ASCAC Panel Report on the Fusion Simulation Project
Final Draft
July 17, 2008

Panel Members:

F. Ronald Bailey
Donald B. Batchelor
David Brown
Stephen C. Jardin
Douglas B. Kothe
Ewing “Rusty” Lusk
Thomas A. Manteuffel
Juan C. Meza
David P. Schissel

Table of Contents

1.0 Executive Summary

The Panel has examined the Fusion Simulation Project (FSP) Workshop Report with regards to critical technologies in applied mathematics, computer science, and computational science, as well as the most beneficial role for Advanced Scientific Computing Research (ASCR) in FSP. In addition, the Panel has examined challenges ASCR is anticipated to face during implementation of FSP.

The Panel finds that FSP will be critical to enabling the U.S. to maximize the benefit from its investment in ITER (International Thermonuclear Experimental Reactor), and to moving toward practical fusion energy production. Furthermore, the panel finds that FSP is well aligned with ASCR goals, and that the project is critically dependent on research, technology, and facilities ASCR will provide. The Panel finds ASCR well poised to make these contributions to the project, in light of the fact that ASCR must also consider a new scope to best serve FSP. ASCR has a long history of technical contributions to fusion energy research and through SciDAC (Scientific Discovery through Advanced Computing), has forged an effective multi-disciplinary and collaborative fusion simulation research model. Through its Leadership Computing Program, ASCR is accelerating the advancement of ultra-scale computing capabilities that will be necessary for the success of FSP.

FSP provides an excellent opportunity to demonstrate fulfillment of the ASCR visionary approach to modeling and simulation. FSP is an application that will drive development of technology that can be applied across a wide spectrum of Office of Science applications, and one that will attract the best and brightest young scientists.

The Panel finds that while FSP is in the early formulation stages and broad scientific and technical issues are reasonably understood, the project objectives, requirements, and approach are not yet well defined. In light of this, the Panel finds the true scope of the research, technology, schedule, and performance required are still unknown. The Panel expects ASCR to participate with Office of Fusion Energy Science (FES) in the anticipated Project Definition Phase, and expects this will bring much needed definition to the project.

After reviewing FSP documentation and holding its own workshop, the Panel finds the FSP Workshop Report provides a reasonably good identification of the critical technology requirements in applied mathematics, computer science, and computational science. The Panel has, however, identified additional areas for consideration. In applied mathematics, additional areas the panel found to be important include methods for incorporating experimental data into simulations, improved methods for sensitivity and uncertainty quantification, the development of new computable models, and methods for analysis of FSP computational and experimental datasets. In computer science, a critical key technology requirement is reaching the estimated sustained performance requirement for FSP on future petaflops and exaflops platforms. In addition, the Panel highlights the need for more robust technology to manage large FSP-generated datasets, visualizing results and sharing data, and for collaborating across the widely distributed FSP research community. Finally, in the area of computational science, the Panel puts additional emphasis on the challenges of developing a suitable FSP software framework, workflow technology, and verification and validation (V&V).

The Panel has identified two principal roles for ASCR in FSP. The first is as a close collaborator with FES. The Panel finds the complex, multi-disciplinary nature of FSP calls for close collaboration among various discipline scientists including working in multi-disciplinary teams to ensure integration of physics, mathematics, and computer science. The Panel finds the SciDAC approach of instituting multi-disciplinary project teams—involving FES and ASCR researchers—will be a highly productive model for FSP to follow.

In the second role, ASCR will act as a provider of science, mathematics, technology, and facilities to FSP. ASCR basic research will provide new knowledge in applied mathematics, algorithms, and programming. The Panel finds that because the technical challenges are so great and the project lifetime relatively long, a strong element of basic research conducted collaboratively among all disciplines will be important to the success of FSP.

ASCR applied research will provide new tools and libraries and contribute to new FSP application software. The Panel envisions that in addition to developing libraries and tools, ASRC will distribute and maintain important software packages. ASCR’s experience in software development will benefit FSP software engineering. ASCR must also work jointly with FSP to focus more attention on software quality assurance (SQA) and define its associated best practices.

Finally, ASCR facilities will provide peta- and exa-scale hardware and software platforms, networks, software environments, and user services. Leadership Computing Centers, the National Energy Research Scientific Computing Center (NERSC), and ESnet are expected to be the primary providers of computational capabilities for FSP, and ASCR needs to begin investigating specific FSP requirements such as allocation and on-demand scheduling, in addition to computational resources of all kinds.

In addition to addressing ASCR’s critical technology challenges and roles, the Panel addressed the challenges ASCR is anticipated to face during execution of FSP. The Panel expects there may be significant sociological issues associated with the FSP effort. In particular, spending time on project duties rather than research leading to publication can potentially have a negative impact on researcher motivation and careers. FES and ASCR must find ways to motivate and reward FSP researchers and make the effort to seek out “hybrid” researchers such as mathematicians and computer scientists that have learned physics, and physicists that have learned mathematics and computer science.

Insertion of ASCR research and technology into FSP presents another challenge. While FSP is active, there will be a chain from basic research knowledge to applied research technology to project engineering product. This process is likely to apply pressure to produce that will, at times, be in conflict with ASCR’s research culture. This is a potential issue that needs to be addressed by defining requirements and expectations for researchers early on.

Good software engineering practices will be essential to FSP's success, and along with this goes its enforcement. The project must have the management structure and delegate authority as-needed to ensure project personnel agree on software design, data structures, and interfaces.

The final implementation issue is the productivity of software developers. Here, ASCR can call on base program and SciDAC projects aimed at improving productivity, on many ASCR libraries and tools that aid in good software productivity, and on the expertise developed in producing ASCR libraries and tools.

2.0 Introduction

2.1 The Charge

Dr. Orbach's charge letter of October 17, 2007 (see Appendix II, Section 8.2) requests the Advanced Scientific Computing Advisory Committee (ASCAC) to consider previous reports and recommendations for an Office of Science Fusion Simulation Project (FSP) and make recommendations as to the most beneficial role for the Advanced Scientific Computing Research (ASCR) Program. The project, to be led by FES with collaborative support from ASCR, will provide the capability to simulate magnetically confined burning plasmas with unprecedented physics fidelity and utility using leadership class computers. This capability is vital to U.S. participation in ITER, for demonstrating fusion energy production beyond ITER (DEMOnstration Power Plant), and for scientific discovery.

ASCAC was specifically asked "to consider what is being proposed with particular attention to the most critical challenges in applied mathematics, computer science, and computational science, and to recommend an appropriate and mutually beneficial role for ASCR in FSP." Consequently, ASCAC formed a Panel (see Appendix III, Section 8.3) to address the charge. The Panel reviewed the two inputs suggested: (a) FSP Workshop Report [http://www.lehigh.edu~infusion/FSP.pdf] chaired by Kritz and Keyes, May 16-17, 2007 and (b) the Fusion Energy Sciences Advisory Committee (FESAC) FSP Panel Final Report [http://www.sc.doe.gov/ofes/fesac.shtml], October 30, 2007. In addition, the panel considered two earlier reports: J. Dahlburg, et al., J. Fusion Energy **20**(4), 135-196 (2002) and D. Post, et al., Fusion Energy **23**(1), 1-26 (2004). Finally, the Panel held a workshop (see Appendix IV for workshop agenda, Section 8.4) April 30, 2008 in Boulder, Colorado to solicit additional views on challenges in applied mathematics, computer science, computational science, and ASCR's role in FSP. The workshop also addressed the topics of software engineering, organization, and conduct of a large, multi-disciplinary software project.

2.2 The Fusion Simulation Project

FSP is envisioned as a 15-year project funded at approximately $25M per year with major milestones at 5-year intervals. According to the 2007 FSP workshop report, at the end of five years, it will be time to prepare for ITER first operation and the vision is “to assemble a new, powerful integrated whole-device integrated framework that uses high-performance computing resources for the simulation of tokamak plasmas.” The five-year goal is “to provide the capability to perform the calculations needed to support ITER diagnostics, plasma control and auxiliary systems design, and review decisions.” The project will focus on a limited number of problems with a framework capable of supporting state-of-the art models interoperating with less demanding ones. In parallel, FSP will develop and/or deploy state-of-the-art scientific software and data management, mining, and visualization capabilities and remote collaboration technologies. The 10-year vision is “to develop a simulation facility required to meet the national scientific and engineering objectives of ITER throughout the remainder of its operational lifetime” with the goal to provide codes that “will be capable of comprehensive integrated time-slice analysis and will be used to develop sophisticated control systems that are actuated by heating, fueling, and current drive systems as well as external 3D magnetic coils.” The project will develop a comprehensive framework that will couple state-of-the-art physical models on multiple time and space scales optimized for the most powerful computer platforms. Finally, the 15-year vision is “to develop a simulation facility that will be sufficiently well validated to extrapolate with confidence to a DEMO reactor based on the tokamak concept or other more advanced magnetic confinement concepts” with the goal of having a “world-class simulation capability that will be used in many ways to get maximum benefit for ITER.”

2.3 Our Report

Our report is presented in four sections followed by a summary. In the first section, we present our overarching findings and recommendations for FSP as a whole. In the second section, we address the specific question of critical technological challenges for ASCR and we take the opportunity to identify those we think are especially challenging. In the third section, we examine ASCR’s role in FSP as a project partner with FES, as a provider of new knowledge and techniques resulting from basic and applied research, and as a provider of infrastructure including ultra-scale computing hardware and software facilities, networks, and user services. Finally, in the fourth section, we present our views on challenges ASCR is likely to face in executing its role in the FSP.

3.0 Panel’s Views on FSP as Proposed

As a result of reviewing the background reports and holding the 2008 Panel Workshop, the Panel has arrived at some overarching findings and recommendations.

3.1 ASCR Alignment

The Panel is very supportive of the FSP and believes it is a critical element in enabling the U.S. to maximize its investment in ITER and in moving forward to demonstrate the practical application of magnetically confined fusion as a practical energy source. FSP is poised at the confluence of three developments as illustrated in Figure 1. The first is the driving requirement for a burning plasma simulation capability within five years. The second is the emergence of petascale computing. The third is the assembly of knowledge and software accumulated under ASCR and FES research programs, including more than six years of richly productive SciDAC collaboration.

ASCR has historically been a strong contributor to magnetically confined fusion simulation and the Panel believes it is well poised to make the major contributions needed for FSP to be successful. The base programs in applied mathematics and computer science have contributed to fusion simulation for many years and continue to make vital contributions. Examples include the widely used Portable, Extensible Toolkit for Scientific computation (PETSc) package of linear solvers and the omnipresent Message Passing Interface Chameleon (MPICH) implementation of the standard message passing interface (MPI) protocol. The SciDAC program has created interdisciplinary research groups that have accelerated development of advanced simulation and fostered the development of an effective collaboration model. The National Fusion Collaboratory Project is one example of an early SciDAC project that used an FES/ASCR interdisciplinary team to develop and deploy computer science technology to benefit both fusion's simulation and experimental communities. ASCR's computer and network facilities have their origin in the magnetic fusion program of the 1970s. Today, ASCR provides some of the world's most advanced computer platforms at Oak Ridge National Laboratory (ORNL), Argonne National Laboratory (ANL), and NERSC and its leadership computing program is advancing the scale of computing to petascale, with exascale on the horizon.

The Panel finds FSP very well aligned with ASCR goals, and that magnetically confined fusion is a prime candidate to drive developments in applied mathematics, computer science, and ultra-scale computing. FSP provides an excellent opportunity to demonstrate the fulfillment of the ASCR visionary approach to modeling and simulation and carry out it programmatic themes. FSP is an application that will drive development of technology that can be applied across a wide spectrum of science applications and one that presents challenges serving to attract the best and brightest young scientists.

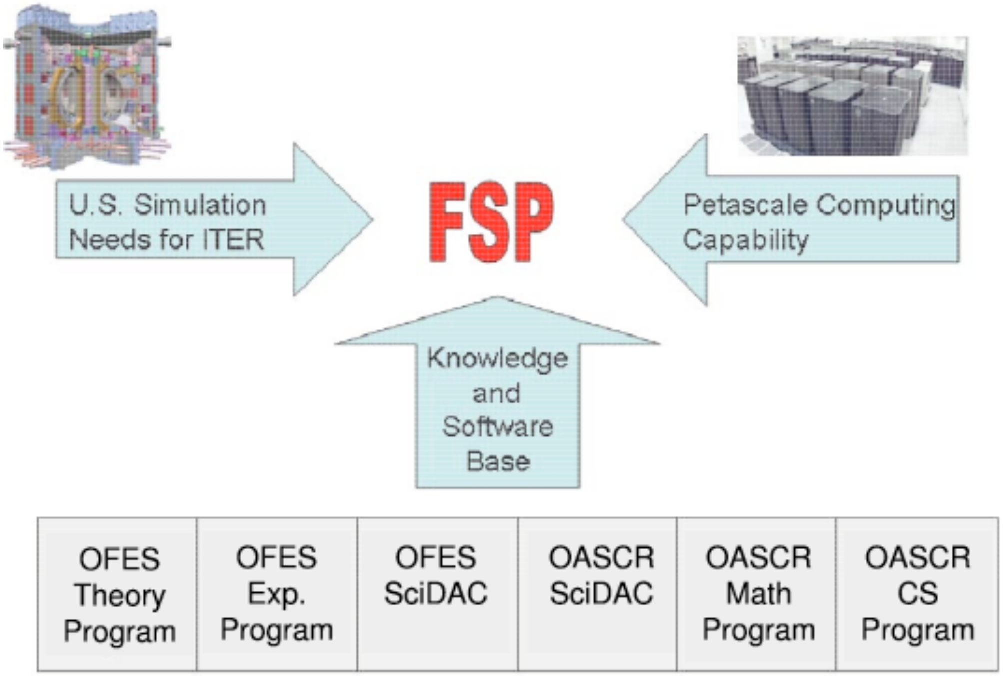

Figure 1: The Fusion Simulation Project is the result of a unique convergence (from 2007 FSP Workshop Report).

As stated in the FSP Workshop report: “the Fusion Simulation Program agenda for applied mathematics and computer science lies squarely on top of the ten-year vision statement ‘Simulation and Modeling at the Exascale for Energy, Ecological Sustainability and Global Security’ prepared in 2007 by the Department of Energy’s (DoE) Office of Advanced Scientific Computing Research.” FSP aligns with the vision’s new approach to modeling and simulation. FSP integration focusing on whole-system behavior of ITER goes beyond traditional reductionism focused on detailed understanding of components. FSP interdisciplinary simulations will spur development of frameworks and semantics that incorporate data from ITER experiments and will treat the whole phenomena extant in confined burning plasma system design, experiment design and analysis, and control. Validated simulations will capitalize on the ability to manage, visualize, and analyze ultra-large datasets.

FSP programmatic themes support the programmatic themes in ASCR’s vision. Top scientists and engineers will be engaged developing the science of complex systems and drive computer architectures and algorithms while providing a challenge to attract the next generation of computational and mathematical scientists. The ASCR and FES co-investment in FSP pioneering science contributes to advancing energy, ecology, and global security. The FSP will spur development of scalable algorithms, visualization, and analysis systems to integrate ultra-scale data with ultra-scale simulation. Finally, the FSP will provide strong justification to build advanced-scale computing facilities and an integrated network computing environment.

FES co-investment will foster developments that will serve a much broader base of computational science as they share many of the needs of FSP. The FSP will help spur development of widely useful software and be exemplar of an "exportable" software philosophy. The FSP will provide a clear example of the contribution of Leadership-Class computers in solving the nation's energy shortfall, lowering the impact of energy production on the environment, and providing energy security.

Finding: The needs and goals of FSP for computer science and applied mathematics support are very well aligned with the capabilities and goals of the ASCR program.

Recommendation: The Panel strongly recommends that ASCR seize the opportunity to participate in this exciting project of national importance in partnership with FES.

3.2 FSP as Defined

FSP was first proposed five years ago and the concept has undergone review by leading experts in fusion energy and computing. This work has resulted in a reasonably good understanding of broad scientific and technical issues. The current proposal, however, is rather vague in many respects as it lacks the specification of requirements, products, and processes needed to define a large software development project. FSP's true scope is not yet defined with regards to technical products, performance, and schedule. So it must be understood that many trade-offs will occur before FSP can be launched, and many more will occur during the project due to its great technical uncertainty. Thus, it is important to begin defining the project's scope now. As pointed out by the FESAC FSP Panel, some of the vagueness in the FSP proposal is a result of differing views as to what it is—research program, integrated computer program, software framework, or software tool suite. Compounding this issue is the lack of precise identification of the stakeholders, customers, and users who pay for FSP product development, pay for the product itself, and use it, respectively. A documented requirements specification defining what is needed, for whom, and at what level of risk; how it is to be used and by whom, is sorely needed. Without taking this additional step, critical technology timelines cannot be precisely identified, nor can precise roles be properly scoped.

Findings: FSP's true scope is not yet fully defined with regards to technical products, performance, and schedule; and stakeholders, customers, and users are not well defined.

Recommendations: A documented requirements specification defining what is needed and at what level of acceptable risk and an identification of the intended use of FSP products and their intended users is needed. We recommend ASCR participate with FES in the Project Definition Phase to begin producing such a document.

4.0 FSP Critical Technology Challenges for ASCR

Our Panel finds the 2007 FSP Workshop did reasonably well in identifying FSP critical technologies. We, however, identified additional areas and areas that need more emphasis.

The Panel views FSP as a very large, complex simulation project that can be characterized as comprising four major technical components: mathematical models, numerical algorithms, high-performance software, and facilities support. All these components present critical technological challenges to be addressed by ASCR programs in applied mathematics, computer science, and computational science. The magnitude of the challenge is great, as is illustrated in Figure 2 (Keyes), which shows a rough estimate of the modeling parameters required to simulate tokamaks on the current university scale (CXD-U), current large-scale (DIII-D), and ITER. The estimate of space-time points, assuming an explicit model and uniform grid, suggest that without algorithm improvements, about a 12-fold increase in scale is required to go from CDX-U to ITER simulations. Anticipated improvements in platform performance fall far short of bridging the gap and a multi-discipline approach is crucial to advancing model performance. The point is, while there are challenges in each program area, it is essential they not be treated in isolation, but rather in a coordinated, multi-disciplinary manner. An example of this is illustrated in Figure 3 (SCaLes report, vol 2). Here, computer platform and algorithmic improvements are combined to improve the effective sustained speed of micro- and macro-scale modeling in magnetic fusion simulations. Clearly, the technology challenges of FSP are such that they cannot be treated by applied mathematics, computer science, or computational science in isolation.

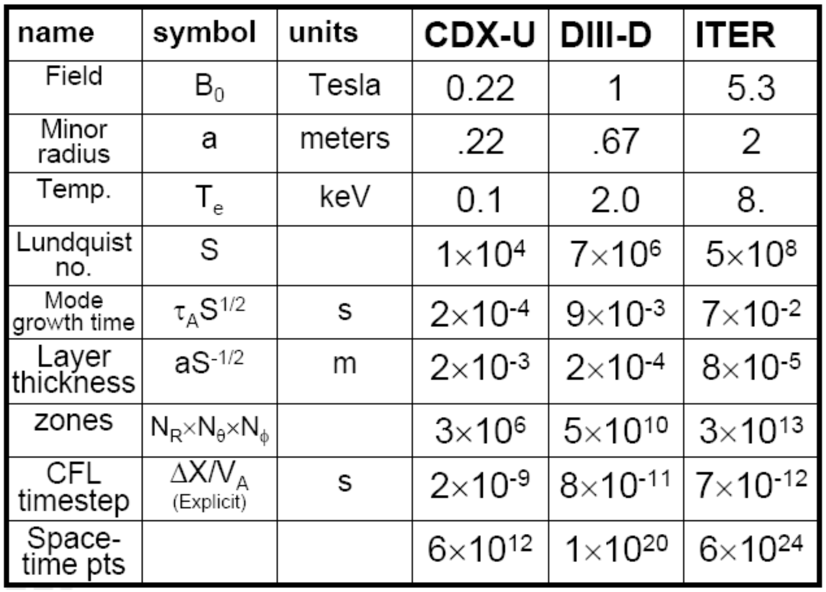

name	symbol	units	CDX-U	DIII-D	ITER
Field	B_0	Tesla	0.22	1	5.3
Minor radius	a	meters	.22	.67	2
Temp.	T_e	keV	0.1	2.0	8.
Lundquist no.	S		1×10^4	7×10^6	5×10^8
Mode growth time	$\tau_A S^{1/2}$	s	2×10^{-4}	9×10^{-3}	7×10^{-2}
Layer thickness	$aS^{-1/2}$	m	2×10^{-3}	2×10^{-4}	8×10^{-5}
zones	$N_R\times N_\theta\times N_\phi$		3×10^6	5×10^{10}	3×10^{13}
CFL timestep	$\Delta X/V_A$ (Explicit)	s	2×10^{-9}	8×10^{-11}	7×10^{-12}
Space-time pts			6×10^{12}	1×10^{20}	6×10^{24}

Figure 2: Estimate for scaling fusion simulations for ITER assuming an explicit, uniform grid (Keyes, 2008 Panel Workshop).

Finding: The 2007 FSP Workshop identification of critical technologies in applied mathematics, computer science, and computational science was reasonably complete. However, experience in large-scale fusion energy simulations shows that coordinating and combining research results from all such areas produces superior results.

Recommendation: **ASCR should coordinate its FSP research efforts in applied mathematics, computer science, and computational science to enhance focus and results for FSP.**

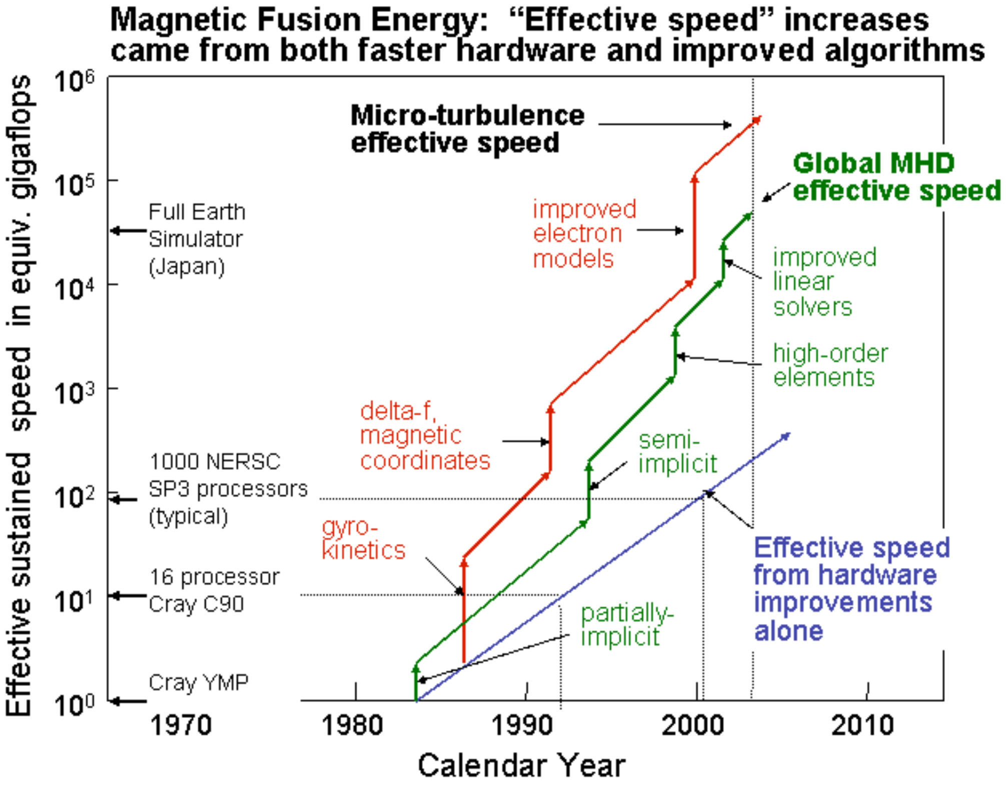

Figure 3: Magnetic fusion energy "effective speed" comes from faster hardware and improved algorithms (SCaLeS report, Vol 2).

4.1 Critical Challenges in Applied Mathematics

FSP anticipates the development of a suite of integrated simulation and analysis codes for modeling the behavior of complex plasmas. These codes will be used for designing and understanding experiments performed on ITER when it comes online. In addition, the codes can play a role in the validation of fusion physics theory. The codes must be targeted for the high-end computing platforms that will be available a decade, hence operating at near exascale speeds. The 2007 FSP Workshop report discusses in some detail many of the applied mathematics challenges FSP will face. These range from identifying the need for new computable physics models for treating some regimes of the plasma (e.g., gyrokinetic turbulence), to development of improved discretization methods, better scalable linear system, and iterative solvers, the mathematics for modeling multi-scale phenomena and multi-physics coupling, and the application of optimization methods. Areas the FSP Workshop report did not mention but which the present Panel found to be important include methods for incorporating experimental data in simulations, improved methods for sensitivity and uncertainty quantification, and methods for analysis of large computational and experimental datasets for the purpose of developing scientific understanding and improving experimental results.

There are a number of areas where the ASCR applied mathematics program can benefit the FSP through development of new mathematical approaches and computational algorithms. The FSP goal of producing a capability for the reliable prediction of plasma discharges in toroidal magnetic fusion devices on all relevant time and space scales will require development of a multi-physics, multi-scale simulation capability that embodies a number of simulation codes focused on distinct physics regimes and spatial and temporal scales. It has been only in recent years that coupling codes in this fashion has been attempted, and there are many open mathematical questions associated with such couplings. Coupling a set of well-posed single physics models together does not necessarily result in a well-posed problem; mathematical analysis is required to establish well-posedness of the resulting multi-physics system. If the single models involve different dominant spatial or temporal scales, this further complicates analysis of well-posedness. Without a well-posed system, a stable and accurate computational model cannot be created. Converting well-posed multi-physics or multi-scale models into practical numerical methods requires additional analysis to assure stability and accuracy of the resulting methods. If discrete simulation methods (e.g., particle-in-cell methods) are coupled with partial differential equation-based simulation methods, this analysis is further complicated by the nonexistence of a general stability theory for particle-in-cell methods. Developing algorithms to solve equations that are multi-scale in nature often requires mathematical analysis and research beyond standard approaches for single-scale problems. Recent advances in adaptive methods that address the multiple time- and space-scale natures of multi-physics problems can also be leveraged by FSP; however, further research may be required.

When simulation codes are to be used in designing and/or understanding experiments, it is natural to consider incorporating experimental observations with the simulations. This process of “data-model” fusion, where data is incorporated with models in an essential way to improve scientific understanding, has received very limited mathematical attention. The need exists to develop robust numerical methods for assimilating data into multi-physics models that are informed by numerical analysis-based error estimates for the simulations and statistics-based error estimates for the assimilated data.

Developing a true predictive capability through simulation requires not only the ability to perform complex simulations, but to understand uncertainties associated with the computations—much in the same way as experimental science assigns error bars to observations. Improvements in the mathematics of sensitivity analysis (understanding the sensitivity of simulations to perturbations in physical parameters) and uncertainty quantification (quantifying the propagation and effects of uncertainty and simulation error on predictions) will have a significant impact on the ability of future FSP simulations to be used reliably as part of the design and understanding of ITER experiments.

For some of the simulation regimes FSP will address, new models must be developed. This presents a challenge both for the fusion physicists who must identify and understand the physics to be modeled, as well as to applied mathematicians, who must help develop models that are realistically computable. An example is the development of gyrokinetic models for edge plasmas far from the Maxwellian regime. Since a complete treatment of edge plasma physics for first principles would be unrealistically expensive to compute, the challenge is to develop a predictive mathematical model of sufficient complexity so as to embody the essential physical processes and interactions at the plasma edge, while at the same time being realistically computable.

The anticipated quantity of data produced by FSP simulations is likely to be extremely large—possibly orders of magnitude larger than present or even future ITER experimental datasets. Current visualization and analysis techniques are unlikely to scale to these data sizes and thus, methods must be developed that can efficiently assist scientists in understanding that data. This understanding must include quantitative comparison of experimental and simulation data that will require the creation of appropriate synthetic diagnostics. Rigorous but computationally feasible methods for meaningful dimensional reduction of data will be an essential element in development of practical methods for data understanding.

4.1.1 Findings and Recommendations

Finding: The FSP is rich with opportunities for the development of new mathematical models, analysis techniques, and algorithms.

Recommendations: FSP should engage the applied mathematics community in many areas including methods for modeling multi-scale phenomena and multi-physics coupling, sensitivity analysis and uncertainty quantification, data-model fusion, the development of computable models, and effective methods for the analysis of computational and experimental datasets.

4.2 Critical Challenges in Computer Science

FSP development will entail taking advances from basic research in science and mathematics and forging them into new applications codes spanning scales many orders greater than today, and that run on exaflop systems. Large-scale tokamak plasma simulations will produce huge datasets to be managed and analyzed. FSP will create a fusion simulation facility supporting a large, distributed community of collaborating fusion scientists. These all present significant computer science technology challenges, many addressed in the 2007 Workshop Report. The Panel considers the issues of sustained performance on large simulation applications and data management, visualization, and collaboration to be especially important and elaborates on them below.

4.2.1 Sustained Performance on FSP Applications

Reaching the sustained performance needed to meet the demands of a burning plasma simulation is a daunting challenge indeed. It is made more challenging by the rapidly changing nature of ultra-scale computer architectures and an anticipated increase in platform diversity.

Power requirements for single-core microprocessors have become the limiting factor in increasing computational capability by simply increasing clock rates. The same overall throughput can be obtained using less power by placing multiple processing cores on a single chip. Quad-core processors are common now, with as many as 80 appearing in experimental processors. The supercomputers of the (near) future consist of multi-core processors. The fact that each core will run at a lower speed than today's desktops implies that such machines will have to have massive parallelism. An example machine delivered this year has 40,960 quad-core processors (IBM BG/P at Argonne).

Further complicating the scene is the use of heterogeneous processors in which the cores on a microprocessor may not all have the same capabilities or instruction set. Such a processor is the IBM/Sony Cell processor, used in the Roadrunner system at Los Alamos. Architecture constraints on power and clock speed will require multi-scale fusion applications to use vastly increased parallelism.

Supercomputing architectures evolve in phases. Periods of evolutionary hardware improvement have relatively little impact on application software development, since the programming model(s) remain relatively stable as the hardware evolves. When significant architectural changes occur however, the programming models can change dramatically resulting in significant, major changes to application software design and implementation. Most (but not all) computational scientists remember when large shared-memory vector supercomputers were replaced by the more scalable distributed-memory architectures. Computer architectures have been relatively stable for more than a decade, as gradually faster single processors and interconnection networks, combined with a standard approach to programming (the MPI standard library interface for message passing, together with advances in C, C++, Fortran, and Fortran 90 (now 2003) compilers, led to a substantial body of scientific software.

The architectural developments described in the preceding sections are putting pressure on the programming model embedded in current applications in which individual processes, each with a relatively large memory footprint, communicate via message passing. As the number of cores per chip increases, with only modest growth in memory size due to cost and power concerns, the amount of memory available to each process will decrease. While some applications may be able to scale down memory requirements per process, most (including fusion) will need to embrace a programming model in which data is shared multiple processes.

No single programming model addressing these new architectural challenges has become dominant. This is a fertile research area. ASCR supports the Center for Programming Models for Scalable Parallel Computing in which a number of approaches are being explored, including extensions to the Global Arrays and MPI libraries, the OpenMP language (with C, C++, and Fortran variations), and the partitioned global address space (PGAS) languages Unified Parallel C (UPC), Co-Array Fortran, and Titanium. To incubate transition of applications to such languages, hybrid approaches are also being studied. Such hybrid approaches may be a natural fit to the hardware architectures consisting of many multi-core and/or heterogeneous processors.

Going beyond these near-term (in the next few years) approaches, the Defense Advanced Research Projects Agency High Productivity Computing Systems program funded a small number of vendors to explore more dynamically parallel languages in which the notion of process is diminished and parallelism is expressed more abstractly. These languages have faded due to lack of community involvement and investment in porting them to multiple architectures. ASCR could adopt this collection of language efforts and encourage collaboration among vendors, university and laboratory computer science communities, and the applications community, with the aim of developing a practical, portable standard implemented on a wide range of architectures. The fusion community would represent an ideal family of applications that could both inform and benefit from such an effort.

We know computer architectures will change radically during the lifetime of FSP and most likely, continue to change when the resulting fusion simulation facility is operational. The near certainty of diverse and changing architectures will make it very challenging to maintain performance of a simulation code as it is ported to diverse platforms. New approaches such as auto-tuning and code generators will be needed, not only for FSP, but for all future large-scale modeling and simulation codes.

4.2.2 Data Management, Visualization, and Collaboration

Technology for managing large FSP-generated datasets, visualizing results, and sharing data and collaborating with geographically dispersed researchers, will be necessary during application testing and when the FSP facility becomes operational.

The FSP will generate data quantities several orders of magnitude greater than present-day fusion experiments or simulations. Furthermore, data requirements for FSP will be greater than most other fields. To be successful, this data needs to be efficiently managed so rapid and easy access is available to both fusion simulation and experimental communities. ASCR has developed expertise over the years in its distributed computing base program and in SciDAC applicable to high-speed file transfer and management of large datasets.

The distributed nature of the FSP team, combined with the anticipated large quantities of data, indicate that present-day visualization techniques are not sufficient to meet the project’s requirements. The output of large-scale simulations can only be fully understood—particularly in the qualitative sense—through scientific visualization. New visualization capabilities are needed to accelerate scientific discovery as well as for real-time sharing of graphics to the distributed team. Data analytics tasks will have to be carried out in part in situ to the actual simulation itself. Fusion simulation data will provide challenges to ASCR's visualization researchers, and their tools will be invaluable in understanding experimental results.

FSP will involve a multi-disciplinary team drawing expertise from within the Office of Science. This team will be geographically distributed across the U.S. since it is not anticipated that all team members will relocate to one institution. To be successful, the FSP needs to be able to support distributed code development and a wide range of coordination activities through shared applications and displays, as well as improvements in interpersonal and group communications integrated with extensible data services. This can only be accomplished through the development and deployment of new tools and technologies placed into the FSP working environment. ASCR, through its work on the previously mentioned Fusion Collaboratory Project, has been involved in development of advanced collaboration technology for FES with the Access Grid as one example. However, it is clear FSP, as well as ITER, require a toolset with significantly improved integration and reach, and ASCR research can have a significant impact on meeting this need.

4.2.3 Findings and Recommendations

Finding: Achieving sustained performance for FSP applications over time and across multiple platforms is made much more complex by the trend of ultra-scale computers toward many-core, heterogeneous architectures.

Recommendation: ASCR should support FSP by developing technology that ensures performance meets expectations as simulation applications move among diverse petascale to exascale platforms.

Findings: To meet its project goals, FSP requires the management of simulation-generated datasets, the visualization of these data, and the facilities and tools for collaboration by a worldwide research community. The panel notes this issue touches other Office of Science projects such as the U.S. work on ITER and the Large Hadron Collider (LHC).

Recommendations: As part of the FSP, FES and ASCR should dedicate resources to developing and deploying distributed data management technologies that will provide rapid and easy access to FSP data, visualization technologies that can be efficiently utilized for the anticipated large-scale data repository as well as for the requirement of real-time graphical information sharing, and collaboration technologies aimed at unifying a distributed scientific team. This finding and recommendation is consistent with that put forth in the FESAC FSP Panel Report, as well as the ASCR report on Visualization and Knowledge Discovery, and the report on Modeling and Simulation at the Exascale for Energy and the Environment.

4.3 Critical Challenges in Computational Science

4.3.1 Frameworks

The development and evolution of a suitable code framework is essential for a project such as FSP. Examples of other community-grown frameworks include Chombo, Cactus, SIERRA, and UPIC. A good framework clearly separates the roles and responsibilities of the expert programmers from those of the scientific domain experts by effectively defining a contract between these two groups.

A good framework enforces and facilitates software engineering style and discipline to help ensure accuracy. For a community code, the framework hides complex domain-specific parallel abstractions from the application scientists, thereby greatly reducing the complexity one must deal with while at the same time, enabling performance. It will allow domain scientists to code nominally serial plug-ins invoked by a parallel "driver" to enhance productivity.

To fit into the framework, a scientific module developer must be willing to relinquish control of the "main" part of his program. Note that for this to work, some standardization on the part of the science community is required. The hardest part will be agreeing upon physics interfaces. This must be initiated and managed by FSP. ASCR personnel can help, but development of a suitable framework will not be solved by a tool or "CS Technology" alone.

A well-designed framework will enable computer and computational scientists to work together in a productive manner. It will virtually eliminate arguments about C++ vs. FORTRAN, and enable easy unit-testing. It will also provide standard functionality such as input/output, communication libraries, and math libraries and solvers. It will enable multidisciplinary collaborations to provide features that would not otherwise emerge in stand-alone codes. It also permits sharing of physics modules among computational scientists.

4.3.2 Workflow

Workflows are presently being used in some fusion simulations, for example, in the Center for Plasma Edge Simulation (CPES) to couple two simulation codes. Scientific workflow is a generic term describing a series of structured activities and computation (workflow components or actors) that arise in scientific problem-solving. This description includes actions performed by the actors, the decisions made (that is, the control flow), and the underlying coordination such as data transfers and scheduling, which are required to execute workflow.

In its simplest case, a workflow is a linear sequence of tasks, each one implemented by an actor. For example, the Kepler workflow tools, developed in collaboration with the SciDAC SDM center, are being used in CPES to run a simulation (M3D) on one machine based on the output of another simulation (XGC) run on a different machine. Scientists use this workflow to submit a job request, and then monitor progress of the workflow as their simulation is running.

The specific tasks performed by the workflow include: submitting a request to the batch system; waiting for the simulation to begin executing; identifying simulation output; transferring output to storage; performing simple analysis on the output; and generating logs that track the current simulation status. Workflows can exhibit and exploit data-, task-, and pipeline-parallelism. In science and engineering, process tasks and computations are often large-scale, complex, and structured with intricate dependencies. Workflows are most useful when a series of tasks must be performed repeatedly.

While current workflow technology is extremely useful, there is still much work to be done before scientists are able to effectively utilize these tools. In particular, better interfaces need to be designed to support quick workflow development and monitoring, the tools need to be extended to better track both data and workflow provenance, and capability-based actors need to be implemented to encapsulate higher-level actions (e.g., a transfer actor instead of ftp, scp, and cp actors). In the area of provenance, workflow environments offer unique advantages over script-based solutions in that they can keep track of the processing history and data dependencies. Provenance information can be used to inform scientists about their results (for example, debugging and re-interpretation of results), or to increase fault tolerance (re-run from checkpoints), or to increase efficiency (smart rerun).

4.3.3 Verification and Validation

Verification ("solve the equations right") and validation ("solve the right equations") is an essential component of a complex simulation package such as that envisioned for FSP. It must take place at many levels, and must be facilitated by the overall software engineering design and framework.

At one level, unit testing and nightly build processes are needed to keep the code base "functional." Strict revision control policies must be developed and enforced. These are good best-practice examples of software verification and configuration management. For "physics verification," or the assurance that numerical solutions are correct and accurate, relatively new techniques such as the method of manufactured solutions (MMS) should be explored within FSP.

Many of the code verification exercises will take place at the module level, but the capability to verify scaled-up models and combinations of modules is also required. The FSP would be well advised to learn some lessons in this area from the DOE NNSA Advanced Simulation and Computing (ASC) Program. Since ASCR does not have many specific tools or capabilities in this area, FSP should have a cooperative effort with ASCR to create tools to facilitate verification and validation (V&V). Visualization and other data analysis tools will play an essential role in this effort.

4.3.4 Findings and Recommendations

Findings: Computational frameworks, utilization of workflow tools, and systematic V&V methodology are all important for the success of FSP.

Recommendation: ASCR has considerable experience and expertise in each of these areas and should take the lead in recommending appropriate solutions to FSP.

5.0 The Role of ASCR in FSP

The Panel has chosen to look into two distinct roles ASCR can play within FSP. The first is as a collaborative member of the FSP team. The Panel understands FES will lead the project and the Panel believes ASCR's most effective and mutually beneficial role is as a close collaborator with FES, as discussed below.

The second distinct role for ASCR is as a provider of knowledge, technology, and facilities to FSP. The scope of that support is broad and includes providing:

- New knowledge, especially in applied mathematics, algorithms, and programming, which will lead to solving problems we cannot currently solve;
- New software—at least in the form of demonstration or pilot codes—that solve specific FSP problems or provide tools that support problem solving;
- High-performance computing and network hardware and software platforms on which to perform FSP simulations and the software environment and user services necessary for efficient, productive use of the platforms;
- Software packages and infrastructure that will be shared with other programs within the Office of Science.

5.1 ASCR Collaborator Role

The Panel envisions FSP will be led by FES with highly collaborative support from ASCR and targeted co-investment. The Panel finds the complex, multi-disciplinary nature of FSP calls for close collaboration among various discipline scientists to ensure the correct integration of physics, mathematics, and computer science. Because FSP requires specified deliverables and schedules, the collaboration must be especially productive. To be most productive, ASCR scientists need intimate knowledge of the day-to-day progress and issues of the task they are supporting. Working on a segment of the problem in isolation and "throwing the solution over the fence" will not be productive, and in fact, may be counter productive. ASCR scientists should, when called for, become dedicated members of FSP project teams and co-located with the team. This also facilitates accountability for any agreed-upon deliverables. When co-location is impracticable, an advanced network-based collaboration environment will be required.

The SciDAC projects in ASCR have led the way in demonstrating that co-investment and direct collaboration between ASCR applied mathematicians and computer scientists on one hand, and application scientists on the other hand, can be successful. Indeed, multiple

SciDAC projects, including centers, Science Applications and associated Partnerships (SAPs), and institutes, have been based on fusion sciences. We anticipate future collaboration between fusion scientists and ASCR scientists will follow similar patterns, whether explicitly organized under SciDAC or not, in the context of an overarching FSP.

FSP will be a complex project involving multiple subprojects, each with a variety of codes or methodologies needing to be improved, scaled, or replaced and interfaced into the overall software architecture of the FSP as a whole. Rather than a large "physics team" and an "applied math/computer science" team, we envision the project management structure will remain on the physics side, with each ASCR scientist interacting closely with a small number of physicists on one or more specific subprojects in the FSP structure.

Such a structure facilitates communication and focused collaboration, as has been demonstrated in the past, both in fusion collaborations and others. The ASCR resources are brought to bear on specific subproject problems. There are often unplanned benefits of close collaboration as knowledge and experience is exchanged informally. For example, this is a conduit for software quality assurance practices, common on the ASCR side, to make their way into the physics subprojects. In the long run, such practices will become part of FSP standards, however adoption of such practices is easier if it is already familiar through infusion into the individual subprojects.

In the Panel's judgment, the collaborator role will be beneficial to FSP and ASCR. For FSP, close collaboration will result in better communication of requirements, more focused knowledge, faster paced implementation, and more successful outcomes. For ASCR, being an effective collaborator will enhance FSP success and contribute to fulfilling its programmatic themes. The FES will be an important driver for new ASCR technology and provide challenges to attract the next generation of ASCR researchers.

5.1.1 Findings and Recommendations

Finding: The technical and project demands of FSP will call for very close teamwork between ASCR and other FSP scientists and engineers.

Recommendation: ASCR should adopt a policy of its scientists participating in integrated task teams when dictated by the needs of FSP. We envision the project management structure will remain on the physics side, with each ASCR scientist interacting closely with a small number of physicists on a specific subproject in the FSP structure.

5.2 The Role of Basic Research in Applied Mathematics and Computer Science

The use of computational models to simulate physical events is one of the most important developments in science and technology of the past century. The basic DOE research programs in applied mathematics and computer science have contributed many essential enabling research advances that have made this possible. Fundamental mathematical developments have translated scientific theory into discrete equations a computer can solve, as well as the mathematical and numerical analysis that provides basic understanding of both the scientific theories and their numerical counterparts. Ambitious projects such as FSP are possible not only due to the exponential growth in speed of high-performance computers over the past decades, but also because of equal or greater increases in capability resulting from research advances in applied mathematics and computer science. An example of coupled research is illustrated in Figure 3 in the case of magnetohydrodynamic simulations described in the SCaLes report, vol 2.

When embarking on a project such as FSP where the goal is to produce a production-quality simulation capability with high-confidence and quantifiable output, milestones on the critical path cannot be strictly dependent on unattained basic research results, which by their very nature, cannot be anticipated in advance. The plan for such a project should be based on results already known in both science and mathematics, or on results that can reasonably be anticipated and produced through applied research. This observation does not negate that fact that new basic research results, when obtained, can provide better solutions than those planned for, with significant enhancements of the resulting capability. In addition, investing in this basic research can help hedge against the risk that some of the planned technology may not function as expected in the final product. Thus, a strong element of basic research, including basic research in applied mathematics and computer science, will be important to the success of FSP. The basic research component will likely be more prevalent in more of the higher-risk areas of FSP.

The role of ASCR basic research in advancing FSP will be to provide new mathematics and computer science results that can enhance the computational simulation and analysis capabilities the fusion community will need to design and understand experiments on ITER and other large fusion devices. Advances supported by basic research in applied mathematics can include development of improved, computable physics models, the development of new algorithms for implementing these models computationally, and algorithms for analyzing and understanding results of the models. Because of the extensive challenges faced as a result of the emergence of multi-core architectures, the distinctions between required applied mathematics and computer science advances are becoming increasingly blurred.FSP should provide a wealth of opportunities for investigation by the ASCR Applied Mathematics-Computer Science Institute proposed for FY2009.

Development of models that provide new or improved descriptions of physics and that are computable can be most effectively achieved through sustained partnerships between researchers in fusion physics, applied mathematics, and computer science. Such partnerships are equally important in the development of numerical algorithms for solving these models on computers. Intimate knowledge of the physics can often be used to develop more effective numerical tools targeted at all stages of a computation.

Continual, sustained partnerships among researchers in the different fields are essential for the success of these efforts. The SciDAC model, where collaboration is encouraged and supported through co-funding of efforts between ASCR and the science offices, might well be employed in developing needed partnerships at the base research program level. In addition to promoting collaboration through an appropriate funding model, mechanisms need to be provided which encourage frequent interaction among the researchers involved. Co-locating mathematicians or computer scientists with application science projects is one way to accomplish this. Providing infrastructure that encourages and enhances remote collaboration can bring larger teams together on a regular basis. It is important to understand that it will be challenging to arrange for these close partnerships, particularly when individuals are resident at different institutions. Close attention should be paid to assuring that math and computer science researchers in academic environments are given equal opportunity to participate in these partnerships as their colleagues at national laboratories.

5.2.1 Findings and Recommendations

Findings: A strong element of basic research including basic research in applied mathematics and computer science will be important to the success of FSP.

Recommendations: The SciDAC model, where collaboration is encouraged and supported through co-funding of efforts between ASCR and the science offices, might well be employed in developing needed partnerships at the base research program level. In addition to promoting collaboration through an appropriate funding model, mechanisms need to be provided which encourage frequent interaction among the researchers involved.

5.3 The Role of Applied Research

ASCR applied research is a necessary step to forge new basic knowledge in applied mathematics and computer science into technology that can be utilized by FSP. For example, it is through applied research that a new, multi-scale algorithm is applied and tested on a real fusion simulation problem or a new concept in many-core, hybrid programming of an FSP model is realized.

Perhaps the most direct role for ASCR applied research is working collaboratively within FSP to take the results of relevant basic research in applied mathematics and computer science from ASCR or other sources, such as the National Science Foundation (NSF), and apply them directly to FSP applications. `This has proven to be effective in SciDAC where co-funded research has produced significant advances in confined fusion simulation, as exemplified by the introduction of both structured and unstructured mesh refinement capabilities and algebraic multi-grid solvers into several of the major fusion codes. It should be noted however that for this mode of applied research to be most effective, researchers must be willing to assist in the transition of their results to production quality software.

Another ASCR applied research role important to FSP is development of software packages that can be applied across a wide spectrum of DOE science domains. There are many examples of important software packages, such as MPICH and PETSc. These software packages are distributed as open source free software and thousands of downloads per month attest to their usefulness to the application development community. This distribution method allows research results to be evaluated for practicality and impact, and users provide feedback, stimulating further research. Thus this relationship between computer science research and software distribution is stable and will prove useful for FSP. Equally important for FSP is that ASCR recognizes the value of this software and supports continued development and maintenance of the most actively used packages. At present, a web page is being prepared for the ASCR website listing the software available together with the software's own web page and e-mail contact information.

ASCR software will be available to FSP in various ways, depending on the individual software itself. For example, it may already be installed at the user's local computer center as there is currently an effort underway to identify an "HPC Software Stack" that will be uniform across large-scale computing centers, possibly including the NSF as well as DOE centers. Some more specialized packages may not be pre-installed but they can be downloaded from a project website and installed locally, either by an individual scientist or his systems support team. Finally, since all ASCR software is distributed as open source, the software may be downloaded and then customized to meet individual requirements.

An additional benefit of ASCR's software development efforts is gained knowledge and experience in developing software. It is generally agreed that a high level of software engineering practices will be necessary if FSP is to succeed. ASCR does not have a specific "general software engineering" program, however many ASCR researchers exhibit excellence in software engineering practices. Such practices not only include code design itself but also quality assurance practices such as nightly builds for portability testing, automated regression tests, overage tests (e.g., is all of the code exercised by the regression tests?), and bug tracking systems. It is likely that through individual collaborations, these practices will be infused into FSP.

5.3.1 Findings and Recommendations

Findings: ASCR applied research will play a critical role in the practical application of new fusion simulation-specific, as well as general knowledge in applied mathematics and computer science as a necessary step to meeting FSP goals.

Recommendation: In addition to playing a strong collaborative role in FSP applied research, ASCR should take steps to ensure applicable results of other applied research supported by ASCR be made available and if necessary, tailored to FSP.

5.4 Role of Facilities

In the FESAC FSP Panel Final Report (Tang et al., Oct. 30, 2007), the FES Advisory Committee recommended that "the computational and software infrastructural requirements for the FSP be communicated early and often to those organizations providing computational and data capabilities for the Office of Science, such as the Leadership Computing Centers, ESnet and NERSC." We concur with this recommendation and would encourage FSP to meet with representatives of these facilities early on, as well as the ASCR staff, to provide as many details as possible. In addition, we would like to suggest a few areas for further investigation.

The ASCAC FSP subcommittee met with representatives of the two Leadership Facilities and NERSC to gather input on how the facilities could support FSP. The consensus was that the leadership facilities would initially have both the computational resources and staff to support an activity such as FSP, but not without redirecting these resources from programs like INCITE for which they are currently fully subscribed. There were, however, several areas where additional resources (appropriately applied) could be of great benefit to both ASCR and FES. These areas include computer allocation resources, in particular, with respect to on-demand computing in support of ITER experiments, software build systems, software package management, and other general software engineering and software quality assurance (SQA) processes. FSP could very well need, for example, one or more dedicated capacity systems for development, V&V, testing, and small-scale simulations. This committee believes the ASCR facilities are equipped to operate these types of systems.

5.4.1 Hardware Infrastructure

Currently at NERSC, 28% of the computer cycles are allocated to fusion energy, spread over 76 fusion codes. Of this workload, over 50% of the cycles are devoted to just 5 codes: OSIRIS, GEM, NIMROD, M3D, and GTC (see Figure 4).

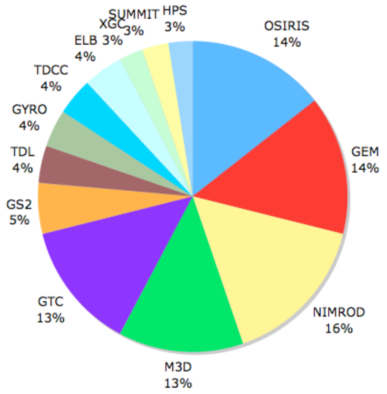

Figure 4: Breakdown of Fusion Codes at NERSC.

Currently within the 2008 INCITE Program allocations at the Oak Ridge National Laboratory (ORNL) Leadership Computing Facility (LCF), four of the 30 INCITE projects are fusion related, comprising 7.2% of the total 145.3 millions processor-hours (see Figure 5). As part of these projects, many fusion codes are present and execute in a reasonably efficient and scalable manner on the Cray X1E and XT4 leadership systems.

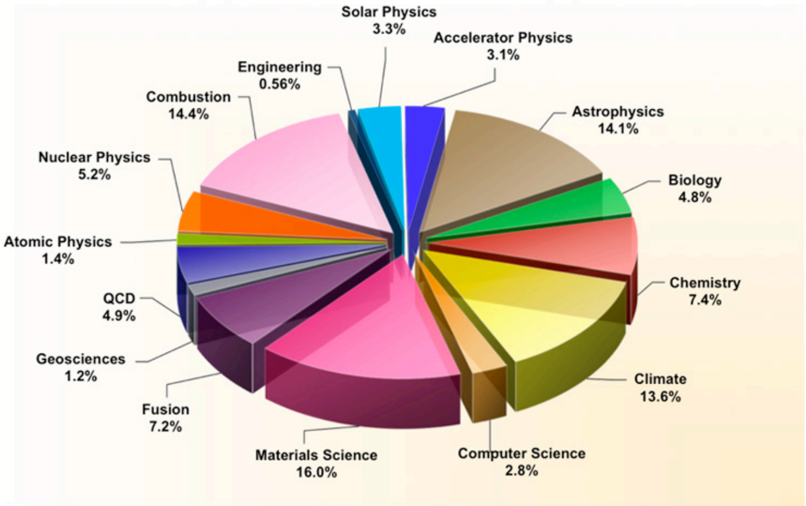

Figure 5: INCITE 2008 Allocation Breakdown by Domain at the ORNL LCF.

An incomplete list includes codes such as AORSA, GYRO, GTC, XGTC, CQL3D, M3D, TORIC, NIMROD, DELTA5D, GEM, and AMRMHD. Two of these fusion codes were selected for the highly visible and high consequence Office of Management and Budget (OMB) Joule metric for ASCR, namely the GTC and GYRO codes in 2007 and 2008, respectively. GTC was also selected as an early access science-at-scale application for the recent Cray XT4 263 TF upgrade, and AORSA achieves a remarkable percent of peak (>50%) on this same system. The existence of this suite of codes demonstrates that the fusion community is capable of building large, high-impact scientific applications capable of effectively utilizing present and future leadership systems.

The committee understands FSP will increase the demand for computational resources ASCR is currently providing to FES. In fact, the current suite of fusion codes run on the LCF systems place some of the highest demands on resources such aggregate compute power and memory, and local storage capacity. The current expected increase in resources (out to 2012) should help to some extent. However, there are specific resource needs for FSP that currently do not fit into the anticipated computational resources. In particular, FSP will probably require a much larger, dedicated allocation over a longer time period than what is currently available. In addition, the Panel believes FSP will require on-demand computing to be able to do “between-pulse” experimental data analysis using reduced physics models. This may take the form of separate queues, dedicated clusters, or some other yet-to-be-determined means.

Another opportunity for both FES and ASCR is in the area of data storage and high-speed networking for all archived FSP runs. The committee believes ASCR should investigate the development—jointly with FES—of a set of data repositories for fusion simulation data. As an example, we point to the repositories established by the Community Climate System Model project. For example, MDSplus, the common data management system for fusion experimental data, has greatly facilitated inter-machine collaboration and comparison. In a similar way, the Panel believes such data repositories for FSP data will lower the barrier for both comparisons of simulation data to experimental results, as well as the comparison of different simulation models. Additionally, such repositories will allow scientists to search the data for patterns and features.

5.4.2 Software Infrastructure

As currently envisioned, FSP would encompass many different codes, components, and/or modules that must be developed quickly and maintained over a long period of time. These simulation codes are comprised of a mix of codes written in different languages and that contain a number of other package dependencies. In addition, the codes must be portable to a number of different architectures. Most likely, the current set of software infrastructure tools will not be adequate in the near future. One possible role for the facilities is in leading the community toward a common set of software management tools for high-end computing. Such a set of tools could contain, for example, build and package tools to aid in the development of complex high-end simulation codes. In addition, the facilities might be able to provide advice on other software management techniques such as regression testing, unit testing, and component test suites. The latter, in particular, should be done at scale whenever possible. Another possibility is to have the facilities provide tutorials on the use of common software management tools such as make, cvs, and svn. The facilities could also benefit FSP with requirements-driven package management (e.g., math library middleware), porting, installing and maintaining visualization tools, and ensuring availability of multiple compiler options. These activities are most easily managed if ASCR were to support liaisons specifically assigned to FSP at the LCFs.

ASCR should also support FSP through their existing visualization and data analysis efforts, both in the Base programs and through SciDAC. Existing SciDAC projects such as VACET are already working closely with fusion scientists to improve visualization of datasets from particle-in-cell simulations and through interactions with other SciDAC projects such as FACETS, CEMM, and GPS-TTBP. These efforts have led to the development of software prototypes that provide fusion scientists with key items for post-query visualization such as the ability to quickly view the path of a particle of interest, and to quickly view the effect of the potential field on the particle over time. In addition, experimental fusion science also places a premium on rapid data analysis in near real-time to support tokamak operation. Other activities such as the current end-to-end workflow automation and “dashboard management” efforts at the LCFs meet FSP needs and would benefit from increased support and focus on behalf of FSP by ASCR.

In the future, FSP will generate significantly larger amounts of simulation and experimental data, all of which will need to be analyzed. The ability to perform “between-pulse” data analysis benefits experimentalists today by making the best use of their allocated experimental time, as it will in the future. FSP will also require the ability to easily find and retrieve data from storage. Techniques to do this are being developed at SciDAC centers such as SDM and VACET. These centers, or similar future efforts, should be encouraged to work with FSP to extend their technologies to the fusion domain. The facilities should then undertake the role of providing production tools based on these technologies.

5.4.3 User Services

While the current level of user services is excellent at all facilities, the Panel believes they do not provide adequate coverage in support of fusion sciences—particularly the "level-two" application support, namely trained and experienced expert fusion computational scientists who engage with the FSP at the model formulation, algorithm development, and software implementation level. User services will need to scale up to match the increase in users and scientific programs requiring new computational power to give FSP the best chance to succeed.

5.4.4 Network Infrastructure

The Panel believes FSP will be implemented by a richly diverse and widely distributed community of researchers, developers, and scientific users. This "virtual FSP community" will depend on the network infrastructure and services provided by ESnet4 to remotely run FSP test and production simulations, and for data access and analysis. ASCR should also accept the role of providing network services and support for collaboration within the community. The fusion community is probably second to the climate community in its need for the sharing and transfer of large datasets. This need is a major driver for high-end network infrastructure.

5.4.5 Benefits to ASCR

Procurement of future machines is guided by benchmarks, which in turn are taken from existing workloads at the facilities. There is a need to establish collaboration between the fusion community and ASCR to define simplified FSP benchmarks that represent fusion requirements for future systems. Areas where it would be particularly useful to gather requirements are the amount of spatial and temporal resolution needed, and performance data related to both weak and strong scaling. A partnership between ASCR and FES to determine an appropriate set of requirements for the FSP would benefit both offices tremendously. The committee suggests that current and future SciDAC performance projects target FSP as a key customer, and that any facility performance tracking and modeling activities also place FSP application as a high priority.

5.4.6 Findings and Recommendations

Findings: The FSP will have special computational needs with regard to capability, capacity, on-demand computation, and data-storage.

Recommendation: A partnership between ASCR and FES to determine an appropriate set of hardware requirements for FSP would benefit both offices. Special solutions will need to be developed such as separate queues, dedicated systems, and special data repositories.

6.0 ASCR Challenges in FSP Execution

A project as complex and technologically challenging as the FSP presents many execution challenges for all participants. The Panel anticipates that ASCR, because of its research culture, will find participation in FSP particularly challenging. Examples include rewarding and motivating researchers, transitioning research results into production codes, implementing software engineering standards, and maintaining high software development productivity.

6.1 FSP Sociology

There are significant sociological issues associated with a large code development effort that need to be taken into account. It is a high-risk proposition for an application scientist to pursue a career as an algorithm/software developer. He or she will be perceived by their peer group as spending a lot of time "not doing physics." It can also be risky for an applied mathematician or computer scientist to dedicate time to code development instead of research. The turnaround time for ambitious code projects does not lead to the steady stream of publications that seem to be required for survival in some institutions. The present practice of rewarding scientists primarily on publication count and citations does not lead to a pool of talent from which to draw the core FSP development team.

The FSP will be judged on production of simulation capabilities and the impact of those capabilities—not publication lists. The kinds of people who can make this work will, of necessity, be multi-disciplinary: math/computer science people who have learned the physics, and physicists who have learned the math and computer science. There are serious questions as to whether we have enough such people, and if not, whether we can train them and give them promising career paths with appropriate professional development.

For the project to succeed, it must develop focused, self-motivated teams. The same metrics should be applied to all team members, and there must be viable career paths for all participants.

Both FES and ASCR must realize that this will be a long-term project, and that they cannot hold the participants to overly restrictive, very short-term accountability. It needs to be recognized that some approaches will fail while others will require a long-term investment to succeed.

6.2 ASCR Technology Insertion into FSP

The FSP will benefit from improved technologies in many areas including models, solvers, programming of multi-core platforms, V&V, and uncertainty quantification, only if these technologies are incorporated into production simulation codes and procedures. Efficient rapid prototyping will reap benefits if it can be done within the FSP software infrastructure. This incorporation can be characterized by a process that first involves gaining new knowledge through basic research, then converting this knowledge into technology through applied research, and lastly, transitioning this technology into a useful product through engineering. Because the FSP is so dependent on new technologies, all three of these processes will be active simultaneously. However, each of these processes has its own, and quite often, different culture that also exist simultaneously. The Panel believes it is important FSP address this issue early and establish requirements and expectations for basic researchers, applied researchers, and developers. Few examples exist, but the Community Climate System Model (CCSM) community, which has heavy involvement by NCAR and DOE Office of Science, has been addressing this issue. CCSM presents a model that should be studied carefully as FSP goes forward.

6.3 Software Engineering

The use of good software engineering practices is absolutely essential for a successful project. FSP can draw upon the software engineering expertise and experience of ASCR personnel through close collaborations and various forms of education. ASCR has developed many software tools that support good software engineering practices; however, these tools alone are not a replacement for disciplined software engineering and management that must be learned, applied, and enforced. ASCR and FSP should also appeal to the best practices established in other advanced simulation programs within DoD and the DOE/NNSA tasked to build large software products to inform high-consequence decisions.

The overall software design must factor the algorithm space into re-useable components and use a layered design to insulate applications from changes in hardware and systems software. High-level organizing constructs will facilitate development by teams of programmers. The overall design structure must also take into account the fact that both models and algorithms will evolve during the project, and that some approaches will fail to perform as expected. Risk mitigation is essential, and there must be a facility for rapid prototyping to test ideas. A good software design will integrate software verification activities and performance tuning into the process.

Development of standard data structures and software practices and documented interfaces between components is a central issue. This is really a sociological issue—not a technological one. ASCR technologies such as the Common Component Architecture can facilitate the definition of interfaces, however the project personnel must ultimately agree on the most appropriate software design, data structures, and interfaces. Uniformity will foster collaboration and go far in leveraging contributions. Uniformity in cross-platform binary file formats will greatly facilitate development of post-processing tools that support feature extraction and detection, fast indexing, and data queries.
The project must have a management structure to break deadlocks when people cannot agree on standards. Some person or group must be empowered to make necessary decisions that are not always going to be popular.

6.4 Application Development Productivity

A challenge facing all applications of significant size is the productivity of software developers engaged in writing actual code. ASCR has a number of both base program and SciDAC projects aimed at improving productivity. Some of these have already been mentioned in the above section on programming models which directly affect productivity. There are other efforts, both in computer science and applied mathematics, that produce mathematics libraries, performance tuning tools, and other libraries and tools that make it easer for programmers to produce correct, efficient, parallel code.

Slightly more speculative are efforts to improve the development environment itself. ASCR computer scientists have experience with build systems, source code management systems, regression testing frameworks, and other elements of modern software development that will be needed by FSP. Transmission of some of this expertise may best occur as a side-effect of direct collaboration recommended by the Panel.

6.5 Findings and Recommendations

Findings: ASCR has a predominately research culture, and participation in FSP will present significant challenges including attracting and retaining researchers, transitioning research results into production software, maintaining good software engineering practices, and maintaining a productive development environment for high-quality software.

Recommendations: ASCR should:

(a) Develop and implement methods to attract, motivate, and reward researchers to participate in FSP;
(b) Early on, address the challenge of fostering a culture in which research creativity exists alongside project engineering discipline, and one that is capable of transferring new applied mathematics and computer science knowledge all the way from research paper to FSP production code;
(c) Establish and enforce good software engineering standards;
(d) Establish a software quality activity tasked to gather and disseminate SQA best practices into FSP (especially in the areas of software verification and testing);
(e) Leverage existing software development productivity tools and further efforts to improve the software development environment.

7.0 Conclusions

The Panel concludes that FSP is well-aligned with ASCR goals and recommends ASCR seize the opportunity to participate in this important project. ASCR has much to contribute to FSP, and its participation will demonstrate ASCR's visionary approach to modeling and simulation. FSP is in its very early stages and the Panel has recommended that ASCR participate in the planned Project Definition Phase.

The Panel, with the assistance of the Panel Workshop participants, concludes that the FSP Workshop Report provides reasonably good identification of critical technologies in applied mathematics, computer science, and computational science. However, the Panel has identified several additional technology areas that need consideration, and areas that need more emphasis.

The Panel concludes that there are two distinct roles for ASCR that need to be addressed. In the first role, ASCR acts as a close collaborator with FES across the project and across the various disciplines to ensure integration of physics, mathematics, and computer science. In the second role, ASCR acts as a provider of science, mathematics, technology, and facilities. In fulfilling this role, ASCR will engage in basic research and applied research, as well as distribute tools, libraries, and experience in software development. ASCR will also provide FSP with peta- and exa-scale hardware and software platforms, networks, software environments, and user services.

The Panel concludes that in addition to the challenges of providing critical technology and fulfilling its FSP role, ASCR must address anticipated challenges in executing FSP. These challenges include sociological challenges in motivating and rewarding researchers in a project environment; new technology insertion into the project; maintaining good software standards across the project; and sustaining good productivity in software development.

8.0 Acknowledgements

The Panel wishes to acknowledge contributions of the participants of the Panel Workshop held in Boulder, CO on April 30, 2008. The Panel offers thanks for the insightful and thoughtful presentations and enlightening discussion.

8.0 Appendices

8.1 Appendix I: Acronyms

ANL – Argonne National Laboratory
ASCAC – Advanced Scientific Computing Advisory Committee
ASCR – Office of Advanced Scientific Computing Research
CCSM – Community Climate System Model
CPES – Center for Plasma Edge Simulation
DARPA – Defense Advanced Research Projects Agency
DOE – Department of Energy
FES – Office of Fusion Energy Science
FESAC – Fusion Energy Science Advisory Committee
FSP – Fusion Simulation Project
HPC – High Performance Computing
HPCS – High Productivity Computing Systems
INCITE – Innovative and Novel Computational Impact on Theory and Experiment
ITER – International Thermonuclear Experimental Reactor (note: this usage has been discontinued)
LCF – Leadership Class Facilities
LHC – Large Hadron Collider
M3D – Multi-Level 3D
MHD – Magneto-Hydrodynamics
MMS – Method of Manufactured Solutions
MPI – Message Passing Interface
MPICH – Message Passing Interface Chameleon
NCAR – National Center for Atmospheric Research
NERSC – National Energy Research Scientific Computing Center
NNSA – National Nuclear Security Administration
NSF – National Science Foundation
OMB – Office of Management and Budget
ORNL – Oak Ridge National Laboratory
PETSc – Portable, Extensible Toolkit for Scientific Computation
PGAS – Partitioned Global Address Space
SAP – Science Applications and associated Partnerships
SCaLES – Science Case for Large-scale Simulation
SciDAC – Scientific Discovery through Advanced Computing
SDM – Scientific Data Management
SQA – Software Quality Assurance
UPC – Unified Parallel C
V&V – Verification and Validation

8.2 *Appendix II: Charge Letter to ASCAC*

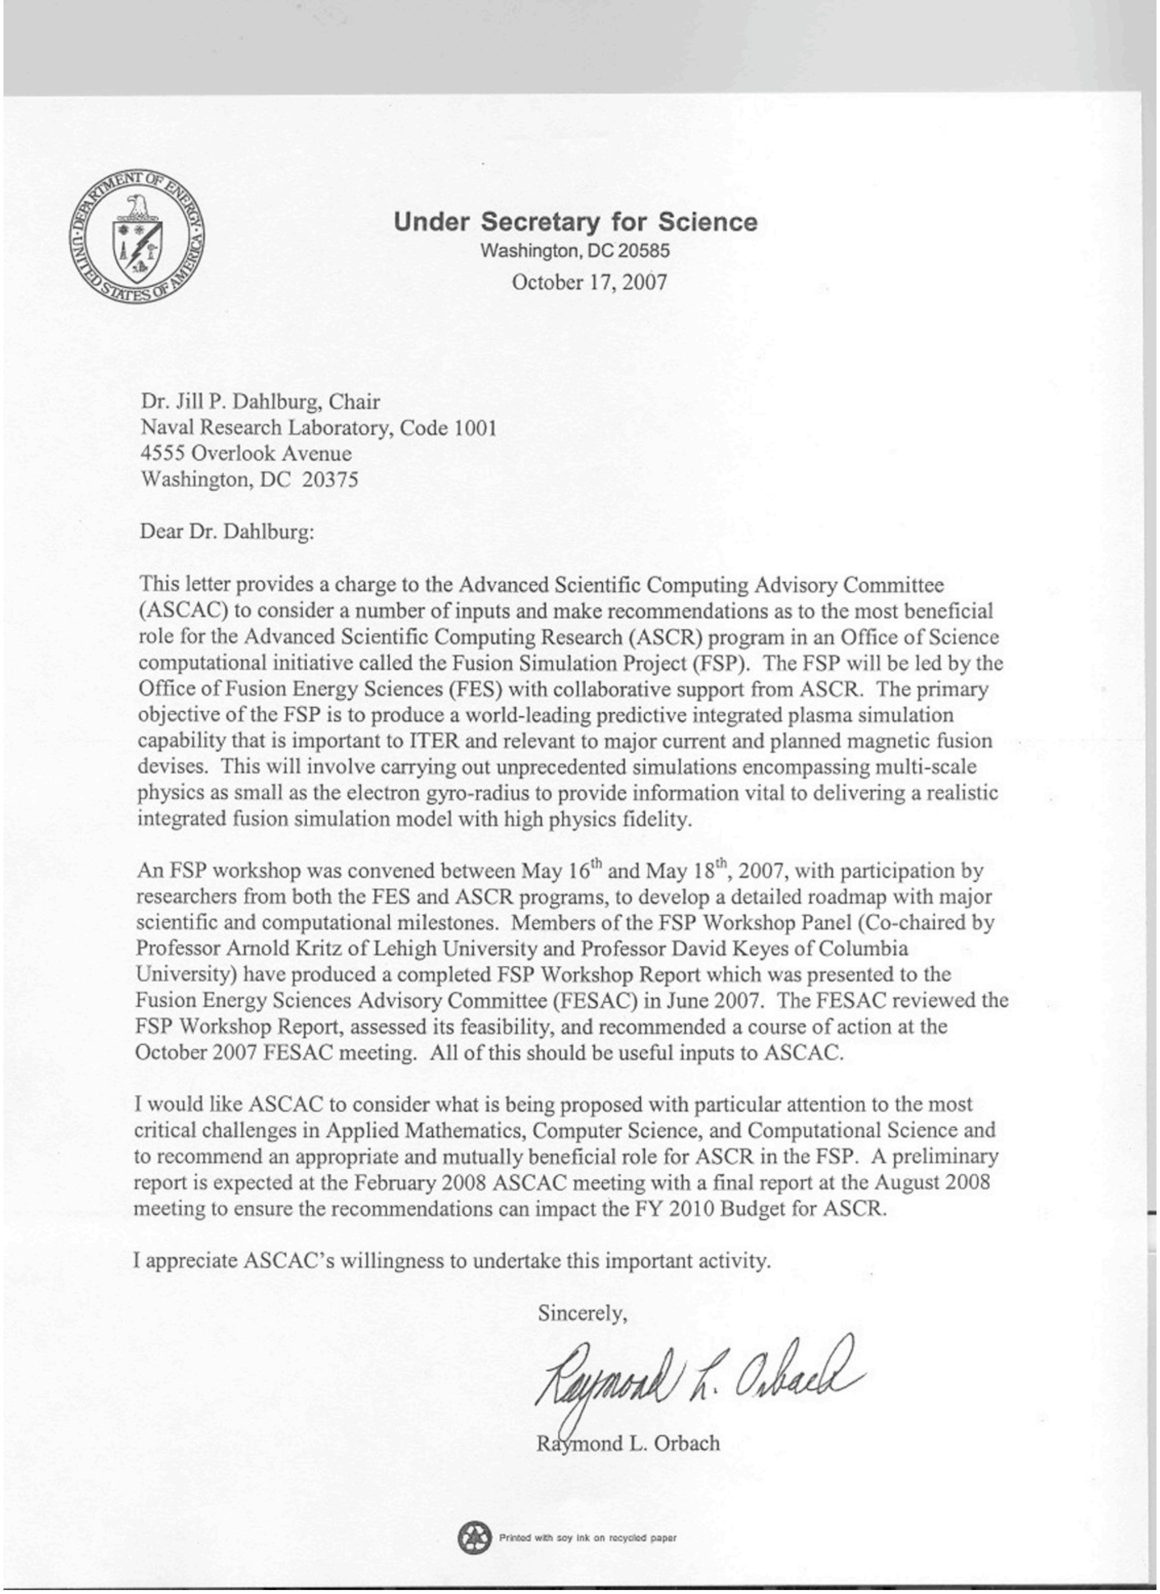

DEPARTMENT OF ENERGY
UNITED STATES OF AMERICA

Under Secretary for Science
Washington, DC 20585

October 17, 2007

Dr. Jill P. Dahlburg, Chair
Naval Research Laboratory, Code 1001
4555 Overlook Avenue
Washington, DC 20375

Dear Dr. Dahlburg:

This letter provides a charge to the Advanced Scientific Computing Advisory Committee (ASCAC) to consider a number of inputs and make recommendations as to the most beneficial role for the Advanced Scientific Computing Research (ASCR) program in an Office of Science computational initiative called the Fusion Simulation Project (FSP). The FSP will be led by the Office of Fusion Energy Sciences (FES) with collaborative support from ASCR. The primary objective of the FSP is to produce a world-leading predictive integrated plasma simulation capability that is important to ITER and relevant to major current and planned magnetic fusion devises. This will involve carrying out unprecedented simulations encompassing multi-scale physics as small as the electron gyro-radius to provide information vital to delivering a realistic integrated fusion simulation model with high physics fidelity.

An FSP workshop was convened between May 16th and May 18th, 2007, with participation by researchers from both the FES and ASCR programs, to develop a detailed roadmap with major scientific and computational milestones. Members of the FSP Workshop Panel (Co-chaired by Professor Arnold Kritz of Lehigh University and Professor David Keyes of Columbia University) have produced a completed FSP Workshop Report which was presented to the Fusion Energy Sciences Advisory Committee (FESAC) in June 2007. The FESAC reviewed the FSP Workshop Report, assessed its feasibility, and recommended a course of action at the October 2007 FESAC meeting. All of this should be useful inputs to ASCAC.

I would like ASCAC to consider what is being proposed with particular attention to the most critical challenges in Applied Mathematics, Computer Science, and Computational Science and to recommend an appropriate and mutually beneficial role for ASCR in the FSP. A preliminary report is expected at the February 2008 ASCAC meeting with a final report at the August 2008 meeting to ensure the recommendations can impact the FY 2010 Budget for ASCR.

I appreciate ASCAC's willingness to undertake this important activity.

Sincerely,

Raymond L. Orbach

Raymond L. Orbach

8.3 Appendix III: ASCAC FSP Panel Members

F. Ronald Bailey* (chair)	Sr. Consultant, Computer Sciences Corp./NASA Ames Research Center
Donald B. Batchelor	Fusion Energy Division, Oak Ridge National Laboratory
David Brown	Deputy Head, Science & Technology, Computing Applications & Research (CAR) Department, Lawrence Livermore National Laboratory
Stephen C. Jardin	Head, Theoretical Magneto-hydrodynamics Division and co-Head, Computational Plasma Physics Division, Princeton University Plasma Physics laboratory
Douglas B. Kothe	Director of Science, National Center for Computational Sciences, Oak Ridge National Laboratory
Ewing “Rusty” Lusk	Director of the Mathematics and Computer Science Division and Distinguished Fellow, Argonne National Laboratory
Thomas A. Manteuffel*	Professor, Applied Mathematics, University of Colorado at Boulder
Juan C. Meza	Head, High Performance Computing Research Department, Lawrence Berkeley National Laboratory
David P. Schissel	Manager, Data Analysis Applications Group and Director Advanced Imagery Laboratory, Theory and Advanced Computing Division, General Atomics Energy Group

*ASCAC member

8.4 Appendix IV: ASCAC FSP Panel Workshop Agenda

ASCAC FSP Workshop Agenda
Boulderado Hotel, Boulder, Colorado
April 30, 2008

TIME	SPEAKER	TOPIC
8:30 - 9:00	Ron Bailey (Chair)	Welcome & Introduction
9:00 - 9:45	David Keyes (Columbia)	Applied math
9:45 - 10:30	Phil Colella (LBNL)	Applied Math
10:30 - 10:45	Break	
10:45 - 11:30	John Cary (Tech-X)	Computer Science
11:30 - 12:15	John Shalf (LBNL)	Infrastructure & Computer Science
12:15 - 1:00	Fred Johnson (OASCR)	Call-in & Lunch
1:00 – 1:45	Open Discussion	
1:45 - 2:30	Ricky Kendall (ORNL)	Infrastructure & Computer Science
2:30 - 2:45	Break	
2:45 - 3:30	Jeff Kiehl (NCAR)	CCM S/W Project Management
3:30 - 4:15	Doug Post (DoD HPCMP)	S/W Engineering
4:15 - 5:00	Robert Harrison (ORNL)	S/W Engineering
5:00 - 6:00	Open discussion	
6:00 - 7:30	Break for Dinner	
7:30 - 9:00	Workshop Wrap-up	