Report of the 2019 RIKEN R-CCS Advisory Council

Meeting of the Advisory Council of the RIKEN Center for Computational Science (R-CCS), July 29 – July 30, 2019, Kobe, Japan

Tasks of the R-CCS AC

The R-CCS AC was tasked with reviewing the activities of R-CCS from 2016 to 2019 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 2019 R-CCS AC.

Members of the R-CCS AC

Jack Dongarra	Prof. Dr.	University of Tennessee, USA	Applied Mathematics
William Gropp	Prof. Dr.	Director, NCSA, USA	Computer Science
Catherine Lambert	Prof. Dr.	Director, CERFACS, France	Fluid Dynamics
Thomas Lippert	Prof. Dr. Dr.	Director, JSC, Germany (Chair)	Particle Physics
Haruki Nakamura	Prof. em. Dr.	Osaka University, Japan	Life Sciences
Hiroshi Niino	Prof. em. Dr.	Atmosphere and Ocean Research Institute, University of Tokyo, Japan	Atmospheric Science
Thomas Schulthess	Prof. Dr.	Director, Swiss National Supercomputing Centre, Switzerland	Materials Science
Satoshi Sekiguchi	Dr.	Vice President, National Institue of Advanced Ind. Science & Tech., Japan	Computer Science
Horst Simon	Prof. Dr.	Deputy Director, LBNL, USA	Applied Mathematics
Shinji Shimojo	Prof. Dr.	Director, CMC, Osaka University, Japan	Computer Science
Shinji Tsuneyuki	Prof. Dr.	University of Tokyo, Japan (Co-Chair)	Computer Science
Katherine Yelick	Prof. Dr.	Associate Director, LBNL, USA	Computer Science

Information provided to the R-CCS AC

- Meeting Program for the R-CCS AC 2019
- White Paper for the R-CCS AC 2019 (Overview, Activities, Responses, Future, SWOT, Appendix)
- Activity Report for the R-CCS AC 2019
- RIKEN 7-Year Plan, April 1, 2018 to March 31
- Presentation Prof. Dr. Shigeo Koyasu, Executive Director RIKEN
- Presentation Prof. Dr. Satoshi Matsuoka, Director R-CCS
- Presentation Shigeo Okaya, Deputy Executive Director RIKEN and Deputy Director R-CCS
- Presentation Dr. Fumiyoshi Shoji, Div. Director R-CCS “Operations and Computer Technologies”
- Presentation Dr. Yutaka Ishikawa, Div. Director R-CCS “Flagship 2020 Project”
- Individual Presentations by 19 Research Teams and 4 Operation and Development Units

Information and Recommendations from Previous Reviews

The RIKEN AICS was founded in July 2010. In 2016 a review of AICS by the AC of AICS under the RAC framework for the period 2010 to 2016 took place, resulting in seven recommendations by the AC. The AICS has been renamed to R-CCS in 2017. Hence, the present review is the second one.

1 Introduction

The reporting period 2016 to 2019 of the meeting of the Advisory Council (AC) of the RIKEN Center for Computational Science (R-CCS) is characterized by the continuous success of the K-Computer, which, despite its long service life, has always occupied leading positions in benchmark rankings, especially on the HPCG and Graph500 lists. In addition, it is the exciting period of the final stages of development of the post-K system, recently named Fugaku.

The evaluation of the R-CCS thus refers to the period just prior to the handover of the baton from the K computer to the Fugaku system and offers, so to speak, the last opportunity to make suggestions on the strategy and direction of the R-CCS research, operation and support activities with regard to the forthcoming operation of Fugaku, its continuous development and its potential "post-Moore" successor in more than half a decade.

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 as well as of the FS2020 project carried out to help co-design Fugaku. Section 3 deals with an in-depth evaluation of the 23 research and operation teams and units. Section 4 provides global conclusions, and section 5 gives the AC's recommendations.

2 General Findings

In the following, the R-CCS' general computing strategy, its operation and infrastructure development strategy, its funding risks, and the FS2020 development project are assessed globally.

2.1 R-CCS Strategy

The R-CCS strategy as presented by its director aims at the tight integration of computer science — the science of computing — with computational science — the science by computing —, thus supporting the advanced computing needs of the digital society through the exploitation of science for computing at the leading edge. The AC welcomes the focus of this strategy towards the convergence of HPC with big data analytics (BDA) and AI, such as the unification of compute-centric processing with data assimilation, the cross-fertilization of first-principle simulations by AI, the acceleration of AI training by HPC and the openness of the center as backend for edge computing. The AC is endorsing this integrative strategic approach as a solid basis for extension of cooperation with institutions, universities and companies in Japan and internationally. The AC has seen many examples of this strategy beginning to work already. This is of particular importance in view of the enormous progress achieved towards installing Fugaku. The AC considers the new strategy together with the upcoming launch of Fugaku to be a very positive development and congratulates RIKEN and its R-CCS on a new R-CCS director with strategic vision, who has as much energy as he has the ability to deal with new situations and integration efforts.

2.2 Operation, Management, Development of Computing Infrastructure

The AC is impressed about the operation and the support of the K-computer at R-CCS over its lifetime. It has produced a large set of scientific results as well as benchmark results including Gordon Bell finalists for demonstrating application results at scale. It has been in the top 10 of the Top500 list, that ranks top supercomputers world-wide, nearly over its lifetime, the HPCG No.1 from November 2016 to November 2017, and at the top of the Graph500 list which similarly ranks international machines but from the standpoint of the more challenging graph analysis problems, being No. 1 from June 2015 to June 2019.

This speaks to the quality and balance of the system, but also to the high quality operation of the facility. Demonstration of performance results at scale can be challenging on any system, but is nearly impossible if the machine has frequent node failures or outages.

Large-scale computing installations like the K computer require several megawatts of power, and a significant fraction of that can be spent on the power distribution and cooling infrastructure, rather than the computer itself. The AC is impressed that the R-CCS operations team was able to significantly reduce this power overhead during the operational lifetime of the system.

The AC acknowledges that management has responded to comments from the previous review regarding user support, and is cooperating with RIST more directly on user issues, has set up a help desk and ticketing system, and conducts workshops to help users adapt their codes to the systems and plan for future installations. The AC considers this as a first positive step to help resolve the problem of separation between users and their support. The AC is convinced, however, that further measures need to be taken to bring users and the center closer together.

2.3 The FS2020 Project, Fugaku and Beyond

R-CCS has an ambitious plan to upgrade to the Fugaku system and is developing research and development plans for a system next to Fugaku. The AC acknowledges that the Fugaku implementation plans are sound and seem feasible. AC is concerned that the user program is delayed and asks R-CCS to intensify user training for Fugaku.

The AC notes that the ratio of network performance to memory performance decreases by a factor of about 7.8 when the K-Computer moves to Fugaku. As a result of co-design, it can be expected that data locality or data-centric programming helps in this respect, i.e., applications with neighbour communications such as stencil will have sufficient network bandwidth, while, as an example, all-to-all collectives are not supported well. The AC considers it important that R-CCS further intensifies its research in this respect and helps users to overcome possible limitations.

The replacement of the K-computer by the Fugaku system leads to a temporary lack of supercomputing resources at the highest level in Japan, i.e. a resource risk, as the flagship supercomputer needs to be stopped for several months to a year in Japan. For designing a new system beyond Fugaku, such a risk should be avoided, and a plan to replace Fugaku with a post-Fugaku in a more smooth manner should be considered when Fugaku is introduced.

2.4 Funding Gap

The AC acknowledges that the investments for 400 cabinets of Fugaku with a total of more than 150,000 nodes are fully funded and that the machine is under construction according to plan. It is planned to be delivered in 2020, and the operational phase is supposed to start in spring 2021. Given that the machine's power consumption will be two times that of the K Computer, the AC was told that there is currently the necessity to allocate about 4 billion Yen p.a. increase in R-CCS' operational budget, or 24 billion Yen over six years in total cost of ownership (TCO). The AC considers this condition to be a major risk for the successful operation of Fugaku. The AC has discussed four options to cover the funding gap:

1. Reduce electricity costs to meet its current budget by reducing the availability of the system. In fact, this will require a partial shutdown of the system;
2. Receive more funding under the nation's Act on the Promotion of Public Utilization of the Specific Advanced Large Research Facilities (Law for promoting the sharing of specified advanced large-scale research facilities);
3. Fund the gap by charging all users, where a credit model could be applied to subsidize strategic projects or the most excellent science;

4. Sell part of the resources to commercial users reducing availability to users who are eligible to access Fugaku through RIST.

Option 1 is considered only a measure of last resort by the AC, when all others fail. Option 2 appears very interesting but certainly is the most risky, and option 3 will likely alienate traditional academic users. Therefore, the AC favours option 4 (favorably in connection with option 2). The stakeholders need to be informed and a sound business model has to be developed. The AC sees this option as a feasible opportunity to expand the user base to a paying customership and to develop a long-term strategy that not only can help in solving the funding gap but might also significantly grow R-CCS' business and its involvement with industry.

3 Summary of Evaluation of Research Groups

In the following section, the AC describes a detailed in-depth evaluation of the individual research teams and operation and development units. The assessment follows the ToRs of the RIKEN President and the Director of the R-CCS. The following criteria have been synthesized from these documents:

1. Is the research unique?
2. Are Co-design activities carried out, contributing holistically to the future of HPC in the respective domain beyond Fugaku?
3. Does the group contribute to the convergence of HPC with Big Data and AI?
4. Does the group contribute to societal issues?
5. Is end of FLOPS-centric performance growth anticipated in the activities?
6. Does the team have global outreach and is cooperating internationally thus representing the R-CCS as top-tier HPC-research centre?
7. Is the promotion of group members and young researchers appropriately realized?
8. Is the team enabled by the directions and governance of the center to make these contributions?

In addition the AC has taken into account the number of publications and citations, other performance metrics as the h-index normalized to postdoc years and domain, as well as the recommendations of the previous AC from 2016, where relevant.

1. Takemasa Miyoshi, Data Assimilation Research Team

Data Assimilation (DA) is a cross-disciplinary science, combining observations and a model to give the best estimate of the state of the system: A reliable initial state is essential for a reliable forecast. The team applies DA mainly in the field of meteorology, and contributes to obtaining world-leading scientific results as well as solving social issues like the improvement of predictions of heavy rainfall and typhoons based on satellite data assimilation that can have direct applications to operational meteorology and thus can be capable of contributing to disaster mitigation and reduction. The team is one of the leading groups of DA in the world, which is confirmed by the number of awards, prizes, publications and media impacts. The team is also active in education of DA in Japan.

The AC welcomes the worldwide leading position of the team. The AC notes that the NICAM-LETKF code developed by the team is one of the co-design target applications for Fugaku, and is estimated to speed up by a factor of 120. The AC further notes the demonstration of a significant improvement in local, short-term (order 30 minutes) weather forecasts with a combination for ultra-fast (order 30 seconds) assimilation of phased-array weather radar data into high-resolution (horizontally order 100 meter) simulations. The AC acknowledges the challenging idea of the team to use these developments during the 2020 Olympic Games. The performance results of the team on the K-Computer are impressive. The team is extending DA to forest simulations and press-forming simulations, and plans also to integrate with Uncertainty Quantification and AI. The AC recommends integration of DA with Uncertainty Quantification and AI as soon as possible. The PI has joint appointments at RIKEN iTHEMS

and RIKEN CPR and contributes to enhance broader collaborations. The team has also international collaborations (USA, UK, France, Germany, …). The AC endorses that the management of the team with around 19 researchers is split into 3 groups of max 5 persons in order to promote young senior researchers as head of the group. The high number of visiting scientists and internships is impressive. The PI considers it very important to keep the motivation of each member high.

The AC acknowledges that the team has an excellent number of refereed publications (80) and a large number of invited talks, press releases, awards and grants.

2. Hirofumi Tomita, Computational Climate Science Research Team

The main objective of the team is to indicate a direction of future climate modelling for high performance computers such as Fugaku. The team used a convection resolving a global model, which they among the first groups developed in the world. This allows them to explore what resolution is necessary to reliably simulate cumulus clouds, which play an important role in climate simulations.

The AC appreciates the efforts of the team to clarify important physical processes, which are parametrized in climate and weather prediction models, and to improve models. which are important for advancing atmospheric science and for better data assimilation. This is important for reliable simulations and also for making better DAs with big data and AI, thus contributing to SDGs. The AC notes that studies of physical processes and those using cloud-resolving global studies will be further advanced with Fugaku. The team contributes to the co-design work for FUGAKU with the fast simulation code named NICAM-LETKF. The team also develops an open source library (SCALE) optimized for HPC. It is a library of weather and climate simulation tools, which can be used as either a global model (SCALE-GM) or a regional model (SCALE-RM) simply by selecting the dynamical core and physical processes and compiling with a driver. Furthermore, so far only a few models except for SCALE exist with truly Open Source Initiative (OSI) compliant licenses. The non-compliance of most model implementations is an impediment for the community, as it inhibits the development of commercial applications of weather forecasting. SCALE-based models would be the first truly scalable codes that could be used in commercial applications of weather forecasting, and could provide a good opportunity for future business development at R-CCS. The AC considers it important to increase the number of users both inside and outside the country and recommends that it may be effective to enhance international collaborations especially in the form of inter-comparison of model performance and reception of foreign visiting scientists. With 11 teams members, the size of the group seems appropriate. The fluctuation rate is seen as positive, as well as the high number of visiting scientists. The policy of the team in fostering young researchers is to develop human resources who not only are productive in scientific outcomes but also are familiar with their models on the methodological level, and thus are capable of developing models.

The AC acknowledges that the team has an excellent number of refereed publications (76) and has a good number of invited talks, press releases, awards and grants.

3. Seiji Yunoki, Computational Materials Science Research Team

The objective is the development of quantum simulators for strongly correlated many body systems and their application for next generation highly functional materials and future revolutionary technologies. Research seems very high level (with highly cited papers). The team is using an implementation of the auxiliary field quantum Monte Carlo method with delayed updates of the Green Function, which transform the vector outer product (BLAS level 2) into matrix multiplies (DGEMM, BLAS level 3). This results in very efficient simulations with a high fraction of peak.

The AC notes that very nice results of the phase diagram were presented for the Hubbard model on a honeycomb lattice with thousands of electrons, and a clear metal-insulator transition could be identified. The AC commends the record simulations with 10,952 electrons allowed the group to

explore the quantum critical region around the phase transition. In the future, the team plans to increase the size of the simulations by an order of magnitude to 100,000 electrons on Fugaku. The AC notes that this is ambitious, because the complexity of the algorithm scales with the cube of the problem size (assuming there is no sign problem), hence we are looking at needed performance increase of three orders of magnitude, which Fugaku alone will not deliver. The team is aware of this challenge and will explore hybrid methods Monte Carlo similar to those used in QCD. Whether this alone will suffice is, however, not clear. Therefore, alternatively, the AC recommends that the team may consider revisiting continuous time version of the auxiliary field (CT-AUX) MC method, that have led to substantial performance improvements in other applications of the Hubbard model. Since the memory footprint of individual MC updates in den CT-AUX algorithm may exceed the available memory on a Fugaku node, the team may have to develop distributed memory version of the update algorithm. A similar challenge to distributed linear algebra algorithms across nodes is faced by the Computational Molecular Science, and interactions between teams should be explored. The team's is approaching models with a sign problems with Density Matrix Renormalization Group (DMRG) methods. The AC commends the large-scale parallel DMRG methods presented. The AC suggests that this DMRG work be continued and maybe look at quantum cluster approaches to effective medium methods that suffer much less from the sign problem. 10 team members are seen as appropriate, the number of visiting scientists is very good. To promote members of the team, the PI proposes them "to show with scientifically significant results". It is a self evident but not a fully sufficient statement.

The AC acknowledges that the team can present an outstanding number of publications (111), has given many invited talks, issued press releases and software suites, and received awards and grants.

4. Takahito Nakajima, Computational Molecular Science Team

Theoretical molecular science that can safely predict the properties and functions of materials and drugs on the molecular levels is a central goal of the team. The team develops NTChem, a highly parallelized code for quantum chemistry (QC) with various unique functions. NTChem is one of the target applications for Fugaku. The stated goal is to put Japanese software back on the map in quantum chemistry. Currently, there is no Japanese software in this domain. The Japanese user community mostly relies on established packages, such as Gaussian, MolPro or GAMESS, which all have poor scalability and for which future development for modern architectures will be expensive or impossible.

NTChem is a full-featured QC package that supports electronic structure methods ranging from different flavours of density functional theory (DFT), to second order perturbation theory (MP2) and beyond, and is applicable to many chemistry problems with small to medium sized molecules. The AC commends the development of NTChem and the resulting performance comparisons of NTChem with CP2K and NWChem that confirm the NTChem code is competitive. An important advancement planned for Fugaku is the development of distributed memory versions of the algorithms, in order to scale the individual quantum chemistry simulations beyond the memory limits of individual nodes. The team also used machine learning for materials development for less toxic solar battery. The NTChem software has a wide number of users: Some of them are industries. Further nice applications of materials design for solar cells and photocatalysis that are based on VASP were presented. These computations comprise a combination of high-performance on-node and distributed high-throughput simulations to scan phase diagrams of binary and ternary alloy compounds: e.g., 30 thousand perovskite oxynitrides and oxide compounds were scanned on the K-Computer to find 42 novel candidate materials for new water-splitting photocatalysts that could work with visible light, six of which are perovskite compounds with inexpensive elements. The AC commends that such computational screening work is very timely and the capability provided by systems such as the K-Computer and in the future Fugaku will provide a decisive advantage. VASP is often used in such studies, because its versatile performance is well recognised in the literature. However, the AC notes

that VASP is not open source, which might result in future reproducibility problems. The AC recommends that the team thus consider using some of the Japanese DFT codes instead of VASP in the future and support the international efforts for development of open source codes for periodic systems. With 8 researchers the teams seems a bit understuffed, the AC recommends for inviting visiting scientists. The PI's objective is to raise the skill and experience of each team member and to produce the excellent researchers who will lead a new research field beyond computational molecular science. He also wants to train the researchers useful to the industrial world. 21 postdocs in 9 years are impressive number. Most of them have obtained positions. We do not have information about international cooperation but the team seems well aware of competing efforts elsewhere. The PI is actively recruiting and employing foreign researchers.

The AC acknowledges that the team can present an impressive number of peer reviewed publications (66) as well as review articles and books (14), has many invited talks, press releases, important software suites, and quite a few awards and grants.

5. Yuji Sugita, Computational Biophysics Research Team

The group develops new methodologies and software in molecular dynamics simulations of biological systems for basic life science, medical science, and in-silico drug design. They have developed new methodologies in MD simulations, and the GENESIS MD software, and tested its performance.

The AC acknowledges that the team carries out MD simulations of biological systems on a high, international level. The team is strongly involved in co-design up to Fugaku. Since computational demands of MD simulations are increasing the AC recommends that the team engages more deeply with developments beyond Fugaku. The team is integrated with experiments for conformational dynamics via Hidden Markov Modelling (HMM). The team certainly has a societal impact as it concentrates on methodologies for life sciences, medical science, or drug discovery. Several applications have been made for actual drug development by simulations of ligand-protein interaction and protein-protein interaction with collaborating industry. By the very nature of the problem studied, a continued focus on flops-centric computing can be expected. The team collaborates internationally with LANL, and competitively compares with NAMD and GROMACS. WIth 13 scientists and 4 visiting scientists, the team's size seems appropriate. The AC commends the focus on the management of group research data and the fostering of group members.

The AC acknowledges that the team can present a good number of peer reviewed publications (34) along with an excellent number of invited talks, and quite a few press releases, software suites, awards and grants.

6. Florence Tama, Computational Structural Biology Team

The goal of the computational structural biology group is to determine the structures of biomolecules that are important to understand biological functions, and for drug development. The team uses cryo EM data as well as data from SPring-8/SACLA and combines the experimental data sets with computational power of K and Fugaku. The team presented some new algorithmic approaches, in particular the use of structural data from the protein data bank.

Biological modeling studies based on cryo-EM are a very competitive field, and thus, the approach of the team is a often-used technology. On the contrary, application to modeling from the dynamic images provided by high-speed AFM (Atomic Force Microscopy) is very unique and promising. The AC very much appreciates the good interaction with experimental groups, in particular the use of imaging data from SPring-8/SACLA light source as well as those from cryoEM, XFEL (X-ray free-electron laser), and AFM, integrated utilizing GMM (Gaussian mixture model). The team's societal impact as its integrative approach can lead to a better understanding of the dynamics essential for drug development. The team is developing new knowledge based structural modeling. The research activities are more data-centric

rather than flop-centric. The AC would have liked to see performance results from K and a discussion on how the group prepares for Fugaku. The team did not report about international collaborations. The AC recommends that the group explores computational approaches that are in use at other light sources internationally. With 5 members, the team seems a bit understaffed. The AC commends the fostering of young researchers via the professor position of the PI at Nagoya University.

The AC notes that the team has a moderate number of peer reviewed publications (23) along with quite some invited talks. One would wish engagement as to press releases, awards and software. The team received 2 research grants.

7. Makoto Tsubokura, Complex Phenomena Unified Simulation Team

The goal of the team is to develop a unified simulation method for meso/macroscopic complex phenomena by utilizing HPC environment, in particular fluid motion, structure deformation, aeroacoustics, chemical reaction, etc. They are developing algorithms for solving multiple partial differential equations, and new approaches to mesh generation and integration of CAD data. They have developed a software framework called CUBE and applied it to many industrial applications.

The AC acknowledges that the team is well recognized by industry, as it has become a member of large consortia. The CAD and mesh generation results in comparison to LS-DYNA, as presented, are considered excellent by the AC. The research is highlighted as FS2020 priority issue 8, i.e. as a digital engineering test bed on the next generation supercomputer within the flagship project. No visible big data or AI activities were presented. As mentioned the industrial involvement is confirmed by being involved in 17 industrial simulation/flop centric vs. the data centric approaches. The team has performed the world's largest full vehicle simulation. Still, the AC would have liked to see validation results with experiment, for example in the wind tunnel. 12 team members are an appropriate number: many visiting scientists and student trainees complement the team. The AC encourages the team to establish a plan for developing and fostering junior scientists.

The AC acknowledges that the team can present a very good number of peer reviewed publications (57) along with many invited talks, and quite a few press releases, awards and software packages as well as research grants.

8. Jun Makino, Particle Simulator Research Team

With particle-based methods one can study phenomena ranging from molecular scale to the Universe. The group is currently developing a "universal" software application for this purpose that runs efficiently on highly parallel computers such as the K computer.

The AC strongly acknowledges the group to be a distinguished research team. Its leader has been Gordon Bell prize winner in the past, and the team has been Gordon Bell finalist with the K computer. It is noted, however, that only scaling (and lines of fraction of peak performances) have been shown and comparisons to other state-of-the-art particle codes have not been made. While universal software applications from molecules to the universe running efficiently on the K computer have been presented, no mention of post-K or Fugaku was made. The AC advises the team to concentrate on Fugaku implementations and to work on systems design beyond Fugaku. The AC notes at this stage that system design beyond Fugaku is an element of the ToR of the director of R-CCS to be followed by all teams. The team has developed processors (MN-Core) in cooperation with the company "Preferred Networks", which are suited for AI applications; however, the group is not involved in work on data science. Here, societal impact might be expected. The research naturally is a flop centric one. Internationally, the PI is very well known as inventor of the series of GRAPE machines and of large-scale N-body simulations codes and has produced important research results since mid of 1995. With 9 team members, the size of the team is appropriate, and an excellent number of visiting scientists is a good complement. The AC identifies different views on tasks, goals and management between the

leadership of this team and the direction of the R-CCS. The AC sees this as an obstacle to the team's contribution to the core research directives of R-CCS and as a risk to the team's future.

The AC notes that the team has a sufficient number of peer reviewed scientific publications (35) in relation to size and previous impact, some invited talks, press releases, awards and software packages, but enjoys an outstanding number of research grants.

9. Mitsuhisa Sato, Programming Environment Research Team

Programming environments for extreme-scale systems are generally of high importance and thus are a important area for the R-CCS. Much of the team's effort is focused around XcalableMP (XMP), a directive-based programming language extension for distributed memory systems that exploits the partitioned global address space (PGAS) programming model. center. It was good to see this work being placed in context.

The AC finds that XcalableMP is a good research vehicle for developing programming tools for extreme scale systems, and the AC applauds the work towards support of other architectures, such as the support for GPUs and FPGAs, and the work in performance modeling to help guide architectural choices, as well as task parallelism. This will be important for gaining adoption for systems beyond Fugaku. Other work described included performance modeling, important for co-design of future processors. This includes opportunities in BD/AI/ML space, particularly for higher level programming systems. The AC notes that there are many efforts to develop programming environments for HPC, but few have succeeded in being adopted. The AC encourages the team to develop an adoption plan and metrics of success. Success might be measured as providing input that influenced OpenMP developments, for example, and need not be establishing a widely adopted programming system that displaces MPI or UPC. The AC acknowledges that several effective international collaborations have been built around work on XMP in the past and encourages the team to increase cooperation internal to R-CCS and outside RIKEN. The number of 17 scientists and company associates is appropriate, there are also quite some visiting scientists in the team. The team is active in working with student summer interns, and conducts annual workshops in XMP.

The AC acknowledges that the team has a good number of publications (58) and notes that, due to the topicality of the publications, computer science tends to publish often in refereed conference proceedings. The team has quite some invited talks, press releases, a software package, many awards and an outstanding number of research grants.

10. Yutaka Ishikawa, System Software Research Team

System Software is essential for the design and the operation of extreme-scale systems and is one of the important areas of research for the R-CCS.

The AC welcomes that the team succeeded developing scalable high performance libraries/ middlewares as well as an operating system already now for the Tofu/FUGAKU environment with highly positive results. Their highlights are the PiP (Process in Process) execution model, the Light-Weight Multi-Kernel Operating System IHK/McKernel and the MPICH implementation for Tofu/FUGAKU. PiP is a new execution model in which processes share the same address space implemented at user-level, achieving better MPI communication performance than OMP or thread-based approaches. McKernel is a light-weight kernel which causes less OS-noise at the application level. McKernel runs on heterogeneous environments including FUGAKU. PiP received a best paper award at HDPC2018, and McKernel received a best paper award at ROSS2017. Both software systems are available as Open Source Software. The team still should consider leveraging existing MPI implementations noting that Fujitsu is providing an Open MPI version or consider closer collaboration. The AC encourages more activity in BDA, AI and ML by the team, for example including more opportunities as to community benchmarks and w.r.t. actual applications. The team has received

substantial third-party funding. The AC notes that there is room for national and international cooperation. With 9 members, the team seems a bit understaffed: in the past it has hosted quite a few visiting scientists and student trainees. The team successfully fosters young international researchers as PostDocs and research associates from industries.

The AC acknowledges that the team has a good number of publications (58) and notes that, due to the topicality of the publications, computer science tends to publish in refereed conference proceedings. The groups enjoys many invited talks, software packages, many awards and quite a number of research grants.

11. Toshiyuki Imamura, Large-Scale Parallel Numerical Computing Research Team

Large-Scale Parallel Numerical Computing is at the heart of simulation on systems like the K Computer, its Successor Fugaku and beyond.

The AC welcomes that the team is developing many necessary components for a high performance numerical libraries for the Fugaku computer. The team focuses on linear systems, symmetric eigenvalue problems, FFTs, random number generator, non-linear equations, and tensor-decomposition. They have strong collaborations with computational scientists, computer scientists, and applied mathematicians. They are focused on algorithms that avoid communication and overlap communication with computation, asynchronous, task-based algorithms, and data-compression. They are exploring novel ideas in this space, including mixed precision, reproducibility, fault-tolerance, and implementations on new hardware devices. The work being done is competitive or is exceeding the performance of leading numerical libraries today. As highlights, they received a best paper award at HPCAsia in 2019 and best poster awards at ISC2016 and ISC2017. 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 in order to help users solve large simulations involving 3-D PDEs. The team has a wide range of experience in matrix computations and should extend its focus on how linear algebra can benefit the big data, machine learning, and artificial intelligence efforts going on in the R-CCS applications. With 9 members, the team is a bit understaffed, still there are many visiting scientists and internships. The AC recommends that the team work out a roadmap to support young researchers.

The AC acknowledges that the team has a good number of publications (44) and notes that, due to the topicality of the publications, computer science tends to publish in refereed conference proceedings. The groups can present quite a few invited talks, a press release, many software packages, as well as many awards and quite a number of research grants.

12. Nobuyasu Ito, Discrete Event Simulation Research Team

Research activities related to discrete event simulation are an important field for large-scale simulations. At R-CCS they extend from traffic simulations, job scheduling on parallel computers, social networks, economics, to quantum computer simulations and more. Results on all of these topics were presented, demonstrating the generality of the work done by the team.

The AC endorses that the applications to transportation analysis in Kobe City, including evacuation plans in the case of a Tsunami in Osaka, are impressive and are important societal applications with particular interest in Japan, and moreover, make effective use of the K computer. The publication record is very strong and the focus on discrete event problems reflects non-traditional HPC applications and moves towards data-intensive computing, although analysis of real-time data was not described in the presentation. The discrete event simulation showed good scalability on the K-Computer. The AC recommends that this project should work with other HPC experts in the architecture space both to address a serious I/O bottleneck in the traffic simulation, and to understand how this type of workload would map onto future architectures beyond Fugaku, representing a workload very different from numerical simulations. Collaborations in China, Finland, Germany,

Hungary, Korea, and the UK were listed, still the engagement of the team as to international collaborations could be increased. The AC recommends that this project should work with others in the architecture space to explore the required system characteristics (including I/O) for good performance, since it represents a very different workload than numerical simulations. WIth 10 scientists, the team is sufficiently staffed, the AC recommends to invite visiting scientists. The AC notes that no substantial account on mentoring of junior researchers was given.

The AC notes that the team has a sufficient number of peer reviewed publications (29) along with many invited talks, a press release, as well as quite some awards, software packages and research grants.

13. Yasumichi Aoki, Field Theory Research Team

The team is concerned with the simulation of the standard model of particle physics, in particular the numerical solution of the dynamical quantum chromodynamics QCD. This comprises the methodology of large-scale simulations especially at low quark mass aka large lattices. Important questions are the order of the QCD phase transition, proton decay form factors, chiral symmetry breaking, and QCD at finite temperature.

The activities of the team are embedded in the activities of many QCD groups worldwide, thus it is not unique in itself. They represent the Post-K priority issue #9 at R-CCS, and it is noted that QCD issues have had much influence on the development of Fugaku. The AC fully endorses the work of this team, as QCD is an important part of usage of the K computer on the 7 to 8 % level and can be expected to play a similar role on Fugaku. The main software packages for QCD generated under post-K priority issues #9 were developed on other platforms like KEK's BG/Q and Oakforest-PACS. The AC notes that the ample software resources for QCD simulations, which have been developed so far in the post-K priority issues #9, can not yet be exploited to full extent on K and Fugaku. This is understandable, given the young age of the team, and the AC welcomes that the collaboration platform is formed to gather the knowledge and experience for the coding and related development for Fugaku and future computers. The team has connections to important Japanese experiments like Kamioka and international experiments, which need the results of such simulations. Here one must also take into account that the team was established quite recently. The team does not directly contribute to societal benefits, however to the generation of fundamental knowledge, and it has an impact on co-designing next generation supercomputers. The AC recommends trying to transfer simulation technologies like the hybrid Monte Carlo to materials science simulations, as for instance done by HMC for graphene elsewhere, to enable societally beneficial outreach. The activities of the team are flop-centric by their very nature, not data-centric. The team is building a platform for collaborations with 15 universities. Increasing outreach to young students and school kids is planned using the R-CCS QCD simulation platform they are in the course to build. Altogether, the team leader showed a good plan for outreach in Japan and internationally. The research group is young, it is supported by internationally well-known senior scientists on visit, young fellows are welcomed on both field, methodology and science. It appears that the number of scientists is insufficient to conduct this ambitious research properly, however, the team is enforced by working together with KEK, and the AC asks the management to increase staffing.

The AC acknowledges that the team has started only recently. Therefore only a few publications are available as well as invited talks. The team is asked to engage in software packages and acquire research grants.

14. Satoru Oishi, Computational Disaster Mitigation and Reduction Research Team

The investigation and anticipation of natural disasters like earthquakes, tsunamis, floods, and inundation is of prime importance for Japan. The research team has the goal to develop large-scale

simulation technologies for this purpose. It especially concentrates on Kobe City and surroundings, the Hyogo Prefecture.

The group is oriented towards large-scale data analytics and has established a patented data processing platform. The Team develops digital twins to simulated complex disaster as digital ensembles. The team seeks to be a bridge between the Science and local Government for the disaster mitigation and reduction. It appears that their research coverage is very broad. this might blur a strong point of uniqueness. The team has built a very fast solver for simulation on Fugaku and is encouraging industries to use Fugaku. The Team intends to develop an integrative Data Processing Platform, DPP. DPP creates the dynamic big data, which will be available for open collaboration with other organizations. It includes an urban structure builder that should become a core system of an intra-data platform designed for the MLIT of the Japanese government (Society 5.0). The AC strongly acknowledges that the team develops a core system of an intra-data platform designed for the MLIT of the Japanese government, thus contributing to the Japanese Government standard. Furthermore, the team is conducting an inter-center activity inside RIKEN with the RIKEN center for Advanced Intelligence Project. The team is addressing societal aspects (natural disasters). The AC asks the team to produce tangible output such as hazard maps. The team is at home on both sides, flop-centric and data-centric activities. The team has received many awards. The AC is surprised that there is no proper dissemination of the important results of this team, e.g. there seems no press release available, and advises the team to increase this activity. The team gives many lectures in fire-fighting departments, river management offices and high schools. It has international cooperation with the USA with PEER. The AC recommends to pay attention to the promotion of group members and young researchers.

The AC notes that the team has a quite moderate number of peer reviewed publications (12) but quite some invited talks. Press releases are missing, despite the high societal relevance of the field. The team enjoys quite some awards and research grants.

15. Kento Sato, High Performance Big Data Research Team

Big data analytics software and AI play a rapidly increasing role in all fields of computational science and engineering, its support on high end HPC system is still a quite unique activity. The team leads research and software developments for accelerating machine learning, deep learning and large scale big data processing (AI techniques) on the K and the post-K computers, i.e., the teams deals with HPC for AI. The team also studies how to accelerate HPC applications and HPC systems by using AI techniques.

The AC welcomes the HPC-for-AI+AI-for-HPC strategy of the team. Since this is a new team launched in October 2018 it is, however, difficult to judge the research results at this stage properly. The team explores new programming techniques, new algorithms as well as new applications of these techniques to AI techniques for the next generation HPC systems, the post K computer and systems beyond it, i.e., what they call “the post-Moore era”. The team does not study AI theory itself. This is the task of the “RIKEN Center for Advanced Intelligence Project (AIP)”, which is in charge of explainable AI in RIKEN. The AC notes that the application of Explainable AI (white-box AI) still could be very useful for the HPC research group in R-CCS and recommends close cooperation. AC recommends to explore more applications of AI and their requirement. This team strongly contributes to the convergence of HPC with Big Data and AI. AC recommends to better define in what context. The AC recommends to further increase the successful collaboration with the High Performance Artificial Intelligence team as well as to include the data assimilation team and other teams. This can provide more diverse big data use cases and will broaden the scope beyond strict I/O problems generated by HPC simulations. The team does not contribute directly to societal issues, They have some collaborations with ANL and LLNL. The number of team members needs a substantial increase. The team has its own policy with long-term career goals. It aims to work with international students as well as domestic students.

Research themes are discussed among the R-CCS research groups for developing new methods and tools.

The AC acknowledges that the team has a sufficient number of peer reviewed publications (30) and invited talks given the very small size of the team.

16. Masaaki Kondo, Next generation High Performance Architecture Research Team

A set of research projects focused on computer architecture and the general problem that future hardware will be energy-constrained, is carried out by the team. Two issues are addressed, heat density on chips, which limit performance improvements, and the need for increased power and cooling infrastructure in the computing center. The team has a track record of multicore chip design and stacked chips.

The AC acknowledges the importance of this work on energy efficient processors and power-aware scheduling. Work in specialized processors to improve performance for particular application domains was described, as well as a power-aware resource manager. The anticipated work on specialized architectures was fairly broad, including support for deep learning, graph algorithms, and graphics. These topics are well-aligned with the overall plan to expand into data and learning applications, although the modest-sized team will need to focus on a specific set of architectures. The team can provide a strong connection with BDA and HPC. The AC recognizes that this is a forward looking team, taking account of hybrid designs with accelerators. The AC notes, however, that the described areas of specialization (deep learning, graphs, graphics, big data, annealing) seem overly broad for a small research team and in areas that are likely to have significant competition from industry. To maximize the impact of the work it should take a stronger co-design approach, starting with important application drivers to aid in design decisions and evaluations. The AC recommends that this work should be focused more narrowly on application problems where the team can expect to have a real advantage, and is aligned with R-CCS system plans. The AC recognizes that the team is orchestrating/collaborating with community-wide efforts like BDEC2, the HPC PowerStack initiative, the IEEE IRDS roadmap, etc., and is fostering international students. Given 2 team members only, it is important that the team has the RIKEN support to grow to an appropriate level.

The AC acknowledges that the team has an outstanding number of peer reviewed publications (44) and invited talks given the very small size of the team. A press release is missing, the team can proudly present quite a few awards, a software release and an outstanding number of research grants.

17. Kentaro Sano, Processor Research Team

The team works on dataflow-based computing models for FPGAs. The team thus might have a broad impact on the whole R-CCS in the future.

The AC welcomes that dataflow based computing models on FPGAs get a basis for a broader impact on the whole R-CCS. Based on his publication record, the team leader is indeed an FPGA expert and has looked at stencils, compression, lattice Boltzmann, and more. FPGAs are likely to be well suited for genome alignment and other integer computations. The work on the C/C++ compiler frontend of the data-flow compiler for custom computing with FPGA is welcomed by the AC. Although this research could well be aligned with BDA, this was not emphasized. The AC recommends that more exploration of the applications including AI/BD is to be carried out. Beyond that, for the future research, coarse-grained reconfigurability is proposed, which the AC acknowledges as a fruitful approach. It is difficult to see a link to societal issues except for energy use of HPC systems. Some of the applications could become such links, however. The team is composed of 8 members (5 post-docs), support of young investigators should be strengthened. The AC acknowledges links with several other R-CCS teams. The AC encourages the team to strengthen its engagement with the next generation high performance architecture team.

The AC acknowledges that the team can enjoy a good number of peer reviewed publications (10) given its short period of existence, and quite some invited talks. There are no press releases or awards, one software release, but a very good number of research grants.

18. Hiroya Matsuba, HPC Usability Research Team

The team aims to develop new techniques and software that contribute to an increase in the number of users or use cases of supercomputers, especially for Fugaku. In order to find new users cases they are focusing on Cyber-physical systems. The challenge is to provide an integrated software environment for Cyber-physical systems.

The AC acknowledges two research achievements: the development of the parallel framework Pyne for easy development of parallel programs (with dynamic code generation) and the development of a cyber-physical simulation model of the cooling facility of the K computer. There is a patent application (2018) devoted to simulation control and computer system. The AC notes that so far the work has not shown significant impact within R-CCS. There is little visible involvement in the development of the cooling of Fugaku and certainly not beyond, nor are there active contributions to the convergence of HPC with Big Data and AI. The team is focusing on cyber-physical systems, which are the systems, in which the information technology interacts with real-world assets or humans to improve economic efficiency or to contribute comfortable life of people. The number of team members is 5, support by the direction and governance is not obvious since all four team members (4) are working with fixed-term employment contracts. The AC recommends to clarify how far the team gives beneficial contributions as well as to increase national and international collaboration and global outreach.

The AC notes that the team has a low number of referred publications (4) even given its short period of existence, and one invited talk. There are no press releases or awards, one software release, nor research grants, but there is a patent application.

19. Satoshi Matsuoka, High Performance Artificial Intelligence Systems Research

Only a few initiatives worldwide focus on the integration of HPC and AI, while it appears to be obvious that the growing size of deep learning systems and other compute-intensive networks will more and more need HPC systems along with software and algorithms research.

This research team, led by the director of the R-CCS in person, has started just recently. The AC much welcomes the establishment of this group extending the portfolio of R-CCS research towards data science and AI, especially the objective of the team, to address the convergence of HPC and AI. The team follows a co-design approach for next-generation AI technology that will utilize state-of-art HPC facilities, especially including Fugaku. AI research is in principle at the center of the solution of many societally relevant challenges that will require increasing amounts of computer power, however, the team does not directly contribute to societal applications yet, but the new devices and software can provide large contribution to AI-related research both inside and outside of the R-CCS. Funded until 2022 by the JST-CREST "DEEP" AI project, the team collaborates with other research institutes in HPC and AI-related research in Japan. Given its recent establishment, the team does not yet show global outreach but certainly is heading for international cooperation given the international standing of its leader. The AC notes that the team leader is one of the busiest people in the worldwide research business, who is obliged to conduct higher level management. The AC recommends that an intermediate level scientist helps take care of the growing team.

The AC acknowledges that the team has just been created and is in the initial phase.

20. Toshiyuki Tsukamoto, Facility Operations and Development Unit

The Facility Operation and Development Unit bears the fundamental responsibility for a smooth and effective operation of the calculation and simulation processes on the large-scale infrastructure as well

as the timely development and preparation for the next generation system. It is evident that the K Computer/Fugaku site has some unique features specific to the site as a power co-generation system or a high-speed current limiting circuit breaker. A major task is to get the power consumption as effective as possible considering the rapidly increasing electricity prices since 2011 and the increasing power demand of future systems.

The AC commends the unit for an overall very stable operation of the K Computer over its 8 years of runtime, a success only possible on the basis of an extremely reliable facility operation. The AC acknowledges the substantial improvements in power consumption of the air handlers by about 40% within the runtime of the system, contributing to a large extent to the reduction of energy consumption by improving the efficiency of the cooling tower through removal of part of the wall based on CFD calculations to predict improvements. The AC applauds the innovative efforts of the unit to reduce energy use of the center, improving the total PUE from 1.43 (initially) to about 1.30 (today). At present the unit is consumed with the Fugaku site preparation and the giant task to get operation of this new system under control. The unit furthermore plans to reduce the PUE down to 1.1. With 11 members, the size of the unit seem appropriate, a clear description of the unit's development or mentoring for this critical operational aspects of the center was not presented, however. The AC recommends the facility operations and development unit to develop a clear and transparent succession plan for continuity of its important activity.

The AC acknowledges that the unit, with a clear focus on its primary task to act as a service unit, has still been able to publish three referred papers.

21. Atsuya Uno, System Operations and Development Unit

The continuous development and improvement of the systems operation and management is again a very important fundamental activity for a world leading large scale supercomputing and data centre. The system utilization is dependent on many factors like the parallel file system stability, but, given the extreme energy demands, measures for emergency job stopping are also important, maybe even more for the operation of Fugaku.

The AC welcomes that the team has developed a tool to monitor power usage and system logs for continuous improvement. While analyzing and improving the systems utilization of the K computer, given its quite heterogeneous job profile, is certainly difficult, the unit managed to arrive at an average utilization of more than 80 % in the last years of operation. It should be mentioned here, that setting the goals for utilization is somehow difficult, however. With these tools, the unit is able to immediately stop a job if the system exceeds its maximum power envelope. The unit is able to anticipate system failures or hardware problems by monitoring system logs, e.g., poorly performing file system accesses may indicate a cabling problem. Also degradation of the file system could be analyzed and traced back to load concentration. The AC notes that operation has measurable KPIs that should be formulated clearly. The AC has the opinion that optimization of scheduling and power consumption can be improved by ML. The AC recommends that the unit looks for the opportunity to apply ML for optimization of scheduling and power consumption in its research portfolio, interacting with the AI activities of the research teams. With 10 members, the size of the unit seem appropriate, development or mentoring of staff for this critical operational aspects of the center was not presented, however.

The AC acknowledges that the unit, which has primarily to act as a service unit, has carried out an important R&D role but still was able to publish four referred papers.

22. Kazuo Minami, Application Tuning Development Unit

The close linkage between system features and application features running on the system is a very important activity for large-scale simulation and data analytics centers. This is a field of activity that,

for the K computer and the forthcoming Fugaku system covers performance evaluation, stable operation, tools development and enhancement, a DL environment and the R-CSS Software Center.

The AC recognizes the high systemwide performances of quite a variety of data intensive codes, in particular the high performance achieved through the TOFU network on the K Computer. The K Computer reached stable operations with a few upgrades only beyond security patches. This makes application maintenance efficient. The AC was impressed by the performance impact of the application tuning efforts on the K Computer and the use of these application benchmarks to prevent performance degradation after software updates. The AC acknowledges that many efforts for performance tuning for Fugaku are on-going covering RSDFT, Adventure, or FrontFlow/blue achieving excellent speed-up on single nodes. The AC would have liked to see in how far the much lower ratio of TOFU performance as to the memory I/O for Fugaku when compared to the K Computer will affect the overall performances of data intensive codes, in particular concerning the positive effects of the "representative process communication". The unit is central to coordinate efforts for Fugaku in this respect. The AC notes that the process for selecting application benchmarks was unclear and should represent the planned workload on the system. It is acknowledged by the AC that investigation for ML and DNN libraries on performance tuning has already been started, however. The AC is concerned about the decreasing size of the group currently equipped with 4 members, formerly 10! This should be reconsidered by the director. The AC recommends the unit to look more at continuous integration / continuous development, since Fugaku will see more updates to the OS/SW stack.

The AC is impressed by the high number of peer reviewed publications (38) and invited talks of the unit.

23. Fumiyoshi Shoji, HPC Usability Development Unit

The enhancement of usability concerns the HPCI shared storage system, high-end visualization tools, workflow tools as well as collaborative environments, the pre environments and post environments of K and Fugaku, and support and development of software at R-CCS (Software Center). In addition the HPCI shared storage system that can be accessed Japan-wide is operated jointly by R-CCS and Tokyo University.

The HPCI environment is essential for carrying out data intensive science. The AC acknowledges that the unit has managed direct data transfer from HPCI to the K parallel file system maintaining important security aspects. The unit developed the web application workflow system WHEEL in collaboration with Kyushu University, as well as the cross platform visualization system HIVE, geared in the direction of interactive in-situ visualization. The AC welcomes the development of the workflow tool WHEEL and the K-specialized visualization system HIVE. The AC has not recognized contacts between the HPC Usability Research Team and this Unit. The AC recommends that the HPC Usability Research Team and the HPC Usability Development Unit show strong interaction. The AC thinks that the unit needs to carry out some benchmarking with corresponding tools so that workloads can be shared risks be reduced. AC recommends a detailed benchmarking including the similar activities from a performance perspective in addition to a functional one.

The AC is impressed by the high number of peer reviewed publications (27) of the unit compared to 4 publications of the Usability Research Team (Team No. 18).

4 Conclusions

The AC shall summarize its findings in the order of the four themes set by the ToR of the R-CCS Director, taking into account the ToR themes of the RIKEN President.

4.1 Overall Contribution to the Core Research Directives of R-CCS

Overall, the AC is very impressed with the high standards of research and development as well as the service quality of the R-CCS teams and units. They are largely at the highest international level. It is evident that the center as a whole has improved substantially since the last AC review in terms of directions and governance. This includes fulfillment of most of the recommendations of the last review. A majority of the teams follows the directives towards a successful operation and usage of Fugaku and the right research beyond Fugaku and into the post-Moore era. The AC acknowledges that the R-CCS has successfully begun to raise funds for the initiation of "Cambrian" computing towards the post-Moore era. Many teams understand the impending end of flop-centric computation growth along with the rise of data science, AI and heterogeneity of computing at the node, rack, system and data center levels.

However, when the research agendas of the many teams are considered individually, the R-CCS does not yet seem to fully coincide with the strategies set out in the ToRs of the RIKEN President and the R-CCS Director. While some groups focus on post-Moore architectures, plan for the post-Fugaku period, think and integrate Big Data and Machine Learning/AI, work directly with experiments or pay very close attention to their societal implications, and are thus very well aligned with R-CCS research guidelines, the AC had the impression that some other teams are only partially on the right track, and a few teams are more or less legacy teams or relics that continue to focus on their traditional research agendas; i.e., the latter are largely simulation-centric in a non-data-intensive sense or focus on previous architectures or alternative architectures that do not really contribute to R-CCS's core research and objectives. The AC has evaluated each team in detail and made specific comments and recommendations on these issues.

The AC intensively discussed whether it might be more appropriate for a centre based on a large computer to better separate those activities, where efforts to develop new methods and results lead the world, and those activities aimed at ensuring competence in areas essential to the operation of the centre. For example, it may be desirable to have smaller teams with a strong understanding and competence in some application areas to ensure excellent user support than a similarly large team trying to compete with excellence in science and other research centres. At the same time, the scale and performance of the supercomputer system offers a number of unique opportunities for truly world-leading research. Part of the problem is that it may be better to focus efforts for global leadership in research on a smaller number of areas and to require sufficient support competence in other areas to fulfil R-CCS's mission. In general, to achieve the goal of world leadership in research, efforts need to know the state of the art and be able to articulate why their work is better or complementary, and there should be a strong tendency to build on the work of others.

This focus proposal goes hand in hand with the fact that there are 19 research teams and all teams report directly to the Director along with 4 units and 4 FS2020 teams. The AC believes that there are too many direct reports. Concentrating on fewer world-leading research teams – coupled with expanding to more supportive teams – could solve this problem, otherwise the AC recommends the introduction of a hierarchical reporting system.

As emphasized in the ToR of the RIKEN President (theme 1), the AC is convinced that the R-CCS meets international standards of research and is world-leading center (1). R-CCS makes a significant contribution to society or is well on the way to increasing its societal impact 1(2). As to 1(3), the AC notes that R-CCS' activities meet RIKEN's current seven year plan in compliance with "Article 5 of the Act of Promotion of Public Utilization" both in terms of the improvement of the shared infrastructure as well as its access by public use. Within the framework of this act, the AC emphasizes that it is imperative to give the director of the R-CCS the flexibility to achieve an optimal balance between the pursuit for global leadership in research and sufficient support competence.

In section 3, the AC has given a detailed evaluation of each of the PIs as to fulfillment of their duties and in accordance with the center's mission and the RKEN seven year plan, following the ToR of the RIKEN President (theme 4).

Finally, the AC recognizes that it will be a big challenge for the R-CCS Director to align all the teams and units with his new strategy and encourages him to be deeply engaged with leading *all* the teams and units to pursue research consistent with the R-CCS strategy. This might invoke a profound transformation of the directions of some teams up to the sunsetting of teams.

4.2 Internal and External Collaborations

The AC recognises that internal cooperation is well developed at R-CCS. This is an important virtue. Where appropriate, the AC has recommended that a range of further links be established between teams and units, as indicated in the team and unit-specific sections and recommendations. Overall, there are few teams with little internal cooperation activity.

The AC recognises that the Centre works well with national and international research in general. This applies to both the global HPC community and domain research. It is reassuring that many of the teams have very good connections to experimental research and theoretically oriented groups. The AC found that the majority of the teams and units are well aligned and integrated into international research, but quite a few are not. This is stated for each team individually.

At the RIKEN level, the AC welcomes the contribution of the R-CCS to the RIKEN cluster of pioneering projects. While the RIKEN representative stated that no additional funding is available, PIs and individual researchers are encouraged to apply for external funding and PIs should encourage their team members strive for external funding in the context of the RIKEN cluster of pioneering projects. It is reassuring that RIKEN does not impose a strict quantifiable expectation on the ratio between the R-CCS' basic duties for a PI and the team vs. their pioneering projects: the RIKEN representative mentioned a ratio up to 30 % vs. 70 % between pioneering projects and basic duties.

To conclude, the AC acknowledges that the R-CCS is continuously enhancing its function as a "science and technology hub", in alignment with the ToR of the RIKEN President (theme 1), (1). The research and development activities of the R-CCS are rightly intertwined with the support and operating competence and thus all elements contribute in a balanced manner to this function.

4.3 Outreach to the HPC Community at Large and to the Industry

The AC notes that the R-CCS, with its world-leading supercomputer operation, is internationally well known and well recognized as a global player. This is reflected in the contribution to many international activities of the HPC community, where R-CCS' expertise is highly recognized. This recognition has been beneficial for many of the connections that the teams have been able to establish around the world, as well as attractive for many international guest scientists, student exchanges, etc. In this respect, it is indeed a stroke of luck that the R-CCS has a director who is one of the world's leaders in the field of HPC.

The AC is less convinced that the R-CCS is active enough concerning its outreach to industry. Certainly, the connection with the builder of the K Computer and of Fugaku is important, but this is not the end user industry as far as HPC is concerned. The AC is well aware that this is a general problem everywhere in the world. Here the AC repeats a suggestion made in section 2.4, i.e., to attempt to sell part of the resources to commercial users to expand the user base and to develop a long-term strategy that might significantly grow R-CCS' business and its involvement with industry.

4.4 Conducting Research Under Proper Ethics

Ethically sound research has become one of the main concerns of governments and donors worldwide. A world-leading supercomputing centre like the R-CCS is obviously and rightly under strong public control. The main concern of proper ethics is the forging of research data, the abuse of personal data, and the misuse of computing facilities for activities that threaten society or the economy.

The AC was pleased to observe ethical awareness at R-CCS concerning data handling and use of supercomputing resources in all teams and units. It is evident the AC is not able to perform deeper control of this issue. The AC sees the necessity that the R-CCS should consult with experts to further develop data policies. This also holds for misuse of supercomputers. There exists a branch of security research, which is concerned with signatures of misuse of HPC system. The AC recommends that the R-CCS should consult with experts in this field.

As regards the ethical responsibility of user research proposals, this is primarily in the area of RIST. The AC recommends that the R-CCS collaborates with the RIST in this respect. The tracking of publications by R-CCS users can, in addition to the positive effect of measuring the scientific impact of the Centre, help to evaluate the misuse of resources. Finally, all software developed to date is open and freely usable. The AC encourages the R-CCS to develop a general licensing policy for released software.

The published work of R-CCS does not appear to be publicly available in general. In order to promote open science and public relations, all published and reviewed works of R-CCS should in principle be openly accessible.

5 Recommendations

1. As a general recommendation, valid for any similar institution, R-CCS should implement succession planning and leadership development activities at all levels of staff.
2. R-CCS should continue activities to further involve user communities of Fugaku, or rather to devote R-CCS for nurturing the HPC communities and also for sustainability of R-CCS.
3. The R-CCS director should reduce the number of reports evaluated by him directly and consider a more hierarchically structured organization of this process.
4. The R-CCS director should evaluate the benefits to focus on a few world-leading research teams and to gear the other teams towards becoming world-class support teams with a partial research component.
5. The R-CCS Director needs to align all the teams and units with his new strategy and is encouraged to invoke profound transformations of the directions of some teams up to the sunsetting of teams.
6. R-CCS should try to increase its engagement for the RIKEN cluster of pioneering projects.
7. The R-CCS should provide open access for all peer reviewed research literature, software components, and data products. A general license policy should be established for released software.
8. While the AC observed ethical awareness at R-CCS concerning data handling, R-CCS should further consult with experts to develop data policies.
9. While the ethical responsibility of research proposals of users is in the domain of RIST at the first place, the R-CCS should be in sync to be enabled for the proper monitoring of programs and the data used and produced.
10. The R-CCS should attempt to sell part of the resources to commercial users to expand the user base and to develop a long-term strategy that might significantly grow R-CCS' business and its involvement with industry.

Signature

Jülich, Germany, October 9 , 2019

Prof. Dr. Dr. Thomas Lippert (Chair),

on behalf of the R-CCS 2019 Advisory Council, Prof. Dr. Jack Dongarra, Prof. Dr. William Gropp, Prof. Dr. Catherine Lambert, Prof. Dr. Haruki Nakamura, Prof. Dr. Hiroshi Niino, Prof. Dr. Thomas Schulthess, Dr. Satoshi Sekiguchi, Prof. Dr. Horst Simon, Prof. Dr. Shinji Shimojo, Prof. Dr. Shinji Tsuneyuki (Vice Chair), Prof. Dr. Katherine Yelick

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
SACLA	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
Top500	Ranking of benchmark results based on a dense matrix computation
ToR	Terms of Reference