Review of the Center for Computational Sciences, University of Tsukuba, 2007

Summary of Conclusions and Recommendations

1. The Review Committee was impressed by the achievements of the Center for Computational Sciences (CCS), its interdisciplinary culture, its leadership and its vision for the future. The Committee believes that the collective expertise in the Center and the research environment it has built are unique worldwide. In continuing the work of the previous Center for Computational Physics, the CCS has established over a period of 15 years a highly successful collaboration between computer scientists and physicists that has built a series of world-class computers and used them to produce results at the forefront of computational physics. This is an example to other areas of computational science and collaborative working has been well integrated into the culture of the Center. The promotion of this interdisciplinary approach to computational science research should remain the core objective of CCS. In order to repeat this success, it is necessary to identify projects in which applications and computer scientists have shared common goals. The CCS should maintain a portfolio of exciting research projects at the forefront of computational and computer science, a talented faculty engaged in this research, and an environment in which shared goals can emerge and be nurtured. It will need continually to seek mechanisms and incentives for participating in interdisciplinary projects. The Review Committee commends the CCS to the University and recommends that the University acknowledges its strategic value, supports the execution of its mission, and encourages the wider academic community to engage with its activities.

2. The reorganisation of the Center in 2004 has been successfully accomplished and the Review Committee congratulates the Center on expanding its faculty and establishing new research groups with a breadth of applications that are at the forefronts of their respective fields and of growing importance in the 21st century. Some of these new groups are sub-critical in size, however, and they will need to build strong collaborative links outside the Center, or be targeted for new hires. The Review Committee also strongly supports the Center's vision for computational science and the collaborative culture that it is fostering between these researchers which was clearly reflected in all the presentations to the Committee.

3. The Review Committee considers the CCS's research in lattice QCD to be world-leading and that the International Lattice Data Grid/Japan Lattice Data Grid (ILDG/JLDG) project is establishing an important global infrastructure. The Division of High Performance Computing (HPC) Systems is also world-leading, being one of the few groups in the world that are successfully building HPC systems in an academic setting. They are highly regarded, productive and innovative. The research in astrophysics is diverse, some of it at the forefront, and has the potential with FIRST to make significant breakthroughs in new application areas. Thus, the Computational Astrophysics Group has the potential to become world-leading if it

can achieve greater coherence around its areas of strength. The condensed matter research focuses on fundamental aspects of future device applications and is excellent. The application of time-dependent density functional theory to both molecular and nuclear physics is a unique combination, and the collaboration between the Computational Condensed Matter Science, Quantum Many-Body Systems and Computational Life Science Groups could soon make the CCS a world-leader in this new approach. The Computational Media Group is also excellent, verging on world-class, but its link to HPC is a future opportunity. The Biological Science Group is excellent, but it is sub-critical in size, and synergy with HPC will also take time to generate fruitful research. The Global Environmental Science Group is producing some excellent work and its focus on downscaling of climate change is appropriate, but it is also sub-critical in size. The Computational Intelligence Group is also excellent and its collaboration with the geo-scientists has created an important opportunity to communicate climate science to the public.

4. The CCS should have an internationally competitive computing facility to support its research projects so that it can focus appropriate computational resources to ensure they have maximum impact in their fields. The Center's strategy to take responsibility for the University's research-oriented supercomputing service will help it to maintain control of appropriate systems and the planned "T2K" acquisition will achieve this for the next 3-4 years. Additionally, CCS should continue to develop new computing technologies in the context of the most demanding applications needs and have the capability to bring these technologies into use to make breakthroughs in specific areas of computational science. The development of novel architectures is high risk but high reward, and although the funding required for any particular development cannot be guaranteed, the Center's successful track record in this area is a major strength which is rare worldwide and should not be lost. There is a rapidly evolving range of new technologies which might be deployed to make significant advances in computational science and the Center is particularly well-positioned to evaluate and exploit these.

5. The CCS has too few technical staff both to support growing use of its computers by its increased range of projects and to develop applications codes that scale efficiently on petascale architectures which will be vital to their continued competitiveness. This situation will be aggravated as the Center becomes responsible for users of the University's supercomputer and for external projects which may not include CCS research staff. However, this also creates the opportunity for the University to transfer the required support staff, or the funding for them, to CCS along with responsibility for its supercomputer service. CCS will need to introduce a more service-oriented user support mechanism for this wider range of users, because the Center will risk its excellent reputation if the service is poor, but it must be careful to ensure that this service activity does not dilute the Center's research focus.

6. The expansion of the CCS's graduate education activities to include a dual degree programme in computational science is strongly endorsed by the Review Committee, because it will help to build a strong computational science community within the

University and to train the next generation of computational scientists. The Committee recognises that there may be challenges in defining a curriculum which is attractive to both students and the wider faculty within the University. We encourage the University leadership to support this initiative by the CCS.

7. The CCS's interdisciplinary research focus, educational activities and pioneering of new technologies should enable it to play a major role in supporting initiatives such as the National Next Generation Supercomputer ("Petaflops") Project. The Center should seek strategic partnerships with national projects such as this, which exploit its complementary expertise in developing new application areas in computational science and its capabilities as an icebreaker for innovative algorithms, high performance software and new hardware technologies.

8. The CCS should actively develop international collaborators for some future core projects, recognising the trend towards multi-national petascale supercomputing facilities and building on the success of ILDG. Today, the CCS has a high international profile. It should leverage this now to participate in a broader range of international projects, so that it continues to be recognised as the center of excellence for interdisciplinary computational science research in Japan.

9. The PACS-CS project was a successful and timely development to provide an appropriate computing resource to support the expanded range of projects in the Center, bridging to the next University supercomputer. The Center's system development expertise created the opportunity to incorporate a more capable (hyper-crossbar) interconnect than was available commercially at the time, and this has already benefited applications on 256 and 512-node partitions. The benefit should be even greater when the system is configured to run applications on up to 2048 nodes. The CCS is now well-positioned to do breakthrough science in QCD and materials science using this full capability.

10. The FIRST project is unique in the world and typifies the Center's ability to use applications needs to drive innovative computer architecture and systems research, which in turn delivers scientific results at the forefront of the applications field. This has enabled the GRAPE technology to be deployed effectively for a wider range of frontier astrophysics applications. Specialised hardware will continue to play a vital role in specific fields. The expertise that CCS has acquired in embedding accelerators in general-purpose clusters will have wider applicability in the future and should be the basis for follow-on projects.

Review Committee, its Operation and Objectives

11. The Review Committee comprised Professor Richard Kenway (University of Edinburgh), Chair, Professor Yoshio Oyanagi (Kogakuin University), Vice-Chair, Professor Masaru Kitsuregawa (University of Tokyo), Professor Michael Norman (University of California, San Diego), Professor Hervé Philippe (University of

Montreal), Professor Horst Simon (Lawrence Berkeley National Laboratory), Professor Akimasa Sumi (University of Tokyo) and Professor Kiyoyuki Terakura (Japan Advanced Institute of Science and Technology).

12. The objective of the External Review was to evaluate the research activities of the Center for Computational Sciences (CCS) in the context of its mission, since its reorganisation in 2004 and at the midterm in its six-year plan. The Review was asked to focus on the work of each research group and the three core projects, PACS-CS, FIRST and ILDG/JLDG, with regard to both their achievements and the effectiveness of their interdisciplinary collaborations. The Committee ranked the research groups, in terms of the originality and impact of their outputs, according to the following categories: "world-leading" (amongst the five best groups in the world); "world-class" (the best group in Japan and internationally recognised for consistently producing excellent work, but falls short of the best in the world); "excellent" (amongst the five best groups in Japan, with international recognition for some work, but lacking the consistency or impact of a world-class group). Finally, the Committee was invited to comment on the future vision and strategies of the CCS.

13. The Review Committee visited the CCS from 30 Oct to 1 Nov 2007. It had available three reports: "PART I, Overview of Center for Computational Sciences Summary of Activities 2004-2007", "PART II, Research Activities, Results, Collaborations and Plan 2004-2007" and "PART III, Vision and Strategies of Center for Computational Sciences". During the visit, the Committee received presentations from the Director, Professor Mitsuhisa Sato, and former Director (now Executive Advisor to the President of the University of Tsukuba), Professor Akira Ukawa, about the activities of the Center since 2004, from each of the research groups (Computational Particle Physics, Computational Astrophysics, Computational Condensed Matter Science, Computational Life Science, Quantum Many-Body Systems, Global Environmental Science, Biological Science, Computational Intelligence, and Computational Media), from the Division of HPC Systems, and from the PACS-CS, FIRST and ILDG/JLDG Projects. In addition, the members of the Committee divided into three subgroups to hear further presentations of research results in their own fields. Finally, Professor Sato presented the future plan and strategies of the CCS. The Committee was able to discuss with members of the CCS on a one-to-one basis and collectively with the senior management. During the latter, it received responses to questions that had arisen during the presentations. The Chair presented the Committee's conclusions and recommendations to a meeting of all the members of the Center at the end of the Committee's visit.

Research Activities

Division of Particle Physics and Astrophysics

14. *Computational Particle Physics Group*. This Group has played a defining role in all major aspects of the Center since its conception. In particular, with the CP-PACS Project, it successfully demonstrated how collaboration between applications and

computer scientists could achieve a world lead by a substantial margin and thereby dominate the field for several years. This project produced a landmark result by falsifying quenched QCD and now the Group is poised for another landmark computation using PACS-CS to carry out the first simulation of lattice QCD at the physical parameter values, removing all need for extrapolation except for the continuum limit. In parallel, the Group continues to produce results for a wide range of physical quantities from lattice QCD, some for the first time and many which are state-of-the-art. The Group has retained coherence in its physics programme over many years and has built an international reputation for excellence in the field which is unsurpassed. In addition, the Group has been a driving force behind the ILDG and JLDG, in the former case promoting international cooperation through its leadership of the International Lattice Field Theory Network. New algorithms which are substantially faster than previous ones, the Group's wealth of experience and the prospect of multi-hundred teraflops years of computing are now accelerating it towards a period of unprecedented physics output. It should continue to be the jewel in the Center's research programme.

15. *Computational Astrophysics Group*. This Group, under the leadership of Professor Masayuki Umemura, is large and diverse. The addition of two junior faculty in 2005 and 2007 represents a healthy investment in the future. The presence of 4 postdoctoral associates, 15 graduate students and 4 undergraduates make it one of the largest computational astrophysics groups in the world. With access to unique computational resources (FIRST, PACS-CS), it has the potential also to be one of the best. A wide variety of research applications in astrophysics and cosmology are being pursued, some of them pioneering. Particularly notable are the algorithm development work, and the applications in radiation hydrodynamic simulations and advanced models of magnetically driven outflows. Not all of the research presented was of uniformly high quality, perhaps because it represented some exploratory forays into new topics and applications by students. There is a danger, with so many codes and abundant computer resources available, of dabbling in many areas. The Group could increase its scientific impact and visibility by concentrating on a few leading problems and developing a coordinated plan of attack, guided by observational data, much as the Computational Particle Physics Group has done. The Committee was happy to hear about several projects that are closely guided by new observations, some of it from Japanese observatories. The Committee endorses the Computational Astrophysics Group's future plans to heavily exploit the FIRST computer, to collaborate with the Division of HPC Systems on the Multiple Architecture Seamless System (MASS), and collaborate with the Division of Material and Life Sciences on *ab initio* calculations of 3-body molecular hydrogen formation in connection with first star formation research. We advise caution on the Computational Observatory Project. While the idea is sound, it represents a substantial effort to implement and maintain, with unclear payoff at the present time. Any such effort should be coordinated with the Theory Virtual Observatory Working Group of the International Virtual Observatory Association (IVOA).

Division of Materials and Life Sciences

16. *Computational Condensed Matter Science Group.* The research activities of this Group are well focused on fundamental studies of next generation device technologies. In this context, the scientific level of their activity is quite high. In their analysis of problems for present and near future Si technology, the identification of structural and electronic properties of divacancy in Si is a good example of the need for large-scale electronic structure calculations. The actual calculation was performed with their newly developed highly parallel real-space density functional theory (RSDFT) code. Another output of technological and scientific importance is their new theory of Schottky barriers. Their search for new properties and phenomena in the construction elements for future nano-scale device technologies has focused on carbon nanotubes (CNT). Some examples are: the interaction between CNT and different substrates, such as Si, Al and Ca; bias voltage dependence of nano-capacitors; and magnetic properties of edge states of CNT. The Group collaborates with the Computational Life Science Group on the analysis of electronic and atomic processes in biological systems. Actual calculations are mostly done using the QM/MM method implemented in the CPMD package, originally developed by Parrinello's community. Professor Mauro Boero was originally a member of Parrinello's group and is still contributing to the development of CPMD. The analysis of the proton diffusion mechanism in Cytochrome c Oxidase is a very important contribution. The Group's development of highly parallel computer codes for density functional theory (DFT) based electronic structure calculations jointly with the Division of HPC Systems is a very important example of collaboration between computational science and computer science, and will promote highly parallel computing within the Japanese community, which is crucial for the National Next Generation Supercomputer Project. The planned attempt to run a 10,000 atom DFT calculation on a 2048-node partition of PACS-CS would be a landmark in the field.

17. *Computational Life Science Group*. This Group studies the function of biological molecules (protein and RNA) using simulations. Several different simulation techniques are used and the most promising one is developed in collaboration with the Condensed Matter Group. The use of DFT is new in Biology and the CCS is therefore well placed to become a leader thanks to the unique interactions that exist within the Center. Simulations are used to study a wide variety of molecules (e.g. from ribozymes to cox). This large diversity of techniques and of applications is potentially useful to tackle the difficult issue of biological function, but great care should be taken to avoid over-dispersion that will hamper research and reduce its international impact.

18. *Quantum Many-Body Systems Group*. In contrast to the Computational Condensed Matter Science Group, this Group mostly treats excited states. Their targets are both electronic and nuclear systems. The central methods used are the time-dependent (TD) Schrödinger equation and time-dependent density functional theory (TD-DFT). The dynamics of excited states is an important subject in materials science and also biological science. This Group has been collaborating with the Computational Condensed Matter Science Group and the Division of HPC Systems. The finite

difference technique to solve the Kohn-Sham equation was first adopted for TD-DFT in this Group. In future, collaboration with Computational Life Science Group will also be fruitful. The Group's work on excitations in both electronic and nuclear systems using the same methodological framework (particularly for TD-DFT) is particularly interesting. Materials under laser irradiation are studied using both TD-DFT and the TD-Schrödinger equation. Both treat tunneling ionization and re-scattering. The former approach can treat larger systems, while the latter may be more accurate. A new approach based on the TD-Schrödinger equation has been developed to study the quantum dynamics of few-body reaction processes for nucleons, such as anti-proton helium production. The TD-DFT method has been applied both to electronic excitations in organic molecules and nuclear excitations in atomic nuclei. From the study of optical response, useful information on the shape of molecules and nuclei can be obtained. Particularly for nuclear systems, the search for a more accurate functional in DFT is now in progress internationally. This Group is contributing to the Universal Nuclear Energy Density Functional (UNEDF) project in the framework of SciDAC in the USA. Based on quantum chemical calculations for clusters, fundamental electronic properties of high-T_c cuprates have been studied. The importance of strong electron-phonon coupling was emphasized.

Division of Global Environment and Biological Sciences

19. *Global Environmental Science Group*. Atmospheric science was a good choice for the first topic of the Global Environmental Science Group, because it has a long history of numerical simulation using high-end computers. As meteorological data and prediction are widely used in society, the JMA/GPV data archive project jointly with the Computational Intelligence Group targets an infrastructural development with wide impact, and is a nice collaboration between computer scientists and atmospheric scientists. The Group has started to develop NICAM (non-hydrostatic global atmospheric model) for petaflops computer systems, and is going to emphasise the down-scaling of global warming to focus on its regional effects. This area is relatively weak in Japan and should be expanded. Considering the scale of the task, the existing Group is very small. Collaboration with other groups, additional supporting staff and postdoctoral fellows will be necessary if it is to be successful.

20. *Biological Science Group*. This Group, although small and recent, is very active, with several publications in high impact journals. It is interested in the evolution of unicellular eukaryotes, especially in the inference of the global eukaryotic phylogeny. A large part of the activity is devoted to data gathering (i.e., to isolate, describe and culture organisms, and finally obtain DNA sequences), which is of prime importance to remain competitive. They are well aware of the difficulties of phylogenetic inference and have initiated collaboration with computational researchers to address one of them (separate model). The result, an MPI version of Tree-Puzzle, is also efficient at avoiding the frequent problem of being trapped in local minima and will certainly be used worldwide. This collaboration is quite promising, but would strongly benefit from an additional professor to help bridge the gap between these distant research fields.

Division of High Performance Computing Systems

21. This Division is led by Professor Taisuke Boku. The Review Committee heard several presentations about its projects, including: Scalable High-Performance and Cluster Network (Boku); Grid-Data Farm, Gfarm, Distributed File System (Tatebe); FFT and Classical Gram-Schmidt (Takahashi); and Power Aware Computing (Kohuta). The Division occupies a central role and is performing excellent research in strategic areas. There are two themes cutting across the projects: one is the focus on applied projects that have an immediate impact on computational science elsewhere in CCS; the other is energy efficient technologies. Both research directions are well chosen and have tangible results.

22. The Gfarm project is addressing the challenges that we will face managing large datasets in the near future, for example, the petascale data expected from the LHC at CERN. It is a commodity-based distributed file system that can be mounted from all cluster nodes and clients, supports fault tolerance and avoids unnecessary duplication. Gfarm is already integrated at CCS, for example, with FIRST. Gfarm has also demonstrated outstanding performance and won an award at SC06, where a total of 52.0 GB/s was measured for parallel read. Gfarm is an open source development, and CCS should be encouraged to continue developing it, since, with Lustre becoming a potentially proprietary product, there is no capable open source alternative out there.

23. The work on FFTE is very impressive. FFTE is a package for complex, mixed-radix, parallel FFTs that supports SSE2/SSE3 instructions, portability, and obtains high performance. The work on Classical Gram-Schmidt (CGS) for the real-space DFT application is a beautiful example of how advanced numerical algorithm work can be integrated with large-scale application codes. This type of success is only possible in an integrated environment such as CCS. By implementing CGS as an iterative scheme with the “twice is mostly enough” Daniel iteration, and then re-writing the orthogonalization as level 3 BLAS with recursive blocking, the team got near theoretical peak speed on a Xeon processor. This was a very impressive accomplishment that was obtained by using clever numerical techniques that are not widely known or appreciated. Professor Takahashi has done some very fine numerical work here.

24. The second theme of the Division focuses on energy efficient technologies. There has been very little work done in power efficient networking, and the research presented on multi-link aggregation with RAID-like ideas is innovative and creative. Similarly, using DVFS (dynamic voltage frequency scaling) for designing a power scaling system on a cluster, and implementing energy aware optimization, is very clever. CCS should continue to work in these areas. In particular, the work in energy efficiency and computer architecture has the potential for far-reaching impact in years to come.

Division of Computational Informatics

25. *Computational Intelligence Group*. This Group is developing tools for integrating heterogeneous databases and various information sources, data mining to discover

knowledge and rules from massive data, and XML-related technologies to deal with various data on the Internet. It is one of the strongest database research groups in Japan. The Group is playing a substantial role, in collaboration with the Global Environmental Science Group, to establish and maintain the GPV/JMA meteorological archive and web site. This is well organized, with sophisticated XML technologies, to integrate a number of data sources using meteorological ontology and is becoming very popular among the research community. The Group's paper on the integrated use of meteorological data was accepted by the IEEE e-Science Conference. The collaboration between the two Groups is a good example of the interdisciplinary work fostered by the CCS. The StreamSpinner stream processing system that the Computational Intelligence Group is developing is a powerful tool for integrating heterogeneous data sources, particularly event driven sensor systems. In future, complex sensor stream mining is likely to be an interesting computationally heavy application domain for the CCS, and an opportunity to cooperate more closely with the Division of HPC Systems.

26. *Computational Media Group*. Professor Ohta's Group aims to integrate sensing, network, grid and visualisation technologies to augment the human senses. It has high international standing and a strong publication record. In addition, the Group has implemented several real systems which show its excellent capabilities – especially the real-time 3D free viewpoint video system and the mixed reality see-through vision system are outstanding. Since the Group joined CCS rather recently, it has not yet exploited the computational power available in CCS. However, its ambitious challenge, 'Real-World Computational Informatics', to process vast amounts of data sensed from human society and the environment so that they can be readily assimilated by humans will very likely create a new demand for huge computational resources. The Division of HPC Systems could therefore make a substantial contribution to this exciting research target newly defined by the Computational Media Group and the project could become an important core component of the CCS activities.

Major CCS Projects

PACS-CS

27. PACS-CS is a custom-designed HPC system for a few specific applications: QCD, nano-material science, astrophysics, environment and biology. The selected target applications required more bandwidth to memory and bisection bandwidth for the interconnect than could be obtained with an entirely commodity cluster. Therefore, following the very successful model of CP-PACS, the team under the direction of Professor Taisuke Boku built an MPP-like system based on commodity technology that was better balanced for the above applications, which are characterized by requiring nearest-neighbour and collective communications. They built a 3D hyper-crossbar interconnect using commodity technology, such as L2 switches with 10 to 20 ports. The resulting 3D configuration is designed for up to 4096 nodes with 16 CPUs

per dimension. The system is configured with 2GB memory per node and, physically, is built in 59 racks, using 550 kW of power.

28. The PACS-CS project was a successful and timely development to provide an appropriate computing resource to support the expanded range of projects in the Center, bridging to the next University supercomputer. The CCS team under the direction of Professor Boku did an excellent job in leveraging the existing expertise and building a more capable technology than was commercially available at the time. Thus, CCS is one of the few centers worldwide that has the expertise to build custom-developed systems that can live up to the expectations of production science users. The speed at which the PACS-CS system became available and the low rate of failures and problems are quite remarkable for a custom-built system. PACS-CS has already benefited applications on 256 and 512-node partitions. The benefit should be even greater when the system is configured to run applications on up to 2048 nodes. The Center is now well-positioned to do breakthrough science on this system with its specific plan for simulating QCD at the physical point and a 10,000 atom density functional theory simulation.

FIRST

29. The FIRST project is unique in the world and typifies the Center’s ability to use applications needs to drive innovative computer architecture and systems research, which in turn delivers scientific results at the forefront of the applications field. This has enabled the GRAPE technology to be deployed effectively for a wider range of frontier astrophysics applications. GRAPE is a hardware accelerator for direct force summation gravitational N-body calculations. GRAPE can be used for direct numerical simulation of collisional N-body systems (e.g., star clusters), or as a hardware accelerator for the PP part of a P^3M collisionless N-body simulation (e.g., dark matter clustering). The integration of GRAPE with a high performance host cluster, better suited to fluid dynamical calculations, enables new types of large-scale coupled fluid/N-body systems to be attacked.

30. The Computational Astrophysics Group is to be congratulated for developing the first such hybrid system, and is encouraged to seek out applications where this architecture can best be exploited and showcased; for example, the formation of star clusters which begin as fluid dynamical systems, pass through a coupled fluid/N-body phase, and end as purely N-body systems. CCS is uniquely positioned by virtue of FIRST to break new ground here.

ILDG/JLDG

31. This is another example of a successful collaboration between the computer scientists and particle physicists at the CCS over the past five years. They have taken the national lead in the international ILDG project to build a global grid of regional data grids for storing and sharing lattice QCD data, as well as developing in parallel the Japanese component grid infrastructure, JLDG, to support the lattice QCD community within Japan. The key features of these projects are their close attention to

user needs, the early introduction into use by the community, and the rapid growth in use across the world.

32. Key to the success of these projects was the definition, and international acceptance, of a standard XML-based mark-up language (QCDml) for the lattice QCD data, developed by the ILDG Metadata Working Group under the leadership of Professor Tomoteru Yoshié. Members of CCS have also played a significant part in the ILDG Middleware Working Group which is defining the global architecture and the interfaces between the various components.

33. The JLDG is being rolled out across five sites in Japan using Gfarm for file sharing. The ILDG project is an on-going development to keep abreast of improvements in grid technologies and to enhance functionality, but has already proved itself to be one of the most successful data grids world-wide and is gaining recognition as such. It is highly likely that it will be quickly accepted by the entire lattice field theory community as the standard data management environment. Credit for this goes to the close collaboration between users and developers from the outset, in which the CCS team has been one of the best examples. The functionality of the middleware is rather generic and opportunities exist for its application in other fields where community-wide agreement on metadata can be achieved.

Interdisciplinary Activities and Collaborations

34. The Review Committee found convincing evidence that the CCS is succeeding in its mission to promote an interdisciplinary approach to computational science research. Collaborative activities and opportunities were highlighted in all the presentations to the Committee, indicating that all members of the CCS attach importance to it. The success of the original CP-PACS collaboration between computer scientists and particle physicists is a powerful example to other projects and continues to influence strategies. The FIRST Project is a recent successful example of this model. Even with the rapid expansion into a broader range of research areas since 2004, it will be a challenge for the CCS to find projects in which two (or more) disciplines have a sufficiently important shared goal, at the same time, to take precedence over their single-discipline efforts and then to nurture them through to successful completion. Potentially exciting opportunities exist for all research groups, on various timescales, although the Biological Sciences and Computational Media Groups are the furthest from the existing competencies of the Division of HPC Systems, whose work has had the most unifying effect within the Center to date. Both this Division and the Division of Computational Informatics hold the keys to future major interdisciplinary projects. Although well placed in its current mix of projects and expertise, and with an enviable track record, the Center will need continually to seek mechanisms and incentives for interdisciplinary working.

35. The Review Committee was impressed by the strong links which the CCS has grown with the University's Graduate School, which embeds it in a much larger research

community, helping its research groups achieve critical mass, and providing it with a vital source of PhD students. The Committee also strongly endorsed the planned dual degree programme in computational science, which will expand the Center's education and training role. Although it will be a challenge to develop a curriculum which is attractive to both students and the wider faculty within the University, the reward will be a stronger computational community and a source of appropriately trained new researchers. These will also benefit the University and the future employers of its students. The Center's links with University researchers will also be broadened through its provision of the University's supercomputing service and this will be beneficial provided the service meets users' needs and does not impose too great a strain on the resources of the Center.

36. The CCS has a variety of external collaborations typical for internationally competitive research in computational science. In addition, it has played a prominent role in the successful ILDG global infrastructure project. Consequently, today, the CCS has a high international profile. As infrastructure and support for computational science, as well as the research projects themselves, become increasingly multinational, CCS will face greater competition for international leadership and so will need to participate in a broader range of international projects if it is to continue to be recognised as the center of excellence for interdisciplinary computational science research in Japan. This strategy will also help to strengthen sub-critical research groups.

Success of the Center in Achieving its Mission

37. The CCS, in continuing the work of the previous Center for Computational Physics, has established over a period of 15 years a highly successful collaboration between computer scientists and physicists that has built a series of world-class computers and used them to produce results at the forefront of computational physics. This is an example to other areas of computational science and collaborative working has been well integrated into the culture of the Center. The Review Committee strongly supports the Center's vision for computational science and the promotion of its interdisciplinary approach should remain the core objective of CCS.

38. The reorganisation of the Center in 2004 has been successfully accomplished with a big expansion in the number of faculty and the establishment of new research groups with a breadth of applications that are at the forefronts of their respective fields. Research themes have been chosen that will be of growing importance in the 21st century. Some of these new groups are sub-critical in size, however, and they will need to build strong collaborative links outside the Center, or be targeted for new hires. Even so, the Review Committee was very impressed by the scientific accomplishments of the new groups in so short a time and the extent to which they have integrated with the interdisciplinary culture of the Center. Most of the research outputs were judged to be world class, or to have the potential to become so. In

several areas, novel approaches were being adopted which are exciting and hold the prospect for major breakthroughs.

39. The PACS-CS and FIRST projects are outstanding examples of major computer developments which were driven by science needs and are now providing unique capabilities for leading-edge research. This ability to design, build and operate novel and competitive computer hardware targeted at challenging applications is at the very heart of the Center and must be preserved. It derives from a mixture of environment and talent that is not replicated anywhere else and gives the Center its competitive edge.

Future Vision and Strategy

40. The Review Committee strongly endorses the Center's overall future vision and strategy to be a worldwide center of excellence in interdisciplinary computational science, taking a global and long-term view. It recognizes that, after the major changes of the past three years, a period of consolidation is needed and it supports the plans of all the research groups and Divisions, while encouraging the Center to expand its participation in international projects, its visitor programme and its support for virtual organisations. Within the current scope of its research groups, the CCS should maintain a portfolio of exciting research projects at the forefront of computational and computer science, and a talented faculty engaged in this research. It should continue to provide an environment which encourages interdisciplinarity and in which shared goals can emerge and be nurtured.

41. The CCS should maintain an internationally competitive computing facility to support its research projects so that it can focus appropriate computational resources to ensure they have maximum impact in their fields. The Center's strategy to take responsibility for the University's research-oriented supercomputing service will help it to maintain control of appropriate systems and the planned "T2K" acquisition will achieve this for the next 3-4 years. However, the CCS has too few technical staff to support the University service as well as use of its computers by a growing range of projects which may no longer always include CCS research staff. The University must transfer support staff, or the funding for them, to the CCS along with responsibility for its supercomputer service. The CCS will need to introduce a more service-oriented user support mechanism for this wider range of users, because the Center will risk its excellent reputation if the service is poor, but it must be careful to ensure that this service activity does not dilute the Center's research focus.

42. The CCS should continue to develop new computing technologies in the context of the most demanding applications needs and have the capability to bring these technologies into use to make breakthroughs in specific areas of computational science. Specialised hardware will continue to play a vital role in specific fields. The expertise that CCS has acquired in embedding accelerators in general-purpose clusters will have wider applicability in the future and should be actively pursued

following-on from FIRST. The Review Committee recognises that the development of novel architectures is high risk, but it is also high reward and, although the funding required for any particular development cannot be guaranteed, the Center's successful track record in this area is a major strength which is rare worldwide and should not be lost. There is a rapidly evolving range of new technologies which might be deployed to make significant advances in computational science and the Center is particularly well-positioned to evaluate and exploit these.

43. The educational activities of the CCS will play a vital role in integrating it with the wider University community and with society. The expanded graduate programme will be time-consuming for staff to develop and deliver, and some resistance to the interdisciplinary concept may need to be overcome. However, the Center's staff should find this work rewarding and the outreach to students should generate new ideas and new recruits. The building of a broadly based computational science community and the promotion of interdisciplinarity should be seen by the University as a major strategic opportunity which the CCS offers, and which could give it a genuine competitive advantage both in terms of its teaching and research. The University should provide appropriate financial support for this expanded teaching programme.

44. The CCS interdisciplinary research focus, educational activities and pioneering of new technologies should enable it to play a major role in supporting initiatives such as the National Next Generation Supercomputer ("Petaflops") Project. The Center should seek strategic partnerships with national projects such as this, which exploit its complementary expertise in developing new application areas in computational science and its capabilities as an icebreaker for innovative algorithms, high performance software and new hardware technologies.

Other Recommendations

45. Overall, the Review Committee was impressed by the achievements of the CCS, its interdisciplinary culture, its leadership and its vision for the future. The Committee believes that the collective expertise in the Center and the research environment it has built are unique worldwide. Already this has delivered world-leading scientific results, giving the Center a high international standing. The expansion of the Center in 2004 has given it the potential for a broader impact in both research and education, and has targeted areas that will be of strategic importance in the 21st century. The Review Committee commends the Center to the University and recommends that the University acknowledges its strategic value, supports the execution of its mission, and encourages the wider academic community to engage with its activities.