Report of the Fusion Simulation Project Steering Committee

Douglass E. Post, chair, (Los Alamos National Laboratory)
Donald B. Batchelor (Oak Ridge National Laboratory)
Randall B. Bramley (University of Indiana)
John R. Cary (University of Colorado and Tech-X Corporation)
Ronald H. Cohen (Lawrence Livermore National Laboratory)
Phillip Colella (Lawrence Berkeley National Laboratory)
Steven C. Jardin (Princeton Plasma Physics Laboratory)

August 18, 2004
(Minor Revisions September 10, 2004)

The Fusion Simulation Project 2004

Table of Contents

Executive Summary

Recommendations of The Fusion Simulation Project Steering Committee

D.E. Post, chair, D.B. Batchelor, R.B. Bramley, J.R. Cary, R.H. Cohen, P. Colella, S.C. Jardin

1. The Fusion Simulation Opportunity
2. What is revolutionary about this opportunity?
3. Scientific Issues and Questions
4. Integration Issues and Questions
5. The Project
6. Evocative examples of the major scientific questions and possible solutions obtainable through simulation
7. Fusion Simulation Project Representative Design Tasks
8. Project Organizational Requirements
9. FSP Steering Committee summary Recommendations
10. Acknowledgements
11. Glossary
12. Appendices
    - Appendix 1 – Charge to FSP Steering Committee from John Willis
    - Appendix 2— Biographical Summaries of FSP Steering Committee Members
    - Appendix 3 - A US-ITER Perspective on Fusion Simulation
    - Appendix 4 - Initial Design of the Fusion Simulation Project
    - Appendix 5 – Candidate Summary Descriptions of Initial Conceptual Design Tasks

Executive Summary

The Fusion Simulation Project (FSP) will unify and greatly accelerate our ability to simulate and predict the physics performance of magnetic fusion experiments. These new capabilities developed by the FSP will maximize the benefit to the United States (U.S.) from the International Thermonuclear Experimental Reactor (ITER), the largest and most complex scientific experiment in the world, and will give the U.S. a competitive edge in the subsequent design of fusion power plants. This ambitious project is now feasible because of the maturity achieved by the DOE fusion and computational sciences programs and the concomitant advances in effective computer power.

Scientific Opportunities: The FSP will enable exciting advances in Fusion Energy Sciences by addressing the overarching question posed in the Department of Energy Office of Science (DOE SC) Strategic Plan, "*How can we create and stably control a high-performance fusion plasma?*" This fundamental question requires addressing a number of outstanding scientific questions, including: (1) What are the conditions under which a sequence of nonlinear magnetohydrodynamic (MHD) events lead to degradation of confinement or a catastrophic disruption of the fusion plasma and how they can be avoided; (2) How do processes in the plasma edge determine its structure, which in turn critically affects global plasma confinement; (3) What are the consequences of plasma fine scale turbulence for global transport length and time scales; and (4) How do we best use heat, particle, and current sources to control plasmas to achieve high levels of fusion power? Answering these questions entails both integrating physics that heretofore has largely been considered in isolation and developing fast mathematical and physics algorithms to handle effects that occur on time scales far too small to model directly. Many examples have been identified, of which two are:

Example 1: When tokamaks are operated at the limits of performance, with high heating power, plasma density and current, disruptive instabilities often terminate the plasma discharge. Computer simulations and experimental evidence indicate that these disruptive instabilities are caused by a combination of large-scale, unstable modes that overlap with one another. Disruptive instabilities in present-day tokamaks are often avoided by controlling the sources of heating power, fuel, and current drive. However, as a result of the fusion heating that occurs within a burning plasma, disruptive instabilities are likely to be a more severe problem in the ITER experiment. Computer codes developed by the Fusion Simulation Project will provide an improved ability to avoid disruptive instabilities while optimizing the performance of ITER.

Example 2: Experiments and modeling both indicate that the plasma density and temperature in the edge region can strongly influence the performance of the core plasma. However a predictive calculation of these quantities remains elusive because of the broad range of space and time scales, the large variation in the rate of collisions between the plasma particles, the complex magnetic field structure, and the multitude of physical processes known to play a role: micro-turbulence, MHD, neutral-particle transport, impurity transport, and plasma-wall interactions. By bringing to bear advanced multiple-scale techniques, solver technologies, hybrid algorithms, and high-performance computing platforms, an FSP initiative could develop a computational capability to address this multi-physics challenge, giving us a powerful tool for optimizing machine performance. Such a tool would also be extraordinarily valuable for predicting the performance of the power exhaust and particle fueling systems.

Mathematics and Algorithms: The required techniques include adaptive and implicit methods for resolving multiple length and time scales in a given mathematical model; hybrid methods for combining different models that are valid for different components of the physics, or on different length and/or time scales; and new analytical methods for coupling between different physics or scales. A broad range of mathematical and algorithmic techniques of these types are sufficiently mature to serve as a basis for designing a new generation of fusion simulation software, and the current robust fundamental research effort in the applied mathematics community will provide a continuous stream of new ideas in the lifetime of a FSP.

Software Design and Project Organization: Design issues include the factorization of the algorithmic and software space into components to maximize reusability, interoperability, and portability; and the development of a flexible user infrastructure of computational, collaborative, and data management tools that supports the development and running of simulations and the integration of data from simulation, experiment, and analytic theory. While the development of scientific software systems of this type for computational physics (i.e. partial differential equations) is of relatively recent provenance, experience obtained in the DOE Office of Science Scientific Discovery through Advanced Computing (SciDAC) program is already providing a useful guide on how to proceed and a source of tools. The FSP will be a software project with many of the same issues as large-scale experimental projects that involve teams of physicists, engineers, data analysts and technicians. It will build on the Office of Fusion Energy Sciences (OFES) project experience.

Design Flexibility: To a far greater extent than many other areas of computational physics, fusion simulation is a moving target. The mathematical models typically are valid only for specific ranges of time and length scales. As we expand the range of scales being represented in a single calculation, the models themselves will change in unexpected ways. This places a premium on a flexible and agile algorithm and software development environment that allows recombining components into new code capabilities in response to these changes, and the need for a continuous pipeline that extends from the research communities through the FSP software development teams and ultimately to the users.

The ITER timeline: Between now and the finalization of the ITER design in 2008, the initial FSP capability will be able to contribute to tokamak system design decisions for plasma heating and current drive, heat removal, plasma control, and diagnostics. The ITER team will then begin developing plans for ITER operation, with the first plasma scheduled in 2017 and the first burning plasma in 2021. The integrated simulation capability described here will provide an essential capability for ensuring that the tokamak will achieve its design goals.

The FSP will create a complete, integrated simulation and modeling capability for ITER-class burning plasmas and beyond. It provides an exciting opportunity for computational science to make key contributions to a program that offers the potential of a plentiful, environmentally acceptable new energy source. The project builds on the successes of the base programs in the Offices of Fusion Energy Sciences and Advanced Scientific Computing and the highly successful SciDAC activities. The fusion SciDAC program has developed several independent 3D plasma models, each with a limited regime of applicability. The FSP will begin with such models to create an interoperable code suite that includes an integrated simulation and then develop the additional capability necessary for a complete high-fidelity simulation of an ITER discharge. The range of different physics regimes, the complex geometry, extreme anisotropy due to the magnetic field, and large range in temporal and spatial scales makes fusion simulations one of the most challenging problems in modern computational physics, requiring a team approach involving application scientists, applied mathematicians, and computer scientists.

Recommendations of The Fusion Simulation Project Steering Committee

D.E. Post (LANL), chair, D.B. Batchelor (ORNL), R.B. Bramley (U. of Indiana), J.R. Cary (U. of Co.), R.H. Cohen (LLNL), P. Colella (LBNL), S.C. Jardin (PPPL)

1. The Fusion Simulation Opportunity

The international magnetic fusion program has formed a partnership—Japan, the EU, the US, Russia, Korea and China—to construct and operate the $ 5B International Thermonuclear Experimental Reactor (ITER). The goal of ITER is to provide the knowledge of the physics of “burning plasmas” needed to design and operate a prototype fusion power reactor. ITER is the highest priority facility for the US DOE Office of Science and is the next step toward the practical realization of fusion energy. Support of the US participation in ITER is a key priority for the DOE Office of Science and the US fusion program. The ITER project has identified crucial contributions that ultra-scale simulation can make to its success (appendix 3). Accurate simulations of the ITER plasma performance would substantially enhance the likelihood of successful operation of ITER, maximize the knowledge gained from ITER operation, and capture that knowledge so that it can be applied to further progress in fusion.

The continued US progress in fusion simulation and computational mathematics together with the continued development of ultra-scale computers makes it possible to develop such a simulation capability. The US fusion simulation and computational mathematical communities —supported by the Office of Fusion Energy Sciences (OFES) and the Office of Advanced Scientific Computing Research (OASCR)—have strong programs in scientific computing in the fusion sciences and in computational mathematics. With the development of ultra-scale computers in the 10 TeraFlops range and the promise of PetaFlops computers within the next five years, computing power has reached the point where it is possible to solve large, multi-physics problems that span very disparate time and distance scales. To successfully meet the challenge of accurately predicting the performance of burning plasma experiments, the existing fusion simulation capability and program will need to be substantially augmented and coordinated in its software development. To accomplish this, we propose that a Fusion Simulation Project (FSP) be launched. This will enable exciting advances in Fusion Energy Sciences by addressing the overarching question from the DOE Office of Science Strategic Plan, “*How can we create and stably control a high-performance fusion plasma?*”

The FSP will have three main elements: a *production component*, a *research and integration component* and a *software infrastructure component*. The production component will consist of a suite of advanced plasma simulation codes and modules that are built around a common software structure allowing results from one code to be used as input to another, as appropriate. An important element of this suite will be an Integrated Plasma Simulation (IPS) capability that is able to simulate the global physics of the whole burning plasma device, although initially at a greatly reduced fidelity. The integrated model for tokamak performance and the suite of advanced simulation tools will be used by the ITER physics community for design, operational planning, plasma control and physics analysis. These tools will be validated by application to present experiments.

The research and integration component has two key elements. The first involves the development of new simulation capabilities that allow the coupling of an ever-wider range of physical phenomena together in the same simulation, and preparing these capabilities for submission to the production component. The second consists of research needed to develop new first principles simulations for highly complex phenomena such as the effects of fine scale turbulence on plasma transport that can then be integrated with simulations of other effects. The project plan envisions staged delivery of capability by the research and integration component. For instance, we estimate that many of the coupling or “integration” projects will require roughly four to five years to complete, but will develop the initial capability to address important issues in much less time than 5 years. The software infrastructure component will allow all of this to take place in a seamless and efficient manner together with visualization capability and access to experimental data as appropriate.

Preliminary estimates of the project scope suggest that development of this capability will require resources of about $20M/year for 15 years in FY 2002 dollars. While the ultimate goal is the delivery of a predictive capability for ITER, the project has intermediate goals to deliver useful computational tools to the ITER and general fusion communities throughout the life of the project.

The physics goals and challenges have similarities with the DOE National Nuclear Security Agency (NNSA) Accelerated Strategic Computing Initiative (ASCI) that was designed to develop the predictive capability required to certify the US nuclear stockpile in the absence of testing, and the Community Climate Systems Model project (CCSM) designed to produce an international capability to accurately predict the Earth’s climate. The FSP will leverage the tools and methods developed by these large-scale projects and will learn from their successes and how they met their challenges.

The FSP is based on the simulation and computational mathematics capability that has been developed by the SciDAC program and the other OFES and OASCR programs. It will build on, incorporate and enlarge this capability, by providing a focus on ITER and the additional resources necessary to accomplish the project. The fusion community has identified a number of fundamental scientific questions that must be addressed. Among them are: (1) What are the conditions under which a sequence of nonlinear magnetohydrodynamic (MHD) events lead to a catastrophic disruption of the fusion plasma and how can they be avoided; (2) How do processes in the plasma edge set its structure, which critically affects global plasma confinement; (3) What are the consequences of plasma fine scale turbulence on global transport length and time scales; and (4) How do we best use heat, particle, and current sources to control plasmas to achieve high levels of fusion power? The FSP will address these questions and others by integrating existing capabilities and developing new capabilities. This integration will also provide leveraging between the FSP projects and other Office of Science programs in computational science.

The fusion program has achieved the level of maturity necessary to both support and utilize the capability that will be developed by the FSP. Present fusion experiments collect and share data (MDSPlus) utilizing an international collaboratory developed with the support of the SciDAC program. Modeling tools are being used to interpret experimental diagnostics, control the plasma and plan experimental scenarios. This provides a capability and a framework for testing and validating all aspects of the FSP as it progresses and prepares for ITER operation. This not only ensures that the FSP will have validated tools with real predictive capability for ITER, but that the FSP will contribute to the ongoing success of the fusion program by aiding the success of the experiments and by increasing the physics understanding of these experiments between now and ITER operation.

ITER experiments will build on this initial collaboratory but will require much more. ITER experiments will be managed remotely via computer networks by distributed teams of physicists and engineers located in the ITER partner countries (China, EU, Japan, Korea, Russia and US). Physics experiments will share the ITER data in a common format, and will use the FSP and other tools to design experiments, control the plasma to ensure successful operation, and analyze the results of the experiments.

The FSP is also a uniquely challenging experiment in software project management and engineering and in the organization and management of scientific research. The FSP will need an organization and governance structure akin to that of an experimental project. The FSP will be accountable directly to DOE and indirectly to ITER for the staged delivery of computational capability over the 15-year project life. At the same time, it must foster sufficient freedom that the individual elements of the project are able to do the research and development needed to solve the immense scientific and technical problems. Finally, the project will consist of non-collocated multi-institutional and multi-disciplinary teams composed not just of computational physicists, but also computational mathematicians, software engineers, and computer scientists.

Building on the initial work of Integrated Simulation & Optimization of Fusion Systems (ISOFS) sub-committee of the Fusion Energy Science Advisory Committee (FESAC)[1] that identified and explored the issues and opportunities, the FSP steering committee has further refined and characterized these issues. The ISOFS report presents a detailed description of the multi-scale issues and challenges, and examines the potential roles and contributions that led it to propose the Fusion Simulation Project.

The committee recommends that the project begin with a “conceptual design phase”, an approach successfully followed for the design of fusion and other large-scale experiments. The conceptual design team would consist of a small, multi-institutional team composed of physicists, computational mathematicians, software engineers and others who will develop the initial design for the project. The initial design team will define a set of realistic goals and deliverables for the project, develop and plan the tasks, develop a schedule and roadmap, and assess the risks and uncertainties. The design team will canvass the OFES and OASCR communities and related groups to assess the presently available capabilities and how they might contribute and be incorporated into the FSP. The design team will work with the US and international ITER community, the similar simulation groups in the European Community and Japan, the US and international experimental communities, the US fusion theory and computational physics communities, the SciDAC program, and the US computational mathematics community to develop the project design and plan. The design team will explore and develop prototypes to the extent feasible and necessary.

2. What is revolutionary about this opportunity?

Wide ranges of physical processes are at play in a high-temperature fusion-grade magnetized plasma. A complete simulation capability must encompass all timescales ranging from the sub-nanosecond electron cyclotron resonance used for RF heating to the hundreds of

---

[1] Dahlburg, J., Corones, J., et al. (2001). "Fusion Simulation Project: integrated simulation and optimization of magnetic fusion systems." Journal of Fusion Energy **20**(4): 135-196.

seconds that it takes for the plasma current to achieve a stationary spatial distribution (Figure 1). A number of multiscale analysis and computational techniques have already been developed to span many of these timescales allowing, for example, the calculation of the evolution of the electron pressure profile over hundreds of milliseconds in the presence of multi-gigahertz electron-cyclotron RF heating. These multiscale simulations rely on the existence of detailed 3D nonlinear models that are each valid on a given timescale, and on multi-scale techniques for bridging the timescales (in this case, that of magnetic surfaces and adiabatic invariants).

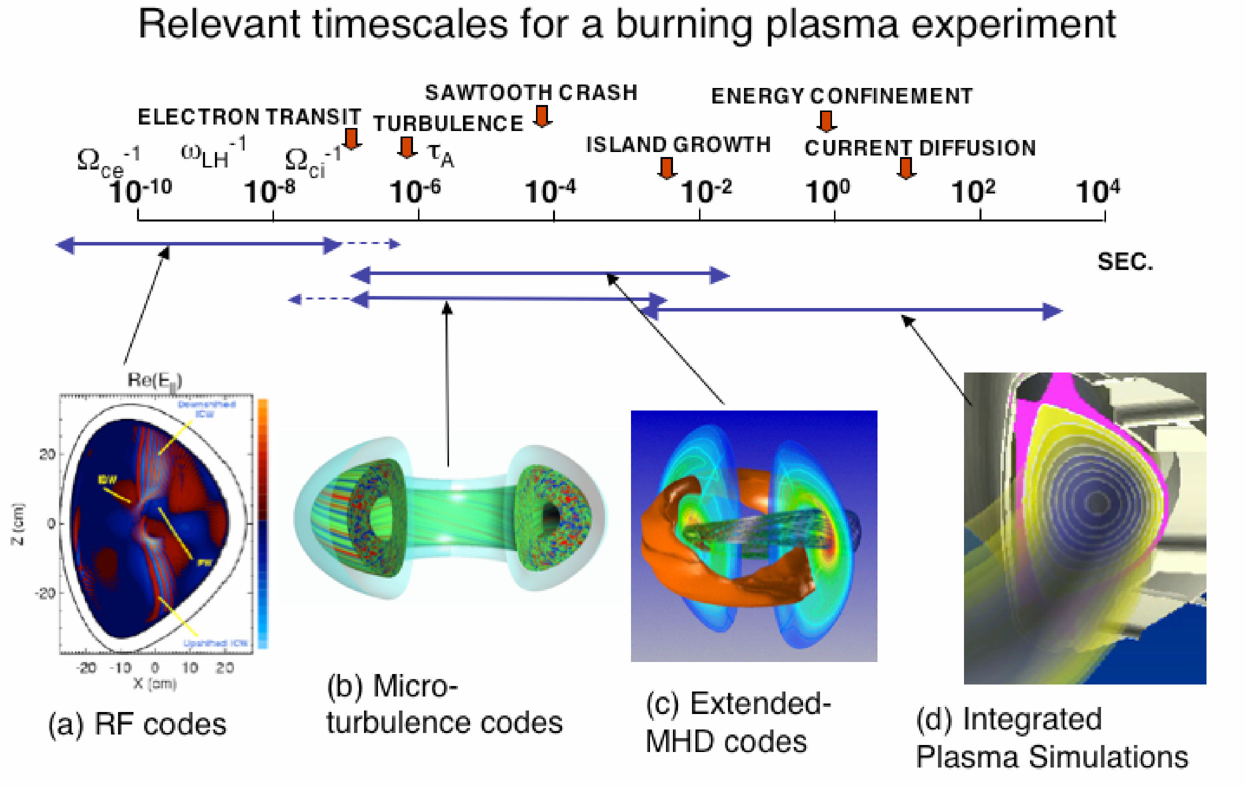

(a) RF codes

(b) Micro-turbulence codes

(c) Extended-MHD codes

(d) Integrated Plasma Simulations

Figure 1. Range of time scales present in a burning plasma.

This initiative will focus on the multiscale coupling of plasma processes that previously have had to be considered in isolation. It will build on the highly successful SciDAC activities that have developed 3D plasma simulation models for extended-MHD, microturbulence, and wave-particle interaction, and will begin to couple these to enable the simulation of important classes of phenomena for which this coupling is essential. A prerequisite for this is that the single-scale nonlinear 3D codes exist and are in a modern and common software framework that enables the coupling. This framework will be developed in this project, as will the coupling algorithms that will allow the simultaneous calculation of an ever-wider range of physical phenomena together in a single simulation. These capabilities will truly represent a turning point in the realism of fusion simulations.

The rapidly increasing computer capability, successes with algorithms bridging multiple scales and treating different types of nonlinearity, and the emergence of software methodologies supporting the coupling of independently developed large-scale scientific simulation codes now offer the promise that a truly comprehensive fusion simulation capability can be developed. The ultimate goal of the FSP is to simulate the behavior of toroidal magnetic fusion devices on all important time and space scales, and to account for the interactions of all relevant processes. It will serve as an integrated "standard model" for the community, permitting the exploration of new multi-scale physics and mathematics. The FSP will fulfill the simulation needs of the experimental program for planning and interpretation of experiments, and for plasma analysis needed for system design. A good analogy can be made to the more familiar phenomena encountered in studying the Earth's climate where very different physical phenomena (e.g. the atmosphere, ocean, land-masses, ice and snow, etc.) must be coupled for reliable predictions of climate evolution. As each piece of the FSP is developed, it will immediately contribute to the success of the fusion program by enlarging our understanding of the fusion science and plasma physics and improving our ability to optimize the performance of ongoing fusion experiments. As these pieces are integrated into a complete and predictive model for burning plasma performance, the FSP will give us the capability to design and successfully operate burning plasma experiments such as ITER.

3. Scientific Issues and Questions

An overall goal of the fusion program is to obtain stable, high-performance plasmas – plasmas that are quiescent, with no large scale, reactor damaging movements, and have sufficiently high pressures and densities so that fusion reactions proceed. To obtain this overall goal, we must answer many fundamental scientific questions, such as:

(1) What are the conditions under which a sequence of nonlinear magnetohydrodynamic (MHD) events lead to a catastrophic disruption of a fusion plasma?

Disruptions, large MHD events, often lead to rapid termination of the plasma. They can result in extreme thermal loads and the transfer of large electrical currents to the confinement vessel wall resulting in major damage to a fusion reactor. The probability of such events is increased by the operation of reactors at conditions of highest performance. For identical plasma performance levels, subtle differences in the plasma shape, profile and current distribution can determine whether the plasma will disrupt or be stable. The ability to determine the precise requirements for disruption-free operation and for controlling the plasma will allow operation of fusion reactors at conditions of maximum possible performance.

(2) What is the physics of the plasma edge that critically affects global confinement of heat and particles?

Conditions at the plasma edge strongly affect the global confinement of the plasma energy and particles. In particular, plasmas with good energy and particle confinement often have sharp gradients in the plasma temperature and density at the plasma edge forming a “pedestal”. The pedestal forms a “transport barrier” there. This barrier often improves the plasma transport in the interior plasma as well. Sufficiently impermeable barriers lead to the highest performance plasmas. An understanding of the physics of the plasma edge will allow prediction of the optimal edge conditions needed to obtain these highest performance plasmas.

(3) What are the consequences of plasma fine scale turbulence on global transport length and time scales?

The central plasma parameters (density, temperature, fusion power production) of fusion plasmas are determined by the leakage of plasma energy and particles from the plasma center to the edge. This leakage across the confining magnetic surfaces is predominately due to fine-scale turbulence. The level of fine-scale turbulence is determined by the plasma profiles, which in turn are determined by the turbulence. Accurate prediction of plasma performance relies on solving this strongly linked system that couples the structure of the fine-scale turbulence with large-scale evolution of the plasma profiles.

(4) How do we best use heat, particle, and current sources to control plasmas to achieve high levels of fusion power?

It has been established experimentally that the sources of heat, particles and plasma current, together with the physics of plasma and current transport and MHD, determine the central plasma conditions, which in turn determine the plasma performance of fusion reactors. These sources can be used to modify and control the fusion power production in fusion reactors. Injection of beams of energetic neutral atoms can increase the plasma pressure and, if done in exactly the right way, will cause the plasma to switch into a high-performance state. Neutral beam injection can also induce plasma rotation, which can stabilize MHD instabilities. Injection of radiofrequency electromagnetic energy can drive currents that can counteract instabilities thought to play a role in the disruptions discussed in (1) above. With sufficient computational capability, the power levels and other parameters of the sources needed for inducing favorable changes in the plasma can be predicted. This will enable active control of burning plasmas in real time to improve performance and eliminate deleterious phenomena.

4. Integration Issues and Questions

The combination of the wide range of spatial and temporal scales, extreme anisotropy, importance of confinement geometry, and the wide variety of physics issues make simulating a fusion plasma one of the most challenging problems in computational physics today (c.f. Figure 1). Although several sophisticated computational models have been developed to treat individual features of magnetically confined plasmas, integrated predictive modeling of fusion plasmas will require major innovations to handle the wide variety of physical processes simultaneously and self-consistently.

(1) What enabling developments are needed in the computational sciences for creating an integrated FSP simulation capability?

An integrated FSP simulation must be simultaneously multi-scale and multi-physics, i.e. treat many different physics effects simultaneously. What are the new algorithms and numerical methods needed to span multiple time and distance scales, and how can these methods be generalized for all of the different discipline physics models in an FSP? Present-day simulations are based on finite elements, spectral elements, Monte Carlo, and Particle in Cell (PIC) techniques. Integrating two or more separate physics models will require developing methods to link models based on different techniques together. It will be necessary to analyze numerically the conservation properties across the boundaries between the different discretization models and to interpolate between meshes that, in some cases, model the same spatial regions at different time scales. New software methods will be needed to improve end-to-end scalable parallel performance when connecting the different codes. Intermediate translation and filtering modules will be required. It will be essential to determine how the configuration of these modules can be dynamically optimized while still retaining the necessary numerical accuracy. The mathematical and software engineering solutions will necessarily be closely entangled.

(2) How can we successfully integrate all of the physics elements necessary to simulate the performance of a burning plasma?

A key goal of the FSP is the development of a predictive capability for all aspects of burning plasma experiments. The predictive capability must include many different physical effects that span multiple time and distance scales. The simulation capability must satisfy multiple goals. At least one configuration must run in real time with sufficient accuracy that it can contribute to the development of a real-time plasma control system (Figure 2). Other configurations must have the accuracy necessary for identifying optimum plasma scenarios for experimental planning, developing and testing models, and analyzing data. Thus each major physics effect must have a hierarchy of models, varying from the simple, small and fast to the sophisticated, large and slower. All levels of the integrated simulation must be validated with experimental data. Due to the large range of multiple time and distance scales and the many different types of physics involved in fusion experiments, developing the physics, computational mathematical and software infrastructure for such an integrated simulation capability is a tremendous challenge. The lessons learned and techniques developed in this endeavor will be applicable to many other fields of computational science that are beginning to address similar issues, such as biological cell and nanostructure simulations.

The customers for the integrated modeling simulation are the community of scientists and engineers who need the ability to accurately simulate burning plasma experiments. The integrated simulation must be useable by all of them with only minimal support from the code development teams. These users must be able to choose the optimum combination of models needed for their task so that almost all models must be able to work together.

Many of the major tokamaks in operation now have simple active plasma control systems that use simple computer simulations that analyze plasma diagnostic information to generate control signals for plasma control “actuators” such as the neutral beam and radio frequency heating and current drive systems and plasma control feedback systems to control the plasma (Figure 2). The ITER team has stated that such a control system will be required to optimize the plasma performance and to avoid exceeding the operational limits (Appendix 3). A reliable and accurate Integrated Plasma Simulation capability with a version that can run in real time would be the basis for developing such a control system.

5. The Project

The FSP will be a major new initiative within the DOE OFES and OASCR. Its ultimate goal is to provide a complete, integrated simulation and modeling capability for ITER-class burning plasmas and other large magnetic fusion confinement experiments. This requires assembling a comprehensive set of theoretical fusion models, combined with the algorithms required to realize them and an architecture and computational infrastructure that enable them to work together. The FSP will build on the successes of both the base programs in OFES and

OASCR and the highly successful SciDAC activities. The fusion SciDAC program had the different aim of producing several efficient and reliable 3D plasma models, each with a limited regime of validity. The FSP will assemble these and other existing models into an interoperable code suite. By combining these heretofore separate activities into a common structure, and extending them as needed with an active research program, we now believe that it is possible to create a powerful new simulation capability that can effectively simulate and model all relevant physics phenomena in an ITER-class fusion device.

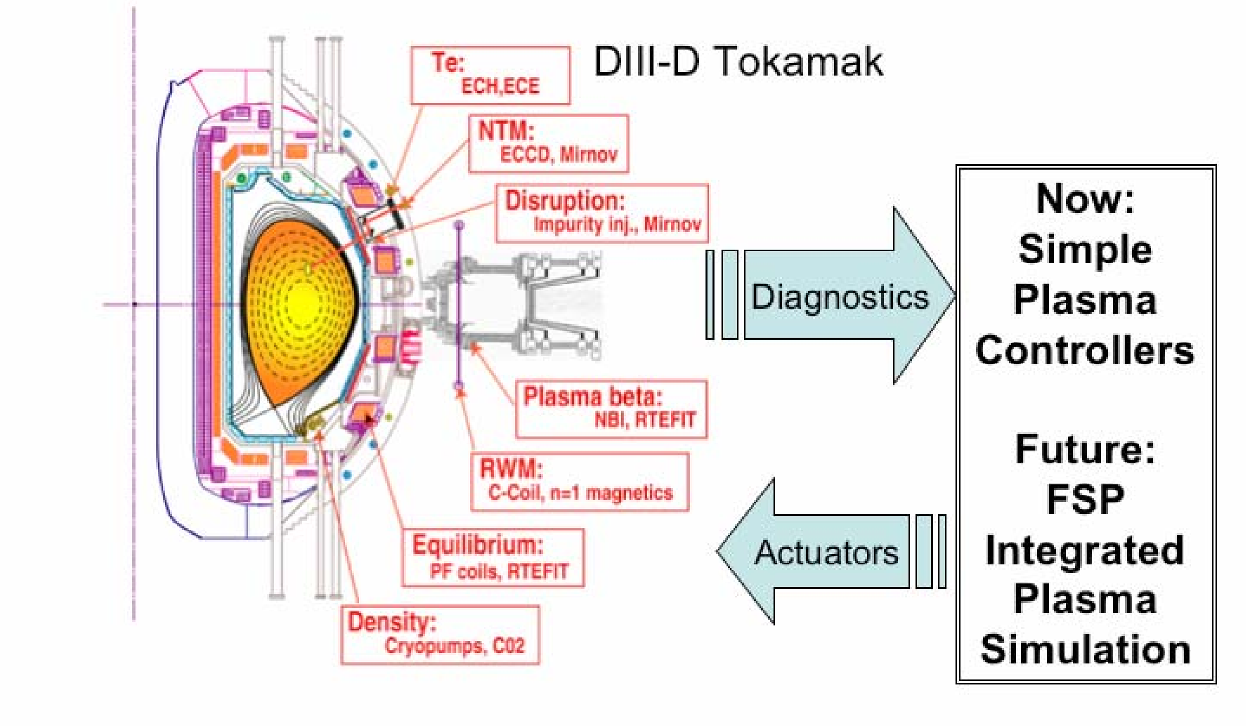

Figure 2 DIII-D tokamak diagnostics and schematic for real time plasma control system.

Combining activities instead of treating them as separate projects will also enable leveraging our investments and efforts. Multi-scale mathematical models for plasma will likely be useful in all FSP projects. The hierarchical, multilevel solvers required (multi-grid, multi-pole, Adaptive Mesh Refinement (AMR)) will enable new capabilities both in the individual discipline simulations and in the new integrated edge, MHD, and turbulent transport models. The development of software in a disciplined fashion will allow reuse of modules for access to online experimental data, visualization of complex 4D structures in toroidal geometries, base libraries for AMR, linear and nonlinear iterative solvers, cross-discretization (finite element, spectral element, toroidal mesh) interpolators, and mesh generation, refinement, and validation tools. An FSP program will also provide greater opportunities for fusion to collaborate with and provide target application for computer science research efforts to develop reusable numerical codes and software frameworks.

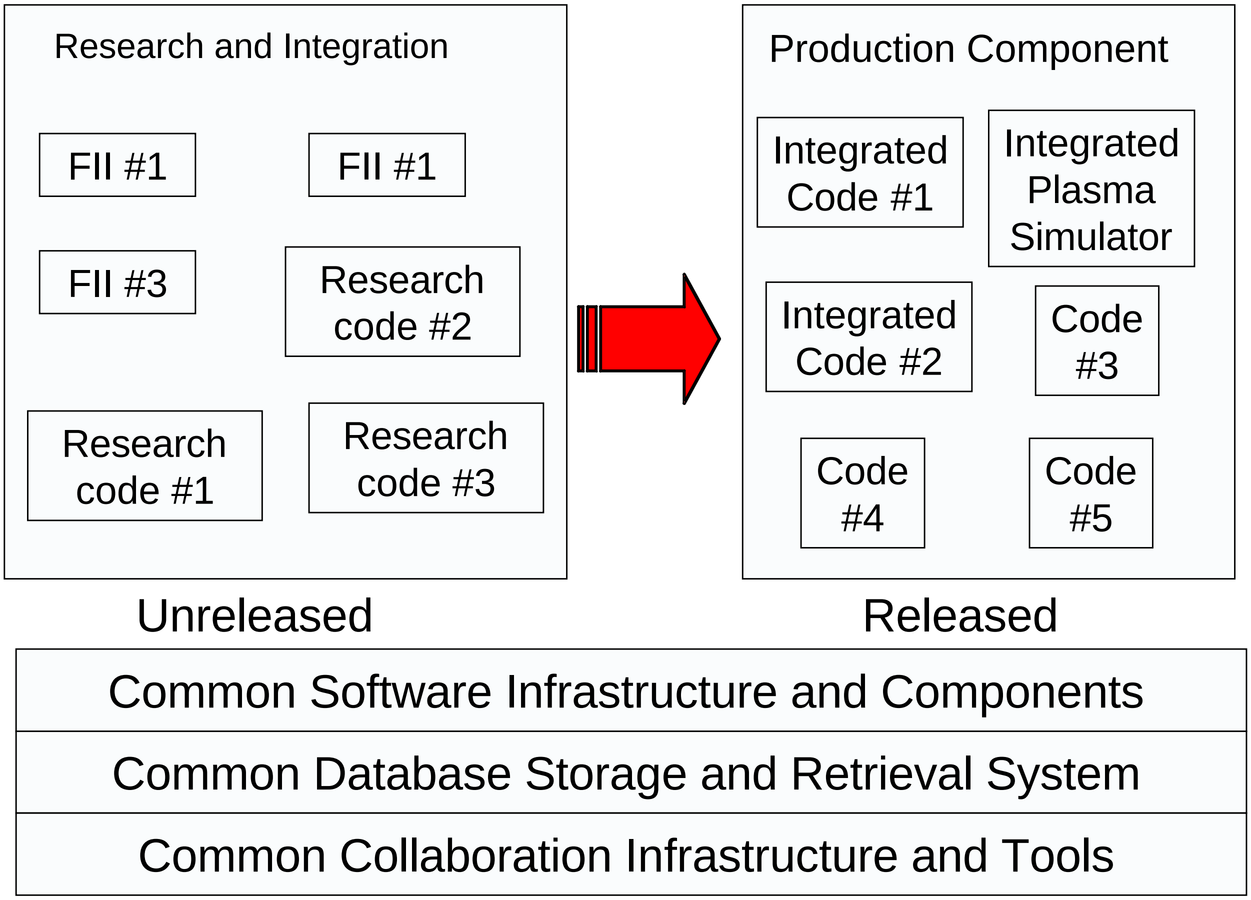

Figure 3. The Three FSP Components.

The project will have a production component, a research and integration component and a software infrastructure component (Figure 3). The production component will contain a common user interaction system that provides access to all production codes and modules in the system. These codes and modules will have a common software infrastructure and share components, and will read and write from a common data base storage and retrieval system. Not only will all of these code analysis packages be readily available to anyone in the community, but a set of common interfaces will be defined and implemented that will allow filters and converters to connect codes in a seamless and convenient manner. This collective suite of codes will enable the burning plasma physics community to simulate all of the important individual and partially integrated aspects of fusion research. In addition, the production component will have a single code, the Integrated Plasma Simulator (IPS), which can simulate the complete integrated behavior of burning plasmas.

Just putting the codes that have already been developed by the fusion community into a unified framework that is readily available to both theoretical and experimental researchers would be a major accomplishment. But the Fusion Simulation Project will also have a research and integration activity the ISOFS committee termed “Focused Integration Initiatives” or FIIs. The goal of each FII is the solution of a compelling problem in fusion physics that requires integrated simulation. This activity will have two major components. The first component will cut across and integrate two or more of the major existing fusion simulation capabilities to provide physics integration both spatially and temporally, with a guiding focus of a single overarching scientific question or topic that is needed for an integrated plasma simulation capability. This integration will not only provide useful tools in a timely fashion to address important problems but will provide important experience in integrating different effects.

The second component of the Research and Integration activity will focus on the refinement of physics models whose integration requires bridging across multiple time and distance scales and the development of new algorithms and methods to accomplish this bridging. This part will emphasize the development of "practical" components and modules that can accurately treat difficult and challenging "single-physics" issues and can be integrated with other components and modules that have much longer time and larger distance scales. For example, present simulations of turbulent transport are far too expensive computationally to be used in an integrated plasma simulation in the foreseeable future. Models must be developed that accurately capture the major features of the effects of turbulent transport and yet are sufficiently computationally efficient that they can be used in an integrated plasma simulation to get solutions in a reasonable time with the computer resources likely to be available soon and by the early stages of ITER operation. This research will rely heavily on the present OFES and OASCR base programs, but will have a more immediate focus on developing "practical" algorithms than the existing OFES and OASCR programs. When these developmental projects reach a production level, they will be incorporated into the production side of the project, thus allowing a means for continual improvement. The FSP will support the use of the production level components by the ITER team and by the general fusion community.

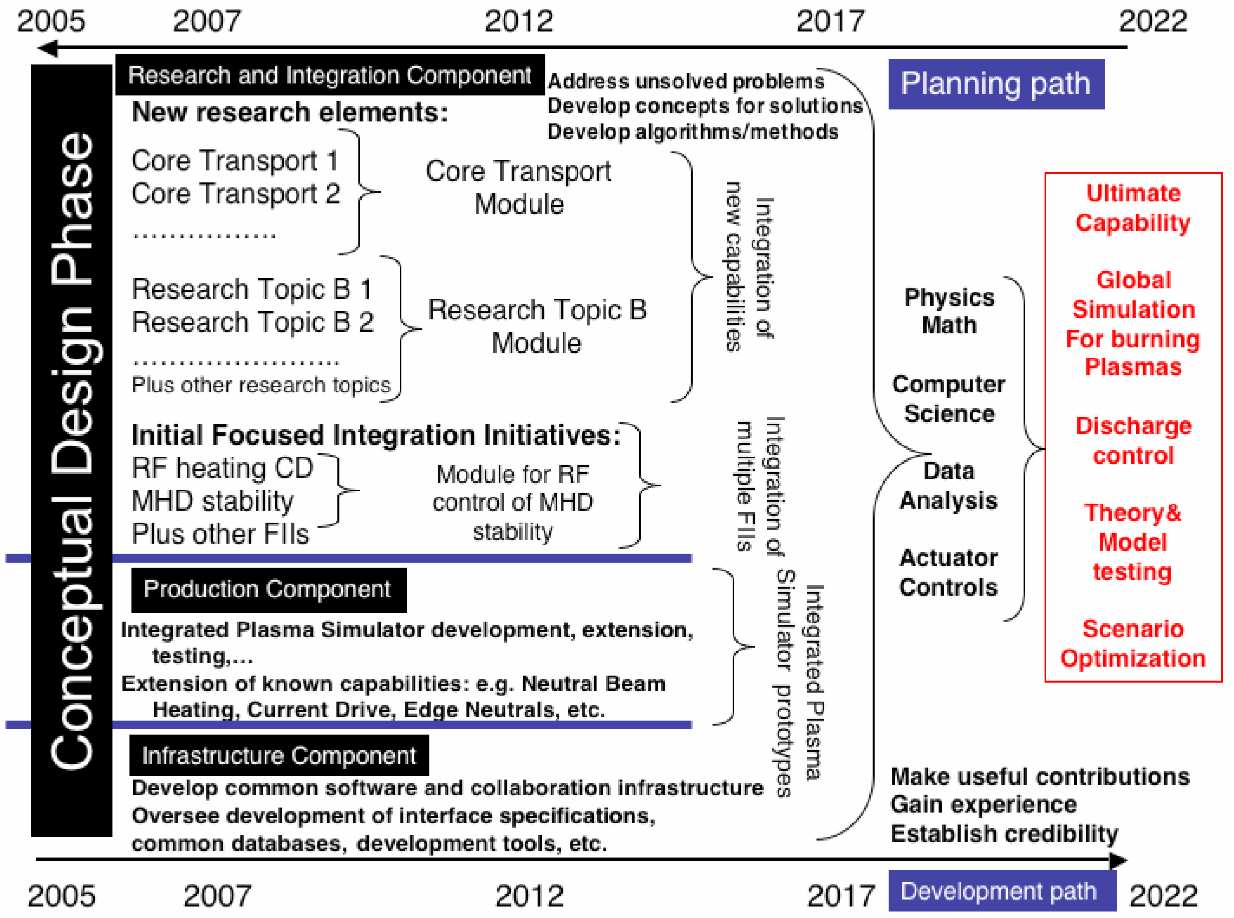

Figure. 4 Illustrative roadmap for the FSP

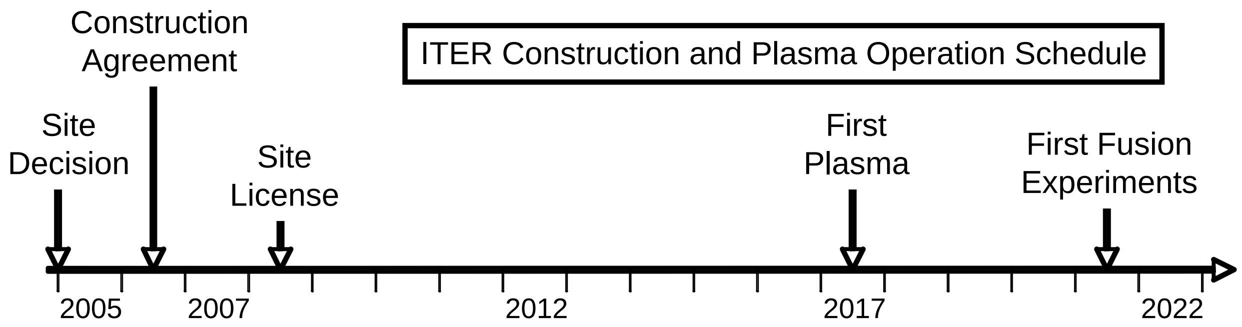

Figure 5. ITER Construction and Plasma Operation Schedule (c.f. www.iter.org)

We suggest that the FSP could occur in approximately 4 stages, similar to the recommendations of the ISOFS report (ref. 4). The project would begin with a 1 to 2 year conceptual design phase, similar to other successful major experimental and computational projects. The conceptual design phase would include development of a project design and plan and include the development of prototypes and exploration of collaboration issues. While the purpose of the conceptual design phase is to develop the project plan, we can illustrate the issues by outlining a candidate project plan (Figure 4).

After the successful completion of the conceptual design, the project would consist of 3 major phases, each lasting about 5 years. A strong effort would be made to deploy capability to the tokamak community as soon as practical, both for verification and validation, and to enhance progress in fusion science and to aid development of the FSP components.

During the first phase (2007-2011) the activities for each component would be:

- Research and Integration: Define and carry out several focused integration initiatives in areas such as the plasma edge and the control of MHD stability with RF. Define and carry out several research efforts to develop practical components for physics effects that presently cannot be run for practical problems for the entire tokamak plasma for time scales of seconds to minutes with foreseeable levels of computer resources. Candidate areas would also include kinetic MHD and turbulent transport.
- Production Component: Develop a common application interface and component infrastructure for many of the existing fusion modeling codes for use by the fusion community and the ITER team. Begin design and development of an integrated plasma simulator. Apply it as soon as feasible to develop the experience that will be necessary for design and development of the production version of the plasma simulator.
- Infrastructure Component: Develop the collaboration infrastructure necessary to carry out multi-institutional and multi-disciplinary code development. Begin development of a common software infrastructure and component interface definition. Begin development of a database storage and data retrieval system, leveraging the SciDAC Fusion Collaboratory and our international ITER partners where useful.

During the Second phase (2012-2016) the activities for each component would be based on an assessment of the results of the first phase and would include:

- Research and Integration: Define and carry out additional focused integration initiatives that will be feasible and necessary for the integrated plasma simulator, such as integration of several of the “single physics” research projects launched during the first phase. Define and carry out development of new research projects indicated by progress in research by the fusion and computational mathematics community and the needs of the ITER team.

- Production Component: Continued development and deployment and support of a capability for integrated plasma simulation. Continued support for providing simulation capability using the codes and suites of codes developed by the FSP and the fusion community.
- Infrastructure Component: Continued support for the FSP and fusion community for collaboration tools, for database and access to tokamak data, and general project support.

The Third phase (2017-2021) would focus the design and development of the integrated plasma simulator based on the prototypes, the "lessons learned" and the capabilities developed during the first two phases.

- Research and Integration: These efforts would continue as necessary to complete and deploy the integrated plasma simulation and to increase the utility of the other production codes and code suites.
- Production Component: Focus would be on development and deployment of the production version integrated plasma simulator.
- Infrastructure Component: Continued support for the FSP and fusion community for collaboration tools, for database and access to tokamak data, and general FSP project support.

This plan is consistent with the likely ITER schedule (Figure 5). Initial tools would be available to help with the late stages of ITER design, operating scenario planning for first plasma conditions, and the initial version of the fully integrated plasma simulator should be working in time for the first fusion plasma experiments. A relatively mature version should be available for full power fusion experiments on ITER.

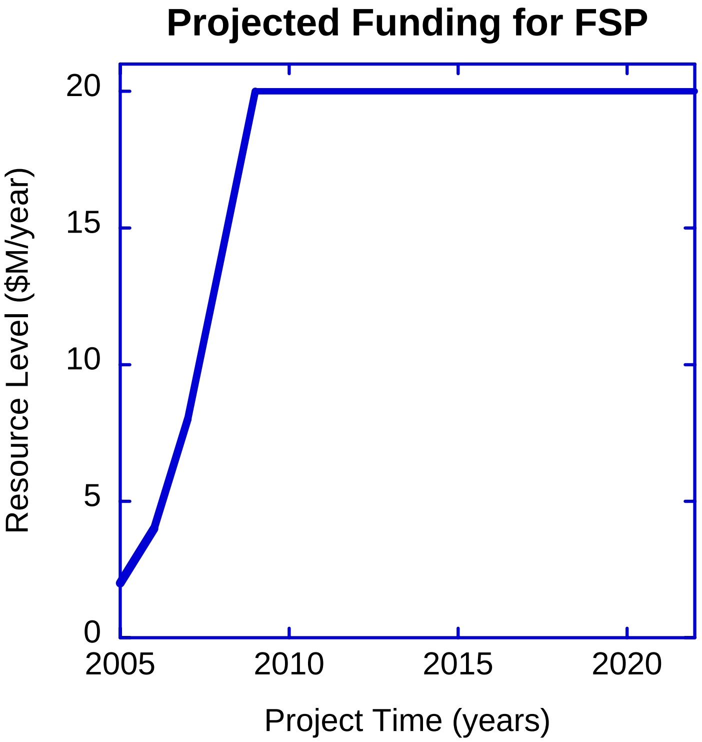

Figure 6. Projected funding profile for the FSP.

The resource plans call for a small initial (~10 to 15%) funding level to launch the project and begin conceptual design (Fig. 6). Then, as the vision and plans mature, the resource levels would increase to the full $20 M during the next 3 to 4 years. This staffing and resource schedule is similar to the resource plans for successful experimental and computational projects. The full level of resources would be allocated only when the project has a clear plan for utilizing the resources efficiently and effectively. For the FSP, it will take time to build up teams to carry out the FIIs and other elements of the project. The history of the ASCI projects, for instance, indicates that staffing up strong code teams takes at least four years when there is not a large reservoir of experienced and knowledgeable staff that can be hired quickly.

Clearly the integration and coupling of all of the multi-scale, multi-physics phenomena required for a comprehensive simulation represents a physics and mathematical challenge of the highest order that will proceed over many years. In addition, to fulfill the programmatic goals of enabling basic research and supporting experiments throughout the fusion research community the integration design must meet some stated goals and fulfill certain requirements:

- It must be extensible. Easy connections are to be made early in the project while more difficult ones, for example those involving very disparate time-scales, can be added as techniques are developed. Its architecture must permit continuous improvements and additions.

- It must be flexible. Only the needed physics modules required for a given study should be interconnected. It must be robust to changes in physics paradigms and introduction of more complex mathematical descriptions

- It must support collaborative research. It should interface well with experimental databases and provide appropriate tools such as synthetic diagnostics to facilitate understanding of output and include protocols for effective communication among geographically and scientifically diverse participants.

- The FSP must run on a variety of platforms, from workstations to massively parallel computers in the multi-PetaFlops range. The numerical methods and software infrastructure will need to enable efficient use of these platforms. Considerable flexibility and agility will be needed to enable the FSP to run on the presently unknown architectures of ultra-scale computers in 2020.

- It must complement existing research. The project must provide value to the individuals involved. Therefore it must not impose undue overhead (computational or human) on the use and development of the separate physics modules. It must provide needed services so as to be of value even to the user of a single module.

- Above all, the integrated capability must technically enable fusion science. It must promote the development of the physics modules and their validation and verification through experimental comparison, beginning in the near term. It must facilitate study of mutual physics interactions presently modeled in separate codes as such interconnections become appropriate. It must increase significantly the depth and breadth of fusion physics compared to today's whole device modeling codes, incrementally, as better modules become available.

6. Evocative examples of the complex scientific questions that can be answered through integrated simulation

To illustrate how the FSP would address the major scientific questions posed above, we describe some evocative examples of the solutions that would be developed as part of the FSP. Each of these examples addresses several of the basic questions posed in sections 3 and 4.

a. How do processes in the plasma edge set its structure, which critically affects global plasma confinement?

Under conditions typical of high-performance tokamak operation, the temperature and density undergo a rapid change over a narrow region in the edge of the plasma, rising steeply across this region. The structure of this region, known as the “edge pedestal”, and the transient MHD events within it (“edge-localized modes”, or ELMs) are found from experiments and modeling to have a strong impact on confinement in the core plasma, and hence on overall device performance. The edge region is also important because the plasma conditions in this region strongly impact the ability to handle power loads, exhaust the helium ash and fuel the plasma.

A predictive model of the edge pedestal remains an outstanding challenge because of the broad range of space and time scales, the large variation in the collisions rates between the plasma particles, the complex magnetic field structure, and the multitude of physical processes known to play a role: micro-turbulence, MHD, neutral-particle transport, impurity transport, and plasma-wall interactions. The edge community has developed a variety of codes that can handle pieces of the problem: two- and three-dimensional plasma transport, three-dimensional plasma fluid turbulence, near-surface chemistry, impurity transport, neutral-particle transport, and MHD stability. These codes are presently mostly run in isolation. There are also gaps in the edge code arsenal which are beginning to be addressed: kinetic edge turbulence codes are beginning to be developed, as are molecular dynamics-based models of plasma-wall interactions. Predictive simulation of the edge region requires an integration of all of the above pieces, completion of those pieces not yet well-developed (e.g. kinetic transport and turbulence), and development of the mathematical and computational advances required to bridge the range of space and time scales. This is a substantial undertaking that will require—and be enabled by—resources of the magnitude envisioned for a topical FSP project.

b. How can the injection of radio-frequency radiation stabilize central MHD oscillations?

Inductively driven tokamaks, like ITER, exhibit a periodic reconnection of the magnetic field lines in the center of the discharge, known as the “sawtooth oscillation”. If the period of the reconnection is short, and the amplitude remains small, this oscillation merely serves to regulate peaking of the plasma current, and does not adversely affect the fusion output of the device. However, in ITER and other burning plasma experiments, fusion produced high energy alpha particles will be present that tend to lengthen the sawtooth period and increase the amplitude of these oscillations. This makes them potentially destructive as these large symmetry breaking oscillations are likely to couple to other plasma instabilities and lead to a major disruption. At the very least, they would lead to large changes in the central plasma conditions and the central plasma fusion power production, and also potentially to large pulsed energy loads on the tokamak structures. Experiments have shown that the application of external radio frequency (RF) fields can control these oscillations, but the detailed computation of the needed powers, frequency, and structure of the RF fields requires the coupling of two of our most advanced plasma models: global MHD stability and Radio Frequency (RF) wave propagation. This coupled calculation would enable us to study the conditions required for sawtooth stabilization, and would give us the tools necessary to design, operate and control a sawtooth stabilization system.

c. How can the validity of MHD be extended beyond the fluid regime by the inclusion of kinetic effects?

The importance of global dynamics in a magnetic fusion experiment cannot be overstated. All power-producing fusion experiments must operate near stability boundaries that are set by the onset of global instabilities, and it is critical that we be able to calculate these boundaries with high-precision. There are presently two mainline thrusts for simulating the basic detailed physics of a hot tokamak plasma: micro-turbulence and macroscopic global dynamics. These two approaches use very different sets of equations to describe the hot, magnetized plasma. The micro-turbulence thrust uses the gyrokinetic equation, which represents the plasma in a 5 dimensional phase space, 3 real space and 2 velocity dimensions. The third velocity dimension, corresponding to the rapid gyration of particles around the magnetic field, is analytically averaged over a gyration period using a local expansion. The macroscopic global dynamic thrust uses 3 dimensional fluid-like equations, the "extended MHD" equations, in which the plasma response has been characterized by a complicated tensor "equation of state" and highly anisotropic turbulent transport coefficients. While these two approaches are each valid in the appropriate limit, it is now clear that there is an important class of phenomena for which neither of these approximate descriptions is valid. A prototypical example of this class of phenomena is the tearing mode, one of the main contributors to plasma disruptions. The tearing mode can best be thought of as containing an "outer region" for which an accurate description of the global plasma is necessary, and a small "inner region" surrounding the potentially unstable surface, where steep gradients and large electric fields occur, and an accurate kinetic description of the plasma is essential. Thus, a promising approach in integrated modeling of tokamaks is to treat the inner region with kinetic equations and the outer region with the extended MHD equations. This is a new research topic and would require writing a new code that builds on the knowledge gained from prior micro-turbulence and macroscopic global dynamics work and integrates that knowledge into a new code.

d. What are the consequences of plasma fine scale turbulence on global transport length and time scales?

Plasma microturbulence is primarily responsible for the transport of heat, momentum, and particles in the core plasma. Through the efforts supported by SciDAC and the base fusion program, there has been considerable progress in simulating plasma turbulence on the relatively short time-scale of turbulence saturation, and extracting from such simulations the fluxes of heat and particles and other statistical quantities. This problem is inherently kinetic in nature, and requires advancing distribution functions in five dimensions (three spatial, and two velocity). It is being addressed with both particle-in-cell codes, which represent distribution functions by discrete particles, and continuum models, which effectively follow the distribution function on a five-dimensional grid. Within each category, there are codes that focus on a local region about a particular field line ("flux-tube" codes) and global codes that can simulate an entire tokamak cross section. These various simulation models are making considerable strides in terms of consistency with one another and with experiments, even though they are still limited in terms of their space and timescale range and even though the physics models still have some restrictions.

Incorporating the results of these simulations into codes that calculate the evolution of the temperature and density profiles is challenging. The plasma temperature and density evolve on much longer length and time scales than found in plasma micro-turbulence. That evolution is currently computed using reduced models of the turbulence-driven fluxes that have a limited range of proven validity, and through transport-timescale equations that make the questionable assumption that the turbulent transport is describable through local diffusion and convection terms. The brute-force bridging of the full range of space and time scales between turbulence and transport is not feasible on current computers and at least unattractive on any computer in the foreseeable future. But there is good reason to believe that this bridging of scales can be achieved via a two-pronged approach. The first approach is to employ multiple-scale techniques to extend turbulence simulations to transport space and timescales. The second approach is to simultaneously perform simulations which encompass the full range of scales and physics that contributes to the turbulence and analyze these simulations to establish their validity and to extract the mathematical nature of the fluxes they drive, and then develop transport-timescale equations consistent with the results. It may be that a combination of these approaches will be successful; for example, one might be able to deduce a partial scaling from a combination of theory and analysis of simulations, and then use simulation to construct a database over a tractable number of remaining parameters. With regard to the first approach (application of multiple-scale techniques), several techniques have been proposed from within the fusion plasma physics community, and there are others under development in the computational sciences communities that could be applicable.

e. How should we design a mathematical and software framework that integrates physics modules containing disparate mathematical models and discretizations to create larger metasimulations?

Numerical analysis rarely treats the problem of multiple discretizations interacting for a single computation, even when each just provides boundary conditions for the others. Current codes that are likely to form initial building blocks for an FSP involve differing mathematical models, discretizations, and initial/boundary assumptions – and in some cases will be modeling the same 3D spatial region rather than interacting only at boundaries. Without care, it is possible to have a highly accurate constituent module while the combined simulation is less than first-order accurate. The analysis of this cross-model problem will likely also be of use in areas such as biological cell and nanostructure simulations. Even when reduced to software components, connecting two models will likely involve a pipeline of intermediate mesh mediators, filters, and interpolators. For the terascale simulations required by FSP, *all* components in the pipeline will need to be scalably parallel. Currently connecting even two parallel programs at runtime is a major research challenge; optimally assembling and connecting the full set of components needed for an integrated FSP will require combined efforts of physicists, mathematicians, and computer scientists.

7. Fusion Simulation Project Representative Design Tasks

The FSP Steering Committee recommends that DOE empower a design team to carry out a conceptual design of the FSP. Just as is the case for a $300M experimental project, the development of a coherent design and project plan for the FSP is an essential prerequisite for beginning the project. Developing that design and plan is an activity requiring dedicated and detailed work by the project initiators. Experimental projects typically allocate about 10% of the total project cost for design development (including initial research and development). The history of the development of computing is replete with examples of poorly designed codes that failed to deliver the required capability due to the lack of a proper design and plan. The project will be ramped up gradually. The DOE has initially allocated about $2M in FY05 for developing the initial conceptual design, $1M from OFES for planning the physics module development and $1M from OASCR for planning the mathematics and software infrastructure. Since $1M will support about 3 Full Time Equivalents (FTEs) staff, this translates to about 3 FTEs for the physics and 3 FTEs for the computer science and mathematics planning. A key part of developing a realistic design is the construction and testing of prototype applications and software infrastructure, and the design activity will include as much prototype development as is useful and feasible.

To accomplish its goal to provide the predictive capability ITER requires, the FSP will need to develop practical algorithms and packages for accurately simulating the major physics elements of ignited tokamaks. This will require research and development in many areas of physics as well as incorporating known physics in an interoperable software framework. It will also be essential to develop the necessary computational mathematics and computational software infrastructure. To give a concrete illustration of one vision for the conceptual design phase, a candidate assignment of tasks by discipline that highlights the areas that need to be addressed by the design team is listed in Table 1. Each area would be addressed by roughly 1/2 FTE. That support would strongly leverage the existing SciDAC and base program activities. The assembled team will have to be multi-institutional to capture the diverse expertise needed for all of the topics. This tentative assignment is only one of several possible set of candidate tasks, but it gives an idea of the level of work that will need to be accomplished during the first year, and the level of work consistent with 6 FTE’s. A key priority is canvassing the fusion science, the computer science and the computational science communities for their perspectives and judgment to help develop the path forward. This is just one example of candidate sets of tasks. Others are also possible.

Table 1 Task areas for conceptual design

Topic	1/2 FTE	1/2 FTE	1/2 FTE	1/2 FTE	1/2 FTE	1 FTE
OFES Physics	Extended MHD	Turbulent Transport	Plasma Edge	Whole Device Modeling	Sources (RF, current, NBI,…)	Central Project Coordination
OASCR Comp. Sci.	Computational Mathematics		Programming Model	Software infrastructure	Collaboration infrastructure	

As the project support level increases during the second year (Figures 4 and 6), the plan developed by the conceptual design team would be implemented. Startup activities would likely involve formation of initial teams for high priority activities, initial prototype code and algorithm development. If progress in some areas is very rapid, prototyping and initial algorithmic develop could begin during the first year.

A candidate summary description of each of the task areas listed above is given in Appendix 5.

8. Project Organizational Requirements

To develop an integrated simulation capability on the timescale needed for ITER will require a large project staff of 60 to 80 physicists, computer scientists, and applied mathematicians. The FSP is thus not only one of the most technically ambitious computational science projects ever undertaken by the Office of Science, but also one that faces new organizational challenges far beyond those faced by past, smaller-scale efforts.

The FSP organization must be accountable to DOE, be flexible to balance project focus with research and development, be able to make timely decisions, and be responsive to the user community, including ITER and the US and international fusion and computational science communities. In addition, it must be able to provide direction and coordination and support for the FSP team and be responsive to DOE requests for planning, budget, and progress documentation.

The FSP requires setting up a focused multi-institutional research organization that is accountable to DOE for timely delivery of a product that meets the needs of the chief customer, ITER. The DOE will require accountability, project planning and tracking, reporting of results and considerable oversight. Accountability can only be achieved if the FSP project management has the authority to make decisions and introduce the necessary level of project discipline. At the same time, most of the FSP efforts will involve research and development of new physics and mathematical algorithms. The need for accountability and a project organization and the need for sufficient freedom by the individual project elements to do the innovative research necessary to develop new ways to solve hard problems pose potential conflicts. In addition, the FSP faces a major challenge in that it must be a multi-institutional project. The expertise needed for the project to succeed is spread across the fusion and computational mathematics community. Typically, almost all large *successful* software projects require close and constant communication. In particular, success of the FSP will require utilizing the OFES and OASCR experimental, theoretical and computational programs, the SciDAC programs, and the experience and expertise of the general computational science community. The FSP leadership must be able to make timely decisions to resolve conflicts, ensure timely progress, allocate and manage resources, and generally function like the project that it is. The FSP must be responsive to all of its stakeholders in the US, in ITER and the international fusion and computational scientific communities. Finally the FSP senior management must be able to provide the technical and organizational leadership and support for the FSP team members that are essential for success.

An important priority for the project will be continual support of and interaction with the ongoing national and international fusion program. That support and interaction should begin at the start of the design phase. Gaining acceptance by and support from the national and international fusion community as rapidly as possible must be one of the main priorities of the FSP. Modeling capability should be deployed into the fusion program as soon as it is developed, both for the purpose of validation and feedback on the correctness and appropriateness of the models, and for feedback on the utility of the capability from the customers. Initial deployment should be within a few years of the beginning of the project

9. FSP Steering Committee Summary Recommendations

We recommend that the FSP proceed in two steps:

1. DOE should empower a design team to develop a conceptual design for the FSP. The DOE should review the conceptual design when it is complete, and decide whether or not to proceed with the project.
2. If DOE decides to proceed with the FSP, the DOE should empower a project team to proceed with the project.

Both the conceptual design and the project efforts should be organized to meet the requirements listed above (accountable, flexible, responsible, decisive, agile, etc.). The members of the conceptual design team should be drawn from many institutions. The conceptual design team and the project team must be accountable to DOE and the OFES and OASCR communities.

Appendix 4 lists a set of criteria that can be used to guide the project management team that DOE will chose to lead the project.

10. Acknowledgements

The FSP Steering Committee gratefully acknowledges the support and assistance from:

- W. Nevins, D. Schissel, S. Scott, C. Kessel, J. Carlsson, N. Sauthoff, S. Kruger, S.Shasharina, S. Hendrickson, L. Pritchard, and A. Sanderson, who participated in the three day planning workshop,
- S. Eckstrand, A. Kritz, W. Miner, J. Willis, A. Davies, M. Strayer and Ed Oliver from DOE OFES and OASCR who provided support and guidance to the committee;
- Members of the FESAC ISOFS committee, especially J. Dahlburg and W. Lokke;
- Members of the fusion and computational mathematics community, especially T.Rognlien, A. Glasser, W. Tang, W. Kruer, J. Callen, G. Bateman, V. Chan, C. Baker and Allen Boozer.
- N. Riebe for editorial and organizational assistance,
- A. Bécolet, P. Strand, H. Wilson, K. Thomsen, G. Cordey, J. Connor, and V. Parail, for discussions about the EU simulation effort,
- A. Fukuyama, M. Honda, A. Akutsu, A. Takabe, Y. Kishimoto, R. More, R. Janev, and T. Kato for discussions about the Japanese simulation effort
- M. Shimada, A. Costley, M. Sugihara, Y. Gribov, Y. Shimomura, N. Sauthoff and A.Kukuskin, for discussions about the ITER requirements for the FSP.
- The Tech-X Corporation, for hosting the May 18-20 planning workshop and for providing computer support.

11. Glossary

AMR—Adaptive Mesh Refinement
CCSM—Community Climate Systems Model
DIII-D—(Doublet III tokamak experiment at General Atomics, San Diego)
DOE—Department of Energy
ELMs—edge-localized modes
FESAC—Fusion Energy Science Advisory Committee
FII—Focused Integration Initiative
FSP—Fusion Simulation Project
FTE—Full Time Equivalent
IPS—Integrated Plasma Simulation
ISOFS—Integrated Simulation & Optimization of Fusion Systems sub-committee of the Fusion Energy Science Advisory Committee
ITER—International Thermonuclear Experimental Reactor
MDSPlus—a set of software tools for data acquisition and storage and a methodology for management of complex scientific data developed by the magnetic fusion community
MHD—magnetohydrodynamics, the science of the force balance between the plasma pressure and the magnetic field pressure
NNSA—National Nuclear Security Agency (a part of DOE that is responsible for nuclear defense and reactor issues)
OASCR—Office of Advanced Scientific Computing Research
OFES—Office of Fusion Energy Sciences
PetaFlops—10^{15} floating point operations per second
PIC—Particle in Cell
R&D-Research and Development
RF—radio frequency
SciDAC—Scientific Discovery through Advanced Computing
SC—Office of Science
TeraFlops—10^{12} floating point operations per second
US—United States

12. Appendices

Appendix 1
Charge to FSP Steering Committee from John Willis

October 7, 2003

Don Batchelor, ORNL
John Cary, Univ. of Colorado
Ron Cohen, LLNL
Steve Jardin, PPPL
Doug Post, LANL
Randy Bramley, Indiana University
Mike Heath, Univ. of Illinois

The fusion community has identified a comprehensive simulation capability as a critical program element for the future. Thus, the Department of Energy is considering initiating a Fusion Simulation Project (FSP) in FY 2005 to provide timely support of present experiments and, ultimately, ITER. Because the FSP will be a very complex project with challenging goals, it is essential to begin planning for the FSP in advance of project initiation. Accordingly, the Office of Fusion Energy Sciences and the Office of Advanced Scientific Computing Research are appointing a committee that will address management issues and provide the initial planning for the implementation of the Fusion Simulation Project.

The task of the committee will be to plan, through an open process that involves community input, the initial phase of the Fusion Simulation Project. In a report, the committee will address the technical scope of the project, the scope during the first five years of the project, as well as the issue of the management structure. In the attached document the nature of the project and the committee's tasks are described in somewhat greater detail.

I am writing to ask if you would be willing to serve on a seven member Fusion Simulation Project planning committee. It is expected that the committee will complete the tasks by end of the first quarter of FY 2005 and that the initial phase of the Fusion Simulation Project will get underway in FY 2005.

Sincerely,

John W. Willis
Director, Research Division
Office of Fusion Energy Sciences
Office of Science

(Mike Heath couldn't serve, but Phillip Colella was able to serve.)

Initiation of the Fusion Simulation Project

Background

The fusion community has identified a comprehensive simulation capability as a critical program element for the future. Thus, the Department of Energy is considering initiating a Fusion Simulation Project (FSP) in FY 2005 to provide timely support of present experiments and, ultimately, ITER. Because the FSP will be a very complex project with challenging goals, it is essential to begin planning for the FSP in advance of project initiation. Accordingly, the Office of Fusion Energy Sciences and the Office of Advanced Scientific Computing Research will appoint a steering committee that will address management issues and provide the initial planning for the implementation of the FSP.

Details

In 2002 the Fusion Energy Sciences Advisory Committee recommended the initiation of a project to develop an integrated simulation capability for fusion systems. A description of the FSP project, that was prepared by a FESAC sub-panel and approved by FESAC, is available at http://ofes.fusion.doe.gov/News/FSP_report_Dec9.pdf FESAC recommended an annual funding level of $20-25M and a 15-year period for the project. The Office of Fusion Energy Sciences (OFES) and the Office of Advanced Scientific Computing Research (OASCR) are considering the initiation of such a project in FY 2005. The initial funding will be a fraction of the recommended funding, but it is anticipated that the annual funding would increase to a significant fraction of the planned level over a 3-4 year period.

The goals of the FSP are very ambitious, and it will be a very challenging project. Much of the required science will have to be developed over the course of the project. In addition, there are fundamental physics/mathematics issues that must be solved to handle the wide range of time and space scales, extreme anisotropy, non-linearity, models of different dimensionality, and complex geometries that must be included in an integrated simulation. Finally, it will be necessary to coordinate wide segments of the fusion science, computer science and applied math communities and interact with two offices in the Office of Science. To succeed the FSP must be managed like a major device fabrication project with specific deliverables and schedules.

Because of the complexity of the FSP, it is likely that up to 15% of the annual funding will be allocated for management tasks. In addition, in order to ensure that the FSP is a multi-institutional project, no institution may receive more than 25% of the total annual funding for both management and technical activities. Since initial funding in FY 2005 is anticipated to be about $4 million, this would mean limits of about $600K for management and $1M total for any one institution. The expectation is that funding would increase moderately in FY 2006, so the corresponding limits would increase similarly. If prospects for the FSP funding remain favorable, OFES and OASCR plan to provide some initial planning funding in mid-FY 2004.

Scope of Tasks to be Carried Out by the Steering Committee

The members of the appointed steering committee will work as a team in investigating technical and management aspects of other integrated scientific simulation projects through a process that includes input from the fusion science, applied math and computer science communities. The appointed committee will prepare a report regarding organization and management structure of a project such as the FSP (for example the role and tasks of committees such as a project steering committee, a project advisory committee, and working groups). The report should also include a description of the responsibilities of key management personnel (for example a project manager, a physics coordinator, and a computer science/math coordinator).

The Steering Committee will also consider and evaluate potential focused integration initiatives (FIIs), for example, the FIIs defined in the FESAC FSP report described above, the computational frameworks that might be used to connect the component physics, and the computer hardware needed to carry out the project. The evaluation will include community input obtained through community workshops organized by the Steering Committee. The Steering Committee will prepare a report that contains detailed plans for implementing one or more focused integration initiatives in the initial phase of the FSP project. The report will describe the physics content and objectives of the focused integration initiative, the approach to integration issues, computational framework to be used in the project, and detailed implementation plans, costs, and schedules. A comparative review process will be used to form funded teams that will begin the work on 1 or 2 FIIs selected for the initial stage of the FSP project.

Appendix 2
Summary Biographies for the Fusion Simulation Project Steering Committee

Douglass E. Post, chair, has had over 30 years of experience developing and applying large-scale multi-physics simulations at the Princeton University Plasma Physics Laboratory (PPPL), the Lawrence Livermore National Laboratory (LLNL) and the Los Alamos National Laboratory (LANL). Dr. Post led the tokamak modeling group at PPPL from 1975 to 1993, and more recently was the AX Associate Division Leader for Simulations at LLNL and then the X Deputy Division Leader for Simulations at LANL. He served as head of International Thermonuclear Experimental Reactor (ITER) Physics Project Unit (1988-1990), and head of ITER In-vessel Physics Group (1993-1998). He has published over 200 papers with 4500 citations in computational, experimental and theoretical physics. He is a Fellow of the American Physical Society and the American Nuclear Society. He received the Outstanding Accomplishment Award from the ANS in 1992 for his leadership of the ITER Physics Project Team. He is presently an Associate Editor-in-Chief of the joint AIP/IEEE publication “Computing in Science and Engineering”. He received a Ph.D. in physics from Stanford University in 1975. His current interests include identifying and quantifying software development management and engineering practices that improve the way the development of “scientific” software is carried out. He leads the multi-institutional DARPA High Productivity Computing Systems Existing Code Analysis team.

Donald D. Batchelor is the Plasma Physics Group Leader in the Fusion Energy Division of the Oak Ridge National Laboratory. He has almost 30 years of experience in plasma and fusion theory and computational physics. He has published extensively on the theory of plasma confinement in a variety of magnetic configurations, and on the interaction of radiofrequency waves and plasmas. The codes developed by his group have utilized to study plasma heating and current drive on a broad range of fusion experiments, including the Tokamak Fusion Test Reactor at Princeton, the DIII-D tokamak at General Atomics, and the Joint European Torus at Abdingdon, UK. He has been a key member of the design teams for a number of proposed fusion experiments including the Burning Plasma Experiment, the Tokamak Physics Experiment, and Quasi-Poloidally Symmetric Stellarator. Donald was the PI for a recent SciDAC project on “The Numerical Calculations of Wave-Plasma Interactions in Multi-dimensional Systems.” He is an active member of the APS Division of Plasma Physics, and is a Fellow of the APS. He received a B.S. in Mathematics from MIT in 1968 and a Ph.D. in Physics from the U. of Maryland in 1976.

Randall B. Bramley is associate professor of computer science at Indiana University, having joined IU in 1992 after three years as senior computer scientist at the University of Illinois Center for Supercomputing Research and Development. Currently he is on the Policy Committee of the IU School of Informatics and has research appointments with the Pervasive Technologies Laboratories and Informatics Research Institute. Over the past five years he has served on committees for all of the major conferences in supercomputing and high-performance computing, and is a member of the Common Component Architecture Forum. Bramley's research interests include large scale scientific and engineering computations, particularly parallel sparse linear solvers and integrating multidisciplinary codes. He works on software component architectures for scientific computing and distributed Grid systems as part of the DoE CCTTSS SciDAC, and scientific data management systems including rapid assembly of data tools, digital libraries, data grids, and portals. He currently is part of the NSF Middleware Initiative, integrating sensors and instruments with web services, data systems, and computations.

John R. Cary is a professor of physics at the University of Colorado at Boulder, where he has served as department chair and science faculty mentor, and he is CEO of Tech-X Corporation, a research corporation in Boulder, CO. His prior positions were at the Los Alamos National Laboratory and the Institute for Fusion Studies at the University of Texas. Professor Cary's interests are in plasma physics, beam physics, and nonlinear dynamics, as well as in the application and development of modern computing techniques to computational physics. He has more than 90 publications on various aspects of physics and computation. His recent work includes architecting and developing VORPAL, an object-oriented, highly flexible plasma simulation application that computes with scaling good to 4000 processors, yet can and does produce significant results on single processors due to its arbitrary dimensionality. This work illustrated that a properly designed, large-scale numerical application could be object oriented and written in C++ without any performance penalty. Already in its short life, VORPAL has been instrumental in one Physical Review Letter and one article in Nature, both dealing with the formation of quality beams via laser-plasma interaction, and it currently being used to study radiofrequency heating of plasmas and the flow of low-density gases. Professor Cary also pioneered the use of remote objects via CORBA for connecting components as needed to create a GUI driven computational server interacting with experimental data. Professor is an associate editor of *Physical Review E,* a member of the National Research Council's committee on plasma physics, and a past associate editor of *Physical Review Letters*. He has chair multiple national committees for the Division of Plasma Physics of the American Physical Society (APS) and served on many others for the APS, of which he is a fellow. He received a B.A. in mathematics and a B.A. in physics from the University of California at Irvine, and an M.A. and a Ph.D. in physics from the University of California at Berkeley.

Ronald H. Cohen is Associate Program Leader for Magnetic-Fusion Theory and Computations at Lawrence Livermore National Laboratory, where he has been a staff member for over 30 years. A Fellow of the American Physical Society, Dr. Cohen has an extensive list of publications in magnetic fusion, particularly in the areas of tokamak edge physics, radio frequency-plasma interactions, magnetic-mirror confinement, and spheromak physics. He has also been active in astrophysics, space plasmas, and computational fluid dynamics, and, currently, heavy-ion fusion. He was co-PI of the team that won the Gordon Bell Prize for computational performance in 1999, and a co-manager of the fusion program's National Transport Code Collaboration. Dr. Cohen also serves on the Advisory Committee for the Plasma Science Advanced Computing Initiative (PSACI), the Fusion SciDAC portfolio. He holds S.B. and PhD degrees from MIT.

Phillip Colella is the leader of the Applied Numerical Algorithms Group at the Lawrence Berkeley National Laboratory. His group develops advanced numerical algorithms and software for partial differential equations and applies the software to problems of independent scientific and engineering interest. The principal focus of Colella's current work is the development of new simulation software tools for multiscale problems in science and engineering. Phillip is the project leader for the Applied Differential Equations Integrated Software Infrastructure Center (APDEC), a SciDAC program. Phillip has been one of the pioneers in the development of adaptive mesh refinement (AMR). AMR is widely used in shock physics, turbulence, astrophysics, flow in porous media and combustion. In 1998, Colella was the recipient of the IEEE Computer Society's Sidney Fernbach Award. Last year, Phillip Colella and John Bell were named as co-recipients of the 2003 SIAM/ACM Prize in Computational Science and Engineering. In 2004 Phillip was elected to the National Academy of Sciences. Phillip received a B.S., M.S and Ph.D in applied mathematics from the University of California at Berkeley.

Stephen C. Jardin is a Principal Research Physicist at the Princeton University Plasma Physics Laboratory. He is presently Co-Head of the Computational Plasma Physics Group at PPPL, Head of the Next Step Options Physics, and Theory Department MHD Coordinator. He has been Lecturer with Rank of Professor in the Princeton University Astrophysics Department since 1986. He was the primary developer of several widely used MHD equilibrium, stability, and transport codes including the Tokamak Simulation Code (TSC). He is the Principal Investigator for the SciDAC "Center for Extended Magnetohydrodynamic Modeling (CEMM)" which involves researchers from over 12 institutions. He holds 4 US Patents, has had over 150 refereed publications in plasma physics, and has supervised 6 Princeton University PhD students. He is a Fellow of the American Physical Society. He has held group-leader positions in several fusion device design teams including those for S-1, PBX-M, CIT, BPX, TPX, ARIES, and ITER. He is presently a member of the NERSC Executive Committee Users Group, a past member of the ESNET Steering Committee, and is Chair of the NERSC Program Advisory Committee and of the National Transport Code Collaboration (NTCC) Program Advisory Committee. He holds a BS in Engineering Physics from the University of California, a MS (Physics) and MS (Nuclear Engineering) from MIT, and a PhD in Astrophysics from Princeton University (1976).

Appendix 3

A US-ITER Perspective on Fusion Simulation
Ned Sauthoff, US ITER Planning Officer

The FSP team should look toward what an experimentalist/ engineer/ designer/ data-analyst should be able to do 10-15 years from now, with intermediate deliverables to enable prototyping and acquisition of experience in such modes of research participation.

I see simulation of the ITER plasma and system as a key tool in 5 application areas:

1. facility design
2. plasma scenario development
3. control system design and tuning
4. experimental shot design
5. data analysis

Particularly in the research phase, leadership in simulation will be a major contributor to winning proposals for ITER run-time. US leadership in simulation will be key to achievement of US research objectives in the environment of the ITER international topical research teams. US effectiveness in ITER research may be tied to leadership in simulation and measurement.

I see the optimum path as one that balances discovery and model-building in topical areas with integration, with a progression of ever-improving and more detailed topical models being integrated into tools suitable for application by non-specialists.

1. Facility Design (from now until ~2008, linked to ITPA design-R&D)

must recognize new burning plasma physics (size-scaling, energetic particles, and strong self-heating with the associated complexity) and include:

- transport prediction, especially the integrated core and pedestal
- power and particle loads, and assessment of mitigation methods
- ELMs and their mitigation
- NTM stability/thresholds/mode-characteristics/saturation techniques (e.g., driven currents in the islands, delta-prime modification)
- rotation drives and damping (including error fields)
- energetic particles modes in positive-, weak-, and reversed-magnetic shear profiles
- heating and current drive (both actuators and plasma responses) for NB, IC, EC, and LH
- more integrated issues like EC-system requirements for NTM stabilization, RWM stabilization by rotation and feedback, disruptions in the 3D vessel structure
- diagnostics specifications (measurement requirements) based on simulations to prototype synthetic diagnostics and data analysis

2. Plasma Scenario Development (from ~2007-8 for heating and current drive, 2012 plasma scenarios)

(Note: the ITPA coordinating committee recently recognized the need for an integrated focus and established the steady-state scenario topical group)

- plasma scenario development will contribute to choices of plasma control tools (heating and current drive, fueling, …) and diagnostics
- prediction and optimization of performance within operational boundaries of stability, confinement, wall interactions, AC losses in cold structures, actuator limits, diagnostic limits
- design of experiments that focus on desired physics and minimize interference from other phenomena (e.g., energetic particle modes with minimum damping from the slowing-down beam distribution; transport experiments that avoid domination by MHD)
- compelling visualizations to enable understanding at many levels

3. Control System Design and Tuning (prototype ~2008-14)

to develop, benchmark, and refine control algorithms based on the standard model of plasma, actuators, and diagnostics

  - full-detail/resolution models will likely not be needed in real-time; reduced models will likely be sufficient for real-time control
  - however, the reduced models will involve parameters to tune algorithms, program state-dependent gains, etc., which could be tuned by off-line use of the full-detail/resolution models

4. Experimental Shot Design (from ~2012)

- ITER shots are too valuable to “play by ear”; a simulation is warranted and may be needed in the more-self-organized burning plasma
- a faithful model of the ITER plasma/actuator/sensor system will be needed to design the shot, to focus on desired objectives and to avoid encounters with limits of the plasma and control tools
- the model should enable tuning of the algorithms as well as derivation of reference waveforms
- US leadership in such modeling would well position the US for ITER research in the environment of international topical teams

5. Data Analysis (from ~2008)

- derive results with flexible application of measured data and simulated parameters, as well as comparisons of measured data and predictions
- benchmark and validate integrated model and embedded sub-models by means of flexible mix of measured data and predicted/extrapolated quantities.

All tools should be compatible with an overall remote participation environment of modeling, data access, results management, run-planning, run-execution, run-assessment, program planning, and team facilitation.

Appendix 4

Initial Design of the Fusion Simulation Project

Department of Energy

SUMMARY: The Office of Fusion Energy Sciences (OFES) and the Office of Advanced Scientific Computing Research (OASCR) of the Office of Science (SC), U. S. Department of Energy (DOE), announces its interest in receiving applications for initial software design of the Fusion Simulation Project software. After a successful review of the initial software design, the successful applicant will be expected to carry forward and execute the project.

The Fusion Simulation Project is to be a long-term activity funded jointly by OFES and OASCR, with the objective of providing an advanced simulation capability for future burning plasma experiments such as the proposed International Thermonuclear Ignition Reactor (ITER), and experimental and power-producing devices beyond ITER. The project will also provide advanced simulation capability for the existing magnetic fusion energy (MFE) experiments. The project will be comprised of a production component, a software infrastructure component, and a research component. The production component will consist of a suite of advanced plasma simulation codes and modules built around a common software structure allowing data exchange between codes, as appropriate. This suite will contain, but will not be limited to, a Whole Device Modeling (WDM) capability that is able to simulate the global physics of the whole burning plasma device, although initially at a greatly reduced level. The research component will entail developing new simulation capabilities allowing the coupling of an ever-wider range of physical phenomena together in the same simulation, and preparing these capabilities for submission to the production component. Each component should be tightly coupled to the OFES and OASCR program elements, and should be applied, as soon as feasible, to the analysis of existing fusion experiments. The software infrastructure will allow all of this to take place in a seamless and efficient manner, and will provide visualization capability and access to experimental data as appropriate.

The nature of this project requires close collaboration among researchers from the disciplines of computational and theoretical plasma physics, computer science, and applied mathematics. It is recognized that the existing legacy software that will form the backbone of this project is spread widely throughout the community, and that the expertise required to combine and to extend this software resides at many laboratories, universities, and industrial sites. Partnerships among universities, national laboratories, and industry are therefore strongly encouraged.

The FSP multi-institutional research organization will be accountable to DOE for timely delivery of a product that meets the needs of the chief customer, ITER. In that respect, it is similar to an experimental project like DIII-D or TFTR. Like an experimental project, the DOE will require accountability, project planning and tracking, reporting of results and oversight of the FSP. The requirements for project accountability and the need to ensure sufficient flexibility to carry out the necessary research and development of new physics and mathematical algorithms pose potential conflicts. The FSP leadership must be able to make timely decisions to resolve conflicts, ensure timely progress, allocate and manage resources, and generally function like the project that it is. The FSP must be responsive to all of its stakeholders in the US, in ITER and the international fusion and computational scientific communities. Finally the FSP senior management must be able to provide the technical and organizational leadership and support for the FSP team members that are essential for success.

Proposals should include the following elements:

- A plan to assess what physics modules are needed for a comprehensive model of a tokamak burning plasma and what are presently available.
- A plan to develop an integrated whole device modeling capability for a burning plasma that would be of near-term value, but would be extendable in a way that will meet the needs of an ITER-class facility in the 10-15 year time frame.
- A plan to develop an appropriate software framework and infrastructure as required so that the different module components will be interoperable. Both serial and parallel operation need to be considered.
- A plan to identify which software interfaces are required between modules, and how these can best be specified in a way that best facilitates code-coupling, yet is extensible in anticipation of future needs.
- A plan to carry out an assessment of each of the subfields of macro-stability, turbulence, plasma edge, and sources to determine what research and development is required in each area, including the identification of promising new computational mathematics algorithms.
- A plan to integrate the existing experimental database seamlessly into the FSP framework to facilitate simulation of and comparison with today's experiments.
- A plan for developing the collaboration infrastructure to support a multi-institutional collaborative project that requires close coordination to succeed with minimal travel among sites.
- A plan for the management that addresses the specific issues and challenges of a multi-institutional project involving a substantial degree of research and development, while ensuring the degree of accountability required of a project of this scope.

It is expected that the proposed planning activity will involve strong community input, and will involve some prototyping of candidate software and approaches. The detailed plans provided are expected to form the basis for a follow-on proposal to begin development of the FSP.

Proposals will be evaluated based on the following criteria:

- The degree to which the proposed work is judged to address the workscope identified in this announcement.
- The competency of the key personnel to carry out the proposed work.
- The degree to which the proposed work takes maximum advantage of appropriate existing expertise and software.
- The degree to which prospective users of the software are brought into the planning process at an early stage and the degree to which user input guides the evolution of the project.

Appendix 5

Candidate Summary Descriptions of Initial Conceptual Design Tasks

As noted in Section 6, the DOE has tentatively allocated $2M for the conceptual design phase of the FSP. This will be $1M from OFES and $1M from OASCR. This will support roughly about 3 FTEs in fusion energy sciences and 3 FTEs in computer science and computational mathematics. A candidate list of tasks is given in Table 1. Other candidate sets of tasks are also feasible.

Table 1 Task areas for conceptual design

Topic	1/2 FTE	1/2 FTE	1/2 FTE	1/2 FTE	1/2 FTE	1 FTE
OFES Physics	Extended MHD	Turbulent Transport	Plasma Edge	Whole Device Modeling	Sources (RF, current, NBI,…)	Central Project Coordination
OASCR Comp. Sci.	Computational Mathematics		Programming Model	Software infrastructure	Collaboration infrastructure	

These tasks are expanded in the following text.

Plasma Physics (OFES):

As described in the FESAC ISOFS report[2], the physics challenge of the FSP is daunting. Phenomena spanning many disparate time and distance scales must be integrated into practical modules that can be combined into a coherent simulation framework. While substantial research is needed for three areas, Extended MHD, Turbulent Transport and the Plasma Edge, as well as in sources (RF, impurities, etc.) as well, there is a base of existing capability in these areas. The key tasks are to develop a detailed strategy for producing a reliable predictive capability in each area beginning with the existing capability. An essential goal is to deliver “incremental” capability to the tokamak community on the short term (a few years after project start) as part of the process of delivering a comprehensive simulation capability by the end of the project.

The Extended MHD, Turbulent Transport and Plasma Edge tasks would involve assessing what’s needed and what is available to use as basis to develop a practical set of predictive modules that could be integrated into a whole device model by the end of the project, and would be available in the integrated model and as stand-alone tools for prediction, analysis and design for ITER. A plan for the required research and development would be written for each area. The ongoing SciDAC programs in Extended MHD and turbulent transport and the similar efforts in the base OFES program would be key elements of these plans. While there is no SciDAC program in edge plasma physics, edge physics is a key area and a strong effort will be needed for the success of FSP. Basic, but focused, research in physics and solution algorithms will be a strong component of each of these topics.

One of the key deliverables of the FSP will be an integrated modeling package for ITER.

---

[2] Dahlburg, J., Corones, J., et al. (2001). "Fusion Simulation Project: integrated simulation and optimization of magnetic fusion systems." Journal of Fusion Energy **20**(4): 135-196.

Whole device models have been useful for tokamaks since the early 1970’s, and have become increasingly sophisticated and complete, but still fall short of being reliable predictive tools. This task will consist of a detailed assessment of the capability and the interoperability of the existing codes and code modules, including those developed and used in the EU and Japan; the identification and characterization of the physics and computational infrastructure improvements needed to improve their predictability; and the development of plans for producing a whole device model by the end of the project that meets the needs of the ITER team and is useful by the international tokamak community.

Modules that can compute the sources for neutral beam and radio frequency heating and current drive are required for all of the tasks, as are models for the physics of neutral hydrogen and impurity transport and their effects on plasma performance. While the level of understanding of these processes is relatively high, practical algorithms still require development and implementation. The existing capability needs to be assessed and tested, and plans need to be developed for how the needed capability will be developed. The SciDAC program on RF heating and the base program efforts in all of these areas will be key resources, but , as in the other areas, the existing level of support and the focus is inadequate to produce the needed modules.

All of these activities will involve prototype development and testing to the extent possible given the resources. The planning tasks will leverage the existing programs (SciDAC and OFES base program) very heavily.

Computational Mathematics and Computer Science (OASCR):

The FSP has many computational challenges. The FSP products must be able to use the state of the art computer platforms fifteen years or more from now, platforms whose architectures are unknown now. It is expected that the FSP will run on a variety of platforms, ranging from massively parallel platforms in the multi-PetaFlops class to small clusters and workstations. Success of the FSP will depend on how well the project can integrate many different types of complex physics algorithms into a coherent set of interoperable codes that can run on the range of computer platforms available to the FSP users and developers, and how well and closely a team with diverse skills located at many different institutions can work together on a daily basis.

Four main tasks are envisaged: 1. Identification and preliminary development of the most promising computational mathematics algorithms to solve the key physics problems, 2. Assessing the most appropriate programming models for the platforms of the future that can be used to solve the major physics problems, 3. Identifying and planning a design concept for the software framework and infrastructure, and 4. Identifying and prototyping the collaboratory infrastructure necessary to carry out the FSP successfully.

1. The computational mathematics task and challenge is to identify and initiate approaches to develop the new underlying mathematics associated with combining multiphysics models and models at disparate time and space scales, and the numerical algorithms that are best for solving the intrinsically nonsymmetric, non-self adjoint, anisotropic, and multiscale problems inherent in federated fusion energy problems. It will be important to determine candidate approaches to produce numerically accurate and re-usable components for mediating between spectral elements, finite differences on toroidal meshes, finite elements, AMR meshes, and the other discretizations required by fusion energy simulations. It will also be important to characterize the likely computer architectures that will be available to the FSP and their impact

on the choices for solution algorithms. These activities will require close collaborations with physicists working in each of the task areas.

2. Approaches for parallelizing the existing fusion codes and identifying approaches for efficient development of new parallel codes will be need to be defined. Massively parallel codes present complex programming challenges. In addition, the continued evolution of the architectures of high performance computers (vector, multiple node shared memory processor clusters, etc.) makes it necessary to assess the various candidate computer architectures that the FSP codes will be required to use, and identify suitable programming strategies. A strategy for performance optimization will also be required.

3. The FSP will be a large, multi-institutional project with many different development efforts. A strategy will need to be developed for the project to successfully integrate all of the components into a set of interoperable codes, and ultimately into a single code structure. The project will need to know how it can evolve disparate research codes under active change at different sites into a federated simulation platform. Currently a base of 50-100 existing codes is used for research in the constituent areas, and it is technically, financially, and sociologically impractical to reproduce their capabilities. It will be important to determine how we can semantically describe and define the interfaces between two parallel components and efficiently connect them at runtime. Current computer science research is addressing the semantics of defining a workflow among a set of serial components, but currently no way exists to even express how codes running differing numbers of processors should interact. It will be important to determine how we can best integrate data from disparate sources (other simulations, experiments, and theory) for use in running simulations, evaluating and comparing results for validation, and supporting tokamak operations. Common infrastructure computational, collaborative, and data management tools that can be widely used and shared in integrated fusion simulations will need to be identified and tested to determine which of those tools can be leveraged from the broader scientific computing community, and which are unique to magnetically confined fusion energy.

4. Experience with other code development efforts shows that close communication and highly coordinated development is essential for successful code development even for collocated code teams. The FSP will be a multi-institutional project composed of non-collocated teams and thus has special challenges. A practical approach for realizing good communication and coordination will be essential. It will be essential to evolve disparate research codes under active change at different sites into a federated simulation platform. The FSP conceptual design team itself will be non-collocated and multi-institutional and will serve as a prototype to explore the collaboration issues. This task will begin by setting up a collaboration infrastructure for the FSP conceptual design team. That experience and the experience of the computational science community including the SciDAC Fusion Collaboratory will then be used to set up the full collaboration infrastructure for the FSP.

Project Coordination

The design team will require full-time leadership and coordination (plus administrative support) to ensure that a coherent design and plan results from the efforts of the non-collocated and multi-institutional design team. The project leadership must be knowledgeable in fusion sciences, the computational science issues and software project management. Experience in coordinating and leading multi-institutional projects is also essential. A strong project leadership will be essential for making the many decisions that the design process will require. Such decisions will need to be based on technical merit but will also have to incorporate institutional issues as well. The project leadership must be able to integrate the concerns of the sponsors (DOE), the stakeholders (the SciDAC and OFES and OASCR communities), the customers (ITER, the US experimental and theoretical communities, and the EU and Japanese “FSP” activities), and the FSP team itself.

The set of tasks described above is not the only possible work breakdown structure for the FSP conceptual design phase, but it gives a sense of the minimum level of planning that will be required before the conceptual design phase can begin.