Science and Technology Facilities Council

Computing Evaluation Panel Report

Executive Summary

Computing is an integral, mission critical, part of all of the projects that are undertaken by Programmes Directorate. If funding for computing is cut it will lead directly to a reduction in the science delivery. Computing provision is essential to the whole spectrum of STFC activities. There is a lot of computing going on that is related to experiment/observations more indirectly.

To maximize science output from the STFC investment in projects, observations and theory simulations it is imperative that the right computing hardware is available to undertake the science. Once the science drivers are identified the most effective e-infrastructure, machines, software and people, to carry out the work has to be put in place. STFC's funding strategy must prepare now for the computing it will require in five years' time and recognise that investment in this area is not a one-off but requires cyclic capital investment to maintain world leadership.

STFC funded computing is evolving continuously with different science questions requiring very diverse computing architectures to be efficient and competitive. STFC needs to guarantee a world-leading computing infrastructure that is heterogeneous in both interconnect (weak, as well as strong, dependence on interconnect) and node-architecture (simple x86, multi/many-core, GPU, accelerators) to cater for STFC's science drivers and maximise physics output (see illustration below).

The world has evolved with traditional boundaries becoming blurred and increasingly emerging computing architectures and the clear needs of the scientific community to require a diverse mix (see the following illustration).

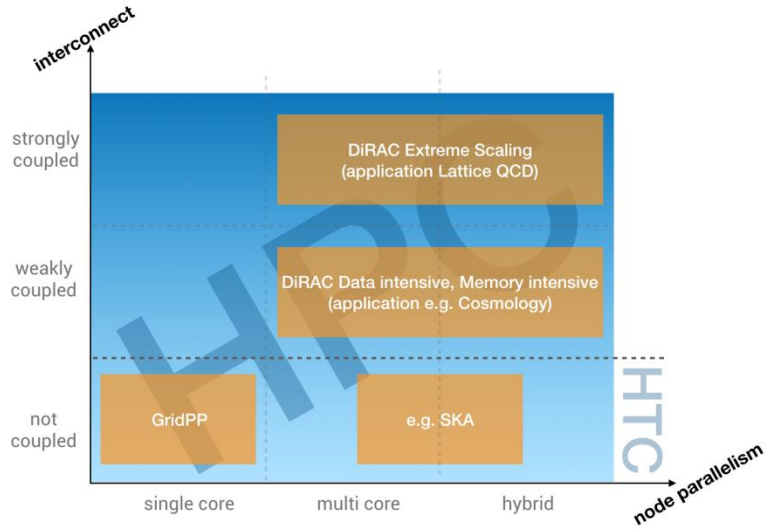

When agreeing that the UK should join an experiment/observation project, STFC should recognise this as a commitment to provide a certain fraction of the computing needed as

computing cannot be separated from the experiment. Computing needs and costs should be reviewed by PPRP at the same time as the project. Specialist reviewers should help with the peer review process to ensure sensible computing plans are in place. STFC could have specialists who help collaborations prepare the computing part of their application.

While computing for experiment and observation cannot be separated from the investment in facilities themselves there is a wide spectrum of activities (such as DiRAC) which are largely independent of capital investment in experimental/observational facilities.

Funding for both DiRAC and GridPP is vital to continue and enhance the current programme. The relationship between hardware provision and ongoing maintenance cannot be ignored and provision is needed to cover on-going costs like electricity. This provision should also include the necessary skilled staff together with a clear career path underpinned by training.

Many “computing people” move to the private sector, which is to be welcomed, however a strong career path within the universities and National Laboratories would ensure that sufficient were retained in scientific research.

Whilst hard to estimate, the computing on university-based systems is significant. It should be noted that in-house capacity is likely to decrease as university funding is stretched, leading to more pressure on regional and national systems, and possible increase in funding requests through Consolidated Grants.

Areas within the STFC programme that have common cutting edge problems should have a mechanism to work together. In a world of restricted funds coordination between different frontier science areas, and other Councils, is important.

The panel agreed with the 2017 UK Industrial Strategy Green Paper which stated “As well as physical and digital infrastructure, we need to make sure that we also have in place an effective data infrastructure. This means the right elements for an economy in which open data drives growth, efficiency and innovation.” The need for data curation in all of STFC’s computing is increasing and should be planned for. Currently it appears that some science communities are being expected to curate data for which they neither have the infrastructure nor the support to successfully carry out. A more strategic approach would be beneficial.

There should be increased coordination across the breadth of STFC’s scientific computing and the recent establishment of the Scientific e-Infrastructure Governance and Oversight Group is to be welcomed efforts should be made to coordinate with Computing Advisory Panel (CAP) and other stakeholders to ensure various groups do not duplicate work and can complement each other’s activities.

New funding streams, such as GCRF or the Newton fund, offer opportunities to bolster the investment in STFC science related computing.

Efforts should be made to increase the diversity (gender/ethnicity) of the community that works in the area of scientific computing.

A. Introduction

1. Computing underpins every aspect of the frontier science programme funded by Programmes Directorate. It has close links with the computing that support activities elsewhere in STFC, which in turn is part of the broader UK computing ecosystem. There is cross Council activity to work towards a National e-

Infrastructure (NeI) under UK Research and Innovation (UKRI). The UK National Infrastructure is shown in diagram 1.

Diagram 1: UK National Infrastructure

2. The computing for STFC's PD programme is largely provided by DiRAC and GridPP with some additional provision coming from universities.

Diagram 2: STFC's e-Infrastructure

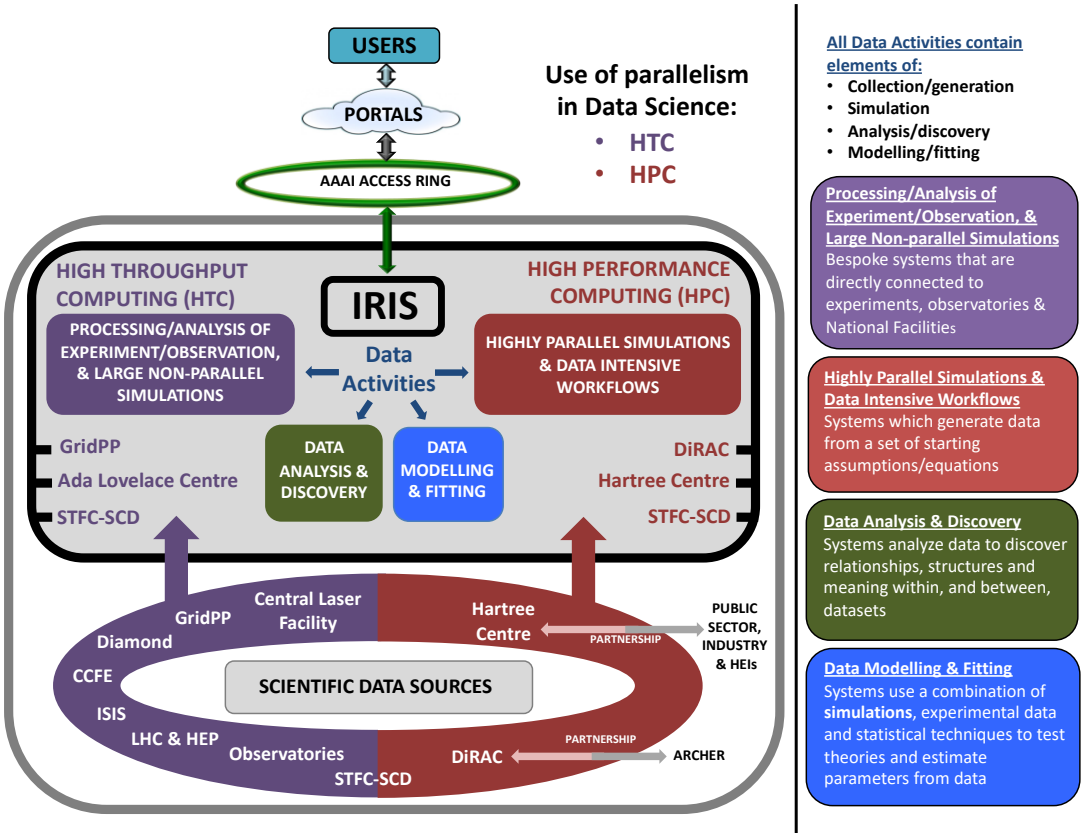

3. To provide the evaluation Panel with an overview of the computing requirements within STFC frontier science pro formas were requested from the UK representatives of Advanced LIGO, ALMA, CASU, CTA, DiRAC, DUNE, ELT, e-Merlin, EUCLID, GAIA, HEP-GridPP, IPPP, LSST, LUX ZEPLIN, SKA, UKT0 (now IRIS) and WFAU.

B. Key science drivers being addressed by the current discipline

4. Computing is driven by the science needs of the frontier science programme of Accelerators, Astronomy, Nuclear Physics, Particle Astrophysics, Particle Physics, Cosmology and Solar Science. The most suitable effective and efficient compute is then sought to solve the defined challenges. The discipline requires a diverse heterogeneous ecosystem, which can accommodate both science output and innovation. Both of these components are essential if STFC is to continue to maximise the scientific return from its experiments.

5. Whilst computing is driven by the scientific needs of the community it serves, the technology itself can also be used to push forward scientific boundaries. Artificial intelligence (AI), machine learning, quantum computing and progress in conventional computing technology can all help advance the scientific expectations and ambitions of the frontier science community. Making the most of these - and future technological advances - is essential as the computing requirement is set to dramatically increase in the near-future as many of the high-profile projects that STFC has invested come on line and produce greatly increased quantities of data, for example LHC Run 3 in 2021 with a step change again for Run 4 HL-LHC in 2024 and SKA around 2026.

Table 1: *The data rates, galaxies and financial and human resources in current and future galaxy surveys.* The date indicates `first light'.

Survey	Data volume per night/day	Galaxies	Cost	Scientists	DiRAC flop/s
DES (2012-)	1 TeraB	~300 Million	~$40M	~400	~0.03 PF
DESI (2019-)	40 GigaB	~35 Million	~$70M	~600	~0.03 PF
LSST (2021-)	15 TeraB	~1 Billion	~$1.0B	~1000	~0.6PF
Euclid (2021-)	850 GigaB	~1 Billion	~$1.5B	~1500	~0.6PF
SKA (2020-)	1 PetaB	~1 Billion	~$1.3B	~1000	~3PF

Table 2: DiRAC computing resources for Theory simulations, capital injections and number of projects and scientists involved. Computing Power, Projects and Capital are UK only. The number of scientists is the number of registered DiRAC users so the majorities are from the UK.

	Year	Computing Power (PFlop/s)	Projects	Capital	Scientists
DiRAC 1	2009	~1	8 (science consortia)	~£12M	393
DiRAC 2	2011/12	1.6	26	~£15M	500
DiRAC 2.5x	2018	2	34	~£9M	~1200
DiRAC 2.5y	2019/20	4.5	~40	~19.5M	~1350
DiRAC 3	2020--	20-30	~50	~£71M	~1550

Table 3: Totals for all four LHC experiments. The table includes the 2016 values and the 2027 projections. 2027 is the end of the first year of data taking for HL-LHC.
LHC (ie the 4 LHC experiment) is servicing the needs of about 10,000 physicists)

Year	2016	2027
Raw data added per year	50 PB	600 PB
Derived data added per year	80 PB	900 PB

Recommendation (i) - Astronomy and Particle Physics experimental compute have entered the era of Big Data. High Performance Computing for the Theory community has entered the multi-Peta Flop era. STFC must ensure it can meet the exploitation demands of the projects that it funds and provide a sustained funding stream for computing that guarantees international competitiveness.

6. The most efficient way to maximise the scientific return is to decide the science that is required and then establish the right e-infrastructure to deliver it as part of the project design. This ensures that the appropriate hardware capable of best delivering the science is in place.

7. Currently the computing trend in experiment is towards convergence with both data and simulation requirements beginning to carry similar weight (such as within the requirements of ATLAS). This should be an intrinsic part of any project, collaborators should work with computing providers and experts to ensure that adequate provision is made, which should be costed in at the proposal stage.

8. PPRP should review the computing needs at the same time as the associated project. Computing experts should be invited to join the Panel to review the quality of the computing plan contained within proposals. In addition, computing experts could be consulted during the planning of experiments. This would ensure that funding requests reviewed by PPRP are as accurate as possible with regard to the computing requirements and financial request. One possibility would be to establish network of experts who can be called upon that can help with the preparation of the computing aspect of applications from the community (this could be built in as a Research Infrastructure and Software Engineer (RISE) Fellowships. RISE fellowship requirement or a short-term 6-month placement within another project, similar to Long Term Attachment at CERN). This type of internal outreach (similar to the Institute for Advanced Architectures and Algorithms (IAA) type model in the US) would make the most of the expertise at our disposal, which could be used to walk applicants through the process, and create a mechanism for increasingly accurate proposals reaching PPRP.

Recommendation (ii) Computing requirements for new projects should be costed at the conception stage. Project proposals to PPRP should include a costed computing plan as part of their application. The computing plans should be reviewed by computing experts and PPRP should ensure that sufficient funds are allocated for computing.

9. DiRAC was established to provide distributed High Performance Computing (HPC) services largely to the STFC theory community. HPC-based modelling is an essential tool for the exploitation and interpretation of observational and experimental data generated by astronomy and particle physics facilities supported by STFC as this technology allows scientists to test their theories and run simulations from the data gathered in experiments.

10. The GridPP Collaboration was formed as part of the UK's contribution to the Large Hadron Collider experiments and represents the UK in the Worldwide LHC Computing Grid (WLCG). The UK has on-going commitments to CERN and the wider international collaborations, which with the growing demands of HL-LHC in

2025 cannot be met with a flat-cash. As the GridPP community has grown and evolved many user communities from a wide range of disciplines, for example SKA, have taken advantage of the computing resources and skills offered by GridPP. At present it is not evident that this approach is sustainable funding for the large-scale computing needs of SKA and other projects.

C. Overview of other related review boards and their conclusions and recommendations

11. In 2014-2015, STFC conducted a strategic review of its current computing capability and future requirements of its scientific and industrial user communities. These communities include STFC-funded researchers; users of STFC National Facilities; other science communities supported by computing facilities or services hosted by STFC; and clients of the Hartree Centre. The Review covered hardware, data services, software, cyber security, networks and computing professionals. The Review recommended actions to meet STFC's future computing needs and improve efficiency.

12. Computing was also addressed in the Balance of Programmes (BOP1) exercise, the STFC e- Infrastructure Strategy Review and the 2017 UK Industrial Strategy Green Paper. The previous reviews reached similar conclusions, stating that funding, skills and collaboration all needed to be increased. It was felt that computing can no longer continue to be seen as an 'add on' and cut from experiments to ease financial pressure and that more investment in this area would lead to a step change.

13. STFC has actively tried to increase coordination of its scientific computing efforts as recommended in the STFC e-infrastructure (eI) strategy with the establishment of Scientific e-Infrastructure Governance and Oversight group. This has met with some success and whilst the group is still developing its role and remit, it will benefit STFC as it will encourage a holistic approach and greater discussion of key issues across the entire STFC eI portfolio.

14. The Computing Advisory Panel (CAP) approaches from a user perspective advises the STFC Science Board on the strategy for, and management of, provision for computing resources (including data handling, data storage, software and hard provisions, skills, and developments in high performance and high throughput computing) in support of programmes either funded or delivered by STFC.

Recommendation (iii) – The Panel endorsed the findings of BOP1 and believed STFC should enhance collaboration and cooperation within the computing discipline.

D. Do the projects/experiments currently sit in the correct scientific discipline?

15. STFC Programmes Directorate computing activities are funded from different sources through project and consolidated grants. The Computing Programme holds the budget for High Performance Computing (HPC) provision through DiRAC and PRACE, as well as EGI membership, and small project awards where possible. A significant part of the programme is strongly linked with the Particle Physics community High Throughput Computing (HTC) facility, GridPP, and is funded through the Particle Physics Programme.

16. DiRAC and GridPP are funded through different mechanisms. GridPP, with its historical association with the Particle Physics programme, goes through panel review on a regular basis (the last review was GridPP5) which enables grants to be allocated and a certain level of stability to exist within the facility. DiRAC has a less stable funding model where development finance is required from large capital government funding. This reliance on spikey finance causes a perpetual challenge together with funding the resource to support and develop the computing facilities. In the past this has caused difficulties in maintaining international leadership.

Recommendation (iv) – STFC needs to improve the funding mechanism for Theory simulation (DiRAC) towards providing for reliable upgrade cycles and covering running cost and operations budget.

17. The UK has an extremely strong HPC community and DiRAC is a powerful computing facility allowing the STFC science community to pursue cutting-edge research on a broad range of topics, from simulating the entire evolution of the universe, from the big bang to the present, to modelling the fundamental structure of matter. DiRAC is both an academic-led and an academic-run facility and has been successful in designing systems specifically to meet the different HPC needs within our scientific community.

18. The technological developments of DiRAC are driven by the science and lead to cutting edge developments that lever additional support from industry. It has been extremely successful in working closely with companies to provide innovative computing architectures which are co-designed with the science need. This has led to positive partnerships with vendors from which other computing infrastructure projects have benefited. The facility is well positioned to take advantage of commercial innovation and higher performance advances as soon as they are available, which benefits the broader science communities.

19. DiRAC and GridPP both lead to International leadership because of the collaborative nature of the ventures (leveraging time elsewhere e.g. lattice QCD). The science driven work that is undertaken would not be possible in the UK without these resources.

20. Both DiRAC and GridPP will need to think differently in the future and find innovative approaches as funding will be unlikely to increase at anything like the same rate as required. New developments in AI and machine learning, as well as new directions in computer hardware, should be fully investigated and adopted where appropriate. Whilst they are both currently embedded in STFC, the boundaries may change in the future as we move to a more integrated ecosystem with UKRI and the NeI.

21. The development of IRIS (formerly known as UKT0) is a cooperative venture that includes both DiRAC and GridPP and brings STFC computing together under a collaborative umbrella (as shown in diagram 2) supporting both the National Laboratory and Programmes Directorate activities. It has recently received a £16M investment from BEIS, which benefits all.

Recommendation (v) – STFC should position itself to take advantage of new technologies. STFC should utilise any opportunities that are generated by this

development. IRIS is a cooperative and shared infrastructure convergence activity evolving and moving towards NeI, this could provide opportunities for STFC.

E. Leadership (world and UK)

22. STFC should aim to deliver computing to a level that is highly regarded internationally. Investment in this key area can help propel the UK into high profile leadership roles. People, depth of expertise and availability of computing capacity are all components of a healthy computing programme. There is an intimate relationship between having the best technology, infrastructure and interoperability with securing scientific leadership.

23. In addition, provision of good computing means that in the future there is also the potential for the UK to obtain work package leadership for the design activities within science projects. Retaining an element of spare capacity in the computing support for the frontier science disciplines is helpful in gaining access to new science projects. This is an essential element is enabling STFC science to buy into a new areas and international collaborations by demonstrating expertise in these areas and undertaking early stage simulations.

24. Both DiRAC and GridPP have leadership with regard to their technical activities. GridPP has enabled the UK to hold positions of leadership within WLCG and broader international collaborations and this mechanism can, and should, be used to influence and ensure efficient and innovative resource allocation procedures are embedded. Adequate Research Software Engineers/Research Infrastructure Engineers (RISES), 5-years analogous to EPSRC scheme, are an essential component of this provision.

25. The UKRI has initiated convergence towards a NeI with roadmap development to produce a coherent infrastructure provision. IRIS is playing a major role in enabling this service.

26. The UK is still at the front in terms of developing areas such as simulation, modelling and the size and speed of network navigation. Robust computing provision is required for the UK to maintain its international presence in these areas, which will continue to be dictated by the science. Frontier science computing providers should look to collaborate with providers in other areas, particularly with regard to data storage and curation. If experiments fail to curate the data that projects generate, they could obtain reputational damage and lose leadership on work packages**.**

Recommendation (vi) –STFC should cooperate and foster enhanced links within UKRI to ensure it has the best opportunities to remain at the forefront of data curation.

F. Overlaps, synergies and gaps with the rest of the PPAN programme (including Education & Training, Global Challenge Research Funding etc)

27. New funding opportunities exist within the UKRI environment. As a science-enabler computing is very well placed to benefit from funding streams such as GCRF, the Strategic Priority Fund and the Newton fund, all of which offer real opportunities to obtain additional funding.

28. In a world of restricted funds, coordination between different frontier science areas is important and IRIS could provide a key role going forward. Further engagement across STFC communities will be essential to meet the data demand. This cooperation can help with the challenge of fully utilising staff effort, sharing best practice and avoiding unnecessary duplication of effort.

29. The returned pro formas demonstrated that there are a number of potential overlaps and activity undertaken where synergies can be found. Multi-messenger observations (such as used in gravitational wave observations) demonstrate this, as do LSST and EUCLID. Community meetings focused on drawing out these synergies could provide a common mechanism for independent groups to solve problems together. This would promote the sharing of best practice and reduce duplication of work being undertaken for example in the area of image analysis. Inter-disciplinarily work can also help establish developing areas such as AI/machine learning/quantum computing for the wider UK community. Overlaps may also exist in the wider STFC landscape, for example Jasmin (which is a NERC-funded facility operated by STFC at RAL) .

Recommendation (vii) – STFC should be stronger in encouraging projects with common computing requirements to work together to foster relations, enhance the sharing of best practice and reduce unnecessary duplication of effort. **Where appropriate, IRIS could be used to help foster relations.**

30. ALMA has used methodology for analysing image data which is applicable to other areas of science and GAIA run workshops to share best practice. Greater collaboration could lead to earlier engagement with industry to development projects and could lead to a national data curation system that is compatible internationally. The UK does not want to build similar but independent infrastructures for SKA and LSST, it would be better to morph an infrastructure to cater for several experiments.

31. The pro formas demonstrate that, on the whole, experiments view their computing requirements very differently. We would strongly encourage all experiments to talk to IRIS to assess their computing needs before submission of SoI rather than leaving computing costs to be thought about at a later date and provided through a separate budget. This change of culture is essential to ensure that the programme is sustainable. There is some progress in this area, e.g. LUX-ZEPLIN used experts from GridPP to help cost their computing needs. If experiments highlight their computing requirements at the SOI stage it will help ensure that the exploitation stage of the experiment can be forecast/financed.

32. GridPP is fully utilised and is a respected part CERN/LHC. It is a mature infrastructure that is beginning to be requested by other particle physics experiments, e.g. DUNE and can reduce the need for duplication of work or overlapping systems. GridPP has been able to lever support from universities in terms of staff time and electricity, this is an increasingly challenging area as running costs are increasing.

G. Science, skills and technology

33. GridPP and DiRAC both develop expert staff members who add to the UK skill base. For some, this transferable learning translates into employment in industry. STFC support for training of students and researchers with computing

infrastructure and software skills (RSE/RIE) is important, however the lack of a clear career path within collaborations and academia means that PDRAs often feel that they do not have visibility and recognition within the science collaborations, which can be detrimental for the research infrastructures

34. Engagement with industry is essential, and the importance of funding CDTs should not be underestimated. The eight well-coordinated CDTs in Data Intensive Science for early career researchers helps to cater for the demand for skilled computational and data analysts in UK industry and academia. The first of the annual CDT summer schools was held at UCL in July 2018[1] and was attended by around 120 students with speakers from 5 companies and academic guest lecturers. This positive approach to training brings together different disciplines and crosses departments as well as industry engagement, which enhances the innovation agenda across all areas of Programmes Directorate and the UK.

Recommendation (viii) - STFC should continue to invest more in CDTs for computational or data-intensive research in our science programme.

35. DiRAC stands out as an example of interacting well with industry. It is actively engaged with industry to help deliver technology innovations that benefit not just its science users but also the wider UK community.

36. A clear career path needs to be put in place to try and retain more of these people in scientific research. RISE Fellowships with an emphasis on computing rather than science (not necessarily post-doctoral but with some relevant post graduate experience) could be a way to promote this, possibly running for five years with a commitment to providing employment at the end, a model which has proved very successful for EPSRC[2]and MRC[3].

Recommendation (ix) – early career fellowships should be established to provide a clear career path for Research Infrastructure & Software engineers.

37. As highlighted in the UK NeI Systems and Service survey 2017 (which will soon be available on the currently migrating HPC-SIG website that hosts the report) the overall gender balance within the discipline is made up of approximately 22% female. Many of these women are in the early stages of their career and many are working in the health & life science areas.

38. The UK does not have a good gender balance in this area and some countries fare better (Source: *UNESCO Science Report: towards 2030*, see the chapter entitled Is the gender gap narrowing in science and engineering). STFC is encouraging better Equality and Diversity balance and should be proactive in trying to improve the situation within this discipline. Improving gender balance on committees, active mentorship and encouraging females at an early age to be interested in the discipline could all help to encourage greater levels of female representation.

[1] https://indico.cern.ch/event/702529/

[2] https://epsrc.ukri.org/funding/calls/rsefellowships/

[3] https://mrc.ukri.org/skills-careers/fellowships/skills-development-fellowships/

39. Further support (perhaps similar to the Dorothy Hodgkin or Daphne Jackson fellowships for people who return to work after a career break) would help increase the number of skilled computational researchers working within this discipline[4] and STFC could also target training to people who re-enter the discipline. This could be in terms of RISE fellowships or enhancing the existing returner schemes in ERF.

Recommendation (x) – STFC should strongly encourage increased female participation within the discipline. Money should be ring fenced for fellowships directly targeting scientists, of either gender, who wish to return to the discipline after an absence from working within the field.

40. **The Research Software Engineer Association[5] now organise an annual conference which focuses exclusively on the issues that affect people who write and use software in research. Organised by the UKRSE, the conference hosts speakers from industry and academia, and a series of talks and workshops from RSEs that investigate everything from the latest tools and engineering practices, to the way the RSE community works. The RSE Conference has attracted contributions from across the world, which will are showcased during an international networking session**. The RSE conference is a positive development that provides good networking and messaging. Additional conferences which promote positive opportunities and messages should be encouraged.

41. RISE help enhance this service provision which is demonstrated by DiRACs investment in RSEs. The facility recognises the importance of centrally investing in developing workable codes, creating models that work elsewhere and allowing collaborations to successfully engage in the facility. More broadly RISE funding and workshops would help as would engagement with Software Sustainability Institute.

H. How the discipline is likely to evolve in the future, including impact of international programme

42. There is a need for local, regional, national and international systems. The interplay between these systems is essential. The future of the discipline is unpredictable, however, there is certain to be an increase in required computing power and in data generation, especially with regard to international collaborations. As many changes will be taking place in the computing discipline, it is important to consider if there is sufficient investment in the R&D aspect and deployment of the underpinning technology. It is no longer sufficient to purchase more of the same compute/disc; innovation is essential in a range of technologies and techniques, such as AI, machine learning and quantum computing to stay at the cutting edge and have the capacity to manage the data challenge. In the area

[4] Suggested names for such a scheme could be Jocelyn Bell, Eleanor Burbidge, Caroline Herschel, Mary Somerville, Annie Russell Maunder

[5] https://rse.ac.uk/

of computing this cannot be done in isolation and is important that the UK maintains its links to international as well as domestic partners.

43. Data is increasing exponentially and curation has to be planned for. Some computing centres are currently going through PPRP and STFC needs to ensure that this is done in a holistic way rather than on a project-by-project basis so that the computing needs and people are in place to efficiently maximize the return from the full portfolio of experiments. Failure to curate the generated data could cause reputational damage, leading to less observational time allocation in the future, and therefore losing leadership.

44. It is hard to assess the contribution from university-based systems, but it is nonetheless important to recognise the significance of this contribution and the resulting reduction in science capability if this contribution is cut. It should be recognised that in-house capacity is likely to reduce as university finances are squeezed, forcing increased pressure on regional and national systems, which could lead to additional computing finance being requested on consolidated grants.

45. Continued investment in innovation of the technology is essential and is an area where small amounts of money could make a large difference to the efficiency level and help generate leadership roles for the UK.

Recommendation (xi) – The computing discipline is evolving fast. Attention must be given to hardware developments (quantum computing as well as AI methods and new directions in conventional hardware)

L. Impact of reduced/increased funding of this discipline

46. Computing is an integral, mission critical, part of all of the projects that are undertaken by Programmes Directorate. If the computing funding is cut on the experiments it will lead directly to a reduction in the science delivery conversely an increase in computing investment will lead to an increase in the science that can be obtained from the funded experiments. However, computing is ~£6.6M resource and currently accounts for around 5% of the frontier science resource spend, this represents a barely sustainable resource level and cuts here could have a disproportionally adverse impact on the experiments and lead directly to a reduction in the science delivery. Many projects have made commitments and there would be reputational damage if not sustained. In addition to resource, £3M is allocated for computing capital.

47. The UK does not currently have enough computing funding or staff for the amount of science that is already being carried out. A 10% cut in computing resource is equivalent to funding for 5-6 PDRAs per year across the entire PD programme so might have a negligible impact when spread across the other programmes but would represent a dramatic loss in a stretched computing programme. In addition it would directly reduce the science exploitation that can be done by those other disciplines.

48. A 10% increase would allow the creation of 5-6 RISE fellows, which would have a dramatic impact. Investment in short-term placements could also be part of this and linking with industry could leverage further resource into the discipline.

Recommendation (xii) – STFC should explore developing and enhancing links with industry through fellowships. This would help in areas such as skills development.

49. If the UK is unable to sufficiently cater for computing needs for experiment, observation and theory simulation it would risk damaging the UKs international reputation. This can lead to losing leadership positions within collaborations and the ability to lever future time on the experiment.

50. Both GridPP and DiRAC are essential to the UK standing in this area. Whilst increased synergy will help and can generate efficiency savings, the need for more staff effort is inescapable and underpins the desire for extra funding on RISE and technical data support. This would help enormously in improving efficiency and sweating assets across the programme.

51. The funding model for compute is patchy not covering all areas of science and not stable, with uncertain capital injections making long term planning difficult and does not match the sustained requirements across the science fields and nor the evolving technology computing needs. The over-demand for DiRAC facility from the community is clearly documented and inevitably leads to loss of science in spite of its recognised efficiency and value for money. The facility is open to innovation and technology foresight and engages with relevant and diverse groups to enhance its level of service provision. The facility caters for specialised needs in many areas and is skilful in the way it matches the physics to the resource level.

Recommendation (xiii) – A decrease in funding in this area would significantly damage all areas of frontier science as all are dependent on computing to make scientific advances. An increase in funding should be used to increase the number of Fellows funded within the discipline.

J. Impact of reduced/increased funding of other disciplines on this one

52. As computing covers all disciplines and is integral to them changes to the other disciplines are felt directly by this one. If spending on the experiments within Programmes Directorate is reduced, it could be argued that the computing costs associated with these experiments could also be reduced however the fragility of the discipline means that this would further damage computing and maintaining computing funding in this scenario would help ameliorate the science reduction

53. An increase in funding elsewhere in the programme should not be at the expense of computing and will necessitate an increase in computing investment.

54. It is not possible to separate the different components of science with theory and experiment both essential and simulation and modelling inherent to both. Consequently cuts in Consolidated Grants and fellowships directly impact

computing facilities. Computing is often wrapped up in PDRA posts for other frontier science disciplines and altering either side of this relationship would have a significant negative impact leaving the UK unable to sufficiently analyse the data form projects it is has invested in. This would jeopardise the international standing of the UK with a consequent loss leadership positions and the ability to leverage future access.

55. Post-Brexit, the picture is uncertain but UK researchers could lose access to funding sources such as European Research Collaboration, H2020 funding and access to facilities such as PRACE (between 2014 and 2017 the frontier science disciplines received over 70,000,000 euros in ERF support). If this proves to be the case it would become increasingly important to fund new fellowships to protect the discipline and to nurture other international collaborations and partnerships.

Recommendation (xiv) - Funding should be evaluated in light of Brexit. Sustainable mechanisms should be put in place to enable the UK to maintain healthy international relationships.

J. Health and vibrancy of the community (Experience pipeline/ Leadership/Scientific Diversity)

56. The funding squeeze has produced innovation in the past but this has more than reached its limit and there is no longer the capacity to do the vibrant research and technology development needed to stay at the forefront of this field.

57. Computing has a diverse ecosystem of both people and architecture and needs to be supported by a diverse population of quality employees.

58. STFC has a strong track-record in this area but the Panel noted that past experience has shown that constrained funding often means that computing is stripped out by grants panels. A good mechanism for sustainably and adequately funding computing facilities is essential as is properly assessing the requirements in the projects and identifying the best funding mechanism.

59. It is essential that the best science is funded irrespective of location. Computing in the UK is flexible with strong host locations spread throughout the country. The hosting organisations can be universities, or at STFC campus sites.

60. There is currently insufficient infrastructure and skills for data curation. This is an essential part of the computing infrastructure and needs to be funded accordingly. This area will become even more important to STFC as several astronomy experiments and particle physics projects begin to generate dramatically increased quantities of data in the next few years. Whilst JISC is currently carrying out some work in this area, it is not on the required scale to cope with the future influx of data from the frontier science fields.

61. There are significant challenges to meet the Pb/PFlop/s era that is on the horizon. This will require preparation not just on the capital side, but also investment in people to ensure a consistent pipeline of experience is in place. Whilst there are many good career opportunities for computational and data scientists, demand for skilled people is not being met in industry or academia. As noted earlier this is further aggravated by the lack of a career path in academia

as without sufficient recognition or reward it cannot compete with industry and other science disciplines. Good opportunities and careers are essential for those who want to follow the technical path.

62. GridPP and DiRAC should undertake more work to raise their public profile and share and celebrate the discipline's successes. Whilst many people are aware of the LHC, very few outside the science community are aware of the computing network that generates the data. The two facilities, due to the nature of their work, have different opportunities to raise their public profile. DiRAC has more opportunity to raise the profile of the facility through engagement with industry. Both should make use of female role models where possible and any fellowship could include a PE component.

Recommendation (xv) – STFC must fully utilise its potential for raising visibility through Public outreach. More should be done to highlight the computing contribution behind many of the success stories associated with STFC science.

63. Whilst DiRAC maintained their own records of the facilities impact, information on this area was hard to obtain comparable data for both DiRAC and GridPP using Research Fish as DiRAC is not included as a facility. STFC could improve the publication data that is available by adding DiRAC to Research Fish.

Recommendation (xvi) - STFC must interact with Research Fish to ensure that publication data for GridPP and DiRAC can be sufficiently recorded and easily accessed.

K. Financial forecast

64. In a time when greater importance is being placed on leadership by the government, it is essential to emphasise how the computing underpins big international collaborations which provide science leadership. If the computing aspect of these collaborations is compromised, so is the science.

65. The financial landscape is changing with the emergence of UKRI and the prospect of a NeI as this develops the Research Councils will need to engage with the aim of improving computing delivery and efficiency for their respective research fields. Within this context the current planning blight and uncertainty of funding of DiRAC3 has the potential to seriously jeopardize the frontier science programme. The potential lack of a sufficient facility if it is not upgraded is a serious threat.

66. Computing is well aligned with the industrial picture as demonstrated by the CDTs and could be a source of funding.

K. Review recommendations (Ranking definitions can be found in Annex 1)

67. The panel gave DiRAC a ranking of:

Exploitation - **g3** (the highest in this category)- A project with high strategic importance in the STFC programme, which has received substantial investment. We would expect to see it adequately funded via grants after peer review.

Impact and Engagement - **i5** (the highest in this category)- - Very exciting impact already under IP management or a close working partnership or exchange with non-academic partners is already in place**.**

68. The panel gave GridPP a ranking of:

Exploitation - **g3** (the highest in this category)- - A project with high strategic importance in the STFC programme, which has received substantial investment. We would expect to see it adequately funded via grants after peer review

Impact and Engagement - **i4 -** Very exciting opportunities proposed, with some first connections made.

L. Concluding remarks

69. Computing is an integral, mission critical, part of all of the projects that are undertaken by Programmes Directorate. If the computing funding is cut on the computing it will lead directly to a reduction in the frontier science delivery. An increase in computing investment will lead to an increase in the science that can be obtained from the funded experiments. You cannot separate the computing from the experiments.

70. To maximize science output from their experiments STFC need to carry out exploitation in the most efficient way. This means that it is imperative that the right hardware is available to undertake the science. STFC have to establish the science needs first and then ensure they put in place the most effective e-infrastructure, machines and people, to carry out the work. STFC need to start to prepare now for the computing science it will require in five years' time

Annex 1

Ranking Scoresheet for Programme Evaluations 2018/19

During the 2017/18 Programme Evaluations, projects/experiments/facilities within each discipline will be ranked. The ranking criteria will cover scientific excellence, exploitation within grants, and impact/industrial engagement. The exercise will look at all funded projects/experiments/facilities and ensure each is considered at whatever its stage of the exploitation cycle.

The panels will consider the merits or otherwise of supporting areas currently receiving STFC investment. This will include consideration of international engagement and subscriptions.

The ranking criteria will be largely based on that previously used by STFC, namely **α** rankings for projects/experiments and “g” rankings for science exploitation themes within grants as used in the last Programmatic Review. In addition a new “i” ranking will be introduced to cover evaluation of impact for the economy and society.

The Panel will be asked to consider the strategic value of the projects/experiments/ facilities that submitted proformas and how highly aligned they are to the mission of STFC. Consideration should also be given to the international standing and the potential for leadership of the area under review. Additional value, such as synergies within the STFC frontier science disciplines (Particle Physics, Astronomy, Nuclear Physics, Particle Astrophysics, Computing, Accelerators) programme should also be taken into account.

The Panel will be asked to score each of the projects/experiments/facilities on the following criteria and submitted 2 days before the meeting.

The Panel member should complete section 1 and 4 below for each proforma. A marking should be given for either section 2 or 3 dependent on which is most appropriate.

The below wording is generic for the six evaluations and may be slightly modified to suit the specific requirements of the individual reviews.

1. **What is the life cycle stage of the Project/Experiment/Facility?**

Early / Developing / Mature

2. **Scientific Excellence of Project/Proposal**

α5 - Highly innovative and very likely to result in seminal changes in knowledge.
α4 - Likely to substantially advance the subject.
α3 - Likely to make an important contribution to the subject.
α2 - Competent, worthy science.
α1 - Interesting science but outcomes considered doubtful.
β - Poor quality, flawed or unlikely to deliver meaningful or interesting results.

3. **Exploitation**

Projects in the science exploitation phase are funded via grant panels. Three categories are defined, intended as strategic guidance to the peer review carried out by grant panels. Please consider the value of exploitation when the area under evaluation reaches maturity.

g3 - A project with high strategic importance in the STFC programme, which has received substantial investment. We would expect to see it adequately funded via grants after peer review
g2 - A project with high potential for excellent science which should be considered via peer review
g1 - A project which is not well matched to the STFC programme, we would be surprised if it were to receive funding via the grants panel.

4. **Impact and Engagement**

Please consider if there is important impact within industry and/or wider society that STFC should be looking to exploit and that will otherwise not happen elsewhere.

i5 - Very exciting impact already under IP management or a close working partnership or exchange with non-academic partners is already in place.
i4 - Very exciting opportunities proposed, with some first connections made.
i3 - Interesting opportunities suggested but needs significant further work.
i2 - Little opportunity, although some could evolve in near future.
i1 - Little opportunity and unlikely to develop significantly in near future.
i0 - No apparent opportunities at all.