Scientific and computational challenges of the Fusion Simulation Project (FSP)

William M. Tang
Princeton University, Princeton Plasma Physics Laboratory, Princeton, NJ

E-mail: wtang@princeton.edu

Abstract. This paper will highlight the scientific and computational challenges facing the Fusion Simulation Project (FSP). The primary objective is to develop advanced software designed to use leadership class computers for carrying out multi-scale physics simulations to provide information vital to delivering a realistic integrated fusion simulation model with unprecedented physics fidelity. This multi-physics capability would be unprecedented in that in the current FES applications domain, the largest scale codes are used to carry out first-principles simulations of mostly individual phenomena in realistic 3D geometry while the integrated models are much smaller scale lower dimensionality codes with significant empirical elements used for modeling and designing experiments. The FSP is expected to be the most up-to-date embodiment of the theoretical and experimental understanding of magnetically-confined thermonuclear plasmas and to provide a living framework for the simulation of such plasmas as the associated physics understanding continues to advance over the next several decades. 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. From a computational perspective, the fusion energy science application goal to produce high fidelity, whole-device modeling capabilities will demand computing resources in the petascale range and beyond together with the associated multi-core algorithmic formulation needed to address burning plasma issues relevant to ITER – a multibillion dollar collaborative device involving seven international partners representing over half the world's population. Even more powerful exascale platforms will be needed to meet the future challenges of designing a demonstration fusion reactor (DEMO). Analogous to other major applied physics modeling projects (e.g., ASCI), the FSP will need to develop software in close collaboration with experimental researchers and validated against experimental data from tokamaks around the world. Specific examples of expected advances which are needed to enable a premier comprehensive integrated modeling capability will be discussed.

1. Introduction

From the perspective of DOE Energy Undersecretary Raymond Orbach, the overarching objective of the Fusion Simulation Project (FSP) [1,2] is 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." Such a capability must be: (i) an important asset for optimizing US participation in ITER – a multibillion dollar collaborative burning plasma fusion experiment involving seven international partners representing over half the world's population; (ii) relevant to major current and planned toroidal fusion devices both nationally and internationally; and (iii) strategically vital to US interests in developing a future demonstration reactor (DEMO). This paper will highlight the associated major scientific and computational challenges which are truly formidable and will demand a strong alliance between DOE's Fusion Energy Science (FES) and Advanced Scientific Computing Research (ASCR) Programs for developing advanced software designed to use leadership class computers for carrying out multi-scale physics simulations to provide information vital to delivering a realistic integrated fusion simulation model with unprecedented physics fidelity. Such a modern multi-physics capability would be unprecedented in that in the current FES applications domain, the largest scale codes are used to carry out first-principles simulations of mostly individual phenomena in realistic 3D geometry while the integrated models are much smaller scale lower dimensionality codes with significant empirical elements used for modeling and designing experiments. The integrated modeling capability of the FSP is expected to be an embodiment of the theoretical and experimental understanding of magnetically-confined thermonuclear plasmas and to provide a living framework for the simulation of such plasmas as the associated physics understanding continues to advance over the next several decades. 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.

If successful, the FSP can be expected to enhance the return on U.S. investments in fusion experiments and help to identify world leadership opportunities for the U.S. fusion program. In particular, the 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 performance enhancements and improvements needed for a demonstration fusion reactor. From a computational perspective, the fusion energy science application goal to produce high fidelity, whole-device modeling capabilities will demand computing resources in the petascale range and beyond together with the associated multi-core algorithmic formulation needed 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. Analogous to other major applied physics modeling challenges [e.g., NNSA's Advanced Strategic Computing Initiative (ASCI)], the FSP software will need to be developed in close collaboration with experimental researchers and validated against experimental data from tokamaks around the world. 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 plasma edge regions.
- 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.

1.1 General Perspective

While many of the technologies used in ITER will be the same as those required in an actual demonstration power plant, 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. Advanced computations in tandem with experiment and theory [3] are essential in this mission – a point well-illustrated in the FSP Workshop Report's first figure [4], 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.

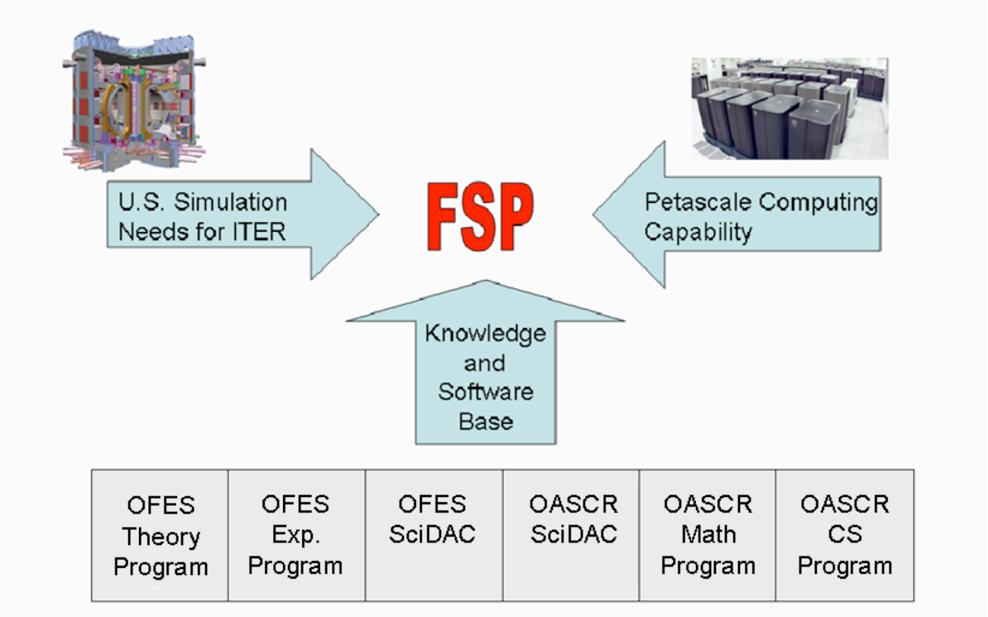

Figure 1. Illustration of expected strong connections to DOE's highest priority missions involving ITER and Leadership Computing as well as to OFES theory and experiment, OASCR CS and Applied Math, and to the SciDAC Program (Ref., DOE FSP Workshop Report, July, 2007)

1.2 Situation Analysis

Experiment and traditional theory are essential components of the fusion energy science research program with a key overarching goal being to reach a level of scientific understanding to enable an accurate predictive capability of a burning tokamak plasma. 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 [3]. However, as noted earlier, the associated challenge of closing the gap between idealized (largely empirically-based) modeling and high-end first-principles-based simulations of mostly individual plasma dynamics is formidable. While the ideal goal is still far away, current progress has shortened the learning curve for developing more complete multi-physics models that are closer to the first–principles equations. Basically, the challenge involves capturing into more realistic reduced modules the most important physics trends produced by the aforementioned large scale first-principles codes that successfully simulate important individual phenomena in realistic 3D geometry. An example is the ongoing development of physics-based transport models [5] that incorporate the predictions of advanced microturbulence simulations that are enabled by access to leadership class computing platforms at the terascale and beyond.

A reasonable approach for the FSP to integrate the most complete physics models available is to begin with binary integration efforts (e.g., SciDAC proto-FSP projects SWIM [6] and CPES [7]) and then extend to higher dimensional integrations to provide a progressively more realistic simulation of experiments. Since many of the physics models are still rapidly evolving, a complementary task involves providing a framework for incorporating new physics as it evolves. An example of this type of effort is the SciDAC proto-FSP FACETS project [8]). There are of course many other examples of smaller efforts aimed at integrating various physics components with a common set of challenges being: (1) keeping up with the constant evolution of validated physics models; and (2) dealing with the incompatibility of various physics components (e.g., with respect to coding structure, dimensionality, scales, etc.). Recognizing that since the relevant scientific areas are not equally developed, it will be necessary to judiciously select an integration strategy working in concert with theory and experiment. An overarching scientific and computational challenge for the FSP is accordingly to provide a living framework for physics as it advances, making it as straightforward as possible to keep the models up-to-date, while at the same time, enforcing standard interfaces to minimize the difficulty of different modules (models) communicating with each other and thereby easing the process for strongly coupled multi-scale integration.

As evident from the DOE FSP Workshop Report (July, 2008), it is fortunate that an excellent collaborative relationship between fusion energy scientists supported by OFES and the computer science/applied math scientists supported by OASCR already exists. 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 J. Dahlburg Report in 2002 (strongly endorsed by FESAC at that time) [1] and the D. Post FSP Steering Committee Report in 2004 [2], but has also been encouraged by US ITER leadership [4]. 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 Accelerated Scientific Computing Initiative (ASCI) Program over the past decade. A lesson learned from major integrated modeling projects in areas such as ASCI is that in order 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 funding are essential. The development of the advanced multi-physics codes targeted by the FSP is expected to take advantage of the ongoing OFES investments in basic theory, the ongoing SciDAC program, and new developments involving joint experiment-theory-modeling 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 to help develop the needed algorithms and improved mathematical models capable of producing high-performance software compatible with the challenging multi-core architectures at the petascale and beyond.

2. Key Scientific Challenges for the FSP

The general consensus from the U. S. fusion energy science research community is that the most prominent and urgent scientific issues impacting the burning plasma program and the successful operation of the ITER experiment include: (1) Disruption effects (large-scale macroscopic events producing rapid termination of plasma discharges), including avoidance and mitigation; (2) Pedestal (steep-spatial gradient) formation and transient divertor heat loads in the plasma periphery; (3) Tritium migration and impurity transport; (4) Performance optimization and scenario modeling; and (5) Plasma feedback control. While undoubtedly a longer list of important scientific questions could be generated (e.g., as found in Table 2.1 of the DOE FSP Workshop Report [4]), these were identified to be the most important and compelling scientific challenges facing the FSP. Independent confirmation of their importance and relevance comes from the fact that the European fusion simulation effort is organized around precisely these five areas [9]. 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 of the DOE FSP Workshop Report [4], 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 illustrates that these physics properties are mutually dependent to a large extent and powerfully makes the case that an integrated multi-physics approach is required. 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 figure 2 (figure 2.2 in Ref. 4). This figure captures a vast amount of physics knowledge, with each topic in its own right requiring detailed physics understanding.

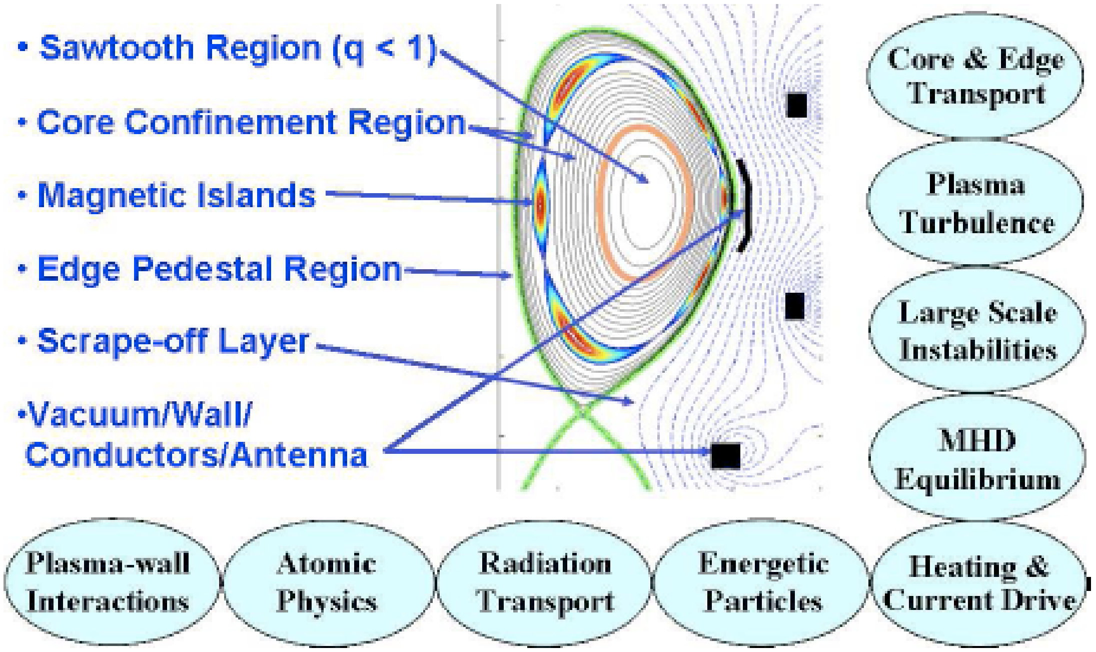

Figure 2. Illustration of the interacting physical processes within a tokamak discharge.

The five issues can be grouped 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. Actually, the science challenges (1-3) require integrated simulations of a different character than the operational challenges (4, 5) – with the former requiring integration of a few "first principles solvers" of high dimensionality and physics fidelity, while the latter requires a larger number of reduced dimensionality models. An associated future modeling challenge for the FSP will be to decide whether it is feasible for these to all be part of a larger, single integrated code.

The five integrated simulation capabilities listed will each also require improved physics modeling components to address: (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. Each of these components has its own set of major scientific and computational challenges. For example, since energetic particles are produced both by auxiliary heating and by fusion reactions in burning plasma experiments such as ITER, a prominent scientific challenge is to carry out self-consistent nonlinear simulations on the transport time-scales of energetic-/alpha-particle-driven Alfven instabilities in the presence of sources and sinks.

The physics modeling components listed here are in fact the traditional fusion topical science areas encompassing everything within the plasma boundary and are currently being addressed at various levels of depth and success in the FES SciDAC projects as well as within the ongoing base FES research program. They are in a sense too high level (and broad) to actually be specific components of the FSP but will, as noted, be separately pursued in a complementary manner in the base and SciDAC FES programs. The expected FSP approach is accordingly to focus on the five critical science issues to deliver reliable predictive tools that would be useful for fusion research in the not too distant future.

Of the five critical areas, the challenge of improved understanding of disruption effects, including avoidance and mitigation, is arguably the most important due to their potential to damage the actual plasma confinement device. When the plasma pressure exceeds linear MHD stability limits, a tokamak plasma is usually observed to disrupt. An important associated scientific challenge is to determine how close to such limits a disruption-free plasma discharge could be confidently expected to be sustainable. From the computational perspective, a formidable but worthy challenge would be to effectively assemble the large (and diffuse) data-base of disruption-relevant events on tokamaks worldwide and then apply the most advanced analysis methodologies to search for key insights into disruption avoidance and mitigation. Other relevant issues include the identification of precursor signals for disruptions and the methods to mitigate the impact of disruption by minimizing the heat load in localized spots. In terms of ITER relevance, an outstanding scientific challenge is the experimental demonstration that a technique or combination of techniques will allow reliable mitigation of disruption effects via control of local heat loads, vertical/horizontal forces, and runaway electron currents. This will demand commencing an associated integrated modeling effort to provide a sound basis for extrapolation to ITER operation. In addition, producing a reliable predictive capability for important performance-limiting MHD instabilities such as neoclassical tearing modes (NTM's) is a major scientific challenge in its own right. This requires the inclusion of long mean-free-path dynamics (e.g., poloidal and toroidal ion flow damping in the presence of 3D magnetic perturbations) which will demand the development of formalisms beyond standard resistive MHD theory.

Pedestal formation and transient divertor heat load impact both the fusion core performance and requirements for materials facing the plasma. Unfortunately, the knowledge in this important area is mostly empirical with very little scientific understanding to guide predictive modeling. 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. In contrast to the years of experience simulating core turbulence and MHD, the simulation of the edge is only beginning, albeit with some significant signs of progress in projects such as the SciDAC proto-FSP CPES [7]. The scientific challenge of better understanding of the edge must be met since information on the edge is needed in order to reliably simulate the performance of the whole device. A practical approach would be to first develop a reduced model based on current knowledge with the expectation that such a model will be improved as more knowledge becomes available.

Development of reliable predictive modeling of tritium migration is another important scientific challenge for the FSP because it can move through the plasma edge region to locations where it is difficult to remove. Effectively dealing with the key problem of tritium retention will require the development of improved capabilities for addressing associated atomic physics and material science issues. In addition, since fusion power production can be degraded by impurity dilution of the deuterium-tritium (DT) fuel, the transport and influx of impurities also needs to be properly predicted – especially with respect to ITER's planned use of carbon facing components for high heat flux regions of the divertor in advance of the DT phase, followed by the installation of tungsten targets before DT operation. The FSP will likely also need to deal with the broader challenge of improved understanding of plasma-surface interactions. In particular, extensive research and development will be needed to establish the physics basis for ITER reference scenarios which are expected to focus on environments with tungsten/beryllium plasma facing components. Associated plasma surface issues include 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.

Since each ITER discharge is expected to cost about $1M, improvement of performance optimization (including sustaining maximum fusion power production) and of scenario modeling capabilities to plan new experiments is important for the FSP. This will require effective tracking and incorporating significant advances in state-of-the-art physics understanding from the four topical science areas. As noted earlier, since some areas are more mature than others, first-principles codes will need to be complemented by reduced models – with both categories of codes needing to be rigorously validated against experiments. Tangible progress in identifying the most reliable models and including them in an integrated multi-physics simulation capability can be expected to provide important insights into the nonlinear coupling between the intertwined physics. Validating the resultant capability against experiments will be an essential and beneficial learning process for the FSP. The associated FSP modeling tools would be most welcome for ITER scenario development activities that include current ramp-up, current ramp-down, vertical stabilization, toroidal field ripple assessment, and various heating and current drive issues. In addition, moving from present-day experiments with tens of seconds duration to ITER discharges dominated by alpha-self-heating and lasting thousands of seconds will present a major challenge for controlling plasma current and pressure.

Plasma feedback control is particularly challenging for ITER because the associated burning plasma regime is a fundamentally new one where self-coupling dynamics are expected to be stronger. This involves discharge start-up, shape control, and disruption avoidance. These applications require mainly integrating MHD physics and have demonstrated progress via the use of modulated heating and current drive and the application of non-axisymmetric fields. Other 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. Another key issue of serious concern to ITER involves controlling the edge localized modes (ELM's) since they can damage the divertor and impact the rapid erosion of plasma facing materials. This will require a good edge pedestal model which does not currently exist. Several schemes are currently under study experimentally to mitigate or eliminate ELMs -- including edge ergodization by introducing external resonant magnetic-field perturbations (RMP's) and modulating the frequency of pellet injection in line with ELM frequencies ("pacemaking"). Transforming what we learn from such experimental studies into a reliable, predictive model constitutes an important scientific and computational challenge for the FSP.

In addition to the multi-physics modeling challenges themselves, the FSP also faces the following key readiness questions:

(1) Which of the emerging or maturing simulation approaches appear most promising in the next 5 years to best make use of multi-core petascale computational resources to accelerate their readiness for integration?

(2) What portions of the integrated model covering the multiple physics features depicted in Fig. 2 have been satisfactorily completed and when might the FSP expect to have an initial working version for experimental validation tests?

At the most generic level, “whole-device” integrated modeling encompasses the entire discharge duration. In particular, 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 damage to the superconducting magnets. Indeed, desired operational scenarios for tokamaks are sensitive to control of the formation phase. It is also important to keep in mind that there are many non-disruptive instabilities which 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. Figure 3 below (as first shown in the DOE FSP Workshop Report [4]) is helpful in placing the integrated system simulation capability within the operational context of actual fusion experiments.

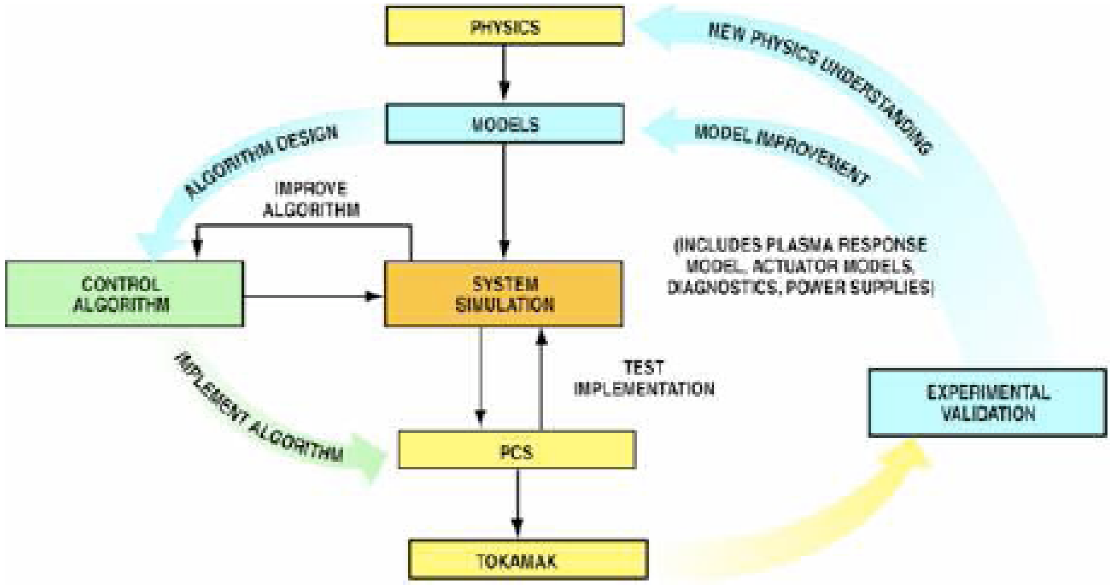

Figure 3. Schematic of integrated plasma control processes which is used by ITER (PCS here stands for a real-time Plasma Control System).

Given the real-time needs of the integrated plasma control process, the orange box labeled “system simulation” represents a reduced “control-level” model capable of rapid execution. It clearly cannot encompass the first-principles whole device simulation model diagrammed in

Figure 2 which would take far too long to run. The plasma control community is already heavily engaged in developing a “Plasma Control System for ITER” with responsibility for this effort led by the Europeans with significant US participation. In order for the FSP to contribute effectively to this ongoing international activity, it must be capable of interfacing with what is already being developed and be able to add value rather than attempt to compete. In particular, an important scientific and computational challenge for the FSP in this context would be to provide a modeling capability to enable effective exploration of ways to address how to control a burning plasma in regimes where the alpha power significantly exceeds external heating and the associated large scale instabilities may not be controllable by the use of modulated heating methods.

More generally, as also noted in the FSP Workshop Report [4], the FSP will need to address the computational issue that the current state-of-the-art software tools for plasma feedback control are limited in their integration with other physical effects and their implementations are characterized as not extensible. Improvements provided by the FSP in this critical area of need would have a most beneficial effect on actual experimental operations. This raises the broader question of priorities when addressing the scientific and computational challenges of the FSP; i.e., which are the most critical and which would most benefit from being addressed earliest? It is apparent that there are at least two integrated simulation software projects required: (1) implementing some version of Figure 2 using a combination of first principles kinetic and fluid solvers for attacking science issues 1-3; and (2) implementing Figure 3 using reduced models. If such a strategy is adopted, an associated challenge is to efficiently connect the first-principles model to the control-level model – including proper validation activities.

3. Verification and Validation Challenges for the FSP

Establishing the physics fidelity of advanced physics modules and the associated integrated software is clearly recognized by the FES and ASCR communities as a major scientific challenge for the FSP [4]. As in the case of the ASCI Project, a plan for verification and validation (V&V) that would rely on ongoing theory and experimental research needs must be properly defined and executed.

Verification is basically the assessment of the degree to which a code correctly implements the chosen physical model. From a purely mathematical perspective, this involves applied math and computer science exercises in the execution of algorithms and touches upon issues associated, for example, with numerical approximations, mesh discretization, temporal discretization, iterative solution of nonlinear equations, and statistical sampling error issues. However, it is important to keep in mind that verification in the context of the FSP must also include systematic comparisons against theoretical physics predictions. 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 gyro-radius corrections, etc.—with the consequence of limitations in the predictive capability of such codes when applied to 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 magneto-hydrodynamic 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 methods, explicit methods, etc.) and/or different mathematical models (Particle-in-Cell, Vlasov or Continuum, Hybrid PIC-Fluid, etc.). In general, the FSP will need to complement conventional mathematical methods for verification of codes with systematic verification through comparisons with theoretical predictions and cross-code benchmarking.

Validation is basically the assessment of the degree to which a code "describes the real world" [4]. The challenge for the FSP is to develop specific strategies with regard to 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 including sensitivity analysis should include data from both the simulation and experiment along with descriptions of data reduction techniques and error analysis, etc. 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 since the coding needs to be sufficiently flexible to accept periodic updating of the physics models. The improved physics models can be assimilated as significant advances in understanding occur. Some examples include a better theoretical understanding of the edge pedestal physics from combining kinetic and fluid models; of wave-heating dynamics from combining Fokker Planck codes and full wave models with finite orbit physics of energetic particles; of transport physics by including self consistent radial electric fields originating from neo-classical theory/simulation on transport time scales; etc. Any one of these processes involves complex nonlinear dynamics – the improved understanding of which should be developed with a combination of experiment, theory, and modeling. Figures 4 and 5 below show examples respectively of a possible approach to the development and validation of transport codes (Figure 4) as well as to RF codes (Figure 5) -- both of which are currently being pursued 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 properly validated, the model would be ready to be assimilated into the FSP project which then integrates and tests multiple physical packages into a "whole" system.

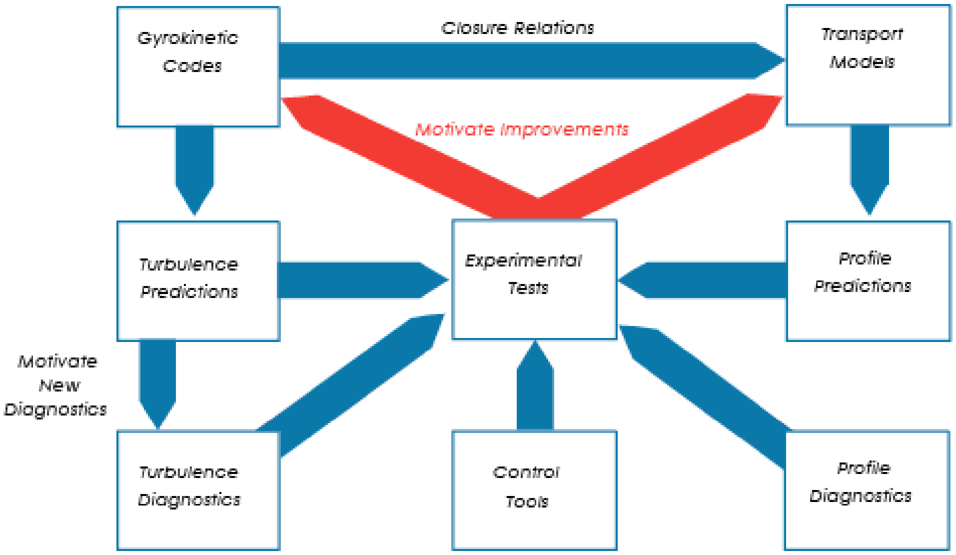

Figure 4. A fully predictive capability is best approached with the combined efforts of modeling and experiment. The above example refers to the development of a predictive understanding of transport in the plasma core.

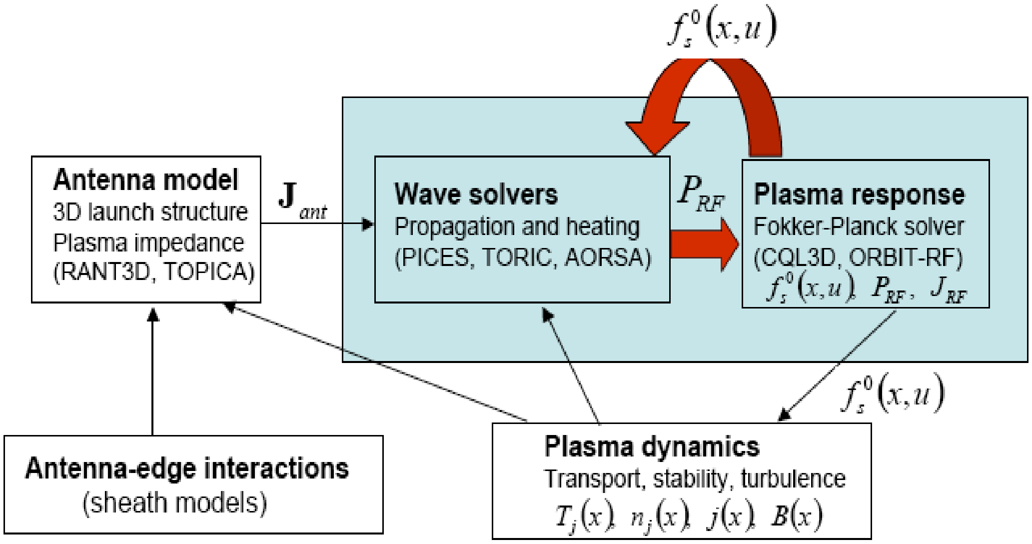

Figure 5. A fully predictive capability in the RF wave-particle interaction area is indicated in the above diagram.

A reliable predictive FSP code can be built only if the individual building blocks have a solid theoretical foundation and careful experimental validation. Basically, theoretical models that capture the targeted science should be implemented and used in the context of self-consistent simulations that can be compared with experimental data – which in turn should be analyzed and organized to facilitate comparisons with the simulation results. Since experiments rely on diagnostics, it is important to build up competence in 2- and 3-D diagnostic capabilities. Furthermore, it is particularly beneficial to validation efforts if each diagnostic is implemented into the codes in a "synthetic" form. It is one of the most powerful tools for validating codes against experiments and has already produced some very impressive results in a number of cases involving RF physics modeling [10] and core turbulence validation studies [5,11]. Development and application of synthetic diagnostics is necessary since almost none of the actual experimental diagnostics can currently provide local data accurate enough to make precise predictions for testing codes.

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. FSP validation efforts will likely depend on the development of specific experimental tools needed to validate the models together with dedicated experiments for validation purposes. In addition, further progress is required in several areas of fusion theory since limitations in current fundamental theoretical understanding will continue to hinder progress in efficient code development. Thus, increased effort for novel development of both laboratory and synthetic diagnostics and for certain aspects of the base theory program will be needed to provide key complementary support to the FSP.

4. Computational Science Challenges for the FSP

The computer science, applied mathematics, and infrastructure (both hardware and its associated software) challenges to produce the needed high-performance software tools for achieving the fusion energy science goals of the FSP are formidable. In order to capture the needed multi-scale (spatial and temporal) fidelity at both the individual physics and the integrated (multi-physics) level, the FSP mathematical models must be well-posed as well as amenable to solutions via efficient numerical methods. Key areas such as simulations of resistive MHD effects, kinetic turbulence dynamics, high-frequency wave-heating, and energetic-particle-driven instabilities in realistic geometry involve mathematical issues encompassing numerical stability, adaptive mesh methodology, space-discretization approaches, efficiency of sparse matrix solvers, accuracy of splitting methods, bifurcation analysis, etc. In particular, the efficient scaling of MHD codes on multi-core platforms at the petascale will need to be aggressively pursued – perhaps via nonlinearly implicit pre-conditioned Jacobian-free Newton-Krylov methods [12]. One of the most daunting challenges facing the FSP is to develop an integrated software product when the scientific basis for a number of key components continues to rapidly evolve. This will obviously 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. In formulating the strategic approach to do so, an associated imperative will be to ensure that national and international best practices and "lessons learned" for large software development projects with strong verification and validation components (e.g., ASCI) are taken into account. 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. Accordingly, a key challenge is for this high-performance FSP software to be developed with an eye toward enabling it to be flexible enough to weather the likely evolution of the hardware architecture as well as that of the associated systems software.

With regard to the overarching technical challenge of dealing with the extreme multi-scale nature of simulating tokamak plasmas for long integration times, a generic approach/vision needs to be properly articulated, for example, for how the "macro/micro" coupling challenge can be achieved even if provided "infinitely powerful" computing power in the future. Although the computational platforms at the petascale and beyond will demand efficient "extreme" parallelization methods, it should nevertheless be kept in mind that the algorithmic challenges are not all parallel in nature. Also, while the scale of the simulation output datasets suggests the likelihood of significant overlap in FSP 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 fusion problems. For example, there will be challenges associated with the development and application of needed tools for improved visualization and pattern recognition of edge plasma events such as edge localized modes (ELMs), turbulent transport structures (e.g., "blobs"), and plasma structures around the magnetic separatrix including magnetic islands.

A clearly dominant trend in the leadership computing hardware development area is the continued addition of 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 less than five years and two orders of magnitude in less than a decade. Together with the high-density, low-power packaging approach to construct large-scale massively parallel distributed memory multi-core computers, this suggests that a petascale or multi-petascale parallel machine in the next decade could reach 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 the FSP as well as in general for fusion energy science and many other application domains. If the fusion energy science applications were only able to effectively utilize a small fraction of the cores on a CPU, major efforts would obviously be needed to accelerate development of innovative new methods for per processor performance.

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

- Efficient scaling of MHD codes beyond terascale levels to enable higher resolution simulations with associated greater physics fidelity.
- Efficient extension of global PIC and continuum 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.
- Development of framework and workflow capabilities for large-scale data management associated with advanced integrated codes.
- Development of innovative data analysis and visualization to deal with increasingly huge amounts of data generated in simulations at the petascale and beyond.

With regard to infrastructure challenges, FSP research enabled by leadership class computing capabilities can be expected to demand much greater computer time and associated support than presently available. For example, a single run using the global particle-in-cell code developed within the SciDAC Gyrokinetic Particle Simulation Center (GPSC) [11] 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) [7] 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 [10]) were also included, then the need for computational resources at the exascale would be a reasonable expectation.

A practical operational challenge for the FSP is to deal with 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). This will give rise to 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 but are expected to be a key element in future requirements planning for ITER.

In general, a requirements and risk analysis for the FSP associated with its computational tools and infrastructure challenges will be needed to determine the appropriate level of direct investment and the expected major increase in capability due to normal developments in the field. Successful examples are evident in the SciDAC Program of joint work funded by ASCR and FES in areas including mathematical techniques, computational libraries, collaboration technology, data analysis, and advanced visualization tools. These efforts will need to be enhanced and resourced with the support levels tightly coupled to the science and engineering goals of the FSP. The associated FSP computational and software infrastructural challenges and requirements will need to 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.

5. Summary

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. It faces the formidable challenge of closing the current gap between idealized (largely empirically-based) modeling and high-end first-principles-based simulations of mostly individual plasma dynamics. This involves capturing into more realistic reduced modules the most important physics trends produced by the high-performance first-principles codes that successfully simulate important individual phenomena in realistic 3D geometry. There are major computer science, applied mathematics, and infrastructure (both hardware and its associated software) challenges for the FSP to produce the needed high-performance software tools for achieving its fusion energy science goals. In order to capture the needed multi-scale (spatial and temporal) fidelity at both the individual physics and the integrated (multi-physics) level, the FSP mathematical models must be well-posed as well as amenable to solutions via efficient numerical methods including the ability to deal with the increasingly complex multi-core systems at the petascale and beyond.

Regarding the scope of the activity, the FSP challenge includes identifying what the deliverables are and their scientific merit. It is necessary, for example, to identify to what extent existing software components can be efficiently integrated now and to what extent “from-scratch” implementations would be required. More specifically, there is a need to better articulate how to selectively 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-modeling efforts as well as the expertise residing in OASCR’s computer science and applied math programs. This will entail identifying the targeted scientific and computational challenges in the form of clear compelling problems with measurable goals. While the FSP Workshop Report [4] is rather generic and “all inclusive” in many respects, it nevertheless contains valuable information for making the case that this project can succeed in answering questions in a timely way that experiment and traditional theory by themselves cannot.

The five critical scientific issues discussed in Section 2 of this paper have generally been identified to be the most important and compelling scientific challenges facing the FSP [4,9]. Each of the five questions is a computational grand challenge in its own right, and requires an integrated simulation capability because the associated physics properties are mutually dependent to a large extent. However, an integration effort encompassing all five of these challenging issues from the beginning is too large a step to be practical. To be practically achievable, the FSP should begin with more modest integration efforts that exhibit a compelling level of verification and validation (V & V). Indeed, the FSP cannot succeed without a viable V & V effort which will likely require enhancement of the existing fusion diagnostic program and an increased synthetic diagnostic development initiative involving personnel with skills in both diagnostic methods and code expertise. Basically, theoretical models that capture the targeted science should be implemented and used in the context of self-consistent simulations that can be compared with experimental data – which in turn should be analyzed and organized to facilitate comparisons with the simulation results.

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-modeling efforts, and the expertise residing in OASCR's computer science and applied math programs. It 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.

In order to be successful, the FSP should not be "everything to everyone." It must be focused and project-driven with well-identified software deliverables that the stakeholders fully support. It is expected to provide great value-added for addressing critical FES programmatic needs, for enhancing return on US fusion investments, and for helping to identify world leadership opportunities for the US fusion program. 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. Increasingly reliable whole-device multi-physics modeling capabilities in Fusion Energy Sciences will 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.

Acknowledgements

A major portion of the source material in this paper came from the FSP Workshop Report [4] led by A. Kritz and D. Keyes and especially from the assessment of the Fusion Simulation Project carried out by the Fusion Energy Science Advisory Committee (FESAC) FSP Panel and summarized in the "FESAC Fusion Simulation Project (FSP) Panel Final Report" (October, 2007) led by the author. He is most grateful for the insightful contributions from his colleagues on the panel: *Riccardo Betti (U. Rochester)***,** *Jeffrey Brooks (Purdue U.), Vincent Chan (General Atomics), Thom Dunning (NCSA, U. Illinois), Charles Greenfield (General Atomics), Brian Gross (GFDL), Michael Norman (UCSD), Miklos Porkolab (MIT), and Rick Stevens (ANL),* and from valuable discussions with David Keyes (Columbia U.), Phil Colella (LBNL), Arnold Kritz (Lehigh U.), Doug Kothe (ORNL), and Ned Sauthoff (US ITER Project Director). The helpful comments and suggestions from the referee of this manuscript are also much appreciated.

References

[1] Dahlburg, J, Corones J, Batchelor D, Bramley R, Greenwald M, Jardin S, Krasheninnikov S, Laub A, Leboeuf J, Lindl J, Lokke W, Rosenbluth M, Ross D, and Schnack D 2001 *J. Fusion Energy* **20** 135

[2] Post D, Batchelor D, Bramley R, Cary J, Cohen R, Colella P, and Jardin S 2004 *J. Fusion Energy* **23** 1
[3] Tang W, Chan V 2005 *Plasma Phys. and Controlled Fusion Research* **47** (2): R1-34
[4] DoE Fusion Simulation Project Workshop Report, ed A Kritz and D Keyes, July 2007, http://www.science.doe.gov/ofes/programdocuments/reports/FSPWorkshopReport.pdf
[5] See, for example, Holland C, 2008 SciDAC 2008 Conference Proceedings, "Validating Simulations of Core Turbulence Simulations: Current Status and Future Directions."
[6] See, for example, Batchelor D, 2008 SciDAC 2008 Conference Proceedings, "Simulation of Wave Interactions with Magnetohydrodynamics."
[7] See, for example, Chang, C S, 2008 SciDAC 2008 Conference Proceedings, "Toward a First-Principles Integration Simulation of Tokamak Edge Plasmas"
[8] See, for example, Cary J, 2008 SciDAC 2008 Conference Proceedings, "First Results from Core-Edge Parallel Composition in FACETS Project."
[9] See, for example, Bécoulet A, Strand P, Wilson H, Romanelli M, and Eriksson L, 2006 International Atomic Energy (IAEA) Conference paper TH/P2-22, "Integrated Tokamak Modelling: The Way Towards Fusion Simulators"
[10] See, for example, Bonoli P, Batchelor D, Berry L, Choi M, D'Ippolito D, Harvey R, Jaeger F, Myra J, Phillips C, Smithe D, Tang V, Valeo E, Wright J, Brambilla M, Bilato R, Lancellotti V, and Maggiora R, 2007 SciDAC 2007 Conference Proceedings, "Evolution of Nonthermal Particle Distributions in Radio Frequency Heating of Fusion Plasmas."
[11] See, for example, Lin Z, 2008 SciDAC 2008 Conference Proceedings, "Verification and Validation of Petascale Simulation of Turbulent Transport in Fusion Plasmas."
[12] See, for example, Chacon L, 2008 SciDAC 2008 Conference Proceedings, "Scalable Parallel Implicit Solvers for 3D Magnetohydrodynamics."