features compared to them, and so the AC appreciates its development and recommends further efforts to explore needs and expand users.

This team is also attempting to develop new materials using Fugaku by high-throughput simulation or the so-called materials informatic approach, and has proposed promising materials for perovskite solar cell materials, heat-resistant polymer materials, biodegradable plastic materials, etc. The combination of NTChem and materials informatics seem to be so powerful that it would be worthwhile to offer them as a package to the community.

3.4 Dr. Yunoki, Computational Materials Science Research Team

The team aims to understand and predict quantum states of condensed matter in the strongly correlated regime. This is one of the most demanding areas of study in condensed matter physics, and in the past has had difficulties finding good applications. However, with the rapid development in quantum engineering and a realistic outlook for significant quantum computing experiments to materialize in the coming decade or two, the simulations of strongly correlated quantum many-particle systems with supercomputers is receiving renewed attention.

The team of Dr. Yunoki has a strong track record in development of large-scale quantum Monte Carlo (QMC) and density matrix renormalization group (DMRG) methods. Both are well known, canonical approaches that require very large-scale computing resources to simulate large quantum many-body systems with high accuracy. The combination of advanced algorithms development and the availability of supercomputers like Fugaku can push the simulated system sizes to scales that are interesting for some of the most advanced quantum computing experiments. More about this in the quantum computing section.

For large-scale QMC, the team is developing a highly scalable auxiliary-field QMC (AFQMC) application for the study of fermionic lattice systems like the Hubbard model. They are using and optimizing the delayed update algorithm, which improves data localization of the updates of the Green's function. They managed to tune the implementation so that it reaches 70% of peak floating-point performance at scale on Fugaku. This enables them to study models with more than ten thousand sites on a honeycomb lattice (graphene), allowing them to make clarifying statements on the universality class of the celebrated Gross-Neveun model in particle physics. They have also been able to investigate models topological Mott insulators

One of the hallmarks of this team is their leading expertise in DMRG, a method that is exact in one dimension, but with its massively parallel implementation on Fugaku, they have been able to apply DRMG to study 2D quantum many-body systems as well. And once again, their algorithms and software implementation reach an impressive ~80% of peak floating-point performance at scale on 82,488 nodes Fugaku. With a kernel polynomial method, they managed to extend the DRMG to time-dependent studies of 2D models.

These are just a few of several algorithmic directions and corresponding software libraries they are developing. The team is quite prolific, producing 78 refereed publications in respected journals over the past 4 years. Their software has been used in 28 HPC projects on R-CCS infrastructure since 2012. Furthermore, the team is fostering young staff, producing tenured research scientists at R-CCS, as well as faculty at two universities.

3.5 Dr. Sugita, Computational Biophysics Research Team

The computational biophysics work at R-CCS is focused on atomistic and coarse grain models of biomolecules. Undoubtedly this is one of the most promising applied domains of physics and chemistry. Almost all research work in molecular studies of life will at some point have to rely on such simulations to determine structure or function of biomolecules. The recent pandemic demonstrated the crucial role of atomistic simulations for the determination of the SARS-CoV-2 spike protein structure and function. This was the basis for the rapid development of vaccines and treatments. The same molecular simulation techniques find applications in many other areas of the pharmaceutical industry as well, and it is thus not surprising that molecular biophysics is well positioned to make significant contributions in a digitalized society. With the development of GENESIS, a free-software suite for scalable atomistic simulations that is co-designed with Fugaku, the team of Dr. Sugita is making a key contribution to this important field.

GENESIS integrates ML methods, Bayesian inference, etc. with state-of-the-art molecular dynamics (MD) techniques to analyze experimental data, in order to determine structure and function of biomolecules. For example, microsecond-long atomistic simulations of SARS-CoV-2 spike proteins with gaussian accelerated MD showed how the inherently flexible protein interacts with glycans to stabilize in two conformations that are linked to the protein's function. There are many studies like this that put Fugaku at the forefront of computational biophysics.

The team recognizes that co-designing GENESIS with Fugaku at scale is particularly well positioning the system for weak scaling to massive numbers of particles. This, however, is a niche with an inherent weakness of reaching very limited simulated time compared to competing codes – tens of simulated nanoseconds per day and reaching microseconds with acceleration techniques, while the relevant timescale for classical molecular systems are in the millisecond range and longer. The team recognizes this and proposes to implement GENESIS for other, competing architectures that have proven to support faster simulations for smaller systems with competing software like GROMACS. While this is definitely a good investment from the point of view of the GENESIS software, we would like to propose that R-CCS consider alternate strategies that leverage the strengths of other teams. Ultimately, highest value will be attained by marrying large scale provided by GENESIS with high performance provided by, e.g., GROMACS.

The team also pointed out the relevance of combining classical molecular mechanics (MM) with quantum mechanical methods (QM) to cover those areas where reactions happen that cannot be described by the force fields used in classical MM. On the very short timescales, such as those studied in pump-probe experiments, such combined QM/MM methods will be necessary. and have been investigated for several decades. Originally the electronic structure part was solved using single particle methods based on density functional theory (DFT). Experience has shown, however, that QM/MM methods applied to biomolecular systems in pump-probe experiments require more accurate many-body methods, such as those based on the Bethe-Salpeter equation. This is challenging and beyond the horizon for most research groups, since it requires access to capability computing and deep knowledge of computational quantum many-body theory. It seems to us that R-CCS is optimally positioned for both challenges – though it would require close collaborations of GENESIS developers in the biophysics teams with colleagues from the quantum chemistry (computational molecular sciences) and the condensed matter physics teams.

3.6 Dr. Tomita, Computational Climate Science Research Team

The objective of this team is to resolve important problems in the modeling of the atmosphere and climate and to indicate a direction of future climate modeling using high performance computers with reliable suggestions. A reliable numerical model is a basis for a better data assimilation that estimates the most reliable current state from big observational data, which in turn gives accurate initial values for the model, thus contributing to society through better weather predictions and climate projections. It also can be used to better understand the mechanism of the atmospheric processes.

The atmospheric motions range from micro-scale phenomena such as turbulent eddies in the atmospheric boundary layer to global-scale phenomena such as planetary waves and El Niño-Southern Oscillation, where the non-linear interactions among the motions at all scales play important roles. This means that, even in the “Fugaku-next” era, small-scale processes such as those occurring among turbulent eddies and among aerosol/cloud/precipitation particles cannot be resolved in the climate models and have to be parameterized. The AC appreciates this team’s efforts to develop reliable weather and climate numerical models and libraries such as SCALE (Scalable Computing for Advanced Library and Environment) through systematic studies of basic atmospheric processes and new attempts of automatic parameter tuning using data assimilation. It is also noted that the execution time was shortened to about 2/3 through their efforts for optimization after Fugaku operation was started.

The team also applied SCALE to clarify important processes in meteorology and climate science. The team has an excellent number of refereed publications (50) since 2018 and has a good number of invited talks, press releases, awards and grants.

The AC appreciate the team’s responses to the 2019 AC’s recommendation: In order to increase the number of users of SCALE (Scalable Computing for Advanced Library and Environment), they started to hold workshops and meetings with the users in Japanese universities, and also initiated international joint researches with ECMWF (European Center for Medium-range Weather Forecast) and the University of Oxford, and model comparison with WRF (weather Research and Forecasting model) in JST-SATREPS project with Argentina. They also plan to study physical processes using a cloud-resolving global model after having studied them in detail in regional models in these years.

The AC acknowledges the team’s attitude to “return to basic research”. It is important to continue their efforts to develop reliable models, and this can be only accomplished by incorporating reliable parameterization schemes. The teams’ plans to use large eddy and direct numerical simulations to clarify the physical processes of turbulence, to combine the big data produced by the simulations to improve the parameterization schemes, and to use automatic parameter tuning are appreciated.

The team accepted 2 Ph. D. course students in a junior research associate program, and 6 internship students (including 2 foreign nationals) for fostering potential modelers.

3.7 Dr. Tsubokura, Complex Phenomena Unified Simulation Research Team

The goal of the team continues to be to develop a unified simulation method for meso/macroscopic complex phenomena by utilizing HPC environment, in particular coupling problems of fluid motion, structure deformation, aeroacoustics, chemical reaction, and others, which is same as the one at the last AC.

CUBE which is a multi-physics solver for fluid dynamics and solid mechanics problems using immersed boundary method to deal with complex geometries is used to efficiently solve various complex industrial and social problems such as efficient aerodynamical design of cars, phonation mechanisms of sibilant, and droplet/aerosol dispersion that significantly contributed to countermeasures for the COVID-19 pandemic in various aspects. These contribute significantly to Society 5.0. Their Fugaku-Based COVID-19 research was not only awarded an ACM Gordon Bell Special prize in 2021, but also was featured by many mass media - more than 400 newspapers, 400 TV shows and 1500 web news all over the world. The team published a good number of peer reviewed publications (34) since 2018 along with many invited talks.

In response to the 2019 AC's recommendation, the team started to apply machine learning to the outputs of CUBE to develop a surrogate model for multi-object shape optimization, and to develop, with an industry company, a digital twin of air conditioning systems using Physics-Informed Neural Networks, both of which reduce drastically the time for the design. Validations of the results of CUBE against laboratory data such as those in wind tunnel experiments were also made successfully. The team carried out their plan for developing and fostering junior scientists: The team accepted 15 student trainees from universities in 2022. The team leader also provided the team members opportunities to give a lecture in his class at Kobe university. As a result, 5 of 6 past members obtained academic positions.

The AC welcomes their future plan to further develop highly integrated HPC simulation of CUBE with data science, to utilize CUBE for the "Fugaku Next" project, and to contribute to Society 5.0 such as policy making and effective manufacturing.

3.8 Dr. Miyoshi, Data Assimilation Research Team

Data assimilation (DA) combines big (observational) data and a numerical model to give the most reliable quantitative estimate of the current state, which also provides reliable initial values indispensable for better prediction of future states. It can be also used to improve numerical models by estimating better model parameter values, and to design time and locations of effective observations or new observational tools that would improve the predictions. The AC acknowledges that this team has been one of the leading groups of DA in the world. The team has an excellent number of refereed publications (85) since 2018 and a large number of invited talks, press releases, awards and grants.

The team made excellent responses to the recommendation of the 2019 AC: Firstly, the team, during the Tokyo 2021 Olympic and Paralympic Games, achieved innovative daily weather predictions that assimilated phased-array weather radar data observed roughly every 30 seconds with a numerical model SCALE developed by the Computational Climate Science Research Team. Secondly, the team started to develop an approach to integrate DA with AI for predicting precipitation, tsunami inundation and so on.

The team has been also developing a particle filter, which is a new ensemble data assimilation method that does not assume Gaussian error distributions, and a mathematical theory that may be applicable to possible weather control with minimal perturbations. The team's works on DA are not confined to meteorological problems, but are now applied to various problems such as predictions of ocean environment, ecosystem control, press-forming, COVID-19 and, glass melting processes all of which contribute to Society 5.0. These activities result from various collaborations with internal and external groups including industries and international research organizations. The team leader has joint appointments at RIKEN iTHEMS (Interdisciplinary Theoretical and Mathematical Science Program) and RIKEN CPR (Cluster for Pioneering Research). These together with the RIKEN President's Initiative for "DA innovation hub" contribute to expanding cross-disciplinary collaborations.

The team considers that they should act as one of the educational centers of DA in Japan. They not only educate young scientists in their group to become experts in DA at an internationally high standard, but also teach DA courses in universities such as Kyoto University and so on. They also regularly hold DA schools and DA workshops.

The AC welcomes their future plan to further develop theory and algorithms for efficient and accurate DA that include a consideration of observation error correlation and an integration with AI technique, and to expand applications of DA to wider simulation fields.

3.9 Dr.Oishi, Computational Disaster Mitigation and Reduction Research Team

Dr. Oishi presented work describing efforts to simulate the effect of disasters, both real and potential, and deliver with these results information of value to policy/decision makers. The development of an Automatic Urban Model Development System allows the rapid creation of models of urban infrastructure (such as bridges or tunnels) that can be used to efficiently estimate failure life-times for these important elements. The group also has built a substantial capability to simulate the effect of earthquakes (at high resolution), including the potential damage to urban infrastructure. This latter code base is highly scalable (demonstrated scalability to over 1M cores) and appears well poised to exploit the Fugaku supercomputer. Dr. Oishi also presented their recent work modeling the effect of flood including the effect of sediment flow. This work exploited the convergent nature of sediment flow as a flood moves downstream to produce an estimate of damage and hazard risk, even with an imprecise knowledge of initiation points. Although impressive in scale, it is nevertheless limited in applicability by both model resolution and assumptions made (such as distribution of initiation times, additional rainfall, diverse sediment types). It would be beneficial to understand the limitations of the existing model, and the group is encouraged to explore the impact of these assumptions by exploiting the scale available with Fugaku.

By developing large-scale analysis methods of urban areas under earthquakes, and ensemble debris flow and flood simulation cooperated with data assimilation team and computational climate research team, this team has been contributing to the society through disaster mitigation and reduction. They have made a laudable effort to disseminate their valuable results and benefit society through disaster mitigation and reduction simulations, including presentations on television and discussions with government officials. Future plans to develop digital twins that capture heat stroke hazard and the extent of potential storm surge (in support of Expo 2025) are exemplary in supporting the mission of R-CCS to provide societal benefit through the use of high performance computing.

The AC notes that the team has been reduced to a single member, which is well below critical mass for this important function. Although there has been a position open for 2 years, there is at this time only a single candidate being evaluated. This is a potential failure mode and while fully recognizing the difficulty of attracting and hiring qualified staff, the AC will emphasize the importance of succeeding. Even with demonstrated substantial collaborations (both internally to R-CCS and externally), the effort is not sustainable with this level of staffing.

3.10 Dr. Tama, Computational Structural Biology Research Team

This team aims to elucidate the structural changes of biomolecular complexes such as proteins and RNAs, which are deeply involved in the expression of biological functions, by combining experimental data with computer simulations. X-ray crystallography has been mainly used for protein structure analysis, but recent real-space and single-molecule measurement techniques are making it possible to investigate the structure of biomolecules under various physiological conditions and to observe the elementary processes of biochemical functions without crystallization. Although these real-space measurement methods do not provide atomic resolution, they may be combined with computational simulations to enable more detailed structural modeling of biomolecules. This team is developing integrative/hybrid modeling methods for precise structural modeling of large biomolecules by combining molecular dynamics simulations with real-space measurement data from cryo-EM and high-speed atomic force microscopy (HS-AFM). These methods reveal not only stable conformational renaming but also dynamic conformational changes of biomolecules. They are also working on structural analysis of nanoparticles using the intense X-ray free electron laser (XFEL).

Biomaterials with complex structures have long relaxation times, making the search for stable structures by computer simulation a difficult task. Moreover, the time scale of dynamical conformational change tends to increase with the size of the molecule, limiting what can be revealed by simulation alone. This team is trying to solve this problem by using experimental data, and is exploring a field that is not simply data science or computational science. The AC appreciates the research approach of this team and looks forward to further development.

This team uses the molecular dynamics code GENESIS developed by Dr. Sugita's team, and in this sense, methodological collaboration is taking place within R-CCS. On the other hand, collaboration and interaction on scientific research content is not evident, and AC recommends that this team, which is small in size, interact with other groups within R-CCS on a daily basis to promote effective collaboration.

3.11 Dr. Sato M., Programming Environment Research Team

Dr. Sato presented work by the Programming Environments Research group, which took advantage of the unique Fugaku system in developing advanced programming models. Their main focus was on XcalableMP (XMP), a directives-based Partitioned Global Address Space model. The group developed and maintained the Omni compiler for XMP, and in recent years significantly redesigned the XMP runtime system to take advantage of features of Fugaku's Tofu-D network. Dr. Sato also authored a book describing the parallelism, communication, and synchronization features of XcalableMP, as well as interoperability with MPI.

Some members of the group, including its leader Dr. Sato, formed the FLAGSHIP2020 Architecture Development Team working on co-design questions in Fugaku. They wrote several high profile papers on this work including an SC20 and IEEE Micro paper. This included work on the A64FX processor architecture and microarchitecture and development of a gem-5 simulator used during co-design. The simulator gave timing as well as memory system information, such as cache misses and access patterns using the high bandwidth memory (HBM).

The team also made significant contributions to Fugaku applications and benchmarking, including: 1) a Particle-in-Cell simulation used for accelerator modeling which as a finalist in the SC22 Gordan Bell competition; 2) the Graph500 benchmark results, where Fugaku continues to be the top machine for one of the kernels (BFS); and 3) the Graph Golf competition to design a minimum diameter graph for a given size and limited degree.

Most of the team members left at the end of FLAGSHIP2020, and the group is being reconfigured to address the end of floating point centric computing. This new group will instead perform research on programming models for hybrid quantum computing with HPC as part of the new QC-HPC Division. They have some initial experience with quantum algorithms based on their work simulating a quantum circuit using Fugaku's A64FX processor. During the review

Dr. Sato presented a vision for a hybrid software stack that leveraged some existing tools and proposed additional ones. The model is somewhat like other accelerators, but at the system scale rather than accelerating an individual node. The AC expects this usage model may change over time as QC evolves, and encourages the team to consider the QC to be more akin to an experimental device integrated with a high speed network to HPC, rather than something accelerating the HPC system.

The AC recognizes the contributions of the Programming Environment team and especially the critical role of the FLAGSHIP2020 Architecture Development effort in making Fugaku a success. While the QC-HPC topic is an interesting one from a research perspective, the AC also notes that broad-based expertise in traditional architectures and software will be needed if any significant level of co-design is expected in a follow-on to Fugaku. In particular, there may be more opportunities for architectural specialization or adaptation, and it will be important for RIKEN to have expertise in the software stack to maintain its position as a world-leading innovator in the design and deployment of HPC systems. This expertise resides in other new and existing groups at CCS.

3.12 Dr. Sano, Processor Research Team

This team, led by Dr. Sano, is investigating HPC processor architectures in the post-Moore era, with an emphasis on reconfigurable architectures implementing data flow and application-specific hardware using FPGAs. The team has been very productive, with 43 refereed publications, 45 invited talks, 13 research grants, and 2 patents applied for. Four projects were presented that build on the team's expertise and experience with FPGAs.

The first project is Elastic and Scalable FPGA-Cluster System for High-Performance Reconfigurable Computing (ESSPER). This was part of the Flagship 2020 project and deployed as an FPGA cluster, connected to sixteen nodes of Fugaku with 100Gbps network links. A full software stack has been developed and explored. With this testbed, a virtual circuit-switching network has been developed and benchmarked. Work is also ongoing on a range of applications, such as stencil, 3D FFT, and breadth-first search of a graph. Many of these investigations involve partnerships with many universities in Japan as well as the Barcelona Supercomputing Center (BSC).

The second project looks at a coarse-grained reconfigurable array (CGRA) for HPC. These architectures have been used in some specialized contexts but their use in general HPC is an open research question. To date, a number of design issues have been identified and the results presented at conferences.

Two examples were presented of using FPGAs for specific applications. These were near-sensor processing for the Riken SPring-8 center and a backend system for quantum error correction. The AC finds both of these to be interesting applications and ones that directly connect the Processor Research Team to full applications.

Dr. Sano presented a response to recommendations from 2019 AC. The AC is pleased with the progress of the dataflow based computing models, as described in the ESSPER and CGRA projects, and with the broad collaborations with researchers at other institutions. Progress has also been made on broadening the applications being considered and in closer collaboration with other next generation architecture teams.

The AC notes that adoption of FPGA techniques in HPC has been slow because of challenges both with software and performance. While benchmarks provide a valuable test, the complexities of applications often introduce features that reduce the benefits of these techniques. The AC recommends that more effort be made to ensure that complete applications will be able to make effective use of the architectures being considered.

3.13 Dr. Imamura, Large-scale Parallel Numerical Computing Technology Research Team

Toshiyuki Imamura presented the work accomplished by the Large-Scale Parallel Numerical Computing Technology Research Team. With a great focus on advancing numerical computing capabilities, this talented group is diligently working towards developing high-performance numerical libraries tailored specifically for the Fugaku computer.

Their research endeavors span a wide array of critical areas within numerical computation, including but not limited to linear systems, symmetric eigenvalue problems, fast Fourier transforms (FFTs), reproducible Basic Linear Algebra Subprograms (BLAS), and iterative refinement methods. By delving deep into these domains, the team aims to unlock new computational efficiency and accuracy levels, essential for tackling complex scientific and engineering challenges.

The team is also actively exploring the potential of leveraging Field-Programmable Gate Arrays (FPGAs) in their computational methodologies. FPGAs offer the prospect of harnessing hardware acceleration to enhance computational performance, thereby pushing the boundaries of what is achievable in numerical computing.

The research team comprises four highly skilled core members, whose expertise is augmented by the invaluable contributions of two visiting scientists. Together, they form a dynamic and interdisciplinary group that thrives on collaboration. By fostering strong partnerships with computational scientists, computer scientists, and applied mathematicians, the team has created an environment conducive to exchanging ideas and developing innovative approaches.

The team's research philosophy revolves around overcoming various challenges of large-scale parallel numerical computing. They actively investigate and implement strategies to minimize communication overhead, maximize computation overlap, and adopt asynchronous and task-based algorithms. Furthermore, they are dedicated to exploring novel techniques for data compression, an area of increasing importance in the era of big data. By pushing the boundaries of these techniques, they strive to achieve optimal data storage and processing efficiency.

The team's pursuit of innovation extends beyond the realm of algorithmic improvements. They deeply explore emerging concepts and technologies, such as mixed precision computing, reproducibility, fault-tolerance, and the development of optimized implementations for cutting-edge hardware devices. By leveraging the unique capabilities of these new hardware architectures, they aim to unlock unprecedented levels of performance and efficiency.

Their numerical libraries consistently demonstrate performance on par with or surpass the leading options available today. The scientific community has recognized this success, as evidenced by the publication of their work in prestigious journals and their participation in esteemed conferences. These accomplishments validate the team's efforts and contribute to the overall advancement of numerical computing.

Furthermore, the team collaborates internationally, reinforcing its commitment to expanding knowledge and fostering global scientific progress. Noteworthy partnerships include their involvement with the Joint-Laboratory for Extreme Computing (JLESC), an esteemed collaborative initiative involving institutions such as NCSA/UIUC, ANL, INRIA, JSC, BSC, RIKEN, and UTK. Additionally, the team collaborates with the Department of Energy (DOE) and the Ministry of Education, Culture, Sports, Science and Technology (MEXT), solidifying their involvement in cutting-edge research projects of global significance.

While the AC notes that the work on dense matrix computations is important, an effort to develop a basic sparse matrix library for Fugaku should be supported to help users solve large simulations involving 3-D PDEs.

Overall, the work conducted by Toshiyuki Imamura and the Large-Scale Parallel Numerical Computing Technology Research Team is a testament to their unwavering commitment to pushing the boundaries of numerical computation. Through their efforts, innovative ideas, and collaborative spirit, they are spearheading advancements that have the potential to revolutionize scientific discovery, engineering solutions, and computational capabilities for years to come.

3.14 Dr. Kondo, Next Generation High Performance Architecture Research Team

Masaaki Kondo, as the Team Leader for the Next-Generation High-Performance Architecture Research Team, presented an extensive feasibility study on the development of next-generation supercomputing infrastructures. This ambitious project was initiated in August 2022 with the primary objective of proposing a comprehensive system image that encompasses architecture, system software, and applications, thereby paving the way for a new era of supercomputing capabilities. The approach taken by the research team involved a holistic analysis of the findings from various studies conducted by different groups, enabling them to gain a comprehensive understanding of the challenges and opportunities in this domain.

Final Version of July 13, 2023

The vision for the next generation of computational infrastructure extends beyond mere technological advancement. It seeks to serve as a powerful platform that contributes to the realization of Sustainable Development Goals (SDGs) and Society 5.0. Central to this vision is the concept of digital twins, advanced virtual models that bridge the gap between scientific research and practical applications. By incorporating digital twins into the scientific landscape, the team aims to revolutionize research by leveraging cutting-edge technologies and computational methods. Their objective is to create an all-encompassing computing infrastructure that enables the seamless execution of complex workflows while harnessing the potential of diverse computational techniques, simulation methodologies, and Big Data at an unprecedented scale.

To achieve this ambitious goal, the research team is engaged in a comprehensive investigation into various aspects of the infrastructure. This includes in-depth analyses of architecture, system software, and library technologies, all of which are carried out in close collaboration with application developers. This co-design approach ensures that the system design is optimized from the ground up, aligning computational capabilities with the unique requirements of diverse applications. At the core of their system design philosophy is the "FLOPS to Byte" concept, which underscores the importance of efficient data transfer and computation within the constraints of power consumption, while maintaining the necessary level of computing accuracy.

The team's research efforts also extend to exploring different system configurations and fundamental technologies that can enhance the overall performance of the next-generation computing infrastructure. By continuously refining and optimizing these components, they aim to create a versatile and powerful computing environment capable of tackling the most demanding computational challenges.

Furthermore, the team is actively involved in developing a roadmap for future system software in Japan. This roadmap takes into consideration the pressing need to enhance data utilization, seamlessly integrate AI technologies with first principles simulation, enable real-time data processing, and ensure the highest levels of security. By addressing these key areas, they aim to lay a solid foundation for the successful implementation and adoption of the next-generation supercomputing infrastructure.

To assess the effectiveness and viability of different architectural choices, the research team is constructing an extensive benchmark set. This set includes comprehensive evaluations of various architectures, allowing them to gain valuable insights into the strengths and weaknesses of each option. Additionally, they conduct exploratory “what if” performance analyses, enabling them to simulate and analyze the potential impact of different algorithms and application parameters. By considering these findings, the team can make informed decisions about the future evolution of algorithms, ensuring that they align with the requirements of the next-generation computing systems.

In summary, the research team's comprehensive approach and careful planning reflect their commitment to the development of next-generation supercomputing infrastructures. With a feasible timeline and a focus on optimizing performance, security, and application compatibility, they are poised to make significant contributions to the advancement of computational capabilities and the realization of transformative scientific research and societal goals.

3.15 Dr. Sato K., High Performance Big Data Research Team

The aim of the High Performance Big Data Research Team is to perform fundamental R&D of system software (with a focus on data), including HPC for AI/Big Data (AI BD workloads on HPC systems) and AI/Big Data for HPC (use AI/BD techniques to improve HPC workload and application performance). Dr. Sato described two research highlights: data compression for the SPring-8 synchrotron and deep learning for Fugaku. The team is also participating in the FugakuNEXT feasibility study.

The SPring-8 synchrotron is projected to generate 1.3EB/year with the next generation detector (CITIUS) starting 2025 (this is the raw output from the sensor). One case is X-ray CT data, which is time evolutionary data. The team devised a novel use of a neural network for data compression that takes into account the time axis, creating a predictor for the evolution of the data. The method was realized in the software TEZip, which provides both lossy and lossless compression. The performance is good: Lossless compression provides 9-15x and lossy (at a few percent error) provides 40-50x compression. The AC is pleased to note that TEZip is open-source and available on github.

DL4Fugaku (Deep Learning for Fugaku) addresses the lack of a tuned DNN numerical library exploiting the ARM SVE. This work extended the Intel OneDNN to generate ARM + SVE instructions instead of x86 instructions. This resulted in a 4-8x speedup in PyTorch and TensorFlow. Combined with 16K node parallelism, Fugaku received the #1 ranking for the MLPerf benchmark in 2020. The team also investigated I/O performance analysis of the Fugaku hierarchical storage system for deep learning applications. This work will inform future storage architectures.

The team has a good publication record, with 26 refereed papers, and is active in the HPC community, with significant participation in HPC conferences.

The response to the recommendations from the 2019 AC were good, with more focus on AI applications and on collaboration with other teams. The AC was pleased to see the participation of student interns and trainees and of visiting researchers in the team. The AC approves of the future plans as presented, particularly with the collaborations described in those plans. One concern is the number of topics, and the AC advises care that the team not undertake too many projects, even with collaborations. The AC encourages consideration of alternative, non-POSIX I/O architectures for future systems, particularly in the context of FugakuNEXT. These should take advantage of the approaches used in big data (but not HPC) systems.

3.16 Dr. Wahib, High Performance Artificial Intelligence Systems Research Team

Dr. Wahib's presentation highlighted the efforts and progress made by the High Performance AI Systems Research Team. The unit was created four years ago and originally led by Dr. Satoshi Matsuoka, with Dr. Wahib assuming the Team Lead role in April 2022. He presented the vision of the unit in three parts: (1) Scaling AI to O(100K) Nodes of Fugaku, (2) HPC in AI, and (3) AI in HPC. The former has been a goal since the inception of the group, while the latter two were started under Dr. Wahib's leadership.

The team has made acceptable progress on the first goal – developing an automated method for generating high performance convolution kernels for the idiosyncratic ARM processor family and demonstrating the ability to train large models using out-of-core memory are two notable results. They are focused on tools that will help the end user, including developing an oracle to guide large-scale training of DNNs and a method for partitioning computational graphs when memory-constrained, both of which will be useful to practitioners on Fugaku.

The two newly defined thrusts are an ideal pair, capturing both the benefit and challenge of AI on Fugaku. It is clear that the team is well poised to address both, as evidenced by their notable publication record touching the whole computing stack for AI. In general, the AC applauds the clear focus on Fugaku by this team; developing AI for (and using AI on) this world-class resource was front and center during the presentation.

A high level of multi-disciplinary collaboration by the team across R-CCS, across RIKEN, and externally, was evident – highly appropriate given the multi-disciplinary nature of AI and its application. This promises to continue as they engage with the US Department of Energy's AI for Science campaign, in collaboration with many leading computational centers in both countries.

The AC supports the team's ambitious future plans. Given the enormous importance of AI for society, it is vital that this team be part of the FugakuNeXT discussions, and we are pleased to see this happening. Anticipated projects such as ultra long sequence lengths for transformers, and the adaptive hiding of samples during DNN training, address some of the most serious challenges facing AI (temporal correlations and the continued scaling of sample count for large models). We look forward to hearing the results of these efforts in their publication record and at the next review. We note continuing collaboration with Fujitsu as they scale towards a GPT-Fugaku through distributed training – addressing another AI challenge. The exercise of developing GPT-Fugaku will be highly instructive, even if the eventual model is overtaken by, e.g., a larger GPT-5.

We are also cheered by future plans to map a mouse brain at high resolution, and the effort to use AI to automate creation of scientific HPC code. The former will move forward neuroscience while showcasing the value of HPC, while the latter would help democratize utilization of these valuable resources – increasing the scope of HPC utilization across the science endeavor in Japan.

Finally, the AC feels that the approach taken by Dr. Wahib for management of human resources is exemplary. Despite other teams at RIKEN reporting difficulty in recruiting or retaining staff, the High Performance AI team grew by 8 members in the last year, and their proactive approach to developing staff (with a focus on growing independent scientists through collaborations) can be taken as a template for success. It is noteworthy that this fully engaged team finds the time in addition to mentor almost a dozen interns and trainees, helping grow the next generation of AI/HPC experts.

3.17 Dr. Domke, Supercomputing Performance Research Team

Dr. Jens Domke presented the work of the Supercomputing Performance Research team, a relatively new group with research in performance modeling and prediction, automatic performance tuning (autotuning), performance tools, architecture evaluation and benchmarking, and hardware-software co-design. As an example of evaluating microarchitecture parameters, the team demonstrated speedups of up to 3x on a finite element benchmark problem from a larger high bandwidth L2 cache using 3D stacked SRAM.

The team has played a key role in evaluation of FugakuNext alternatives as part of both the Application and Architecture Groups. On the applications side, they worked on the design of a new Octopode benchmarking set that identifies a set of kernels and impact metrics from existing applications. In contrast with fixed proxy applications, this approach allows additional flexibility in co-design, and may yield more realistic performance measures. On the architecture side, they worked with the interconnect networking subgroup and the CPU/GPU evaluation subgroup. They are also looking to the future with papers on a benchmark suite for machine learning and another on optimizing the communication pattern in deep learning to better support data partitioning instead of replication.

In spite of the relatively junior level of the group they have strong international collaborations. For example, a paper describing a translator to map GPU code to CPU code has collaborators from MIT, Lawrence Livermore National Laboratory, Google (in France), as well as Tokyo Tech, and a paper on communication for deep learning is with other researchers in Japan along with others in France and China. Along with R/CCS Director Matsuoka, the Octopode position paper is coauthored with leading researchers at Argonne, Oak Ridge, and Berkeley National Laboratories, a paper on the myths and legends of HPC with a leader at ETH, and the ML benchmarking paper with a wide range of international researchers in ML and HPC. The team also has a number of other formal and informal international collaborations.
The group is small and beyond the group lead by another junior researcher, it is mostly populated by visitors and summer interns. While much of the work in the group is well-suited to standalone student projects and dynamically built within RIKEN, across Japan, and internationally, the AC notes that sustained software projects, even benchmarking suites or

simulation tools, will require a more permanent workforce. Given the importance of this activity to the success of FugakuNext and the future of R/CCS, the AC recommends that the group continues to leverage students and visitors, but also adds some additional permanent staff.

3.18 Dr. Okuno, Biomedical Computational Intelligence Unit

The development of successful therapeutics has become more and more difficult and costly, with the current cost estimated to be in excess of $2B per approved drug. The Biomedical Computational Intelligence Unit is one of four units in the new HPC-AI-driven Drug Development Platform Division which has a stated aim of developing a new platform (combining HPC, AI and experiment) to accelerate (and substantially lower the cost of) the development of novel therapeutics. The BCI Unit has two unit-specific aims:

a. Exploring new approaches that integrate simulation, AI, and real-world data, such as wet experiments and clinical settings, in biomedical research
b. Developing new computational methods to understand disease mechanisms and to identify drug-targeting proteins/genes.

The presentation by Unit (and Division) Leader Dr. Okuno described some of the recent successes from this newly-created Division. Understandably (given the short existence of this Division) much of the progress involved involves work that overlaps significantly with Prof. Okuno's well-established group at Univ.Tokyo. They describe progress on many fronts, including the use of Bayesian networks to extract biomolecular networks (in the process identifying disease-related genes in SARS-CoV-2 and gastric adenocarcinoma) and the analysis of cryo-EM images. This novel approach exploits molecular dynamic simulations and AI to analyze cryo-EM data, extracting an experimental measure of residue dynamics from large proteins that is difficult to acquire otherwise.

The AC applauds the efforts begun to port various capabilities to Fugaku, such as the conformational search pipeline and PyAutoFP. The AC expects that as time goes on, the unit (and Division) will take even more advantage of Fugaku, highlighting their utilization of this world-class resource. We are excited by the future plans that were presented by Dr. Okuno – including their strong involvement with FugakuNeXT. Expanding the use of computational resources at scale to solve societal issues (including health) is a fundamental pillar of the RIKEN philosophy, and the AC feels that this focus is well placed.

As time goes on, the AC would expect the center of activity to move from Tokyo to R-CCS, where there are ample opportunities for collaboration with world-class experts in computational structural biology and cryo-EM image analysis. In addition, we note that this unit in particular is described as operating at the end of the drug discovery pipeline. We believe that the unit would benefit from a greater focus on later-stage development of therapeutics, including more contact with wet-lab biology, pharmacokinetics, toxicity, manufacturability, and clinical trials.

We also note that there was scant mention of gender or national diversity – this is a center-wide priority for RIKEN, and the AC would like to hear more about how the new Division and the individual units are addressing these issues.

3.19 Dr. Okuno, AI-driven Drug Discovery Collaborative Unit

The AI-driven Drug Discovery Collaborative Unit plays a key role in the HPC- and AI-driven Drug Development Platform Division, aiming 1) to develop an HPC- and AI-driven platform for drug discovery, and 2) to apply the platform to practical drug discovery processes. The presentation by the Unit Leader Yasushi Okuno (also the Division Leader) demonstrated a clear overall vision for developing such a platform, and exciting new results in areas such as the development of kMoL, a machine learning library for molecular systems, and the application of structure-based generative AI models. As for the second aim, a number of collaborative research projects are on-going, and specific applications were presented such as the screening of anti-SARS-CoV2 molecules on Fugaku and the prediction of the impact of mutations in cancer genomes. The latter research can contribute to precision medicine and is a new example of how HPC/AI addresses societal challenges.

While those results are highly impressive, the AC recognizes that building an AI-driven drug discovery pipeline is still an ambitious goal. Good collaborations exist with other divisions of RIKEN and external organizations (including pharmaceutical companies) but more will be needed with clear division of labor and better coordination. The AC recommends that a more regular communication be established with other related elements across R-CCS (and Riken) to maximize opportunities for collaboration and contribution.

Related to this point, the AC notes that some of the component predictive models will require more experimental data. It is, therefore, recommended that communication be with both dry and wet-lab communities to discuss how to collect additional data in a systematic manner. Another issue is that the presentation was not always clear about the effective use of HPC resources. For example, it is unclear whether kMoL and other deep learning tools developed can run effectively on Fugaku. Finally, there was no mention of gender or nationality diversity, an issue common to all the units in the AI-driven Drug Discovery Collaborative Unit. The AC recommends more efforts be made in all those areas.

3.20 Dr. Honma, Medicinal Chemistry Applied AI Unit

The Medicinal Chemistry Applied AI Unit, led by Teruki Honma, aims to 1) develop AI systems focusing on medicinal chemistry for drug discovery, and 2) utilize those AI systems to aid in the design and optimization of drug candidates. Within the HPC/AI-driven drug development platform division, this Unit is uniquely positioned in contributing to the generation of hit compounds and optimizing them by considering numerous factors including drug-likeness and synthetic accessibility.

Notable achievements include the development of AI models for generating novel and diverse chemical structures, and the application of multi-objective optimization techniques to lead optimization. It is also impressive that this Unit played a key role in establishing a consortium of a large number of pharmaceutical companies and academic institutions in Japan.

While the AC recognizes that this consortium makes important contributions in building a large number of predictive AI models, the presentation was somewhat unclear about how those attempts were related to other activities in the Division, in particular, the development of an AI-driven drug discovery platform explained by the Division Leader. A discussion with the Unit Leader after his presentation made it clear that the Drug Discovery AI platform mentioned here should be a part of the larger AI platform pursued as the ultimate objective of the whole Division. The Medicinal Chemistry Applied AI Unit is, indeed, developing critical pieces of the AI-driven drug discovery platform, i.e., molecular structure generation and optimization, but this point should be better clarified.

The Unit also developed a virtual screening system that can deal with “ultra-large scale” compound libraries. Screening billions of synthesizable molecules has attracted much attention recently in drug discovery research, but dealing with such ultra-large databases will pose significant challenges in computing. Utilizing the power of Fugaku is a natural option; the Unit is well placed as a new addition to R-CCS, and its contributions in this research area should be applauded.

Another application area that the Medicinal Chemistry Applied AI Unit pursues is a technique known as the fragment molecular orbital method (FMO). This method allows for the application of quantum mechanical calculations to large molecular systems such as proteins. This approach is potentially helpful in producing more accurate interaction energies, but the AC recommends better integration of the FMO-related work with other activities in the Unit.

The Unit’s second aim is to apply the new methods being developed to specific drug discovery projects. Here, it is crucial to collaborate with other Divisions of Riken or external partners. The AC encourages further collaborations and expects to see more results, but the exact roles played by this Unit and its collaborators should be explained properly.

3.21 Dr. Ikeguchi, Molecular Design Computational Intelligence Unit

The Molecular Design Computational Design Unit aims at 1) developing computational methods for designing drugs, including small molecules, middle-sized molecules, and macromolecules, and 2) integrating molecular simulation and AI techniques to enhance drug design processes. While conventional drug platforms such as small-molecule compounds are well established, new formats (known as modalities) including middle-sized molecules and biologics have emerged recently. The presentation by the Unit Leader Mitsunori Ikeguchi showed considerable potential for accelerating the development of new modality drugs by combining AI and HPC techniques.

Specific achievements included a method for predicting the membrane permeability of cyclic peptides, antibody design protocols, and tools for designing antisense oligonucleotides. Unlike small-molecule drugs, conformational variability needs to be considered when dealing with those larger molecules. In the case of cyclic peptides, efficient conformational sampling was achieved by utilizing GENESIS and novel sampling techniques developed by the Computational Biophysics Research Team. This is a nice piece of collaboration within the R-CCS, and the AC encourages further attempts along those lines.

The Unit exhibited an impressive list of publications to report a wide range of applications for protein dynamics and molecular interactions. Overall, the Molecular Design Computational Design Unit plays a significant role in the HPC- and AI-driven Drug Development Platform Division and contributes to the R-CCS. On the other hand, the AC notes the following issues: 1) no clear descriptions of the availability of the software and tools developed, 2) the utilization of Fugaku not obvious in some applications, and 3) further clarification needed on how the Unit can contribute to the development of an AI-driven Drug Development Platform, the ultimate objective of the Division. The AC suggests the Unit focus more on the integration of HPC and AI technique, and expects to see explicit results in this area, as it is one of the stated objectives of the Unit.

3.22 Dr. Miura, Facility Operations and Development Unit

The team's most crucial challenge was the installation of the Fugaku system. However, they completed the building as planned, and Riken operated the Fugaku system successfully. They carefully considered power constraints and the air conditioning system's capacity, achieving the same level of Power Usage Effectiveness (PUE) as in the past. This achievement helped minimize the overall increase in operating costs. The successful installation demonstrates the team's expertise and ability to handle complex technological tasks. They can leverage this experience to tackle future challenges effectively and further enhance their reputation.

Despite successfully installing and operating the Fugaku system, the team encountered a significant challenge in staff recruitment. As the previous Advisory Committee highlighted, the team has yet to accomplish the necessary recruitment to augment its workforce. Increasing the number of personnel engaged in the team and supplementing the existing staff is crucial to establish a more sustainable and robust workforce.

To address this challenge, the Advisory Committee suggests exploring various recruitment methods instead of solely relying on the conventional open recruitment process. This recommendation reflects the need for innovative approaches to attract qualified candidates and expand the talent pool. While it is understandable that the team is currently facing profound difficulties regarding recruitment, prompt action is necessary due to the urgency of the situation. Efforts should be made to explore alternative recruitment strategies, such as targeted recruitment campaigns, collaborations with academic institutions or research organizations, and leveraging professional networks to identify potential candidates. Additionally, the team could consider offering attractive incentives, benefits, or professional development opportunities to attract and retain top talent.

The AC (Advisory Committee) emphasizes the importance of taking immediate action in response to recruitment challenges. By prioritizing this issue, the team can ensure its operations' continuity and long-term sustainability. Addressing the recruitment needs will alleviate the current strain on the team and provide the necessary resources to tackle future projects and initiatives effectively.

3.23 Dr. Uno, System Operations and Development Unit

Among these efforts, the Advisory Committee (AC) highlights the remarkable contributions of this team, particularly in successfully achieving a smooth transition from the K computer to Fugaku amidst the challenges posed by the onset of the Covid-19 pandemic and semiconductor shortages. It is worth noting that they effectively addressed software and job scheduling issues during this transition. Additionally, the AC recognizes the exceptional contribution of this team in collaborating with users to meet their demands and assisting in the acquisition of the Gordon Bell Prize.

Furthermore, during the energy crisis in 2022, this team skillfully utilized Fugaku's energy-saving feature, known as "PowerNob," to ensure stable operation without depleting the energy budget. They have also successfully responded to the previous Advisory Committee's comment regarding setting Key Performance Indicators (KPIs).

In response to the real-time processing demands of the Data Assimilation Research Team's "Big Data Assimilation," this team has ingeniously devised job scheduling techniques to minimize the impact on other users while effectively addressing these requirements. This has further earned substantial trust from the Japan Meteorological Agency (JMA) regarding the Fugaku architecture. The AC acknowledges that the unit, which has primarily acted as a service unit, has carried out an important R&D role but still was able to publish four referred papers. The collection of various log data, which serves as the foundation for automated operations using Machine Learning, makes a significant contribution to ensuring the continuous operation of this field, which is constantly faced with a shortage of skilled personnel. Therefore, we anticipate further advancements in this area.

3.24 Dr. Murai, Software Development Technology Unit

Hitoshi Murai presented the work accomplished by the Software Development Technology (SDT) Unit, which has been organized since July 2022. The team comprises nine people with a wide range of research backgrounds. The main aim of SDT is to maintain and enhance the software environment of Fugaku for its more efficient and effective operation. SDT also works for the feasibility study of the next supercomputer.

The Fujitsu compiler is a source of poor optimization, and the SDT has made an effort to develop a Clang/LLVM compiler for A64FX.

They are working to provide Fugaku users with the following tools for software development to complement the proprietary ones.

- Score-P/Scalasca, an open-source performance profiler developed by Julich Supercomputing Center. They are trying to port Score-P/Scalasca to Fugaku under a joint research agreement.
- TAU, an open-source performance profiler developed by the University of Oregon. TAU has already been ported to Fugaku by the developer under a joint research agreement, and we are testing it before releasing it to the users.
- VeloC, an open-source checkpointing tool developed by Argonne National Laboratory. VeloC is already ported and available on Fugaku.

About 100 packages are available on the Fugaku compute nodes as of April 2023.

In addition, they have been working to create/modify Spack recipes for A64FX, many of which have been merged into the Spack mainstream and are publicly available.

In addition, they are providing user support for AI training on the Fugaku.

3.25 Dr. Shoji, HPC Usability Development Unit

Dr. Shoji presented the work of the Usability Development Unit, which has a number of activities related to data and workflow services, pre and post processing on Fugaku, and connections with commercial cloud providers. The work is focused on the operational side of running a supercomputing center and less on basic research, but the work is critical to the success of science projects using Fugaku. The team has 5 permanent technical staff, which is more than some of the research groups, but necessary to stability of operations.

To support data intensive computing and preserve the results of scientific simulations, the team supports HPCI (High Performance Computing Infrastructure), with 50 Petabytes of storage and servers that form the western hub of this distributed national resources. It supports 78 different science projects and has an impressive uptime of close to 3.5 years. The total storage is expected to double in 2023 at each of the two sites, with the University of Tokyo providing the eastern hub. Beyond traditional simulation problems, the team is supporting a tight integration between large scientific experiments and the HPC facilities such as the national radiation facilities (SPring-8/SACLA) which send data into HPCI-west where it can be analyzed on Fugaku. The unit also supports a key international collaboration with a shared data service agreement with the Singapore National Laboratory (A*Star). They also support both in-house

(HIVE) and externally developed (ParaView, POVRay, GMT, VisIT, OpenViSUS, Vampir, AVS/Express) tools for scientific visualization.

Many scientific applications require complex workflows that tie together data analysis steps with a set of scale simulations, which can be hard to manage on a batch-queued HPC system like Fugaku. The team developed WHEEL (Workflow in Hierarchical distributEd parallEL), an HPC workflow management system that is an open source tool providing a broad range of services and--along with the visualization tools--an easy “click through” setup process. They have also optimized pre and post processing steps to improve the user experience.

Another important and forward-looking partnership for R/CCS is with cloud providers, which allows users to move computations between Fugaku and the Cloud. The Usability Development Unit worked to establish a free and security data connection to Oracle Cloud.

The AC acknowledges the importance of the work performed by this unit, which is not measured in publication count but by a satisfied user community, and advises R/CCS continue to retain and replace-as-needed the staff in the group. The AC also commends the group for its high quality services and innovations in cloud access, data, and workflow tools, and recommends that additional integration of experimental facilities with Fugaku will help to cement even further the national importance of Fugaku to the scientific research enterprise.

3.26 Dr. Yamamoto, Advanced Operation Technologies Unit

(Lead: Sekiguchi, Help: Shimojo, Yelick)
The Advanced Operations Technology Unit, established in FY2020, is crucial in conducting research and development for advanced operations within the data center. This effort encompasses the Fugaku supercomputer and various infrastructure components, such as power supplies, cooling facilities, and networks. The unit closely collaborates with the System Operations Development Unit to ensure seamless integration of their efforts.

One of the critical areas of focus for the Advanced Operations Technology Unit is the development of operational data analysis infrastructure. By analyzing data from this infrastructure, the team gains valuable insights that inform the optimization of high-performance computing (HPC) and facility operations, particularly emphasizing power considerations. The team has made notable progress in improving energy-saving processes by leveraging operational data analysis. They have successfully implemented an advanced power-saving mechanism, the Power Knob, which the team and stakeholders highly appreciate. Furthermore, the recent expansion of network bandwidth from 100 Gbps to 400 Gbps has been a welcome development, enhancing overall system performance and efficiency.

In addition to infrastructure and operations, the unit has also focused on enhancing user experience and usability. They have developed Tomitake's Software-as-a-Service (SaaS) environment, simplifying the utilization of HPC resources for users. This software development aspect contributes to technological sophistication and combines operational technology to establish an efficient operating environment that provides users with stable and reliable services. To ensure effective management and performance evaluation, the Advisory Committee (AC) recommended that the unit leader defines key performance indicators (KPIs) that align with the unit's objectives and goals.

In line with advancing automation and streamlining operations, the AC recommends further research and development of a data analysis infrastructure incorporating artificial intelligence (AI). The utilization of AI in automating various operational aspects holds great potential for enhancing efficiency and reducing human effort. However, developing AI capabilities necessitates substantial investment in human resource development. The AC acknowledges the challenges and supports the unit's ongoing efforts in this critical area.

By continuously refining its operational processes, leveraging advanced technologies, and investing in human resource development, the Advanced Operations Technology Unit is expected to drive innovation, improve operational efficiency, and provide a stable and user-friendly environment for the data center's diverse stakeholders.

4 Conclusions

R-CCS is a world-leading research and operations organization. Fugaku is an outstanding success. Through both Fugaku and the research teams and units, R-CCS is delivering on the science and is positioned to continue into the future. The leadership of R-CCS is to be commended on their strategic vision and the implementation of that vision. R-CCS is effectively combining computer science, computational science, and an aggressive yet well-grounded approach to the design of Fugaku next, taking into account the changes in HPC such as AI and new application areas for the benefit of society.

5 Recommendations

1. Stay the course with FugakuNEXT plans and fully fund development.
2. The co-design strategy employed in the development of Fugaku has proven to be highly effective and should be replicated for FugakuNEXT.
3. The co-design strategy for FugakuNEXT would benefit from better coordination with RIST and HPCI-C.
4. Encourage stronger collaborations across groups (which will lead to FugakuNEXT inputs). For example, division leads working with team PI's could develop a shared research vision. This would help identify opportunities to share ideas and resources and especially to identify gaps needed to achieve the vision.
5. Consider best practices for hiring diverse talent. These include salaries that are competitive for disciplines; spousal hirings incentives; promoting a family-friendly (https://myfamilyberkeleylab.lbl.gov) work environment.

Acknowledgments

The AC members would like to thank the R-CCS scientists and administration, in particular Nobuo Sugiura and his team, for a very effective organization of the review and the friendly, constructive and very pleasant atmosphere during the review.

Glossary

AC	Advisory Council
BLAS	Basic Linear Algebra Subprograms
Cryo EM	Cryo-electron microscopy
DA	Data Assimilation
DGEMM	Double precision GEneral Matrix Multiply
FS2020	Flagship 2020 project for post-K computer
Fugaku	Flagship 2020 computer (post-K computer)
Graph500	Benchmark based on data-intensive loads
HPCG	Benchmark based on sparse matrix computation
HPCI	High Performance Computing Infrastructure
JLESC	Joint Laboratory on Extreme Scale Computing
PUE	Power Usage Effectiveness
QCD	Quantum ChromoDynamics
RAC	RIKEN Advisory Council
R-CCS	RIKEN Center for Computational Science
RIST	Research Organization for Information, Science and Technology
SCALA	RIKEN X-ray Free Electron Laser (XFEL), "SPring-8 Angstrom Compact Free Electron Laser"
SDGs	Sustainable Development Goals
SPring-8	RIKEN large synchrotron radiation facility,

	"Super Photon ring-8 GeV"
SWOT	Strength, Weakness, Opportunity, and Threats
TCO	Total Cost of Ownership
ToR	Terms of Reference