Astrophysics theory, high performance computing and astronomical data handling

Ad-hoc consultation panel: report to PPAN

Carlos S. Frenk, Chair

13 May 2008

Contents

Astrophysics theory, high performance computing and astronomical data handling
Executive summary

1. The outcome of the Programmatic Review (PR) is deeply disappointing for theoretical research in particle physics and astrophysics and threatens the viability of some of the most productive, highest impact areas of mainstream science in the UK. The community deplores the lack of astrophysics theorists on high-level STFC committees, including PPAN, which should be redressed at the first opportunity.

2. HPC-based modelling is an essential tool for the exploitation of observational and experimental facilities in astronomy and particle physics. The UK has an extremely strong HPC community which includes genuine world leaders in QCD, cosmology, MHD, and star and planet formation. These points were forcefully made in the community input.

3. Following a decision in 2006 by the DTI to make approximately £13M available for HPC hardware, PPARC issued a call for proposals. After thorough peer review, the PPRP ranked 9 consortia as "highly fundable' (world-leading, or potentially world-leading, research) and 'fundable' (internationally competitive) and awarded £12.3M from the capital budget and £4.2M from the non-capital budget. However, this £4.2M for 'HPC operations' is ranked as 'lower priority' in the PR.

4. PPAN's decision to rank HPC operations as lower priority was driven by tactical considerations in the expectation that universities would be willing to cover operations costs in exchange for hardware. The community – and this panel – object in the strongest possible terms to this tactic and to the resulting separation of the hardware from operations, both for reasons of principle and fairness, as well as for the practical implications. PPAN's decision to place HPC operations in the 'lower priority category' is risky and could have unintended consequences. It endangers the £13M offered by DIUS which will be released only if a realistic business plan for operations can be formulated. Such a loss would be catastrophic for UK science. It also sends completely the wrong message about the sustainability of HPC to Vice-chancellors and peer review bodies.

5. It is illogical and unfair to lump into a single package projects which STFC's own rigorous peer review process has given a range of rankings, from 'world-leading,' through 'internationally competitive,' to 'unfundable'. By choosing not to evaluate individual HPC projects using the same criteria which it used to rank other projects and facilities, PPAN dismissed the outcome of the peer review process. This approach runs counter to PPAN's stated and commendable aim of supporting world-leading components of STFC's programme. Applying PPAN's own published criteria, this panel finds that the 5 HPC projects ranked by PPRP in their highest category (Exeter, MIRACLE, UKMHD, UKQCD (I) and VIRGO) would fall in the 'high' or 'medium/high' priority groups in the PR, with the remaining 4 (HORIZON, UKQCD (II), COSMOS and Leicester) falling in the 'medium/high' or 'medium/low' priority groups. The panel recommends that the 'HPC operations' package be disassembled and the individual projects ranked in the PR according to their scientific merit and the published grading criteria.

6. Although fundamentally disagreeing with PPAN's attitude to HPC, the panel reluctantly decided to explore whether university funds might help to rescue UK HPC. After consulting with consortia PIs, the panel believes that a viable package for HPC operations can be put together at a total cost to STFC of £1.7M over 3 years, with universities contributing much of the remaining £2.5M. Since PPRP allocated £0.7M less than the expected DIUS hardware award, this balance could be used to enhance hardware allocations in exchange for 'in kind' university contributions. STFC should enter immediately into negotiations with consortia to try to leverage operations funding from their universities. These negotiations should take into account the PPRP ranking of the individual HPC projects and STFC should be prepared to contribute to the costs of operations in a way that reflects

those rankings. The amount sought from STFC represents 1% of the sum allocated to projects and facilities in the 'high' and 'medium/high' priority lists in the PR. Yet, HPC projects would have been ranked in these two lists had PPAN applied to them the same criteria that they applied to the rest.

7. AstroGrid was set up in 2002 with the aim of designing, developing and deploying a Virtual Observatory (VO) framework to promote astronomical discoveries through Grid-enabled data mining of current and future astronomical data centres across the world. Through an investment so far of £7.7M, AstroGrid has secured UK leadership in the International Virtual Observatory Alliance (IVOA). AstroGrid software was released on 1/04/08 at the UK National Astronomy Meeting. The services it provides are expected to be at the heart of future astronomical research at a cost which is a small fraction of the huge investments needed to generate the databases in the first place.

8. AstroGrid has been scrutinised by the PPRP and the PPARC Science Committee and, in March 2007, PPARC Council approved the third phase of the project over the period Jan 2008 to Dec 2009, at a cost of £2.72M. However, in the PR AstroGrid has been ranked near the bottom of the priority list on the grounds of (i) delayed delivery and (ii) limited impact. AstroGrid strongly contests these criticisms and points out that delays have been minor and, in any case, agreed by their Oversight Committee. Impact cannot yet be gauged since the software has just been released.

9. There is strong support in the community for AstroGrid. Although the panel ranks some of the HPC projects requiring operational support at a higher priority than AstroGrid, it believes that, in view of the continuing strategic importance of this area of e-Science and the large targetted investment to date, to discontinue core funding of AstroGrid before its usefulness has been established would be a false economy. At some point in 2009 it will be possible to assess quantitatively the impact of an operational VO service for UK astronomers. The funding approved by PPARC Council has yet to be released so AstroGrid is operating with bridging funds provided by the universities involved. While the panel believes that there is room for STFC to negotiate with AstroGrid on the committed costs, this anomalous situation is unacceptable. STFC must ascertain the minimum funding level required for AstroGrid to be in a position to be reviewed in 2009. Such funding could come from a levy on ground- and space-based projects which will use the AstroGrid infrastructure. In the longer term, the future of AstroGrid will rely on European collaborations through ESO and ESA.

10. The panel recognises the important contribution of the CASU and the WFAU to the analysis of data from current (UKIRT) and upcoming (VISTA, GAIA) facilities and missions. It recommends that future funding of the service elements of the CASU and the WFAU should be identified directly with specific project or facility grants, and that future funding of the research elements should be competitively assessed against other STFC-supported research. This would effectively remove CASU/WFAU from the PR, although some interim funding may be needed to maintain their capabilities.

Astrophysics theory, high performance computing and astronomical data handling

1. Introduction

1. The panel was asked to cover the areas of astronomy theory, high performance computing (HPC) (including astronomy and particle physics) and astronomical data handling in the Programmatic Review (PR). While we believe that the PR fell short of accepted standards of peer review, we welcome this more broadly based consultative exercise and are participating in it on the understanding that our report will be made public. The low ranking awarded to theoretical activities such as HPC was greeted with disbelief by the community, both nationally and internationally, forcefully expressed through the formal input into the PR invited by STFC and by spontaneous comments from abroad. We trust that the results of this exercise and the atmosphere of inclusivity and openness with which it is being conducted will help dispel the widespread disenchantment with the review process that followed the publication of the PR.

2. Nevertheless, the panel feels that the narrow remit of each subpanel is over restrictive. Our panel has been asked to review the relative rankings of three projects which have all been classed at 'lower priority.' Two of these projects, HPC and AstroGrid, have important ramifications across the rest of the programme.

3. Most of the comments received from the community under the rubric of 'astronomy theory' mentioned also HPC. Although there is no specific 'astronomy theory' project or facility included in the PR, the decision by PPAN to impose a 25% cut on the grants line has a direct impact on theoretical research and so we include theory explicitly as a separate section in this report. The astronomical data handling item covers AstroGrid as well as the Cambridge Astronomical Survey Unit (CASU) and the Wide Field Astronomy Unit (WAFU) in Edinburgh. For each of the three topics considered in this report, we first summarize the input from the community and then add the panel's own commentary. The community input on HPC is summarized separately for astronomy and particle physics. The composition of the panel and procedural details are given at Appendix A. Appendix B contains a summary of recent reports on HPC computing and provides useful background to this topic. Finally, Appendix C gives an overview of recent UK achievements in theoretical astrophysics and solar system physics and Appendix D a brief overview of HPC requirements in particle and astroparticle physics.

2. Astrophysics theory

2.1 Summary of community input: astrophysics theory

4. The comments from the community on astrophysics theory are virtually unanimous and centre on four main points:

   (a) UK astrophysics theory is of outstanding quality. The theory community exercises world-leadership across broad swathes of the subject.

   (b) There is deep concern that the cuts to the grants line have a disproportionately large effect on the theory community, as grants are its main, and in some fields only, source of funding.

   (c) The continued strength of the UK astrophysics theory community is utterly reliant on access to world-class high-performance computing resources; hence theory and HPC cannot be decoupled.

   (d) The lack of theorists on high-level STFC committees, including PPAN, alienates the sizeable UK theory community and vitiates decisions that are potentially catastrophic for a vital part of the UK's science programme.

2.2 Panel commentary: astrophysics theory

5. UK theoretical astrophysics is widely recognized as one of the very best areas, not just of UK physics, but of UK science as a whole. Astrophysics theory is a major contributor to the UK's international standing in physics. This point has been repeatedly made in reviews of the international perception of UK physics, most recently in 2005, and there is a wealth of evidence to support it. For example, the UK contributes 5 of the top 15 most cited astronomers of all time; 3 of them are theorists. The rest of the UK's effort in astronomy benefits from this international leadership, since theoretical work provides the framework for observing programmes and the planning of new facilities.

6. Despite its outstanding reputation, freely recognized by the rest of the astronomical community, theory has suffered very badly in the review process. Theory uses essentially only two resources, HPC and postdocs, and both have been severely hit by the PR and the associated 25% cut in grants. HPC operations is rated at the lowest priority, and the grant cut is not even up for discussion. Although significantly less expensive than most other areas of STFC science, theory is probably more vulnerable to cuts in skilled manpower (postdocs) and, because of its high level and reputation, particularly at risk of losing talented people to other countries.

7. There is a strong feeling that the poor treatment of astrophysics theory reflects a consistent blind spot in STFC perception, starting from its early mission statements. Many people have pointed out the weakness of the STFC peer review structure, including its lack of genuine input from international experts at a high level. PPAN contains no theorist and was patently more severe on theory than any other area. This is a potentially disastrous way to manage a strategically vital part of the UK science programme. As a minimum, STFC must ensure that there is proper representation of astrophysics theorists in its high-level structure.

3. High Performance Computing (HPC)

3.1 Background

8. We briefly discuss the events that led to the formulation of STFC's current policy on HPC and present the background to PPAN's ranking of HPC operations as 'low priority.' Further details of the reports discussed in this section are given in Appendix B.

9. Following an in-depth review in 2003, the Webber report, PPARC's Science Committee formally recognized the vital importance of HPC for astronomy and particle physics. The Webber report concluded that "HPC is now a crucial investigative tool across all areas of astronomy and large parts of particle physics" and that "UK HPC work is of high scientific merit and world-leading in many areas across the PPARC remit". Commenting on the frustration expressed by the community regarding the difficulty of obtaining maintenance and staff costs, the report says: "The Panel is unanimous that PPARC must provide the necessary supplementary and supporting costs". The Panel estimated that "continuation of the current level of support would require annual investment of around £6.0M, of which approximately £1M would provide for maintenance and running costs".

10. In 2005, the International Review of Research using HPC in the UK published a report in which it recognized that "computation has now become essential for the advancement of research across science and engineering and HPC enables the tackling of problems and opportunities that cannot be approached in any other way. Of the PPARC consortia selected for detailed review, the report noted that UKQCD and VIRGO were playing a leading role in setting international standards.

11. A coherent strategy for sustaining and enhancing the leading role of HPC in the UK was developed in the 2005 Kenway report. The Kenway report was accepted by Science Committee in 2006 and remains the policy of the STFC. It recommended that "all aspects of HPC provision are reviewed

together" and, in particular that "[s]ystem and software support should be treated as an integral part of funding HPC." It revised the estimated cost of the "HPC facilities and services required to address the immediate science goals of the PPARC research community" to £6.5M pa.

12. STFC accepted the Kenway working group recommendations and issued the first of a series of three-yearly calls for proposals for HPC resources - computer hardware, system and software support and operating costs, with a closing date in April 2007. Proposals were invited for up to five years for staff support and for up to three years for equipment and consumables. The budget for capital items was approximately £13M earmarked in 2006 by the then Department of Trade and Industry from the previous Comprehensive Spending Review, with non-capital costs coming from a 'new-HPC' planning line. Eleven proposals requesting £34.5M were received. They were extensively reviewed by PPRP augmented by additional experts.

13. The PPRP ranked the proposals as "highly fundable on the basis that the programme of work was internationally competitive and had aspects of world-leading research" (Exeter, MIRACLE, UKMHD, UKQCD (I), VIRGO); "highly fundable on the basis that the programme of work has the potential to be world-leading or produce a step change", though of "higher risk because they were speculative or because they were proposed by new groups" (HORIZON, UKQCD (II)); "fundable on the basis that the programmes were internationally competitive" (COSMOS, Leicester); and "unfundable" (UKAFF, UKQCD (III)). The total recommended cost of the highly fundable/fundable proposals consisted of £12.3M from the capital budget and £4.2M from the non-capital budget.

14. The PPRP report was provided to PPAN as input for the PR. The outcome is that the non-capital expenditure of £4.2M for 'HPC Operations' has been ranked as 'low priority'.

3.2 Summary of community input: HPC in astrophysics

15. Forty-four sets of comments were received from a broad cross-section of the community, including observers. They unanimously condemn the ranking of HPC operations as 'lower priority' in the PR and express a sense of despair at the perceived damage that will be inflicted upon the highly rated theoretical astrophysics research in the UK, much of which relies on HPC. PPAN is urged to reconsider its ranking of HPC operations.

16. Every single comment emphasizes that the UK has an extremely strong HPC community which includes genuine world leaders in the development and applications of HPC in a number of areas such as cosmology, galaxy formation, MHD, star and planet formation. The importance and world standing of HPC-based research has been widely recognized in a number of recent reviews, including two international ones (the 2005 'Review on International Perceptions of UK Research in Physics and Astronomy' and the 2005 'International Review of Research using HPC in the UK'). Some of the most successful and highest profile projects in UK astronomy involve HPC. It is difficult to understand how a country with such a rich theoretical community can allow its one theory facility to rank so low compared to the list of observational facilities.

17. HPC is vital for the exploitation of many astronomical facilities because it provides essential tools for data interpretation and modelling. An example is the 2-degree field galaxy redshift survey which provided evidence for the existence of dark energy and discovered baryon acoustic oscillations, relying heavily on cosmological simulations by the Virgo consortium. This is recognized even in STFC's own roadmap which, under 'Strategic Priority' states: "HPC is a crucial investigative tool for the theory community across all areas of astronomy (including astrophysics, cosmology and solar system science) and large parts of particle physics. Most of the research activity of the UK theory community relies on access to HPC facilities." More generally, within the STFC remit, HPC makes the largest contribution to training skilled individuals with expertise relevant to the UK knowledge and skills-based economy.

18. Ranking HPC operations as lower priority dismisses the carefully crafted strategy for HPC which STFC developed through two major reviews (the Webber and Kenway reports) which were accepted by PPARC's Science Committee. Even if this ranking is tactical, it sends completely the wrong message.

19. The HPC community has done remarkably well in the past at securing funding from HEFCE initiatives (e.g. JIF, JREI, and SRIF) at no cost to PPARC. For STFC to disown operations costs is incomprehensible. The descoping of the HEFCE initiatives now makes it difficult to obtain HPC funding from RCUK. STFC must restore an HPC funding line to cover hardware, running costs and support staff (as expected under fEC).

20. It makes no sense to lump together HPC operations for a variety of projects which have been extensively peer reviewed and given different rankings by PPRP. PPAN should unravel the package and award each project the priority determined by expert peer review. How many of the highly ranked projects in the PR have the same level of international recognition as the top rated HPC projects?

21. It would be disastrous to lose the £13M provisionally allocated for hardware by DIUS for lack of support for operations. Placing capital costs in jeopardy has the potential to destroy theoretical astrophysics research in the UK utterly. The damage would be irreversible.

22. The lack of expertise in theoretical astrophysics and HPC on high level STFC committees, particularly PPAN, is alarming and casts doubt on the reliability of the low ranking awarded to HPC.

3.3 Summary of community input: HPC in particle physics

23. The UKQCD collaboration welcomes the judgement of the PPRP that a significant fraction of the science projects submitted as part of its HPC bid are 'highly fundable' and accordingly a high priority for STFC. The funding recommended by PPRP is £6.8M capital and £1.5M recurrent. The capital is understood to be available from DIUS, but the recurrent costs were included in the PR and awarded low priority by it, despite the PPRP's high rating of UKQCD science and the likely negative impact on the ability to leverage the capital from DIUS. Although UKQCD has not had any formal feedback from STFC, UKQCD understands that its science case was not included in the PR. Certainly, at no stage was UKQCD invited to make a submission to this exercise. UKQCD is concerned about the potential negative consequences of having 'HPC Operations' explicitly listed as low priority. This is unhelpful for at least three reasons: firstly, it lumps together entirely distinct areas of science proposed by distinct communities, which runs counter to the stated aim of the list reflecting science priorities; secondly, it creates a misleading impression that the science cases did not fare well at peer review, which may prejudice the treatment of related bids by subsequent peer review panels; and thirdly, it makes little sense to give different priorities to the capital and recurrent costs when both the Webber and Kenway reports stress the importance of tying them together. So UKQCD proposes that 'HPC Operations' be removed from the priority list.

24. UKQCD accepts that the funds available for recurrent costs are severely limited and is prepared to reduce its recurrent costs and/or find money from elsewhere in order to be able to unlock the capital funding. The approach UKQCD is proposing would offer ownership of the high-value capital items for accepting the risk on recurrent costs. It claims to know of organizations that are prepared to bid to do this for the Blue Gene systems and the Tightly Coupled Cluster. It also expects that most of the universities hosting the small clusters and data stores required for distributed data analysis will be able to manage without operating costs from STFC. UKQCD estimates the absolute minimum recurrent funding required from STFC to be in the region of £0.5M (for QCDOC electricity, and 2FTE of staff for system and software support) and requests the opportunity to discuss these possible solutions with the STFC. While the reduced resources will inevitably raise issues regarding scheduling and throughput, and hence have some impact on the science, UKQCD is confident that it can deliver the scientific objectives detailed in the original proposal.

3.4 Panel commentary: HPC in astrophysics and particle physics

25. Over the past decade, HPC has revolutionized astrophysical theory and modelling and this, in turn, has had a huge impact on the way in which astronomical data are analyzed and interpreted. HPC-based modelling has now become an essential tool for the exploitation of observational facilities dedicated to solar physics, star and planet formation, galaxy evolution and cosmology. For example, mock catalogues which mimic the details of particular observational setups are essential for developing statistical analysis tools, controlling systematics and, crucially, for relating the data to an underlying physical model. Without a solid theoretical and modelling base, it is difficult to imagine that large-scale observational facilities will produce cutting-edge science.

26. In particle physics, HPC plays an equally fundamental role. Computer simulation of QCD has offered the only rigorous and systematically improvable way of extracting, from first principles, predictions of the theory of the strong force. This is important because the confinement of quarks cannot be handled analytically and so limits our ability to confront Standard Model predictions with measurements at the level of precision now possible experimentally. UKQCD's current proposals should permit the computation of a wide range of physically interesting quantities in which all theoretical uncertainties are controlled at the few percent level, matching the experimental precision. As LHC begins to make discoveries, it is very likely that computer simulation will play a crucial role in understanding the emerging new physics and that the HPC requirements for this work will exceed those of QCD.

27. As its Chairman has acknowledged, PPAN recognizes that a number of the HPC proposals reviewed by the PPRP contain world-leading science which have consistently been ranked at the highest priority by various peer review bodies and which PPAN wishes to support. PPAN's decision to rank HPC operations in its lower priority category is purely tactical: PPAN's perceptions is that with £13M from DIUS earmarked for hardware on offer, universities will be prepared to cover in full, or at least contribute to, the operations costs. While it is indeed likely that some external funds can be leveraged, this panel objects in the strongest possible terms to the tactic followed by PPAN both for reasons of principle and fairness, as well as for its practical implications.

28. On the face of it, PPAN's approach to HPC operations is illogical and unfair. Firstly, the 'HPC operations' rubric encompasses projects to which STFC's own rigorous peer review process has given a range of rankings, from 'world-leading,' through 'internationally competitive,' to 'unfundable'. Lumping all these projects together is as illogical and unfair as lumping all ground-based astronomy or all particle physics experiments into a single package. Secondly, PPAN's decision to accept in full the PPRP's recommended funding for astronomy and particle physics projects classified as 'high' and 'medium high' priority in the PR is justified on the grounds that PPAN commendably wishes properly to support its high priority projects. Yet, the criteria used by PPAN to rank other projects and facilities has not been applied to HPC operations, not even to those that PPRP has classed as 'world-leading.' Similarly, PPAN's tactical approach to leveraging funds from universities has not been applied to any other project. Would the same logic not apply, for example, to universities that are awarded grants for large experimental equipment?

29. PPAN's strategy of placing HPC operations in the 'lower priority category' is risky and could have unintended consequences. In the first place it endangers some of the highest profile, internationally recognized, world-leading projects within STFC's remit. In the short term, this strategy risks losing the £13 million offered by DIUS which will only be released if a realistic business plan for operating the hardware can be formulated. It goes without saying that failure to secure the £13 million would be a disaster which would spell the death of HPC in the UK and wipe out at a stroke a large fraction of theory activity in the country. Surely this is not PPAN's intention. Finally PPAN's strategy risks sending the wrong message: in spite of the hardware carrot, would a university vice-chancellor be willing to commit funds to a project deemed by STFC to be low priority? They would certainly be entitled to ask whether such activity is sustainable in the long run.

30. PPAN's strategy for funding HPC operations is not only risky but it is also impractical. Except for Exeter and Leicester, the HPC proposals are by consortia of several universities (a model hailed as a major UK strength by the 2005 'International Review of Research using HPC in the UK'). In the majority of cases, however, the requested hardware will be sited at a single university. The £13M carrot argument will hold little sway with consortia universities which will not be hosting hardware. On the other hand, those universities which will host hardware may find it difficult to justify a subsidy to other universities in spite of the prestige associated with hosting the hardware. This potential for conflict can jeopardize the very existence of the consortia. A further disincentive for universities is the fact that other Research Councils cover the full fEC rate for HPC facilities. Being out of line with other councils, STFC's approach is liable to be questioned by universities.

31. For all the reasons just given this panel strongly believes that PPAN was wrong to lump together the operations of all the HPC projects which had been individually and rigorously peer-reviewed by the PPRP and to assign a low priority to the entire package. In the interests of fairness, each HPC project should have been evaluated individually using the same criteria that were applied to all other projects and facilities in the PR. Based on the published evaluation criteria and scores, this panel has attempted this exercise. Unavoidably, our results are only indicative because the weightings given by PPAN to different criteria have not been released to the panel. We also warn that 4 of the 5 panel members are part of HPC consortia. However, our scores, which were moderated by the non-conflicted panel member, Prof T. Hartquist, follow rather straightforwardly from the comments made in the PPRP peer review report on HPC.

32. We find that of the 5 HPC projects ranked by PPRP in their highest category, UKMHD, UKQCD (I) and VIRGO would have each scored maximum points in all criteria except the fourth (uniqueness of the facility – all our competing nations provide access to HPC facilities), for which we assign a score of 2, and industrial benefit where UKQCD scores 3, VIRGO 2 and UKMHD 1. The other 2 projects in the highest rated PPRP category, Exeter and MIRACLE, are smaller and would have scored 2 rather than 3 in the scientific user base criterion (5) and 1 on industrial benefit. Independently of the unknown weightings, we feel we can safely assume that these 5 projects would have fallen in the 'high' or 'medium/high' priority groups in the PR. For the remaining 4 (HORIZON, UKQCD (II), COSMOS and Leicester), we subtracted one point from the first (strategic importance/uniqueness) and fifth (scientific user base/need) criteria and concluded that these projects would fall somewhere in the 'medium/high' or 'medium/low' priority groups. The panel recommends that these rankings for HPC operations be adopted by PPAN, or alternatively, that PPAN produce their own rankings using their published evaluation criteria and the results of the PPRP peer review process.

33. In spite of our concerns about PPAN's tactics, in view of the limited funding available, the panel feels it cannot oppose PPAN's strategy of trying to leverage funding from universities in exchange for hardware. We recommend that STFC enters into negotiations with the consortia immediately. These negotiations should take into account PPRP's ranking of the individual HPC projects and STFC should be prepared to contribute to the costs of operations in a way that reflects these rankings. As discussed below, the panel believes that a viable, but limiting package for operations can be put together at a cost to STFC substantially below the £4.2M recommended by PPRP.

34. After consulting with representatives of the 9 HPC consortia, the panel has been able to identify, in broad terms, areas where there is some scope for reduction in the PPRP's recommended level of STFC funding for HPC operations. While this panel has neither the authority nor the desire to enter into negotiations with the consortia, we hope that the following remarks will inform STFC and help them approach negotiations with each consortia. In summary, the panel feels that a ballpark figure of £1.7M for operations over 3 years would enable the consortia to deliver most of the science they propose. Some of this money might come from any balance in the amount provisionally allocated by DIUS for hardware (PPRP allocated only £12.3M out of an expected £13M) since, as discussed below, some consortia might be able to use an enhanced hardware allocation in exchange for 'in kind' university contributions towards operations. If STFC is unable to secure enough additional

funds to cover its part of the operations package, funding should be obtained by levying a small percentage from other projects funded in the PR, some of which will benefit from associated HPC activity. A mere 1% of the sum allocated to projects and facilities in the high and medium high priority lists in the PR would finance a realistic contribution to HPC operations. The intellectual case for such a levy is that several HPC projects would have likely been ranked higher than many projects in this list had PPAN applied uniform criteria across the board. PPAN should be under no illusion that such a drastic reduction in operations support will not impact on the quality of the science that will be delivered. It is imperative that negotiations with consortia be conducted immediately. The HPC funding process has already suffered numerous delays and the start date of 1/10/07 anticipated in the call for proposals is well behind us. These delays are causing serious staffing problems and eroding the competitiveness of the world-leading consortia.

A. Highly fundable – world-leading research

In this summary the consortia are grouped into the three categories recommended by PPRP, arranged in alphabetical order within each.

Exeter: Exeter are in a unique position since they applied only for access charges (i.e. operating costs) for an existing supercomputer at Exeter. A possibility might be for STFC to purchase hardware from the Exeter supercomputer for the Astrophysics group in exchange for operating costs. Another possibility might be for Exeter to team up with a consortium which is able to provide operations costs in exchange for an enhanced hardware allocation to that consortium.

MIRACLE: There is scope for negotiation with the main consortium partner, UCL, regarding staffing costs and for trading off operations costs for an enhanced hardware allocation.

UKMHD: Since PPRP recommended hardware at two sites for this consortium, negotiations to reduce operations costs with the host universities are likely to place less of a strain on the consortium than in the single-host case. There is room for negotiation with the proposed hardware hosts, St. Andrews and either Leeds or Warwick.

UKQCD (I): See paragraph 24 above

Virgo: The main consortium partner, Durham University, is prepared to negotiate a reduction of STFC's contribution to operations costs, specifically the machine running costs (£300k), in exchange for support for essential staff and an increased hardware allocation.

B. Highly fundable – with potential for world-leading research

HORIZON: There is scope for negotiating with the main consortium partner, Oxford, in particular for reducing STFC's contribution to operations in exchange for opening access to non-STFC researchers.

UKQCD (II)): See paragraph 25 above

C. Highly fundable – internationally competitive

COSMOS: The lead institution, Cambridge, is prepared to negotiate their operations costs with STFC. In addition, unlike other consortia, there is a prospect that other members of the collaboration could make some contribution to the staffing costs.

Leicester: Some of the running costs would be absorbed by the university, but staffing remains a problem.

35. There is an indication from DIUS that the hardware funds should be used for a 'centrally organized delivery facility.' All consortia have reiterated the point made very strongly in the Kenway report and in the individual consortia bids that there are sound scientific and economic reasons for the consortia model operating a variety of computer architectures. The panel fully concurs with the universally held view that a single machine is not what is needed to deliver the science. This does not, of course, preclude a degree of central organization and coordination of geographically distributed hardware and, indeed, this panel strongly recommends such an approach.

4. AstroGrid

4.1 Background

36. AstroGrid was set up in 2002 with the aim of designing, developing and deploying a Virtual Observatory (VO) framework to promote astronomical discoveries through Grid-enabled data mining of current and future astronomical data centres across the world. AstroGrid products conform to global standards agreed by the International Virtual Observatory Alliance (IVOA). As one of three founder members of this world-wide campaign, AstroGrid has been critical in securing UK leadership in the IVOA. The project currently includes development effort in eight UK institutions (Bristol, Cambridge, Edinburgh, Central Lancashire, Leicester, UCL, Manchester, RAL).

37. AstroGrid software (V2008.1) was released on 1/April/08 at the National Astronomy Meeting. The AstroGrid VODesktop is an IT tool which allows astronomers to find, analyse and visualize data that they can store in a remote VOSpace. This will greatly help astronomers generate publishable science. The delivery of this system was the main aim of the first two phases of AstroGrid. AstroGrid has begun to work directly with a rapidly engaging community to maximize the take-up of these new services.

38. In March 2007, PPARC Council approved the third phase of AstroGrid over the period Jan 2008 to Dec 2009 at the cost of £2.72M (including working allowance and contingency) and bringing the total investment to £10.39M. Much of this money came from earmarked e-Science funding awarded to PPARC in the 2000 and 2002 Spending Reviews (at the same time as GridPP was funded). The main goal of AstroGrid-3 over the 2008-2009 period is to work with the astronomical community to optimise scientific return from the VO by ingesting key data centres into the system and by distributing, maintaining and developing core AstroGrid software products. About 15 people are primarily employed by AstroGrid. The implied ongoing commitment to AstroGrid (as a virtual facility rather than a project) was contingent on a late-2008 review of implementation and community uptake of AstroGrid products, as well as the strategic priority of the programme.

4.2 Summary of community input: AstroGrid

39. Ten positive comments on AstroGrid were received from around the world (e.g. Canada, Europe, Japan, South Africa, USA), emphasising the global standing of the project, and from key researchers across the astronomical wavebands (from X-ray, through optical, to radio). The international inputs clearly identified AstroGrid as the leading international VO project, and judged it 'years ahead' of the main competition in Europe and in the USA. Two major points were common to the responses received: (i) the clear need within the UK community for VO tools, and (ii) the world leadership secured by AstroGrid for the UK within the IVOA. Supporting letters were received from groups active in the IVOA and from astronomy projects planning on distributing their data using AstroGrid tools. It was pointed out that the UK is required to make VISTA data products available in VO format and, at present, AstroGrid is the only way of achieving that. Responses also pointed out the importance of developing e-Science activities like AstroGrid on the pathway to future major facilities such as ALMA, E-ELT and SKA.

4.3 Panel Commentary: AstroGrid

40. Modern astronomy relies increasingly on combining very large databases. The costs of doing this efficiently are very small compared to the huge investments needed to generate the databases in the first place through the design, construction and operation of new space and ground-based astronomical facilities. This situation looks certain to remain true as new facilities like VISTA, ALMA, Planck and GAIA begin to generate data, and astronomers begin to think seriously about the challenges posed by the huge datasets that will be generated by future facilities such as E-ELT

and the SKA. The IVOA is therefore likely to be at the heart of all future astronomical research, and PPARC is to be commended on securing earmarked funds from previous Spending Reviews that have placed UK researchers in world-leadership positions. The AstroGrid team is also to be congratulated on its recent delivery of functioning software to the community, and the downloading of 743 copies of the software in April is evidence that the community is starting to assess its capabilities seriously.

41. The key questions are now whether or not AstroGrid provides competitive tools for astronomical data mining, and whether or not these secure scientific leadership for the UK. Ultimately these questions will be most simply answered by considering evidence on the take-up and use of AstroGrid products by the UK and international communities. As AstroGrid software has only just been released to the community, the success of the investment in AstroGrid can only be judged after the planned late-2008 major review of AstroGrid.

42. The reasons given for placing AstroGrid in the 'lower priority' group were (i) delayed delivery and (ii) limited impact. We agree with the response of the AstroGrid Board to these comments, namely that (i) delays have been relatively minor and, in any case, agreed with the Oversight Committee and (ii) prototype software systems have generated significant interest, but true impact can only be assessed once the recently released software has been available to the community for an extended period. We endorse the points made by the AstroGrid Board concerning the strategic importance of UK leadership within the Euro-VO and the Global IVOA, the potential damage to EU funding, the potential damage to future facilities and missions, and the potential loss of future collaborations with other academic and industrial programmes.

43. The reasons for placing AstroGrid at the very bottom of the PR ranking are questionable. STFC evidently recognizes the importance of e-Science infrastructures through its support of GridPP (ranked as medium-high priority) and the RAL e-Science centre (although this does not appear to be under consideration at all in the PR). e-Science infrastructures are certainly at least as important to astrophysics as they are to the science areas covered by GridPP and the RAL e-Science centre.

44. AstroGrid has been carefully scrutinised by the PPRP and the PPARC Science Committee and approved at the minimum funding level that would allow delivery of all key elements of the programme. Given the large targetted investment to date, the continuing strategic importance of this area of e-Science, and the world lead secured for the UK by AstroGrid, our panel considers that it would be a false economy to discontinue core funding of AstroGrid before its usefulness has been tested. At some point in 2009 it should be possible to assess quantitatively the impact of an operational VO service provided by AstroGrid for UK astronomers and hence the optimum mix of funding routes needed to exploit the VO for current and future projects. Funding to enable AstroGrid to prove itself and identify its future requirements could come from a levy on other astronomy ground and space-based projects ranked in the higher categories which will make use of the AstroGrid infrastructure.

45. There is room for STFC to negotiate with AstroGrid regarding the level of costs committed. One possibility is a reduction in the period over which a commitment is currently secured, e.g. perhaps committing funding to mid-2009, rather than late 2009. In any case, STFC should honour the commitment to AstroGrid approved by PPARC council and make good the bridging funds that universities have disbursed while awaiting the PR. STFC should then work with AstroGrid to identify the minimum level of funding required to enable the project to prove its worth to the community. Support for the continued integration of new databases should form part of this package. There may also be some possibilities for minor cost savings by sharing costs of staff with University-based e-Science activities. In the longer term, the funding of AstroGrid will have to rely heavily on European collaborations through ESO and ESA.

5. CASU/WFAU

5.1 Background

46. The Cambridge Astronomical Survey Unit (CASU) at Cambridge (PI Mike Irwin) and the Wide-Field Astronomy Unit (WFAU) at Edinburgh (PI Bob Mann) have both established impressive track records in the provision of data-processing services to optical and near-infrared astronomers in the UK and Europe. They have been central to the success of projects like the UKIRT Infrared Deep Sky Survey (UKIDSS) and are critical for upcoming missions like GAIA and facilities like VISTA. They have been supported by separate grants with similar structures incorporating: infrastructure; design, operation and development of data processing systems and archives; and survey-related astronomical research. Through their close links with the Virtual Observatory, e.g. through AstroGrid, both the CASU and the WFAU are well placed to extend their remit to include a broader range of datasets, and to play a critical part in developing strategies for the processing of data streams to be generated by next-generation facilities like the E-ELT and the SKA, and for providing UK Access to datasets from external facilities like the LSST.

5.2 Summary of community input: CASU/WAFU

47. The only additional input to the PR from the consultation exercise came from the WFAU group. They noted that PPAN's concerns about the impact of the CASU/WFAU facilities was at odds with community feedback. They also noted the developing links between the WFAU and ESO, future space missions like the JWST and Euclid, and the US Large Synoptic Sky Telescope (LSST).

5.3 Panel commentary: CASU/WAFU

48. Our panel broadly endorses the views expressed in the additional consultation comments from the WFAU. The high international profile of projects like UKIDSS, and services like the UKIRT WFCAM Science Archive, mean we can be confident in the delivery of high-impact science from the CASU/WFAU work on data streams from upcoming missions and facilities like GAIA and VISTA. It is evident that UK-led world-class data management and access facilities, such as those at both the CASU and the WFAU, are of great benefit to UK leadership in world astronomy. The lack of any direct evidence of community support for the CASU and the WFAU via the consultation exercise was unfortunate. Nevertheless, our group believes that the CASU and the WFAU continue to provide essential services to the UK and European astronomical communities, generating data products both for direct use and for legacy use via the Virtual Observatory. We recognise the strength of having research-active astronomers central to both facilities.

49. Our first recommendation is that the service elements be written directly into grants for those projects or facilities requiring CASU or WFAU support, extending the models provided by the the Gaia Data Flow System and the Vista Data Flow System. In this manner, the true costs of data processing for facilities and projects will become transparent, and the CASU and WFAU capabilities will be encouraged to develop along the lines needed to support the most highly rated future facilities and projects. Our second recommendation is that the research elements of these groups are supported via research grants so that their impact can be judged alongside other STFC-supported research. We note that some central funding of both the CASU and the WFAU might need to be maintained over the transition period between direct and project/research-grant support if damage to essential infrastructure and loss of key personel is to be avoided.

Appendix A: panel composition and methodology

A.1 Panel composition

Professor Carlos Frenk (Durham University) – Chairman

Professor Tom Hartquist (Leeds University)

Professor Richard Kenway (Edinburgh University)

Professor Andrew King (Leicester University)

Professor Andrew Liddle (Sussex University)

Professor Steve Rawlings (Oxford University)

A.2 Terms of reference

Your panels have been set up to provide independent advice to PPAN and PALS, as well as to the Science Board and the STFC executive. The panels will use material generated by the Programmatic Review, and comments received during the consultation.

The information that will be made available to you will be

• The inputs received as part of the consultation process

• The project information made available to PPAN and PALS on your projects

• The financial information made available to PPAN and PALS on your projects

• Feedback from PPAN and PALS on the projects

• Any additional information on the PPAN and PALS rankings which the chairs of these committees agree to provide

The output of your panel should be a report that digests and summarises the community's input, adds your panel's commentary about the prioritisation of the projects and recommends a course of action. This course of action can include adjusting priorities, recommending reductions in funding on some projects to support increases in others, or changes in time scale for projects. Panels are also invited to explain what could be done with a modest increase in funding for their area of science in order to allow us to tension the opportunities that exist in different areas.

Normal rules relating to conflict of interest will apply during your panel's meetings.

STFC will cover reasonable expenses relating to the activities of these panels but cannot undertake to provide administrative support.

The panels' reports will be delivered to PPAN or PALS and to the Director, Science Programmes. The panel chairs will make a presentation of the contents of the report to the PPAN or PALS committees. The Science Board will receive the reports together with PPAN and PALS committees' recommendations on the route forward. The panel reports will be made public after STFC council has considered the recommendations on the future science programme (currently foreseen for 1 July).

A.3 Meetings

Face to face meetings were not deemed necessary. The panel met by teleconference on 28/March, 9/April, 22/April, 5/May and 12/May.

Appendix B: summary of recent reports on high performance computing

B.1 International standing of UK research using HPC

1. The International Review of Research using HPC in the UK published its report in December 2005 and, as part of that review, the international panel visited the UKQCD, VIRGO and UKAFF consortia. Written input was also provided on the research activities of the other PPARC consortia. The review concluded the following:

2. Computation has now become essential for the advancement of research across science and engineering and HPC enables the tackling of problems and opportunities that cannot be approached in any other way. The three PPARC consortia selected for detailed review were producing scientific results that compared with the highest international standards, and UKQCD and VIRGO were playing a leading role in setting international standards.

3. The consortium model is an organisational asset of HPC in the UK and has led to benefits in the efficient and cost-effective use of HPC systems, in software development and the training of young researchers. Consortia should be allocated coordination funds on presentation of an agenda for the management and coordination of their consortium.

4. The level of integration of computational science with computer science in the UK has not reached international standards and many research groups in the UK appeared to be deficient in adapting to modern programming practices, anticipating the need for data management and visualisation and developing interoperable software and algorithm developments. Consortia should generate plans for handling and analysing data in their disciplines. The UKQCD consortium was cited as a model of how this should be done. The training of computational scientists in the UK is too narrow and too short for interdisciplinary training.

5. To best realise the scientific pay-off for funds already invested, a diversity of HPC resources are required in the UK, spanning the spectrum from leading-edge facilities to university-based departmental or research group facilities, including both capability and capacity systems and offering a diversity of architectures. There was an 'unnatural divide' between the ability of PPARC-funded and other Research Council researchers to access the high-end national facilities. While there were historical reasons for this, the divide should be re-examined. Local HPC facilities have many advantages over national facilities and also serve as development platforms for work that will migrate to the national facilities. It is crucial that university-based facilities are sustained and upgraded under fEC.

B.2 Review of recent HPC policy

B.2.1 The Webber report 2003

6. In 2001, PPARC's Science Committee initiated a review of its approach to HPC by a panel chaired by Professor Bryan Webber. The report, published in 2003, concluded that "HPC is now a crucial investigative tool across all areas of astronomy and large parts of particle physics" and that "UK HPC work is of high scientific merit and world-leading in many areas across the PPARC remit". The report recommended that "HPC should be considered holistically and not as a separate item like a particle experiment or a space mission", by which it meant that "computation/simulation should be seen as a 'third discipline"'.

7. The Webber report concluded that "[l]arge externally-managed HPC facilities have not proved attractive for PPARC science" due to poor cost-effectiveness and inflexibility of operation. It also noted that "many groups and individuals have been commendably successful at fund raising for

equipment" while flagging the risk that "PPARC may need to make significant investment" if the alternative sources of funding cease. However, the Webber Panel did not encourage an 'isolationist' view in respect to central facilities and advised the PPARC Executive "to exploit viable cross-council infrastructure options which would offer value, but without compromising the current investment in PPARC community facilities".

8. The Webber report states "responses received from the community overwhelmingly indicate frustration that many of the funding sources accessed for major infrastructure in recent years have not provided for maintenance and staff costs within the overall award". It goes on to say "[t]he Panel is unanimous that PPARC must provide the necessary supplementary and supporting costs". The Panel estimated "that continuation of the current level of support would require annual investment of around £6.0M, of which approximately £1M would provide for maintenance and running costs".

B.2.2. The Kenway report 2006

9. In 2004, PPARC asked Professor Richard Kenway to chair a working group to advise on PPARC's involvement in national HPC facilities and to develop a strategy for HPC provision to support PPARC's science programme. The funding context was an annual expenditure by PPARC of around £1M pa, plus the availability of a capital amount of £13M from the Large Facilities Capital Fund for six years from 2006.

10. The Kenway report was accepted by Science Committee in 2006 and remains the policy of the STFC. It confirmed many of the conclusions of the Webber report. Thus, it reiterated that "[c]omputer simulation has become an essential part of PPARC science and is growing in importance" and "[HPC] is a crucial investigative tool for the theory community across all areas of astronomy and large parts of particle physics". It noted that "t]he International Review of HPC found that PPARC's HPC programme is producing scientific results that compare with the highest international standards and that some consortia are playing a leading role in setting international standards". It repeated the need for a holistic approach, because "[t]he distinctiveness of HPC and commonality of requirements across PPARC applications calls for a programmatic approach in which all aspects of HPC provision are reviewed together". On the basis of the inputs received from consortia, the Working Group revised upwards the estimated cost of the "HPC facilities and services required to address the immediate science goals of the PPARC research community" compared with the Webber report to £6.5M pa.

11. The Kenway report concluded that a one-size-fits-all approach to machine provision was inappropriate, stating that "[t]he consortia require access to a wide range of HPC systems, mostly commodity clusters, but, significantly, also shared-memory systems and a topical machine for lattice QCD. Some of the consortia additionally need occasional access to national class facilities". Thus, it recommended that "PPARC partnership in the new national HPC service, HECToR, should be delayed until after the service has been procured and/or the outcome of the HPC peer review is known. PPARC should continue to explore with EPSRC options for bringing together some PPARC HPC facilities and HECToR under a common service framework, in order to benefit from shared experience and the opportunity to trade resources, and to position the PPARC community to participate in future national and international HPC developments.

12. The Kenway report recommended that "[s]ystem and software support should be treated as an integral part of funding HPC", observing that "[t]he funding of software development, in particular the development and implementation of novel algorithms, has been largely neglected to date. This threatens scientific productivity and narrows the hardware options available to consortia, raising costs". Thus, it recommended that "[t]he cost-effectiveness, competitiveness and sustainability of the total package of hardware, system and applications software should be a key part of the peer review".

B.3 STFC call for proposals for HPC resources 2007

13. STFC accepted the Kenway working group recommendation that it should take a programmatic approach to the provision of HPC facilities, within which proposals for major investments in HPC provision for the STFC Theory Programme are reviewed together, establishing consistency in the peer review process and combining the evaluation of science quality with the technical evaluation of the choice of hardware, software, algorithm and implementation strategy. Therefore, it issued the first of a series of three-yearly calls for proposals for HPC resources - computer hardware, system and software support and operating costs, or alternately for access charges to local or national computing facilities, with a closing date in April 2007. Proposals were invited for up to five years for staff support and for up to three years for equipment and consumables. The budget for capital items was the earmarked sum of £13M in the Large Facilities Capital Fund, with non-capital costs coming from a 'new-HPC' planning line. Eleven proposals requesting £34.5M were received. They were reviewed by PPRP augmented by additional experts.

14. The PPRP agreed that while "[topical] centres were not a viable option for this call", because "setting up such a facility would take at least a year and this would result in groups losing their international competitiveness", it recommended that "planning should start as soon as possible to develop options for funding such centres at the next call". The Panel "focused the [limited] non-capital funding on providing staff to allow the hardware resources requested to be exploited in full, at the expense of requests not directly connected with user support" and it "reluctantly restricted the period of these appointments to three years maximum".

15. The PPRP ranked the proposals as "highly fundable on the basis that the programme of work was internationally competitive and had aspects of world-leading research" (Exeter, MIRACLE, UKMHD, UKQCD (I), VIRGO); "highly fundable on the basis that the programme of work has the potential to be world-leading or produce a step change", though of "higher risk because they were speculative or because they were proposed by new groups" (HORIZON, UKQCD(II)); "fundable on the basis that the programmes were internationally competitive" (COSMOS, Leicester); and "unfundable" (UKAFF, UKQCD(III)). The total recommended cost of the highly fundable/fundable proposals consisted of £12.3M from the capital budget and £4.2M from the non-capital budget.

16. The PPRP report was provided to PPAN as input for the PR. This resulted in the non-capital expenditure of £4.2M for "HPC Operations" being ranked as "low priority".

Appendix C: recent UK astrophysics and solar system physics theoretical activity

In each section of this appendix examples of major results are given before some problems to be addressed in future HPC studies are described.

I. Cosmology and Galaxy Formation

The UK is widely recognized as one of the strongest, if not the strongest, country in theoretical cosmology worldwide. For example, the UK contributes five scientists to the list of the fifteen most cited astronomers of all time. Three of them (Efstathiou, Frenk and Rees) are cosmology theorists. The distinction of UK cosmology theory was explicitly recognized by the 2005 Review on International Perceptions of UK Research in Physics and Astronomy which stated: "UK scientists have participated centrally in major advances in cosmology, defining the Standard Model and measuring the average density of dark matter and galaxy formation and evolution... Theoretical groups continue to lead the way in fundamental research."

The UK owes its international reputation in theoretical cosmology to contributions in four main areas: numerical simulations; theory and modelling of galaxy formation; early Universe and the microwave background radiation; analysis of galaxy redshift surveys. UK cosmologists and their colleagues are responsible for the N-body simulations that established the viability of the now broadly-accepted cold dark matter (CDM) cosmological model with which the UK is strongly identified. Simulations of the structure of dark matter halos have led to some of the most cited papers in the astronomical literature (Navarro, Frenk & White 1996, ApJ, 462, 563; 1997, ApJ, 490, 493). The UK is the base of the Virgo consortium, the world-leading collaboration in the subject. Virgo has consistently produced cosmological simulations of such importance as to lead the report of the 2005 "International Review of Research Using High Performance Computing in the UK " (ICC review sheet) to state: "The Virgo computations have redefined the standard for large-scale HPC computations in cosmology..." "With these simulations they have regained the world lead in cosmological simulations from the US ...". A recent example of the impact of the work of the Virgo consortium is the "Millennium simulation," in which Durham, Edinburgh, Sussex and Nottingham played a major role alongside the Max Planck Institute for Astrophysics in Garching (Springel et al. 2005, Nature, 435, 629). This simulation employed a record-breaking 10 billion particles to calculate the formation of structure in a region comparable in size to those covered by the largest galaxy surveys with enough resolution to follow small galaxies, leading to significant scientific advances (the first paper was featured on the front cover of Nature) and generating intense public interest.

The second main strand of the UK's research strength is galaxy formation theory and modelling. The basic ideas behind what has become the standard theory of galaxy formation can be traced back to the seminal paper by White & Rees (1978, MN, 183, 341), which was reformulated in the context of the CDM model by White & Frenk (1991, ApJ, 379, 52). These ideas led to the technique of semi-analytic modelling, which was developed in the UK (Kauffmann et al. 1993, MN, 264, 201; Cole et al. 2000, MN, 319, 168) and has been emulated by many groups worldwide. This technique allows modelling of the evolution of galaxies in sufficient detail to permit meaningful comparisons between theory and observations. It provides a vital point of reference for observational studies of galaxy formation, one of the highest priorities of STFC research (e.g. Lacey et al. 2008, MN, 385, 1155; Baugh 2006, Reps Prog Phys, 69, 3101).

The UK has developed world-leading strength in studies of the physics of the early Universe, and in particular in the testing of these models with astrophysical data including the cosmic microwave background (CMB) anisotropies. A particular focus has been the COSMOS consortium, comprising researchers at the Universities of Cambridge, Central Lancashire, Manchester, Nottingham, Oxford, Portsmouth and Sussex, and Imperial College and UCL, which for over a decade has operated the main high-performance computing facility for such work. The UK can point to leading work in exploiting new ideas in fundamental physics, for instance the braneworld, in cosmological models. Most of the innovation, for instance the curvaton scenario for the origin of structure, in inflationary cosmology in recent years has come from the UK, and UK theorists have also led the way in developing theoretical calculations of cosmic non-gaussianity, which may well be the next observational frontier in cosmology. The UK is providing world leadership in applications of advanced (and computationally intensive) data analysis techniques, particularly of the Bayesian variety, to cosmological data; most determinations of cosmological parameters worldwide use the CosmoMC package of Lewis and Bridle. Many of the community will be actively involved in theory exploitation of Planck Satellite data from within the Planck consortium.

The fourth main strand of the cosmological research in the UK is the analysis of large galaxy surveys. The UK has played a leading role in the major surveys of the past two decades, starting with QDOT and PSCz in the 1980s and 1990s and culminating with the "2-degree field galaxy redshift survey" (2dFGRS) in the 2000s. The 2dGRS was a collaboration between Australia and the UK (Cambridge, Edinburgh, Durham, Nottingham and Oxford). It is a major highlight of UK astronomy in the past decade and, in the area of large-scale structure, has overshadowed the US-based Sloan Digital Sky Survey (SDSS). These two surveys have revolutionised the study of the local large-scale structure and played a pivotal role in bringing about the era of 'precision cosmology'. The question of the energy budget of the Universe was addressed and solved by the 2dFGRS which, in combination with WMAP, showed that the mean cosmic density corresponds to $\Omega_{\rm m} = 0.25 \pm 0.03$ (Peacock et al. 2001, Nature, 410, 169), and, independently of SN distance scale measurements, that the energy density of the Universe is dominated by dark energy (Efstathiou et al. 2002, MN, 330, 29; Cole et al. 2005, MN, 362, 505; Sanchez et al. 2006, MN, 366, 189). The 2dFGRS also provided the first accurate estimation of the galaxy stellar mass and luminosity functions (Norberg et al. 2002, MN, 336, 907; Cole et al. 2001, MN, 326, 255), the construction and exploitation of the first modern catalogue of groups and clusters (Eke et al. 2004, MN, 348, 866), the first accurate determination of the galaxy power spectrum (Percival et al. 2001, MN, 327, 1297) and the discovery of baryon acoustic oscillations in the galaxy distribution (Percival et al. 2001, MN, 327, 1297; Cole et al. 2005, MN, 362, 505) One of the reasons for the resounding success of the 2dFGRS was the extensive use of realistic mock catalogues constructed from the Virgo consortium simulations, which gave the 2dFGRS a crucial advantage over the far better resourced SDSS.

The last twenty years have seen the emergence of a standard paradigm for the geometry and material content of our Universe, as well as for the origin and evolution of all structure within it. According to this ΛCDM model, we live in a flat universe where at least two thirds of all mass-energy is now in the form of a dark energy field that drives the observed acceleration of the cosmic expansion. About a quarter is dark matter, most probably a new weakly interacting elementary particle yet to be detected directly on Earth. Only about 5 percent is ordinary or baryonic matter of which only about a tenth is in stars today and the rest resides mostly in the intergalactic medium. Structure was seeded by quantum fluctuations of the vacuum in the very earliest instants of the Big Bang. These produced weak sound waves in the near-uniform primordial plasma that left observable imprints on the surface of recombination, seen when the Universe was only 400,000 years old. These tiny ripples, mapped by imaging the CMB radiation, grew into the full richness of structure we see around us today. It is this transformation from a near-uniform primordial soup to a cosmic web of galaxies, clusters and larger structures that is the focus of Virgo's simulation programme.

Computer simulations are the main theoretical tool for calculating the formation history of cosmic structure given a set of assumptions regarding the generation of inhomogeneities in the early Universe, the identity of the dark matter, and the values of the cosmological parameters. Simulations provide the best means for testing theories of the early Universe against observational data. Over the past 25 years, computer simulations by the Virgo consortium and the precursor collaboration that led to the formation of the consortium have played a central role in defining the modern approach to physical cosmology and establishing ΛCDM as the standard model of cosmology.

II. Star and Planet Formation; Accretion Discs

The UK has produced important computational studies of different aspects of star formation. They include investigations of the effects of magnetic fields on the creation of spiral distributions of molecular clouds (Dobbs & Price 2008, MN, 383, 497) and of the roles of fast-mode shocks and thermal instability in producing magnetically dominated regions of low-mass star formation and of slow-mode shocks in triggering high-mass star formation (Van Loo et al. 2007, A&A, 471, 213). Work by Bate, Bonnell and their collaborators on the gravitationally-driven collapse of structures in molecular clouds to form clusters and brown drawfs (Bate et al. 2003, MN, 339, 577) and high-mass stars (Bonnell et al. 2004, MN, 349, 735) has stimulated a huge international response. The relative roles of turbulently induced compression, gravitationally-driven collapse and ambipolar resistivity in the creation and evolution of dense cores in which low-mass stars form have been at the center of much international debate. The UK has a clear lead in the numerical techniques necessary for a definitive identification of the role of ambipolar resistivity (Van Loo et al. 2008, arXiv:0803.4422). Collaboration between those performing such MHD simulations and theoretical, as well as observational, molecular astrophysicists in UCL, Leeds, Belfast and Manchester performing complementary investigations (e.g. Tsamis et al. 2008, arXiv:0803.0519) will ensure that ALMA will be exploited to the full in star formation research.

The country excels in computational studies of protoplanetary discs. Stamallos & Whitworth (2008, A&A, 480, 879) have employed a radiative transfer scheme and a treatment of the internal energy of molecular hydrogen that has resulted in the most reliable and important study of the possible role of gravitational instability in the initiation of giant planet formation. Rice et al. (2004, MN, 355, 543) have stimulated world-leading activity, pursued in Edinburgh, St Andrews, Queen Mary and Cambridge, in the investigation of dust dynamics in gravitationally unstable discs in promising efforts to solve the classic and vital problem of planetismal growth. The work of Nelson & Papaloizou (2004, MN, 350, 849) brought together their pioneering efforts in studies of hydromagnetic effects on disc evolution and the effect of disc dynamics on planet migration. They and their collaborators continue to lead the world in computational studies of planetary system formation. Computational models of the chemistry of protoplanetary discs developed by Millar and his collaborators (e.g. Ilgner et al. 2004, A&A, 2004, 415, 643) prepare the way for ALMA to have the maximum impact in observations of protoplanetary discs. A UCL – Leeds collaboration on chemistry in dynamically more complicated disc models will enhance the preparation.

Accretion discs are studied extensively in a number of other contexts. The groups at Leicester (King, Nayakshin and Lodato) and Cambridge (Pringle and Clarke) have considered the effects of self-gravity on large discs (Clarke et al. 2007, MN, 381, 1543) and disc warping (Lodato & Pringle 2006, MN, 368, 1196). While these studies are relevant to protoplanetary discs, they are also important for understanding how the ultraluminous X-ray sources represent a short but spectacular stage of high-mass X-ray binary evolution (Begelman et al. 2006, MN, 370, 399). Nayakshin et al. (2007, MN, 379, 21) have performed simulations of star formation in a massive disc around the supermassive black hole (SMBH) at the Galactic Centre.

Ambipolar resistivity plays major roles in the evolution of star forming regions. It affects the propagation of waves that may trigger collapse and waves that carry angular momentum. In some stages of low-mass star formation, it may regulate the rate at which gravity causes further compression. It must be included in studies of the generation of 0.1 parsec-scale dense cores that are precursors to protostars. Little reliable computational work has been done on the formation of protoplanetary systems from such objects. This is partly a consequence of the need for a detailed treatment of the dynamics of dust grains, which are important charge carriers and determine the ambipolar resistivity in dense cores as they evolve. A three-dimensional, time-dependent MHD treatment incorporating a minimum of three gas-phase and two grain fluids is required to investigate the birth of protoplanetary systems around low-mass stars. The combination of computational power and efficiency of the numerical method must allow the study of density differences of at least eight orders of magnitude. The models must include chemical networks for the fractional ionisation and distribution of charge on grains.

Ambipolar resistivity is also of great importance in protoplanetary discs and determines the locations of “dead-zones”, regions in which the Magnetic Rotational Instability does not operate and the effective viscosity is very low. Dust dynamics, chemistry and radiative transfer are all necessary elements of any numerical calculations that will advance the understanding of planet formation. The problem of the evolution of a protoplanetary system is even more computationally challenging than the problem described in the previous paragraph because of the coupling between radiative transfer and dynamics.

III. Quantum Data, Stellar and Exoplanet Atmospheres; Stellar and Starburst Nebulae and Supernova Ejecta; Galactic and Extragalactic Interstellar Media and Starburst- and AGN-Driven Winds

Groups in Belfast, Sheffield Hallam, Strathclyde and UCL continue the UK's outstanding tradition in the computation of atomic and molecular data without which spectra obtained with many ground-based and space-borne facilities costing billions would remain unexploited. UCL physicists produced the BT2 water line list (Barber et al. 2006, MN, 368, 1086) and have used it with Hertfordshire and Liverpool John Moores colleagues to construct models of cool stars' atmospheres. All very old, low metalicity stars studied to probe the elemental abundances during galaxy formation are cool, and the interpretation of their spectra and calculation of their evolution (Harris et al. 2007, MN, 377, 1520) require good molecular data. Most exoplanets to be discovered will orbit cool stars. Excellent molecular data will be necessary for modelling the stellar spectra in sufficient detail that spectroscopic studies, including those designed to detect signatures of life, of the atmospheres of the exoplanets will be possible. As a result of their quantal studies of water, the UCL group has already been able to find water in a transiting exoplanet (Tinetti et al. 2007, Nature, 448, 169).

Due to the UK's strength in the calculation of relevant quantum data, it has long been at the forefront of research on photoionised nebulae, which include young supernova ejecta. Multidimensional radiative transfer calculations have become important in the study of such nebulae. The production and survival of dust in supernovae are key issues in the early evolution of galaxies because the presence of such dust would have induced a qualitative change in the mode of star formation from that producing only stars so massive that they would become supernovae. Using the UCL MOCASSIN radiative transfer code, Sugarman et al. (2007, Science, 313, 196) have shown that up to 0.02 solar masses of dust formed in the ejecta of a Type II supernova within 250 days of outburst.

Supernovae power the large-scale dynamics of the Milky Way's interstellar medium, which affect how the Galaxy evolves and where star formation occurs. Many models of supernova remnants and of the global interstellar medium have been based on single fluid nonmagnetic descriptions. The need to include magnetic fields and cosmic rays, which make contributions to the total pressure at least comparable to and sometimes exceeding that of the thermal pressure, is clear. The models must be self-consistent in the sense that a flow in the thermal fluid modifies the magnetic field and the cosmic ray pressure just as they modify that flow. The Newcastle group has shown that a galactic fountain efficiently removes magnetic helicity from galactic discs, which allows the mean magnetic field to saturate at a strength comparable to equipartition with the turbulent kinetic energy produced by supernovae (Shukurov et al. 2006, A&A, 448, L33). Tony Bell, now in Oxford, independently discovered the diffusive shock cosmic-ray acceleration mechanism in 1978. His fundamental work on the turbulent amplification of the magnetic field during diffusive acceleration (Bell 2004, MN, 353, 550) has major implications for the energies to which cosmic rays are accelerated. The Leeds group has developed a model of cosmic-ray modified flows including nonequilibrium radiative cooling (Wagner et al. 2007, A&A, 463, 195). It has been used to interpret optical emission line profiles for a shock in the Tycho remnant in order to infer the cosmic ray injection rate, diffusion coefficient and energy transfer rate from cosmic rays to the thermal fluid. The code can be applied to more global studies informed by the results obtained for specific shocks.

Galactic winds driven by supernovae in starburst galaxies and by Active Galactic Nuclei (AGNs) are central to feedback during the era of galaxy formation. The comparison of observations with the results of hydrodynamic models of the effects of AGN outflows on accretion flows at the centers of clusters of galaxies developed in Durham, Oxford and Southampton (e.g. Dalla Vecchia et al. 2004, MN, 355, 995; Omma et al. 2004, MN, 348, 1105) is providing great insight into processes important in such feedback. The Cambridge and Leicester groups study the formation and growth of SMBHs by accretion and the feedback of the SMBHs on the surrounding galaxies. Lodato & Natarajan (2007, MN, 377, L64) have investigated the formation of SMBH seeds at high redshift (about 15) due to the self-gravity of accretion discs in collapsed haloes. King et al. (2008, MN, 385, 1621) have studied how randomly oriented discs around SMBHs in AGNs evolve and how their evolution affects the SMBHs' spins. The self-gravity of such a disc has major implications for the accretion efficiency, mass growth timescale and other important parameters. King (2003, ApJ, 596, L27; 2005, ApJ, 635, L121) showed that radiation-pressure induced outflows from hyperaccreting SMBHs have a determining effect on the host galaxies, and are probably responsible for the well-known relations between the SMBH mass and the velocity dispersion and between the SMBH and bulge masses.

The maintenance of a molecular data programme, which will become increasingly important and broader as Darwin observations of exoplanets approach, requires roughly 660K CPU-hours and 7 Tbytes storage in the short term. The short term goals include the calculation of data for deuterated water (important for the development of a test to determine whether ultra-low mass stars have reached the deuterium burning phase), ammonia (important for models of brown dwarfs) and acetylene (important for the study of cool carbon stars and terrestrial atmospheric evolution).

Additional future work on exoplanet atmospheres is mentioned in the Planets and Space Plasmas section of this appendix.

Extragalactic photoionized nebulae associated with starbursts will constitute an important area of application of three-dimensional radiative transfer codes. Data from Spitzer and Herschel, together with the models, will allow the derivation of the metalicities in each nebula and the distribution of masses of the stars that power them. This work will bear on the important question of how the initial mass function varies between environments. In the short term 1000K CPU-hours and 3 Tbytes memory will be required for these and related studies.

The inclusion of magnetic fields and cosmic rays into more advanced models of global interstellar dynamics, starburst superwinds and AGN outflow - accretion interactions will have major effects on the global flows. Dissipation in shocks will be much more distributed leading to most shocks being much weaker than in analogous hydrodynamic calculations. The diminished strengths of model shocks will have major implications for the interpretation of observations of active regions. The inclusion of magnetic fields and cosmic rays will make what have already proved computationally challenging problems even more demanding.

IV. Galactic Dynamics

The study of galactic dynamics started in the UK with the work of Jeans and Eddington, and UK astronomers are still in the vanguard of work on galaxies. As a consequence, Monthly Notices of the Royal Astronomical Society is the premier journal for galactic dynamics, and an inspection made in mid-April showed that, according to ADS, 51 of the 100 most recently published papers with "galactic dynamics" in the abstracts appeared in that journal. In addition, graduate courses on galactic structure and dynamics at most leading graduate schools (including Harvard, Yale, MIT, Princeton and Caltech) are based on the two volumes by Binney (Oxford) and Merrifield (Nottingham) and by Binney and Tremaine (Cambridge-educated).

Gilmore & Reid (1983, MN, 202, 1025) discovered the Milky Way's thick disc and Binney et al. (1991, MN, 249, 678) wrote one of the two papers that convinced the community that the Milky Way is a barred galaxy. Dehnen (1998, AJ, 115, 2384) discovered the dynamical streams in the distribution of stars near the Sun. Ibata et al. (1995, MN, 277, 781) discovered the nearest major satellite of the Milky Way, and the Cambridge group has led the world in studies of the outer reaches of Local Group galaxies (e.g. Belokurov et al. 2007, ApJ, 658, 337). Raha et al. (1991, Nature, 352, 411) uncovered what is probably the dominant mechanism by which spiral structure heats stellar discs. Bedregal et al. (2006, MN, 373, 1125) first securely established that spiral galaxies in clusters passively fade to S0 galaxies. Heggie (1975, MN, 173, 729) laid the foundations of our understanding of the profound role played by binary stars in the evolution of star clusters. Several major surveys of the properties of nearby galaxies (SAURON, SINGS and THINGS) are led by UK groups.

N-body modelling has played a central role in the development of galactic dynamics. The observed parts of galaxies involve $< 10\%$ of the total masses of the galaxies, which are dominated by their dark halos. Consequently, in N-body simulations the majority of the particles represent the unobservable dark halo, which severely restricts the resolution achievable in the observable inner regions. Recently we have begun to observe the chaotic outer low surface-brightness regions of galaxies. To interpret these data we need dynamical models that extend much further out, and include environmental disturbances from neighbouring galaxies. This requirement dramatically increases the minimum computational scale of simulations.

Particle noise is invariably a major issue in the interpretation of simulations, especially in connection with disc phenomena, such as spiral structure and warps, neither of which can be considered well understood. The only sure way to beat down particle noise is to increase the number of particles, and simulations now routinely employ $> 10^7$ particles. Substantially larger particle numbers are required to address questions relating to either the dynamics of nuclei, such as the dynamics in the vicinity of central black holes, or the dynamics of extended discs.

Gas, which is key for the dynamics of spiral galaxies, is hard to simulate with adequate resolution. We will not have a good understanding of optical and radio-frequency observations of the Milky Way until gas dynamics is well studied with very large-scale simulations.

Massive surveys of stars in the Milky Way are underway. They will culminate in the Gaia mission, which will catalogue 10^9 stars. On account of observational biases, science can be extracted from these data only by fitting them to dynamical models. Each star is associated with a non-negligible error ellipse in the space of observables, so likelihoods must be evaluated by integrating over a billion ellipses for each candidate model. Significant challenges are posed by existing catalogues, which contain millions of stars and have not yet been the subject of maximum-likelihood analysis.

V. Relativistic Astrophysics

Noncosmological theoretical relativistic astrophysics research is pursued in a number of UK institutions including Cambridge, Cardiff, Edinburgh, Exeter, Leeds, Liverpool John Moores, Portsmouth, Southampton and UCL.

High quality UK activity contributes significantly to advancing the state of the art modelling of neutron stars and black holes. Much of this work is driven by the effort to detect gravitational waves and provides support for the major facilities GEO600, LIGO and LISA. The theoretical developments are also beginning to impact on X-ray and radio observations, as a result of the models becoming more sophisticated and including more detailed physics. The UK is world leading in the development of these models, as demonstrated by recent studies of neutron star oscillations and magnetar QPOs (Samuelsson & Andersson 2007, MN, 374, 256), neutron star free precession and superfluid turbulence (Glampedakis et al. 2008, PRL, 100, 081101; see also Haberl et al. 2006, A&A, 451, L17) and gravitational waves from neutron stars with quark cores (Haskell et al. 2007, PRL, 99, 231101). Detailed models of dynamical black holes are also being developed in order to meet the requirements for LISA. UK groups are particularly strong on the modelling of extreme-mass-ratio binaries. Activity in this area ranges from investigations of detailed issues in general relativity (Barack & Sago 2007, PRD,75, 064021) to considerations of problems concerning spacetime mapping (Barack & Cutler 2007, PRD, 75, 042003; Gair et al. 2008, PRD, 77, 024035) and stellar dynamics in galaxy cores (Hopman et al. 2007, MN, 378, 129).

Important investigations involving HPC have included fully general relativistic studies of core-collapse supernovae (Ott et al. 2007, PRL, 98, 261101) and black hole formation (Baiotti et al. 2007, CQG, 24, S187). Major work is also done on gamma-ray bursts and jets (Barkov & Komissarov, 2008, MN, 385, L28), neutron star mergers (Price & Rosswog 2006, Science, 312, 719; Domainko & Ruffert 2008, Adv Sp Res 41, 518) and black-hole neutron star mergers (Rantsiou et al. 2007, arXiv:astro-ph/0703599). In the area of numerical simulations of black hole collisions, the UK involvement has mainly been in the development of techniques (Booth & Fairhurst 2008, PRD, 77, 084005; Garfinkle et al. 2008, CQG, 25, 075007).

The demand for HPC in relativistic astrophysics is increasing as the models become more realistic. As stated above, the UK has contributed very significantly in the development of relativistic models. UK groups also play prominent roles in numerical modelling, even though the investments in HPC have not been optimal. Fully general relativistic numerical studies (e.g. of black hole mergers) are particularly demanding due to the number of coupled equations to be solved. Adequate investment in UK HPC is absolutely essential to enable full exploition of the UK's world leading contributions to the incorporation of physics in the models and the development of powerful numerical tools.

VI. Solar MHD and Plasma and Helioseismology

The UK is world-leading in the theory of the generation of magnetic fields via dynamo action, which gives rise to the eleven-year solar cycle and the instability and emergence of the magnetic field to form active regions at the solar surface. The UK has made fundamental contributions to the theory of mean-field electrodynamics, which underpins much solar dynamo theory (Cattaneo & Hughes 2006, JFM, 553, 401; Courvoisier et al. 1996, PRL, 96, 34503). The important question of the interaction of dynamo action with shear and turbulence has been addressed by Ponty et al. (2001, JFM, 435, 261), Zhang et al. (2006, JFM, 546, 25) and Proctor (2007, MN, 382, L39). The UK has also made major contributions to answering the vital question of the formation and stability of the tachocline, the region of strong shear at the base of the solar convection zone. Gough & McIntyre (1998, Nature, 394, 755) demonstrated the inevitability of a magnetic field in the solar interior, and Tobias et al. (2008, ApJ, 667, L113) and LeProvost & Kim (2007, ApJ, 654, 1166) have investigated angular momentum transport in this stably stratified sheared environment. Once the dynamo generated magnetic field has emerged at the solar photosphere via magnetic buoyancy (Kersalé et al. 2008, ApJ, 663, L113), it forms sunspots. Thomas et al. (2002, Nature, 420, 390) proposed a new mechanism for the formation of the filamentary structure of the penumbrae of sunspots. This work has had a major impact internationally with a number of groups investigating this process.

UK scientists have played central roles in the establishment and theoretical development of helioseismology as a powerful tool to probe the solar interior (Thompson et al. 2003, ARA&A, 41, 599).

Because of its huge scale compared to laboratory plasma devices, solar plasma provides an unequalled opportunity for the investigation of basic plasma phenomenon, including magnetic reconnection and particle acceleration, which are ubiquitous in the cosmos and crucial to laboratory and fusion plasma studies. The outflow of solar matter and radiation is also important in the climate change debate and for hazards posed to powerlines, communications and space flight. The UK solar physics community is very strong in plasma physics and diagnostic theory and in HPC simulations, as well as in research at the all important interface of data and theory. The key to the solution of long standing problems concerning magnetic energy release and particle acceleration (coronal heating, flares and CMEs) lies at the challenging interface of MHD and kinetics, which was the subject of the first workshop (at RAL) organised by the UK Solar Theory Steering Group founded in 2007. This workshop played a part in the crucial drive to narrow the gap between kinetic theory and MHD in solar theory and establish new collaborations across the solar theory, space plasma and laboratory plasma communities. All three areas of research are very demanding of analytic ingenuity and intense computation because of the hugely diverse range of characteristic length and time scales involved. The UK also has a leading position in the modelling of emission processes and is actively involved in the continuous improvement of CHIANTI, the code used to calculate synthetic spectra and to carry out plasma diagnostic studies (Landi, Del Zanna, Young & Mason 2006, ApJ, 162, 26), and in modelling non-thermal particle transport. The acceleration of non-thermal particles in solar flares is a flourishing area in solar physics with strong participation from Glasgow, Manchester, St Andrews and elsewhere. Comparable theoretical challenges are involved in the inference of solar flows and fields (Potts & Diver 2007, Solar Phys, Nov, 175), the distribution of nonthermal particles (Brown et al. 2006, ApJ, 643, 523) and plasma parameters from the copious high resolution data coming from current space missions and to be obtained with advanced future facilities like the Solar Orbiter. These non-equilibrium transient/kinetic plasma phenomena and intense beam electrodynamics lie far outside the MHD regime but are ubiquitous throughout astrophysical sources from the Sun to pulsars, AGNs and the early Universe. Solar observations with future space- and ground- based facilities like Lofar and SKA will stimulate further theoretical efforts.

Studies of solar surface phenomena based on MHD models continue to be of great value. The rise of active-region flux through the convection zone and its emergence at the solar surface has been studied over the years with increasingly detailed simulations. Adaptive 3D compressible MHD codes with shock capturing (e.g. Archontis & Hood 2008, ApJ, 674, 113) now can treat partially-ionised plasmas (Leake et al. 2005, A&A, 442, 1091). Weiss et al. (2004, ApJ, 641, 577) developed a leading model of how emerged flux forms into sunspots. The overall picture of the evolution of the global corona is studied with magnetofrictional codes (Mackay & van Ballegooijen 2006, ApJ, 641, 577). Pioneering work in the theory and MHD simulation of magnetic reconnection, and its role in forming observed coronal structures, has long been a strength in the UK (e.g. Priest et al. 2005, ApJ, 624, 1057), as have the theory of MHD wave propagation (e.g. De Pontieu et al. 2005, ApJ, 624, L61) and the development of coronal seismology. Work based on the latter technique, uses ducted waves as coronal diagnostics and recently has been focussed on slow- and fast- mode wave trains in realistic coronal loops excited by impulsive bursts of energy (Taroyan et al. 2005, A&A, 438, 713; Nakariakov et al. 2004, MN, 349, 705).

The dynamics of the solar interior are dominated by the interaction of extremely turbulent convection with stratification, rotation and magnetic fields. Turbulence is characterised by the interaction of processes on a wide range of spatial and temporal scales and provides one of the stiffest tests for High Performance Computing. Solar convective turbulence operates at the extreme end of the parameter space with the parameter regime of interest one of extremely large fluid and magnetic Reynolds numbers and extremely small Prandtl numbers. It is now believed that, for a complete description of the dynamics, over 10^{18} degrees of freedom need to be included in a turbulence calculation. The calculations will be particularly challenging because it takes many dynamical times to determine whether or not the flow really does act as a mean field dynamo.

The release of magnetic energy during a solar flare or a sub-storm in the geomagnetic tail requires the collapse of scale-lengths as a current sheet form means that the fluid approximation breaks down and any attempt at explaining effects such as particle acceleration in flares must address the coupling of the large scale trigger, often described by MHD, to the collisionless, kinetic acceleration of particles. The trigger mechanism for flares may be due to the emergence of new magnetic flux from the solar interior. This requires the coupling of the solar interior to the corona. Hence, small-scale photospheric features must be resolved as input to the large-scale features of the diffuse corona. Note that in flux emergence simulations the density varies by over 8 orders of magnitude. It is essential, therefore, that the large scale MHD simulations contain enough resolution to allow the separation of scales. In 3D in order to resolve the length scales varying by 3 orders of magnitude the minimum computational domain is approximately 1024^3 grid points. Thus, whether we are trying to couple different physical regions of the solar system, e.g. convection zone to chromosphere, or different plasma models for flares, we are driven to a computational constraint of 1024^3 grid points to allow scale separation between domains and models.

VII. Planets and Space Plasmas

The UK community has built a world-class reputation over the past 25 years in theoretical modelling of the Earth's upper atmosphere and its coupling to the magnetosphere beyond. The most recent investigations have concerned the global distribution of nitric oxide, an important species in thermosphere and stratosphere physics (Dobbin et al. 2006, J Geophys Res, 111, A07314). Over the past 10 years these studies have also extended to the upper atmospheres of the gas giant planets and their moons. The UK community has the only published global models of Saturn's thermosphere and ionosphere (Mueller-Wodarg et al. 2006, Icarus, 180, 147; Smith et al. 2007, Nature, 445, 399), and of Titan's thermosphere (Mueller-Wodarg et al. 2003, J Geophys Res, 108, A121453; Galand et al. 2006, Geophys Res Lett, 33, L21101). The Saturn models have shown how rapid rotation strongly inhibits the global transport of auroral energy input, while the Titan models are of particular importance in relation to the ongoing Cassini/Huygens space mission. Most recently these codes have also been applied to the upper atmospheres of exoplanets, uncovering the pivotal role of H_3^+ in atmospheric stability (Koskinen et al. 2007, Nature, 450, 845). Models of the latitudinally-banded jet flows, resulting from the interaction of rotation and convection dynamics, at lower altitudes in Saturn's and Jupiter's atmospheres have been developed for simple geometries (Jones et al. 2003, Geophys Res Lett, 30, 1731).

Theory and computation in solar system plasma physics in the UK spans a broad range of topics from processes taking place on kinetic scales to large-scale modelling of global systems. Particular examples illustrating the range of recent work include (a) PIC simulations of high Mach number collisionless shocks that have shown for the first time that quasi-perpendicular shocks have intrinsic time-dependence (Chapman et al. 2006, Space Sci Rev, 121, 5), (b) models of turbulence in Earth's ionospheric plasma demonstrating energy transfer between Langmuir and ion-acoustic modes (Guio et al. 2008, PRL, submitted), (c) comprehensive kinetic models of the effects of proton precipitation on the Earth's upper atmosphere and consequent auroral output (Galand & Chakrabarti 2006, J Atmos Terr Phys, 68, 1488), (d) studies of anomalous scaling in MHD simulations of turbulence with application to the solar wind (Merrifield et al. 2007, Phys Plasmas, 14, 012301), (e) PIC simulations of the lunar wake in the solar wind flow (Birch & Chapman 2002, Phys Plasmas, 9, 1785) and (f) global models of magnetosphere-ionosphere coupling at Jupiter and Saturn with application to the origins of auroral emissions (Cowley et al. 2005, J Geophys Res, 110, A11209). Overall, this represents a substantial body of work that has significant impact and influence at an international level.

The UK programme in planetary atmosphere modelling relies heavily on HPC, since models must simultaneously calculate a variety of processes ranging from atmospheric thermodynamics to plasma physics, electrodynamics and chemistry, in general in three spatial dimensions and with time. Enhanced computing power is required to provide these models with the necessary scope to grow in complexity and sophistication. For instance, important physical processes, such as the coupling to the external magnetic and plasma environment, will be incorporated and larger height ranges and finer grid spacing will be adopted. Given adequate resources, future work will become possible on e.g. (a) extending the parameter ranges and generalising the geometries of models of jet flows in Jupiter's and Saturn's atmospheres that include the effects of compressibility, (b) developing time-dependent models of magnetosphere-ionosphere coupling at the giant planets, (c) investigating the dependence of exoplanet atmospheric properties on orbital phase and (d) constructing Mars exosphere models that include the effects of solar wind interactions.

Development of the broad programme of theory and modelling in solar system plasma physics outlined above to more realistic scenarios requires access to HPC with increasing power. For example, although much has been learned about collisionless shock reformation and ion acceleration from existing PIC calculations, key remaining problems require the simulation of multiple ion species as well as the consideration of processes in three dimensions and time. HPC is also required in simulations of global aspects of space plasma systems, not only in relation to the interaction of solar system bodies with the solar wind where models of increasing sophistication have been developed over recent years, but also particularly in the interaction of moons with magnetospheric environments, particularly that of Titan inside Saturn's magnetosphere and Io at Jupiter.

Appendix D: HPC in particle astrophysics and theoretical particle physics

Particle Astrophysics research requiring HPC is summarised in an annex of the Kenway report. Some is mentioned in the section on cosmology and galaxy formation in the above appendix. Substantial HPC resources will be required for the processing and interpretation of data from a wide variety of particle astrophysics facilities including Planck, GEO600, Advanced LIGO, LISA and the Pierre Auger Observatory.

The Theoretical Particle Physics research for which HPC is required is also addressed in an annex of the Kenway report. Key areas requiring HPC support include Lattice QCD and phenomenology.