FESAC FUSION SIMULATION PROJECT (FSP) PANEL REPORT

William Tang[1], Riccardo Betti[2], Jeffrey Brooks[3], Vincent Chan[4], Thom Dunning[5], Charles Greenfield[6], Brian Gross[7], Michael Norman[8], Miklos Porkolab[9], and Rick Stevens[10]

October 18, 2007

I. INTRODUCTION

In his letter of charge (see Appendix 1) to FESAC, Dr. Raymond Orbach clearly identifies the overarching objective of the proposed Fusion Simulation Project (FSP) as being to "produce a world-leading predictive simulation capability that will be of major benefit to the overall science and mission goals of the US Fusion Energy Science Program." The expectation is that such a capability must be: (i) an important asset for optimizing US participation in ITER; (ii) relevant to major current and planned toroidal fusion devices; and (iii) strategically vital to US interests in developing DEMO. The associated major challenge for this project, which demands a strong alliance between DOE's Fusion Energy Science (FES) and Advanced Scientific Computing Research (ASCR) Programs, is to develop advanced software designed to use leadership class computers for carrying out unprecedented multi-scale physics simulations to provide information vital to delivering a realistic integrated fusion simulation model with high physics fidelity. Accordingly, our FESAC FSP Subcommittee was appointed to address this vitally important subject and respond specifically to five questions posed in Dr. Orbach's letter of charge. This Report, which contains our response to this task, is organized as follows:

---

[1] Chief Scientist, Princeton Plasma Physics Laboratory & Associate Director, Princeton Institute for Computational Science and Engineering, Princeton University

[2] Professor of Mechanical Engineering and Physics, University of Rochester

[3] Senior Computational Nuclear Engineer, Argonne National Laboratory

[4] Director of Theory & Computational Science, General Atomics, San Diego, CA

[5] Distinguished Professor of Computational Chemistry, University of Illinois @ Urbana Champaign, & Director of NSF's National Center for Supercomputing Applications (NCSA)

[6] Deputy Director, Experimental Science Division, General Atomics, San Diego, CA

[7] Deputy Director, NOAA's Geophysical Fluid Dynamics Laboratory, Princeton, NJ

[8] Professor of Physics and Center for Astrophysics & Space Sciences, University of California @ San Diego

[9] Professor of Physics & Director of Plasma Science Fusion Center (PSFC), MIT

[10] Professor of Computer Science, University of Chicago & Associate Director for Computational Sciences, Argonne National Laboratory

- Section (I) Introduction: This portion of the Report contains the Subcommittee's general perspectives together with a situation analysis that includes connection to the two earlier FSP reports led by Dr. Jill Dahlburg [J. Dahlburg, et al., J. Fusion Energy, **20**(4), 135-196 (2001)] and by Dr. Douglass Post [D. Post, et al., J. Fusion Energy, **23**(1), 1-26 (2004)].

- Sections (II) through (VI): These portions of the Report contain the Subcommittee's specific responses to the five questions posed in the letter of charge from Dr Orbach regarding a critical review of the Workshop Report [Reference: http://www.lehigh.edu/~infusion/FSP_report.pdf] resulting from the FSP Workshop (May 16-17, 2007) – led by Professor Arnold Kritz of Lehigh University and Professor David Keyes of Columbia University.

- Section (VII): This portion of the Report contains the Subcommittee's summary assessment of the feasibility of the proposed FSP and recommendations for a course of action.

- Appendices
  1. Letter of Charge from Dr. Raymond Orbach
  2. Example of Work Breakdown Structure (WBS) from Combustion System Simulation
  3. Timeline of Activities for FESAC FSP Subcommittee

1.1 General Perspective

While many of the technologies used in ITER will be the same as those required in an actual demonstration power plant (DEMO), further science and technology is needed to achieve the 2500 MW of continuous power with a gain of 25 in a device of similar size and field. Accordingly, strong R & D programs are needed to harvest the scientific knowledge from ITER and leverage its results. As emphasized in the FSP Workshop Report, advanced computations in tandem with experiment and theory are essential in this mission – a point well-illustrated in the Report's Figure 1, which depicts the imperative for the FSP to leverage ongoing investments in OFES' base theory and experimental programs, in OASCR's computer science and applied math programs, and in the interdisciplinary SciDAC Program. The associated research demands the accelerated development of computational tools and techniques that aid the acquisition of the scientific understanding needed to develop predictive models which can prove superior to extrapolations of experimental results. This is made possible by access to leadership class computing resources which allow simulations of increasingly complex phenomena with greater physics fidelity. With ITER and leadership class computing being two of the most prominent missions of the DOE Office of Science, whole device integrated modeling, which can achieve the highest possible physics fidelity, is a most worthy exascale-relevant project for producing a world-leading realistic predictive capability for fusion. This should prove to be of major benefit to U.S. strategic considerations for Energy, Ecological Sustainability, and Global Security.

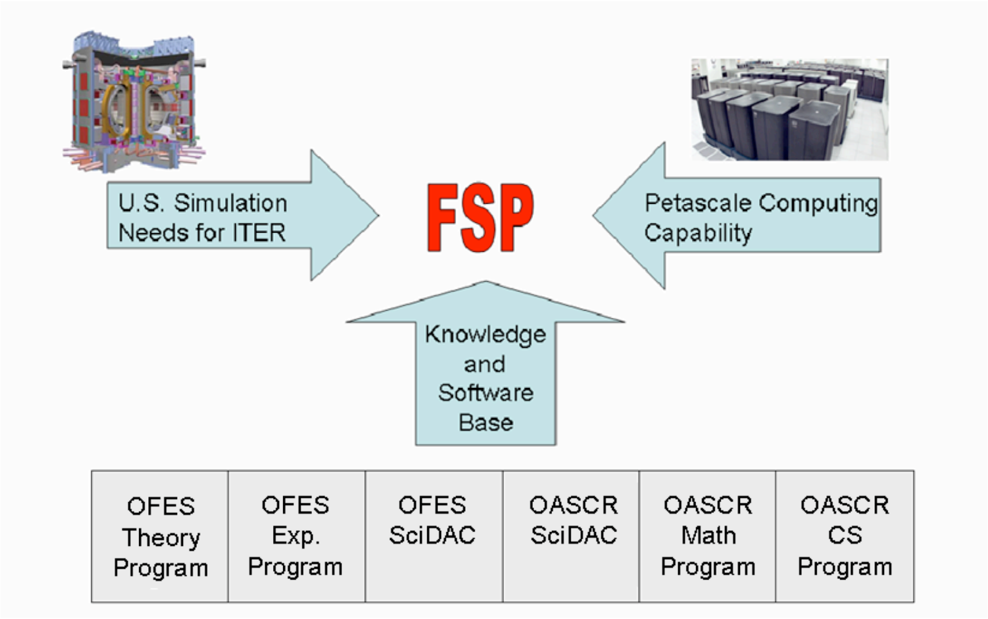

Figure 1 from FSP Workshop Report illustrating expected strong connections to DOE's highest priority missions involving ITER and Leadership Computing as well as to OFES theory and experiment, OASCR CS and Applied Math, and to the SciDAC Program.

The FSP under consideration is being charged with the primary objective of producing a world-leading predictive integrated plasma simulation capability that is vitally important to U.S investments in ITER and is also relevant to major current and planned toroidal fusion devices. This will involve the development over the next decade of advanced software designed to use leadership class computers (at the petascale and beyond) for carrying out multi-scale physics simulations to provide scientific insights vital to improved understanding of magnetically-confined burning plasmas. This activity should result in the delivery of a realistic integrated fusion simulation modeling tool with unprecedented scientific predictive capability. Modules with much improved physics fidelity will enable integrated modeling of fusion plasmas in which the simultaneous interactions of multiple physical processes are treated in a self-consistent manner. The associated comprehensive modeling capability must be developed in close collaboration with experimental researchers and validated against experimental data from tokamaks around the world. Since each long-pulse shot in ITER is expected to cost over $1M, this promises to be a most valuable new tool for discharge scenario modeling and for the design of control techniques under burning plasma conditions.

Some specific examples of expected advances which are needed to enable a comprehensive integrated modeling capability include:

- The effective coupling of state-of-the-art codes for the plasma core and the plasma edge region.
- The effective coupling of state-of-the-art codes for MHD dynamics and auxiliary heating of the plasma via RF waves.
- The development of more realistic reduced models based on results obtained from the DNS-type (direct numerical simulation) major codes which use petascale capabilities.
- The development of advanced frameworks and workflow management methods needed for code coupling.
- The development of an appropriate verification and validation effort to ensure reliable predictive capability.

1.2 Situation Analysis

The FSP Workshop Report was a compelling testimonial to the excellent collaborative relationship that currently exists between fusion energy scientists supported by OFES and the computer science/applied math scientists supported by OASCR. Much of the admirable depth of such alliances is due to DOE's truly interdisciplinary Scientific Discovery through Advanced Computing (SciDAC) program which has now been in place for over 6 years. The FSP builds upon and updates not only the original ISOFS/FSP Dahlburg Report in 2001 (strongly endorsed by FESAC at that time) and the D. Post FSP Steering Committee Report in 2004, but also has benefited from direct input from US spokesmen from ITER (N. Sauthoff and W. Houlberg) and from the Burning Plasma Organization (D. Batchelor). Consistent with the recommended levels in the Dahlburg Report, the targeted budget for the proposed FSP is around \$25M per year with a 15 year timeline. As noted in the earlier reports, this is in line, for example, with the \$25M per year allocated to just the University Alliances portion of the ASCI Program over the past decade. The development of the *advanced physics modules* targeted by the FSP is expected to take advantage of the ongoing OFES investments in basic theory, the ongoing SciDAC program, the SciDAC proto-FSP integration projects (including SWIM, CPES, and FACETS), and new developments involving joint experiment-theory-modelling efforts to predict and improve tokamak performance as time progresses. It will also be complemented by the expertise residing in OASCR's computer science and applied math programs together with access to leadership class computing resources for both "capability and capacity" computing applications.

1.3 Preview of Recommendations

As a preview of the final Section (VII) of our Report, the following bullets provide a representative picture of the FESAC FSP Subcommittee's summary assessment of the feasibility of the proposed FSP and recommendations for a course of action:

- While it was felt that the FSP Workshop document came across as too generic and "all inclusive," the FESAC Subcommittee believes that it contains sufficient information for making the case that the FSP can succeed in answering questions in a timely way that experiment and traditional theory by themselves cannot.

- In order to be successful, the FSP should not be "everything to everyone." It must be focused and project-driven with well-identified deliverables that the stakeholders fully support.

- The FESAC FSP Subcommittee agrees with the five critical scientific issues identified in the Workshop Report as important areas of focus appropriate for the FSP. However, an integration effort encompassing all five of these challenging issues from the beginning looks to be too large a step. To be practically achievable, the FSP should begin with more modest integration efforts that exhibit a compelling level of verification and validation. This recommendation is in line with a similar position taken in the original FSP Report from Dahlburg, et al.

- The FSP should be a repository of the latest physics as it evolves. In this sense it cannot be a "stand-alone" project. It must be properly coordinated with theory, experiment and fundamental simulation. More specifically, a proper implementation of the FSP will demand an effective plan for developing "advanced scientific modules" via utilization of results from the OFES base theory program, the SciDAC FES program, new insights from joint experiment-theory-modelling efforts, and the expertise residing in OASCR's computer science and applied math programs.

- The FSP cannot succeed without a viable validation and verification effort, and this will imply expanding the diagnostic effort and linking it better to the FSP, for example through an increased synthetic diagnostic development effort. This will require special personnel with an appreciation of both diagnostic methods and code expertise.

- The management of the FSP should be organized with clear accountability and oversight and work out a clear and compelling work-breakdown-structure (WBS). It should also seek advice and guidance from a broad community of stakeholders, experimentalists, analytic theorists, fusion engineering scientists, applied mathematicians and computer scientists.

- The FSP should establish and maintain strong connections with relevant international projects and also draw on the large experience base from existing scientific software development projects from other fields.

- The DOE should properly launch a true FSP only if a sufficient critical funding level can be realistically met and sustained.

II. FESAC FSP Charge Question 1:

“Has the report identified key scientific issues and grand challenges that can be addressed by this approach to linking the scientific knowledge base for fusion energy?”

The “Scientific Issues” section in Chapter 2 of the FSP Workshop Report was both interesting and informative and did a reasonably good job of handling Question 1 of the Charge. We are accordingly answering this first charge question with a conditional “YES.” The report identifies five critical science issues: 1) Disruption effects, including avoidance and mitigation; 2) Pedestal formation and transient divertor heat loads; 3) Tritium migration and impurity transport; 4) Performance optimization and scenario modeling; and 5) Plasma feedback control. These issues were identified as the most urgent for the burning plasma program and for the successful operation of the ITER experiment. While undoubtedly a longer list of scientific questions could be generated (cf. Table 2.1 of the Workshop Report), we found these to be both important and compelling. Independent confirmation of their importance and relevance comes from the fact (noted in the workshop report) that the European fusion simulation effort is organized around precisely these five areas [see, for example, International Atomic Energy (IAEA) Conference paper TH/P2-22, “Integrated Tokamak Modelling: The Way Towards Fusion Simulators,” A. Bécoulet, et al. (2006)]. Each of the five questions is a computational grand challenge in its own right, and requires an integrated simulation capability. This is made clear by Table 2.1 in the Workshop Report, which is essentially a matrix of how the properties of tokamak plasmas depend on a variety of physical processes. The fact that the matrix is nearly full tells you that everything depends on everything else, and powerfully makes the case that an integrated approach is required. We note that a more helpful picture of what is actually included in item (4) above (the key topic of performance optimization and scenario modeling) is depicted below in Fig.1.1 (Fig. 2.2 in the original workshop report). This figure captures a vast amount of physics knowledge, with each topic in its own right requiring detailed physics understanding.

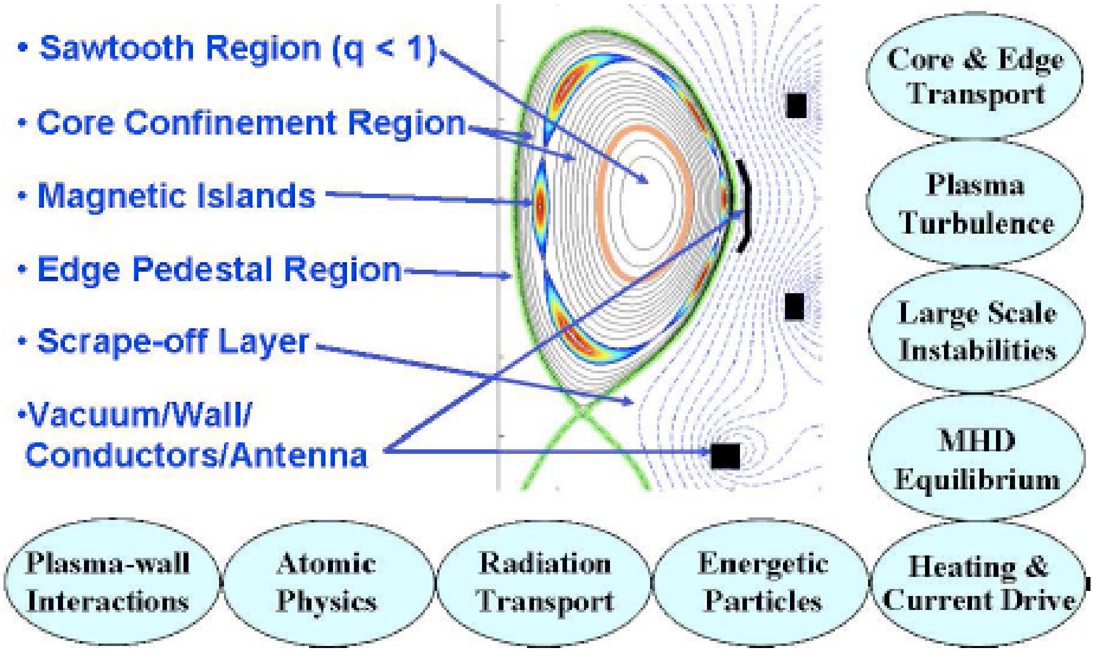

Fig 1.1. Illustration of the interacting physical processes within a tokamak discharge.

The five issues are split into three topics (1-3) focusing on improved scientific understanding of physical processes, and two (4, 5) focusing on new tools for operational control of ITER experiments. In reality, item (4) overlaps scientific understanding and operational control. Although this is not clearly brought out, it would seem that the science challenges require integrated simulations of a different character than the operational challenges. The former requires the integration of a few "first principles solvers" of high dimensionality and physics fidelity, while the latter requires a larger number of reduced dimensionality models. Whether these are all part of a larger, single integrated code was not clearly explained.

In addition to the five grand challenges, four physics modeling components needed to address these challenges were discussed in some detail: 1) core and edge turbulence transport; 2) large scale (MHD) instabilities; 3) sources and sinks of heat, momentum, current and particles; and 4) energetic particle effects. The report states that these are "examples of physical processes that affect the optimization of burning plasma performance", and goes on to list several others. It is not stated how these four were chosen, or whether they are the four most important components. By calling them components, the strong implication is that they would be four key components in the integrated simulation capability required to address the five science issues.

We were generally satisfied with the three-question format used in both the five science question sections 2.1.1-2.1.5 and the four physics components sections 2.2.1-2.2.4. For non-specialists, this provided a good sense of the scientific importance, current state-of-the-art, and new capabilities needed in each area. Some sections were more complete and comprehensible than others. However, we believe that the report writers missed an opportunity to ask and address the following key readiness questions:

1. Which emerging or maturing simulation approaches appear most promising in the next 5 years that are now or soon will be ready for integration?
2. How close are we to begin implementing the integrated model shown in Fig. 1.1?
3. Do we already have part of it completed?

Figure 1.2 (Fig. 2.3 in the original workshop report) was helpful in placing the integrated system simulation capability within the operational context of the fusion experiment.

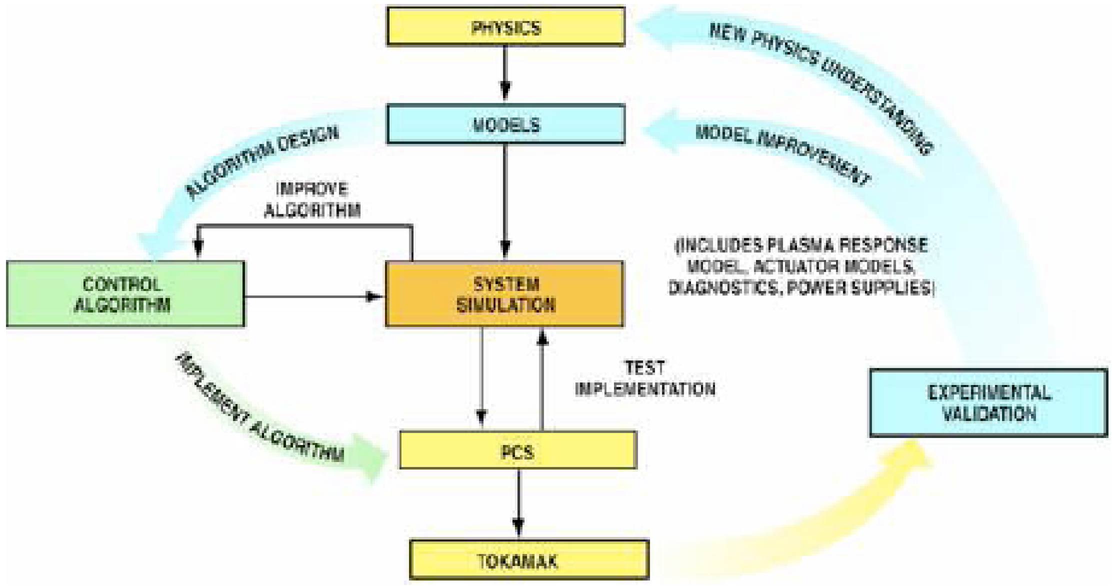

Fig. 1.2 Schematic of integrated plasma control processes which is used by ITER (PCS stands for a real-time Plasma Control System).

Given the real-time needs of the integrated plasma control process, the orange box labeled "system simulation" could not be the first principles whole device simulation model diagrammed in Fig. 1.1 which would take days to weeks to run, but rather some reduced "control-level" model that could execute quickly. The current state-of-the-art for plasma feedback control is described in Sec. 2.1.5 of the original workshop report; the software tools are limited in their integration with other physical effects, and their implementations characterized as not extensible. It would seem that a FSP is most critically needed and would be most beneficial in this area. This raises the question of priorities. Of all the needs expressed in Chapter 2, which are the most critical and which would most benefit from a FSP? It is apparent that there are at least two integrated simulation software projects required; one implementing some version of Fig. 1.1 using a combination of first principles kinetic and fluid solvers for attacking science issues 1-3, and one implementing Fig. 1.2 using reduced models. The report alludes to this strategy in Chapter 2 but states it most succinctly on page 67, Sec. 4.7. If this is the strategy to be taken, which appears to be sensible, the following questions then arise:

(1) How will these two efforts be connected?
(2) How will insights from the first principles approach be folded into the control-level models?
(3) How will the proposed validation activities connect to these two efforts?

In addressing these issues, we make the following observations.

The four topics of core and edge turbulence, large scale MHD instabilities, sources and sinks of heat, etc, and energetic particle effects are the traditional fusion topical science areas encompassing everything within the plasma boundary. The FES community has spent years studying them and will continue to do so. They are in a sense too high level (and broad) to actually be components of the FSP. Fortunately, we do not need to understand everything in these four areas to proceed with integrated simulation. Focusing on the five critical science issues is the right strategy if the FSP is to deliver a tool that would be useful for fusion research in the not too distant future.

Of the five critical areas, disruption effects, including avoidance and mitigation, is driven dominantly by MHD physics with appropriate kinetic modifications. We know, when the plasma pressure exceeds linear MHD stability limits, the plasma is usually observed to disrupt. One question is how close to the stability limit we can operate and still have confidence the plasma will not disrupt. Associated key issues include the identification of precursor signals for disruptions and how to mitigate the impact of disruption by minimizing the heat load in localized spots. While we do not currently know the answers to these questions, progress in nonlinear MHD simulations (that include more realistic physics) together with systematic validation against experimental results can be expected to get us there in reasonable time.

Pedestal formation and transient divertor heat load impact both the fusion core performance and requirements for materials facing the plasma. Unfortunately, we know very little about this topic. Theoretically, it is very challenging because: (i) it is a multi-scale regime; (ii) there is no obvious separation of scale for MHD and transport here; (iii) atomic physics is important; and (iv) the geometry is truly 3-dimensional. While simulation of the edge is just starting, there has been years of experience simulating core turbulence and MHD. Nevertheless, information on the edge is needed in order to reliably simulate the performance of the whole device. The most promising approach would be to develop a reduced model based on current knowledge with the understanding that the model will improve as we learn more.

The report lists "tritium migration and impurity transport" as a high-level, key scientific issue. Regarding tritium migration (and the relevant high association with use of carbon divertor material in ITER), while we agree that this is an important matter, we recommend that FSP include a much more general focus on the plasma surface interaction challenge, including issues that are equally critical as "tritium migration" if not more so. These plasma surface issues include such topics as sputtering erosion/re-deposition (and sputtered impurity transport/plasma-contamination), erosion-dominated component lifetime, dust formation, flaking, as well as tritium migration and *trapping* via co-deposition in beryllium, and tritium *trapping* in bulk tungsten.

Performance optimization and scenario modeling capabilities will require integration of state-of-the-art physics from the four topical science areas. Since some areas are more mature than other areas, it is unlikely that we can integrate first-principles codes in a practical way in the foreseeable future. While there are many reduced models in existence, they have not for the most part been rigorously validated against experiments. Selecting the best models and integrating them into a simulation capability will allow us to learn about the nonlinear coupling between the intertwined physics. Using the resultant capability to validate against experiments will also be a beneficial learning process.

The focus of plasma feedback control for ITER is start-up, shape control, and disruption avoidance. These applications require mainly integrating MHD physics, and is something that we can make good progress on. Other control capabilities, such as profile control to get to advanced performance regimes, will require knowledge of particle and momentum transport – about which there is currently very little scientific understanding. Controlling the edge localized modes (ELM's) will require a good edge pedestal model which does not currently exist.

III. FESAC FSP Charge Question 2:

"Have all the critical technical challenges been identified for which predictive integrated simulation modeling has a unique potential for providing answers in a timely fashion, in a way that traditional theory or experiment by themselves cannot?"

The Fusion Simulation Project report identifies five critical scientific issues:

1. Disruption effects, including avoidance and mitigation
2. Pedestal formation and transient divertor heat loads
3. Tritium migration and impurity transport
4. Performance optimization and scenario modeling
5. Plasma feedback control

These are to be addressed with physics models in four areas

1. Core and edge turbulence and transport
2. Large-scale instabilities
3. Sources and sinks of heat, momentum, current, and particles
4. Energetic particle effects

Are these the critical technical challenges?

The four topical science areas are very general, and arguably could be taken to include all areas relevant to a tokamak reactor for which our understanding can be advanced through integrated simulation. So the short answer is "yes." However, as suggested in the earlier Dahlberg report, attempting to simultaneously integrate all of these areas from the start would be an impossible challenge. A more realistic approach would be to choose a subset of these issues to start with so as to ensure there will be useful deliverables in the not too distant future, and integrate additional issues as warranted by progress and experience. An example would be to start the scenario modeling task by coupling transport and MHD stability. This would have the most immediate impact, and other elements could be added later.

Although the above set of critical scientific issues is well-recognized as important issues for ITER and fusion power plants, it would strengthen the FSP initiative by calling out specific examples of their importance.

- Non-disruptive instabilities: The FSP report correctly identifies disruptive instabilities as a critical issue due to their potential to damage the tokamak. However, there are many MHD instabilities that can cause the plasma to rapidly lose much of its stored energy, often coincident with a rearrangement of the current profile. Although there is little or no danger of machine damage, the plasma control system will likely not be able to return the plasma to the desired operating scenario following such an event.
- Edge Localized Modes (ELM) mitigation schemes: This falls under transient divertor heat load issues and is a serious concern for ITER because of the rapid erosion of plasma facing materials. Several schemes are currently under study experimentally to mitigate or eliminate ELMs. They include development of plasma operating scenarios without ELMs, and the use of external magnetic field perturbations to minimize ELMs. Our physics understanding of these is incomplete, and evolving. Transforming what we learn from the experiments into a reliable, predictive model would thus be of very high priority.
- "Whole-device" integrated modeling really needs to include the entire discharge duration. The plasma can be particularly delicate during the formation and shutdown periods as parameters change quickly and sudden changes in electromagnetic conditions may cause significant damages to the superconducting magnets. Also, careful control of the formation phase may be critical to obtaining the desired operational scenario.

Does predictive integrated simulation modeling have a unique potential for providing answers in a timely fashion, in a way that traditional theory or experiment by themselves cannot?

The goal of much of the current fusion energy science research in the United States is to reach a level of scientific understanding of a burning tokamak plasma to allow for accurate prediction. Experiment and traditional theory are essential components of this. An experiment encompasses all the realistic physics but is limited in its scalability by the hardware. Traditional theory makes simplifications to first-principles equations to enable analytical solutions in special limits. Simulation ideally bridges the gap of experiment and traditional theory by taking advantage of state-of-the-art development in applied mathematics, computer science, and high performance computers. While that ideal goal is still far away, simulation has already shortened the learning curve for developing more complete physics models that are closer to the first–principles equations. An example is the recent development of physics-based transport models that nicely reproduce the predictions of microturbulence simulations, without adjustable normalization parameters.

The goal of FSP is to integrate the most complete physics models available, beginning with binary integration, and extending to higher dimensional integrations, providing a progressively more realistic simulation of an experiment. It should be aware that many physics models are still rapidly evolving, so FSP should provide a framework for incorporating new physics as they evolve. Not all scientific areas are equally developed. By judiciously choosing its integration strategy working in concert with theory and experiment, FSP should be able to produce answers in a timely way that cannot be obtained in its absence.

There are many examples of smaller efforts aimed at integrating various pieces. There are several transport codes that integrate sources and transport. There have been some efforts aimed at combining MHD stability and transport. All of these efforts are characterized by (1) rapidly evolving physics understanding, necessitating constant evolution of the physics models, and (2) incompatibility (including both coding structure and physics components; i.e., dimensionality, scales, etc.) of the different pieces requiring significant effort to join them. A lesson learned is that for the FSP to succeed, a partnership of all the stakeholders including fusion physicists, applied mathematicians and computer scientists, with fully supported computing resources and stable critical funding are essential.

The key contribution of the FSP could be in addressing these two concerns, and less so in the development of the physics models themselves, since that is the goal of the DOE fusion science program. As the FSP moves forward, equally significant advances will take place in experiments, especially as diagnostics become more and more sophisticated. Traditional theory and fundamental simulations will also advance. The FSP could and should provide a living framework for physics as it advances, making it as straightforward as possible to keep the models up-to-date. At the same time, standard interfaces could minimize the difficulty of different modules (models) communicating with each other, thereby easing the integration process.

What should be the focus on plasma feedback control in FSP?

The control community is already heavily engaged in developing a Plasma Control System for ITER. Responsibility for this effort has mainly been assigned to the Europeans, but there is significant US participation. The FSP report appears to imply that it could provide the control system for ITER. While this is a laudable goal, timeliness and international politics might well stand in its way. Nevertheless, FSP does have much to contribute to plasma control.

Any usefulness of the FSP in plasma control will require flexibility on the FSP development side. Since the ITER PCS team might have to follow a more rigid schedule to deliver a working system, the FSP effort should be capable of interfacing with what is already being developed and be able to add value rather than attempt to compete.

It is also important to note that in a regime where any external heating tools are overwhelmed by alpha power, large scale instabilities will not, as the FSP Workshop document claims, be controlled by the use of modulated heating. As is already recognized by the control community, this is a major challenge for plasma control. In general, the capability to explore effective ways to control a burning plasma is a desirable product of the FSP.

IV. FESAC FSP Charge Question 3:

"Is there a clear plan to establish the fidelity of the advanced physics modules, including a sound plan for validation and verification?"

Establishing the physics fidelity of advanced physics modules in the FSP is definitely recognized in the Workshop Report, and the associated essential role of verification & validation (V&V) in enabling a successful FSP is clearly emphasized ("*A verification and* *validation (V&V) program is essential to the role envisioned for FSP")*. For the purposes of this discussion, we assume that the *advanced physics modules* include the outcome of ongoing basic theory, the ongoing SciDAC program, the SWIM projects, and new developments to predict and improve tokamak performance as time progresses. While the FSP provides an appropriate project framework/mechanism to move in the right direction, as it stands now it does not provide a clear path toward establishing the fidelity of advanced physics modules, neither does it offer a plan for validation or verification (V&V) that would rely on ongoing theory and experimental research. In fact, part of the vagueness of what is being proposed stems from a diversity of opinions of the authors of the Workshop Report of what the FSP should be. Some authors referred to it as a research program, while others referred to it as an integrated computer program. Still others referred to it as a software framework which would allow the integration of simulation components, while still others referred to it as a software tool suite (in the sense of a workbench) with improved user and data interfaces. Clearly, given the gap between the current state-of-the-art and the needs expressed, the FSP project needs all of these items in a properly phased and coordinated way. Some of these elements will be supplied by activities funded in the base program and through the SciDAC program. As emphasized earlier in this FESAC Subcommittee Report, the FSP is expected to interface with these activities before it can effectively move forward as a unique program within the US fusion program. One of the challenges of the FSP project will be how to deal with the potential problem of different algorithms being used by the various SciDAC projects which may not be easily combined.

Verification.

The report states ***"****Verification assesses the degree to which a code correctly implements the chosen physical model",* and we agree with this. However, the sentence ends with the statement *"and is essentially a mathematical problem"* and we cannot agree with this simplistic conclusion. The authors then proceed to describe applied mathematics and computer science "technicalities" as to how verification is to be carried out. At this level, connection to the real physical world is lost to a large extent and verification becomes an exercise in the execution of algorithms, issues involving numerical approximations, mesh discretization, temporal discretization, iterative solution of nonlinear equations, and statistical sampling error issues as well as resource management issues, just to mention a few examples.

While this may be acceptable from a mathematics perspective, in our opinion **verification** must also emphasize a comparison against theoretical physics predictions in a more intimate manner than described in the document. For example, codes are often developed to study highly nonlinear or even turbulent stages of plasma instabilities. The linear and weakly nonlinear phases of such instabilities are usually well described by analytic or semi-analytic theories that can be used to *verify* the code accuracy in the same regimes. Benchmarking codes with theoretical predictions is an essential tool to *verify* the code convergence in the limits where the theories are valid. In cases when numerical difficulties prevent the solution of the most comprehensive mathematical models, alternate theoretical approximations may be necessary, or a reduced set of equations may have to be formulated to expedite numerical solutions. Other examples include replacement of kinetic equations with fluid equations, or reductions of the range of finite Larmor radius corrections, etc. Of course, the consequence will be limitations in the predictive capability of codes to predict actual experimental situations. Another fundamental issue of code verification concerns the limitations on physical parameters set by the numerical solvability of the mathematical model. For example, computational considerations often limit the magnitude of the magnetic Reynolds number (MRN) used in resistive magnetohydrodynamic simulations. Typical values of MRN used in the simulations are orders of magnitude less than the actual values in relevant fusion plasmas. Thus, it becomes crucially important to develop a strategy for assessing the effects of such limitations on the validity and relevance of the numerical results.

Though theoretical predictions can help in verifying code accuracy, analytic or semi-analytic solutions are often not available to verify codes in highly nonlinear and turbulent regimes. In such cases, a "cross-code" verification can be a satisfactory approach. Cross-code verification requires the systematic comparison of multiple codes which use different numerical algorithms (finite difference, finite elements, spectral methods, implicit, explicit …..) and/or different mathematical models (Vlasov, Particle-in-Cell, Hybrid PIC-Fluid, Fluid…..). Such different codes can be used to benchmark each other in those regimes where analytic solutions are not available.

In summary, we recommend that special emphasis be placed on code verification through cross-code benchmarking and comparisons with theoretical predictions.

Validation is correctly defined by the authors as a highly desirable goal, and it "*assesses the degree to which a code describes the real world.*" However, the goals are not developed into actual action items and the text becomes void of specific recommendations of what needs to be validated and how it would be done. Clearly, simulations can never perfectly model physical reality, but can nevertheless be superior to empirical extrapolation if they can demonstrate a reasonable level of agreement with reliable results from systematic experimental measurements. In particular, documentation from such validation tests should include data from both the simulation and experiment along with descriptions of data reduction techniques and error analysis, etc. Regarding the question of validation, the description in Chapter 3 is also lacking somewhat in specifics. In particular, validation needs to be more than just carrying out experimental tests of the model. Writing code with today's version of the model is only the first step (note that this implies that the coding needs to be flexible enough to accept frequent updating of the physics models). The improved physics models will come in periodically as major advances in fusion theory occur, for example, in a better understanding of the edge pedestal physics by combining kinetic and fluid theoretical models; in RF theory by combining Fokker Planck codes and full wave models with finite orbit physics of energetic particles; in transport codes by including self consistent radial electric fields originating from neo-classical theory on transport time scales; in developing a theoretical understanding of the density limit in tokamaks; in a better understanding of magnetic reconnection processes in sawtooth and perhaps ELM triggering mechanisms; in better understanding of alpha particle physics and their impact on MHD stability, and so forth. Any one of these processes involves complex nonlinear dynamics – the improved understanding of which has to be developed with a combination of experiment, theory, and modeling. Figures 3.1 and 3.2 below show examples respectively of a possible approach to development and validation of transport codes (Fig.3.1) as well as to RF codes (Fig. 3.2) -- both of which are currently being developed under the SciDAC FES program. The benefit of such approaches is that the accuracy and value of both the models and experiments improve together over time and are integrated with ongoing experiments and theory. Once understood, the model will be ready to be assimilated into the FSP project which then integrates and tests multiple physical packages into a "whole" system.

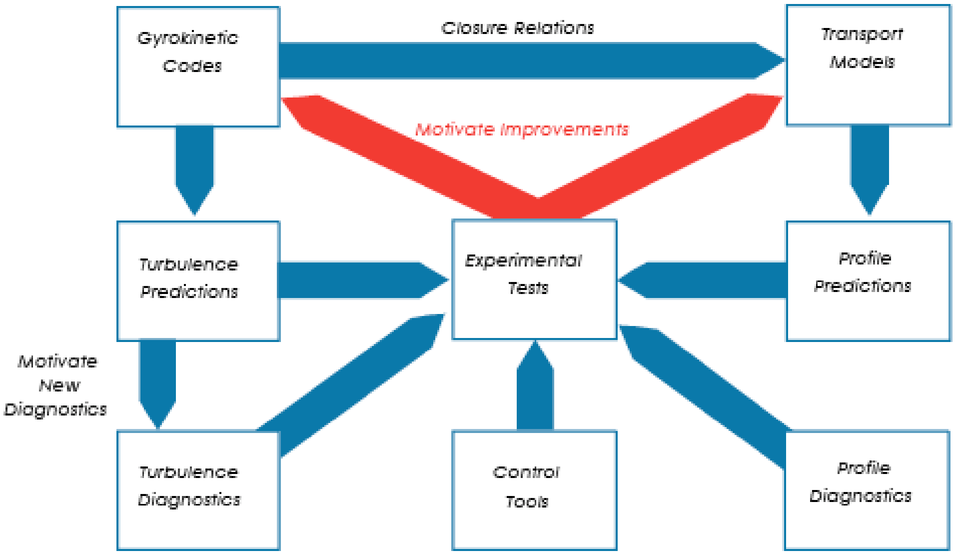

Figure 3.1. A fully predictive capability is best approached with the combined efforts of modeling and experiment. The above example refers to the development of a predictive understanding of transport in the plasma core.

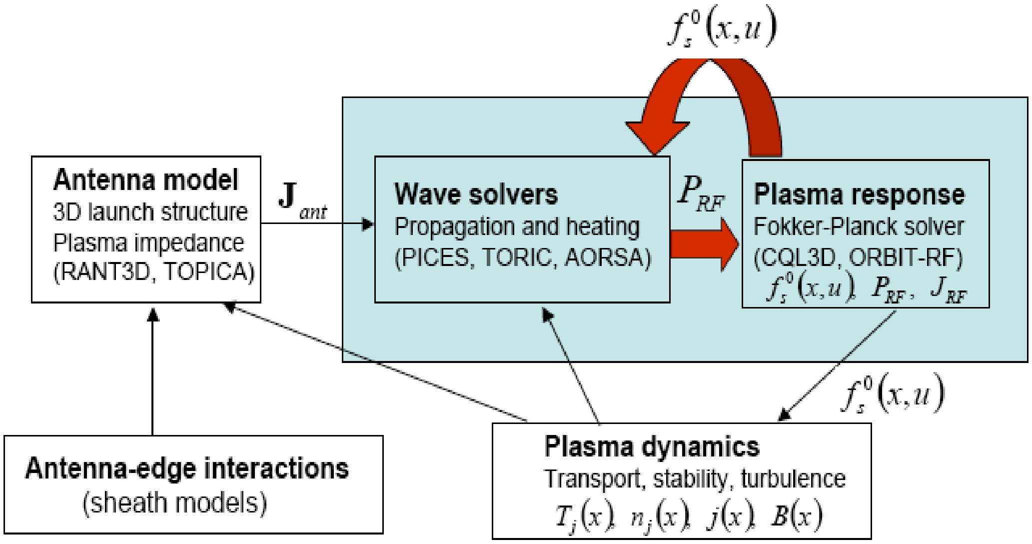

Fig. 3.2. A fully predictive capability in the RF wave-particle interaction area is indicated in the above diagram.

We emphasize that a reliable and predictive FSP code can be built only if the individual building blocks have a solid theoretical foundation and careful experimental validation. Since experiments rely on diagnostics, we must build up our competence in 2 and 3 D diagnostic capabilities, and furthermore, each diagnostic will have to be implemented into the codes in the form of a "synthetic" diagnostic. This is necessary since almost none of the diagnostics provide local data accurate enough to make precise predictions for testing codes. Thus, increased funding for novel diagnostic development and certain aspects of the base theory program will have to be supported in a more aggressive manner to provide adequate support to the FSP.

In summary, our primary observations/recommendations are:

1. Code validation is successful only to the degree to which a code describes the real world. Therefore the code needs to be flexible to accept advances made as a result of comparisons between experiment and new theoretical developments.
2. Validation and model development should not be regarded as separate activities.
3. The value of synthetic diagnostics cannot be overlooked. This is often done as an afterthought, but is the most powerful tool for comparison between model and experiment. This area of research is in its infancy and it has already produced some spectacular results in a few limited cases such as RF physics.

4. Current devices are capable of performing experiments in all five of the critical scientific areas identified earlier. However, a properly designed validation exercise may require resources not already available, such as diagnostics and/or control actuators. The FSP should advocate development of the experimental tools needed to validate the models. In addition, further progress is required in

several areas of fusion theory that are presently not well supported. Limitations in fundamental theoretical understanding will seriously limit the usefulness of code development.

V. FESAC FSP Charge Question 4:

"***Does the FSP Workshop Report clearly identify the critical areas of computational science and infrastructure in which investments would likely produce the tools required for the FSP to achieve its goals?"***

Computational science issues including enabling mathematical techniques and infrastructure are covered in Chapters 3-6 of the FSP Workshop report. The report discusses applied mathematics/numerical methods generally and with application to plasma simulation. Such key areas as resistive MHD simulation, fluid plasma calculations, and turbulence simulation are discussed. Mathematical issues covered include adaptive meshes, space-discretization, sparse matrix solvers, and bifurcation analysis.

The report recognizes the need for specific development and application of computational techniques to several unique features of ITER and plasma simulations in general. These include the visualization of edge localized mode (ELM) plasma events, pattern recognition of plasma structures such as the magnetic separatrix, magnetic island detection/visualization, turbulent transport structures (blobs, etc.), and the like.

In general terms, the FSP Workshop report does a reasonable job in describing the computational science methodologies needed to produce the tools required for the FSP to achieve its goals. However, there is a lack of specificity on what the software deliverables would be (even the number and kind) and how they might best be applied to FSP codes. Specificity with respect to bibliographic or other references to actual code names and algorithms relevant to the FSP is actually provided in the earlier Dahlburg and Post reports – not so much in the present FSP Workshop Report. With regard to the overarching technical challenge of dealing with the extreme multi-scale nature of simulating tokamak plasmas for long integration times, the Workshop Report's description of how this will actually be done is also lacking in specifics. On the other hand, the Workshop Report does concentrate on the key petascale parallel computing issues. While clearly petascale computers will be required, it should nevertheless be noted that the algorithmic challenges are not all parallel in nature. It is suggested that future work provide a more substantive description of the vision for how the formidable macro/micro coupling challenge can be achieved even if provided "infinitely powerful" computing power in the future.

In areas such as data management, visualization, and analysis the arguments are, likewise, less specific to fusion and more based on the scale of the simulation output datasets. While there is likely to be significant overlap in requirements for data analysis, visualization and data management with other disciplines, it is also likely that the FSP will need specific tools and new algorithms specifically devised for the characteristics of the fusion problem.

Furthermore, it is also likely that the need for simulations to be performed in a time critical fashion (for the interpretation of shot data or for experimental planning against a timeline) will drive some new requirements for deadline-driven data assimilation methods, and other quasi real-time methods that will offer new challenges to both systems architecture and operational infrastructure. The precise nature of the deadline driven requirements are not yet fully articulated and these should feature in future requirements planning.

It is recommended that one of the first tasks of the FSP should be to conduct a requirements and risk analysis associated with the computational tools and infrastructure to determine the appropriate level of direct investment and the expected increase in capability due to normal developments in the field (which are expected to be considerable).

It is further recommended that the FSP engage in SciDAC-like joint partnerships to develop the FSP specific capabilities for computational tools and infrastructure. There do exist successful examples of joint work funded by ASCR and FES in the areas of mathematical techniques, computational libraries, collaboration technology, data analysis and advanced visualization tools. These efforts should be encouraged, resourced, and the support levels tightly coupled to the science and engineering goals of the FSP.

Finally, it is recommended that the computational and software infrastructural requirements for the FSP be communicated early and often to those organizations providing computational and data capabilities for the Office of Science, such as the Leadership Computing Centers, ESnet and NERSC.

Comments on Multi-core Challenge:

A dominant trend emerging in the ultrascale hardware development area is to continue to add more and more multiple CPU cores onto the same chip to deliver high aggregate computing performance. Current estimates are that the number of cores per chip is expected to increase by an order of magnitude in five years and two orders of magnitude in a decade. Together with the high-density, low-power packaging approach to construct large-scale parallel computers, this trend indicates that a petascale or multi-petascale parallel machine in the next decade could reach as many as 10 million CPU cores. The formidable challenge here is of course to develop new methods to effectively utilize such dramatically increased parallel computing power. This will be necessary to achieve accelerated scientific discovery in fusion energy science as well as many other application domains.

Some examples of outstanding challenges in the fusion energy science application area are:

- The efficient scaling of MHD codes beyond terascale levels to enable higher resolution simulations with associated greater physics fidelity.
- The efficient extension of global PIC codes into fully electromagnetic regimes to capture the fine-scale dynamics relevant not only to transport but also to help verify the physics fidelity of MHD codes in the long-mean-free-path regimes appropriate for fusion reactors.
- The mastery of Data Management to help with the development (including debugging) of advanced integrated codes.
- The development of innovative data analysis and visualization to deal with increasingly huge amounts of data generated in simulations at the petascale and beyond.

Required Investment: The FSP will require on the order of $25M/year over the course of the next 15 years and more. In addition, research enabled by ultrascale compute power will also demand much greater computer time. For example, a single global particle-in-cell code [developed within the SciDAC Gyrokinetic Particle Simulation Center (GPSC)] run carried out at present to investigate the long-time evolution of turbulent transport requires around 100K cores * 240 hours = 24M CPU hours. Since the current version of a plasma edge code [developed within the SciDAC Center for Plasma Edge Simulations (CPES)] requires roughly the same amount of time, the actual coupled simulations of the core and edge regions noted earlier could demand approximately 50M CPU hours. If additional dynamics (such as the modeling of the RF auxiliary heating) were also included, then the need for computational resources at the exascale would be a reasonable expectation.

Major Risks: If dedicated investments are not made, the ability of the leading fusion codes (e.g., the MHD codes with their scaling challenges) would run the risk of not being able to effectively utilize the large number of processors at the exascale. It should also be noted that the development of effective mathematical algorithms for integration/coupling is very difficult and could well be difficult to achieve within the next decade. Finally, if the fusion energy science applications were only able to effectively utilize a small fraction of the cores on a CPU, major efforts would be needed to develop innovative new methods for per processor performance.

VI. FESAC FSP Charge Question 5:

"Have the issues associated with project structure and management of the proposed FSP been properly addressed?"

The FSP Workshop Report provided a list of 13 management issues (Section 6.1) along with a "sample FSP structure" (Section 6.2). While the list is useful and reasonably comprehensive, the "sample structure" was not adequately compelling or clear. In order to more effectively address the project structure and management of such a challenging project, we believe that it is important to first identify the ***most*** critically important management issues and then develop and apply key guiding principles to address and resolve these problems. The following are specific examples to illustrate our point.

1) Developing an integrated product when the scientific basis for some components is still rapidly evolving.

Guiding principles:
Best practices and interdisciplinary integration – The FSP will require that some of the most advanced methods of computer science and applied mathematics be effectively integrated with advances in computational and theoretical plasma physics. The project structure should ensure that national and international best practices for similar types of projects be incorporated in the management plans. The previously proposed DOE SC Scientific Simulation Initiative (SSI) for Combustion Systems Simulation is appended as an illustrative example (see Appendix 2).
Risk assessment and mitigation – Software projects are unique in the risks associated with them, to a degree that arguably exceeds even experimental device construction. A project of this magnitude and complexity needs to be able to quantify the risk associated with each key part of the software project and to have appropriate backup solutions and/or recovery methods identified.
Technical decision making – The FSP should have a detailed Work Breakdown Structure (WBS). The WBS organizes and structures the project by work elements at several, distinct levels. A successful WBS will identify all of the activities that need to be accomplished to achieve the project objectives. The proposal for each work element in the WBS should contain goals and technical justifications, cost and resources required, schedule and deliverables, and risk assessment. The WBS should be reviewed by an appropriate technical/management committee.

2) Managing a geographically diverse team, whose members have technically diverse interests (reward system issue).

Guiding principles:
Communication – It is expected that the FSP will be a large, multi-institutional, and geographically distributed project. Requirements, schedules, progress, and issues must be efficiently disseminated throughout the project. Difficulties encountered by sub-teams must be made known in a clear and timely manner in order to facilitate the development of solutions.
Motivation and Evaluation – The project management and structure will need to ensure that the project scientists and supporting staff members are highly motivated by recognition within the project, within their home institutions, within the scientific community at large, and by appropriate compensation. It is important to establish mechanisms for ensuring that accomplishments are appropriately rewarded. This should be accomplished by implementation of a performance-based management system which can efficiently evaluate and identify areas where productivity is problematic and then move forward in a timely way to get these elements of the project back on track.

3) Identifying deliverables (short and long-term) that have the potential to answer questions that are clearly beyond existing capability and focusing resources to maximize success.

Guiding principles:

Value-added – Since the FSP deliverables must be useful to the stakeholders, there should be clear paths for obtaining input, both solicited and unsolicited, from them. Stakeholders include the software users from the theoretical, modeling, and experimental communities. They also include those planning future experiments and, ultimately, future reactors. Mechanisms must be in place to effectively assess the usefulness of the FSP project, in whole and in part. Improved predictive capability of key phenomena supported by Verification and Validation is a major component here. At the highest level, the accelerated delivery of improved predictive capabilities enabled by truly interdisciplinary advances from FES together with ASC must be clearly evident to OFES, OASCR, DOE-SC Headquarters, and OMB.

Delivery and quality – The FSP should identify the mechanisms by which it will ensure that its deliverables are provided on time and that all quality standards are enforced. Quality standards include basic ones, such as portability across computational environments, as well as reproducibility and the ability to predict well-studied cases. An aspect of overall quality assurance falls under the Verification and Validation area.

Staffing and resource management – From the initiation through evaluation processes, the FSP will require the effective assignment of staff resources (such as access to computers and auxiliary staff). Of primary importance is the identification of the responsibility for making such decisions. The project should explicitly delineate the mechanisms for this key aspect of the management plan. In addition, if the project has a diffuse funding mechanism (such as a research grant), it will need a proper prescription for reassignment of tasks, in partnership with the Department of Energy.

4) Ensuring oversight and accountability when all the computational experts in the fusion community are likely participants in the project.

Guiding principles:

Accountability – The FSP will be a large software and scientific development project, as opposed to a conventional research project. Since there will be scheduled deliverables, the project structure needs to make clear who is ultimately responsible for project deliverables as a whole as well as for the individual parts of the project.

Expertise, advice, and oversight – Success of this project will rely on obtaining needed expertise from the communities of fusion science, applied mathematics, and computer science – all working together in a productive interdisciplinary sense. The project structure should include identified mechanisms, such as advisory committees and/or panels, by which the required advisory and oversight functions can be effectively served.

Conflict resolution – In the event of disputes, which might include decisions on technical approaches, task and resource assignments, differential recognition, and priorities, it is important that the management plan identify the responsible person and/or mechanism by which conflicts will be resolved.

5) Enfranchise the fusion community (i.e. theorists, experimentalists, technologists), as well as the applied mathematics and computer science community in this project.

Guiding principles:
Utility – In common with the "value-added" component of the guiding principles for identifying attractive deliverables (item #3 on this list), the scientific deliverables must be very useful – in an interdisciplinary sense -- to the stakeholders in the FES and ASC communities. Again, at the highest level, the accelerated delivery of improved predictive capabilities enabled by truly interdisciplinary advances from FES together with ASC must be clearly evident to OFES, OASCR, DOE-SC Headquarters, and OMB.
Mentoring and education – FSP will be a long term project -- one that is expected to last through the ITER period and beyond into DEMO phase, ultimately helping to deliver sustainable, environmentally friendly energy solutions to the U.S. and to the world. Its human resources will need to be replenished through highly competitive recruitment. Management will need to ensure that there exist mechanisms for educating and bringing into the project scientifically capable personnel from other fields, as well as establishing and encouraging liaisons with training and education institutions, especially universities.

Illustrative Organizational Chart

The sample organization chart (Figure 6.1) in the FSP workshop report, while containing the generic structure, does not fully delineate the functionality and responsibility of the organizational elements in addressing the five most challenging issues listed above. By reorganizing the reporting structure and clarifying the level of detail that should be in the technical task groups, the illustrative organization chart shown below (Fig. 5.1) would more effectively respond to the guiding principles put forth to address the five challenges.

We retain the model that there is a lead institution for the Fusion Simulation Project that is chosen by an open, competitive process. The lead institution will be responsible to DOE for meeting the project goals and milestones. A proposal from a prospective lead institution should identify a Project Director, who is the key interface with the funding agencies OFES and OASCR on an operational basis. The Project Director should assemble a strong management team in the form of a Management Coordinating Council that will coordinate and ensure the success of all facets of the project, including technical and managerial elements. The Project Director is a permanent member (possibly the chair) of the Management Coordinating Council. The other members should be chosen to reflect a broad representation from the institutions participating in the project. This is necessary for managing a geographically diverse team, to ensure fair and equitable recognition for contributions, and to reallocate resources when the need arises.

A Program Advisory Committee, consisting of senior managers with experience in managing large R&D projects, will advise the director of the lead institution on high-level programmatic issues such as the continuing support of the FSP in the DOE Office of Science portfolio. A Scientific Steering Committee consisting of experts in relevant fields who are not part of the project will advise the Project Director on technical execution plans, goals and deliverables, and resource utilization among other issues. It will assist the project in identifying deliverables that have the potential to answer questions that are clearly beyond existing capabilities. The two committees together will ensure oversight and accountability in a project that involves such a broad participation.

The Management Coordinating Council will rely on a Software Standards Committee, a Users Advisory Committee and a Verification & Validation Committee for guidance in developing a technical execution plan, including milestones and deliverables that meet the requirements of the stakeholders. For this reason, the members of these committees should be selected from the stakeholders in the FES and ASC communities.

An essential element for success of an R&D project of this size is a clearly defined Work Breakdown Structure (WBS) for each technical task. This is especially the case for the FSP when the scientific basis for some components is still rapidly evolving. The WBS for several illustrative technical tasks is given in the chart below. The WBS should be developed to a level of detail such that the work flow leading to meeting the task milestones is transparent, and the risks and resource requirements are clearly identified.

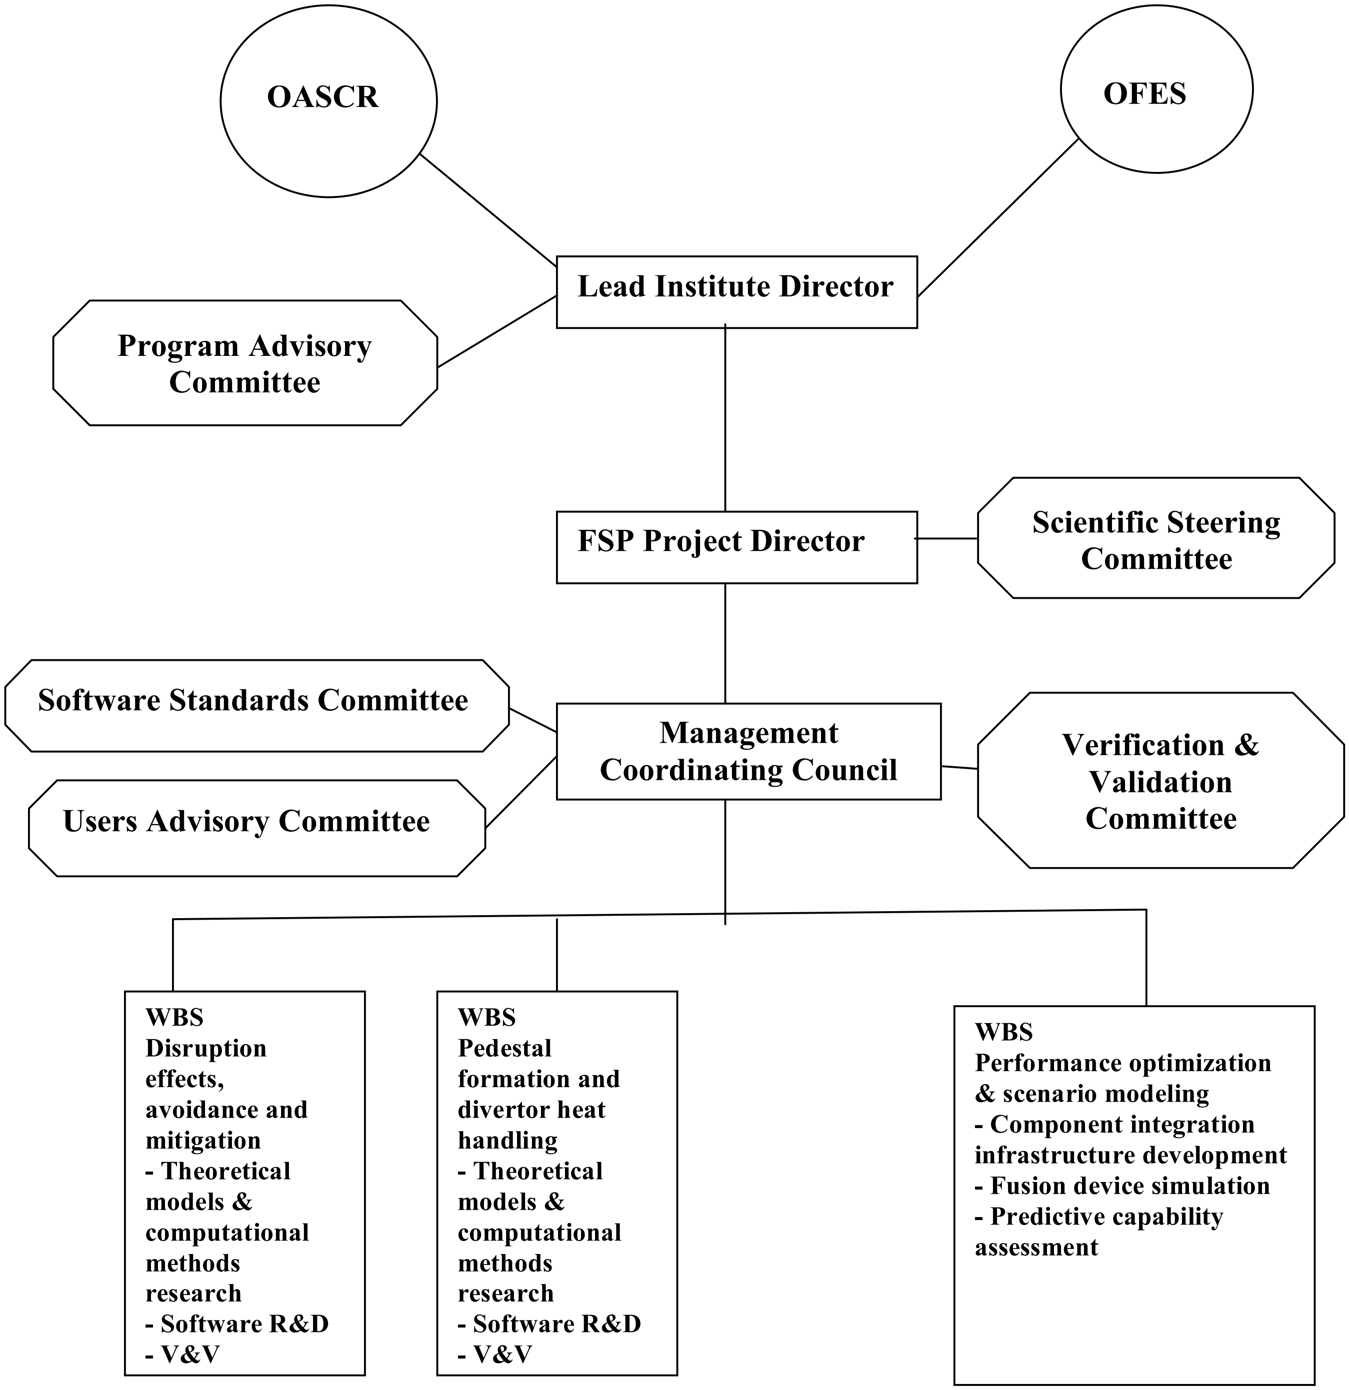

Fig. 5.1 FSP Organization Chart

VII. CONCLUSIONS & RECOMMENDATIONS

This section of the Report contains the Subcommittee's summary assessment of the feasibility of the proposed FSP and recommendations for a course of action.

7.1 General Observations

The integrated modeling capability developed through the FSP should be an embodiment of the theoretical and experimental understanding of confined thermonuclear plasmas. As such, substantive progress on answering the outstanding scientific questions in the field will drive the FSP toward its ultimate goal of developing a reliable ability to predict the behavior of plasma discharges in toroidal magnetic fusion devices on all relevant time and space scales. The FSP should also provide great value-added for addressing critical FES programmatic needs, for enhancing return on U.S. fusion investments, and for helping to identify world leadership opportunities for the U. S. fusion program. In particular, a successful FSP will better enable the study of burning plasmas and greatly aid the U. S. role in the operation of the ITER experiment and in harvesting the associated scientific knowledge. This could enable discovery of new modes of operation, with possible extensions of performance enhancements and improvements needed for a demonstration fusion reactor (DEMO). In general, increasingly reliable whole-device modeling capabilities in Fusion Energy Sciences will surely demand computing resources at the petascale range and beyond to address ITER burning plasma issues. Even more powerful exascale platforms will be needed to meet the future challenges of designing a demonstration fusion reactor (DEMO).

While we find the report convincing on the need for and the potential benefits of an integrated simulation capability (ISC) for the US fusion community, the plan for attaining it is not as clearly spelled out as it could be. The FSP Workshop Report would benefit from more specificity as to the scope of the activity, what the deliverables are and their scientific merit, and how the work would be divided and coordinated. It was not clear to what extent existing software components would be integrated into an ISC, and to what extent "from-scratch" implementations would be required. More specifically, there is a need to better articulate how to integrate into FSP the ongoing SciDAC program, the SciDAC proto-FSP integration projects, including SWIM, CPES, and FACETS, and new developments involving joint experiment-theory-modelling efforts as well as by the expertise residing in OASCR's computer science and applied math programs. These considerations have helped motivate the following FESAC FSP Subcommittee recommendations.

7.2 Recommendations

The FESAC FSP Subcommittee assessment of the FSP Workshop Report has led to the following recommendations:

- While it was felt that the FSP Workshop document came across as too generic and "all inclusive," the FESAC Subcommittee believes that it contains sufficient information for making the case that the FSP can succeed in answering questions in a timely way that experiment and traditional theory by themselves cannot.

- In order to be successful, the FSP should not be "everything to everyone." It must be focused and project-driven with well-identified deliverables that the stakeholders fully support.

- The FESAC FSP Subcommittee agrees with the five critical scientific issues identified in the Workshop Report as important areas of focus appropriate for the FSP. However, an integration effort encompassing all five of these challenging issues from the beginning looks to be too large a step. To be practically achievable, the FSP should begin with more modest integration efforts that exhibit a compelling level of verification and validation. This recommendation is in line with a similar position taken in the original FSP Report from Dahlburg, et al.

- The FSP should be a repository of the latest physics as it evolves. In this sense it cannot be a "stand-alone" project. It must be properly coordinated with theory, experiment and fundamental simulation. More specifically, a proper implementation of the FSP will demand an effective plan for developing "advanced scientific modules" via utilization of results from the OFES base theory program, the SciDAC FES program, new insights from joint experiment-theory-modelling efforts, and the expertise residing in OASCR's computer science and applied math programs.

- The FSP cannot succeed without a viable validation and verification effort, and this will imply expanding the diagnostic effort and linking it better to the FSP, for example through an increased synthetic diagnostic development effort. This will require special personnel with an appreciation of both diagnostic methods and code expertise.

- The management of the FSP should be organized with clear accountability and oversight and work out a clear and compelling work-breakdown-structure (WBS). It should also seek advice and guidance from a broad community of stakeholders, experimentalists, analytic theorists, fusion engineering scientists, applied mathematicians and computer scientists.

- The FSP should establish and maintain strong connections with relevant international projects and also draw on the large experience base from existing scientific software development projects from other fields.

- The DOE should properly launch a true FSP only if a sufficient critical funding level can be realistically met and sustained.

As a final observation, we note that the effective "enfranchising" of more of the fusion community -- especially the experimentalists and technologists as well as analytic theorists -- in the FSP will require that this program produces first-rate scientific capabilities that help advance the research of a large user base of scientists working in these areas, particularly as their work relates to ITER and burning plasmas.

APPENDIX 1: Letter of Charge from Dr. Raymond Orbach

DEPARTMENT OF ENERGY
UNITED STATES OF AMERICA

Under Secretary for Science
Washington, DC 20585

June 8, 2007

Professor Stewart C. Prager, Chair
Fusion Energy Sciences Advisory Committee
Department of Physics
University of Wisconsin
1150 University Avenue
Madison, Wisconsin 53706

Dear Professor Prager:

This letter provides a charge to the Fusion Energy Sciences Advisory Committee (FESAC) to assist in the evaluation of a computational initiative called the **Fusion Simulation Project (FSP)**, which will be led by the Office of Fusion Energy Sciences (OFES) with collaborative support from the Office of Advanced Scientific Computing Research (OASCR). The primary objective of the FSP is to produce a world-leading predictive integrated plasma simulation capability that is important to ITER and relevant to major current and planned magnetic fusion devices. This will involve the development of advanced software designed to use leadership class computers for carrying out unprecedented simulations encompassing multi-scale physics as small as the electron gyro-radius to provide information vital to delivering a realistic integrated fusion simulation model with high physics fidelity.

The intention of the FSP is to use high performance computers to develop modules with much improved physics fidelity that will enable integrated modeling of fusion plasmas in which the simultaneous interactions of multiple physical processes are treated in a self-consistent manner. This comprehensive modeling capability will be developed in close consultation with experimental researchers and validated against experimental data from tokamaks around the world. The FSP will be valuable for discharge scenario modeling and for the design of control techniques in ITER.

A FSP workshop was convened between May 16th and May 18th, 2007, to develop a detailed road map with major scientific and computational milestones. Members of the FSP Workshop Panel (co-chaired by Professor Arnold Kritz of Lehigh University and Professor David Keyes of Columbia University) are expected to produce a completed FSP Workshop Report by the end of June 2007.

I am requesting that FESAC critically review the FSP Workshop Report, assess its feasibility, and recommend a course of action. In examining this report, FESAC should respond to the following questions:

1. Has the report identified key scientific issues and grand challenges that can be addressed by this approach to linking the scientific knowledge base for fusion energy?
2. Have all the critical technical challenges been identified for which predictive integrated simulation modeling has a unique potential for providing answers in a timely fashion, in a way that traditional theory or experiment by themselves cannot?
3. Is there a clear plan to establish the fidelity of the advanced physics modules, including a sound plan for validation and verification?
4. Does the FSP Workshop Report clearly identify the critical areas of computational science and infrastructure in which investments would likely produce the tools required for the FSP to achieve its goals?
5. Have the issues associated with project structure and management of the proposed FSP been properly addressed?

In summary, FESAC should evaluate the prospects for the proposed FSP to produce a world-leading realistic predictive capability that will be of major benefit to the overall science and mission goals of the U.S. Fusion Energy Sciences program. A preliminary FESAC report is expected by the end of July 2007 and a final report by mid-October 2007.

Sincerely,

Raymond L. Orbach

Raymond L. Orbach
Under Secretary for Science

APPENDIX 2: <u>Example of Work Breakdown Structure (WBS) from Combustion Systems Simulation</u>

1 COMBUSTION SYSTEMS SIMULATION (CSS)

1.1 Technical Management and Integration

1.1.1 APPOINT COMBUSTION SYSTEMS SIMULATION ADVISORY COMMITTEE

1.1.2 INTEGRATION OF COMBUSTION SYSTEMS REGIMES

- 1.1.2.1 Develop CSS Project Plan
  - 1.1.2.1.1 Prepare CSS Work Breakdown Structure
  - 1.1.2.1.2 Prepare CSS Network Diagram
  - 1.1.2.1.3 Prepare CSS Three-Year Outlook
- 1.1.2.2 Establish CSS Regimes Integration Group

1.1.3 INTEGRATION OF COMBUSTION SYSTEMS SIMULATION AND COMPUTER SCIENCE & ENABLING TECHNOLOGIES

- 1.1.3.1 Participate in CSS-CSET Working Group

1.1.4 INTEGRATION OF COMBUSTION SYSTEMS SIMULATION AND NATIONAL TERASCALE COMPUTING INFRASTRUCTURE

- 1.1.4.1 Participate in CSS-NTCE Working Group
- 1.1.4.2 Prepare CSS Computing System User Requirements Document
- 1.1.4.3 Participate in Computing Systems Procurements

1.2 Device Simulation

1.2.1 THEORETICAL MODELS AND COMPUTATIONAL METHODS RESEARCH

- 1.2.1.1 Specify Critical Research Problems in Device Simulation

1.2.2 SOFTWARE RESEARCH AND DEVELOPMENT

- 1.2.2.1 Unstructured-grid Device Code
  - 1.2.2.1.1 Develop Collaborative Problem-solving Environment
  - 1.2.2.1.2 Unstructured-grid Device Code Software Development
- 1.2.2.2 Structured-grid Device Code
  - 1.2.2.2.1 Develop Collaborative Problem-solving Environment
  - 1.2.2.2.2 Structured-grid Device Code Software Development

1.2.3 VALIDATION OF DEVICE SOFTWARE AND MODELS

- 1.2.3.1 Define Validation Suite
- 1.2.3.2 Run Validation Suite
- 1.2.3.3 Prepare Validation Report

1.3 Combustion Device Applications, Analysis, and Assessment

1.3.1 IDENTIFY COMBUSTION DEVICES AND SCENARIOS

- 1.3.1.1 Prepare Report on DOE needs
- 1.3.1.2 Prepare Report on National Needs

1.3.2 DEVELOP DEVICE SIMULATION RESULTS DATABASE INFRASTRUCTURE

1.3.3 MODEL SCENARIOS

- 1.3.3.1 Execute Identified Model Scenarios
- 1.3.3.2 Populate Device Simulation Results Database

1.4 Sub-grid Models

1.4.1 THEORETICAL AND COMPUTATIONAL METHODS RESEARCH

- 1.4.1.1 Specify Critical Research Problems in Sub-grid Models

1.4.2 ESTABLISH COMPUTATIONAL TESTBEDS

- 1.4.2.1 Establish Structured-Grid Surrogate Testbed
- 1.4.2.2 Establish Unstructured-Grid Surrogate Testbed

1.4.3 SUB-GRID MODEL RESEARCH AND DEVELOPMENT

- 1.4.3.1 Prepare Report on Sub-grid Model Needs for Selected Combustion Devices
- 1.4.3.2 Research on Subgrid Models

1.4.4 VALIDATION OF SUB-GRID SOFTWARE AND MODELS

- 1.4.4.1 Define Canonical Flows
  - 1.4.4.1.1 Prepare Report on Canonical Flows
- 1.4.4.2 Turbulent Momentum Transport
- 1.4.4.3 Turbulent Mixing
- 1.4.4.4 Turbulence-Chemistry Interactions
- 1.4.4.5 Turbulent Multiphase Sprays and Soot
- 1.4.4.6 Turbulence-Radiation Interactions

1.5 Fine-continuum Simulation

1.5.1 THEORETICAL MODELS AND COMPUTATIONAL METHODS RESEARCH

1.5.1.1 Specify Critical Research Problems in Fine-continuum Simulation

1.5.2 SOFTWARE RESEARCH AND DEVELOPMENT

1.5.2.1 DNS Collaborative Problem-solving Environment

1.5.2.2 Low Mach Number Direct Numerical Simulation Code

1.5.2.3 Compressible Direct Numerical Simulation Code

1.5.2.4 Spray Code Suite

1.5.3 VALIDATION OF FINE-CONTINUUM SOFTWARE

1.5.3.1 Low Mach Number Direct Numerical Simulation Code

1.5.3.2 Compressible Direct Numerical Simulations

1.5.3.3 Spray Code Suite

1.5.4 PREPARATION OF FINE-CONTINUUM ARCHIVAL DATABASES

1.5.4.1 Report on Database Needs for Validating Subgrid Models

1.5.4.2 Fine-continuum Archival Databases
1.5.4.2.1 Develop Archival Database Infrastructure
1.5.4.2.2 Populate Archival Databases

1.6 Combustion Reaction Mechanism Simulation

1.6.1 THEORETICAL MODELS AND COMPUTATIONAL METHODS RESEARCH

1.6.1.1 Specify Critical Research Problems in Combustion Reaction Mechanism Simulation

1.6.2 SOFTWARE RESEARCH AND DEVELOPMENT

1.6.2.1 Collaborative Problem-solving Environments

1.6.2.2 Reaction Mechanism Development Code Suite

1.6.2.3 Chemical Energetics Code

1.6.2.4 Chemical Kinetics and Dynamics Code

1.6.2.5 Chemical Mechanism Databases

1.6.3 APPLICATIONS

1.6.3.1 Natural Gas

  - 1.6.3.2 Gasoline
  - 1.6.3.3 Diesel Fuel
  - 1.6.3.4 Aromatics Oxidation Submechanism
  - 1.6.3.5 NOx Formation Submechanism
  - 1.6.3.6 SOx Formation Submechanism
  - 1.6.3.7 Soot Formation Submechanism

- 1.6.4 VALIDATION OF CHEMISTRY MODELS
  - 1.6.4.1 Validate Chemical Mechanism Approaches
  - 1.6.4.2 Validate Chemical Energetics Models
  - 1.6.4.3 Validate Chemical Kinetics and Dynamics Models

1.7 Materials Simulation

- 1.7.1 PREPARE “CRITICAL MATERIALS NEEDS” DOCUMENT

- 1.7.2 PREPARE MATERIALS COMBUSTION SIMULATION ROADMAP
  - 1.7.2.1 Define Materials Science Software Requirements

- 1.7.3 MATERIALS SIMULATION PROJECT PLAN
  - 1.7.3.1 Work Breakdown Structure
  - 1.7.3.2 Network Diagram
  - 1.7.3.3 Three-Year Outlook

1.8 Combustion Systems Simulation Software Deployment and User Support

- 1.8.1 DOCUMENTATION OF CSS SIMULATION SOFTWARE

- 1.8.2 DEPLOYMENT OF CSS SOFTWARE ON NTCE

- 1.8.3 USER SUPPORT & TRAINING

APPENDIX 3: Timeline of Activities for FESAC FSP Subcommittee

TIMELINE FOR FESAC FSP SUBCOMMITTEE ACTIVITIES

• MAY 16, 17 '07 -- FSP Workshop – Prospective FESAC FSP Subcommittee members invited to attend as observers – first meeting of this panel during Workshop

• JUNE 7, 8 '07 -- Briefing on FSP Workshop to Plasma Science Advanced Computing Institute (PSACI) by A. Kritz – significant number of prospective FESAC FSP Subcommittee members in attendance

• JUNE 8 '07 -- Final Version of FSP Charge Letter to FESAC released

• JUNE 15 '07 -- Full Membership of FESAC FSP Subcommittee announced

• JULY 3 '07 -- FSP Workshop Final Report distributed to FESAC FSP Subcommittee

• JULY 16 '07 -- FESAC Meeting – FESAC FSP Subcommittee Chair makes presentation on: *Discussion of Charge & Plans*

• JULY 23 through OCTOBER 12 -- Series of FESAC FSP Subcommittee Teleconferences:

-- Responsibility for development of written response to the Charge Questions distributed among panel members with at leads for each of the 5 questions assigned

-- Series of 8 full panel Teleconferences held over this time-frame (supplemented by a significant number of additional discussions involving smaller segments of the panel)

• OCTOBER 19 -- *FINAL REPORT from FESAC FSP Subcommittee submitted to full FESAC*

• OCTOBER 23 -- Discussion of FESAC FSP Subcommittee Final Report at the October 23-24 FESAC Meeting

-- FESAC FSP Subcommittee Chair makes presentation on: *Discussion of Findings/Recommendations of FESAC FSP Subcommittee*

-- Associated discussion of formal FESAC final response to Dr. R. Orbach's FSP Charge