Fusion Simulation Program Execution Plan

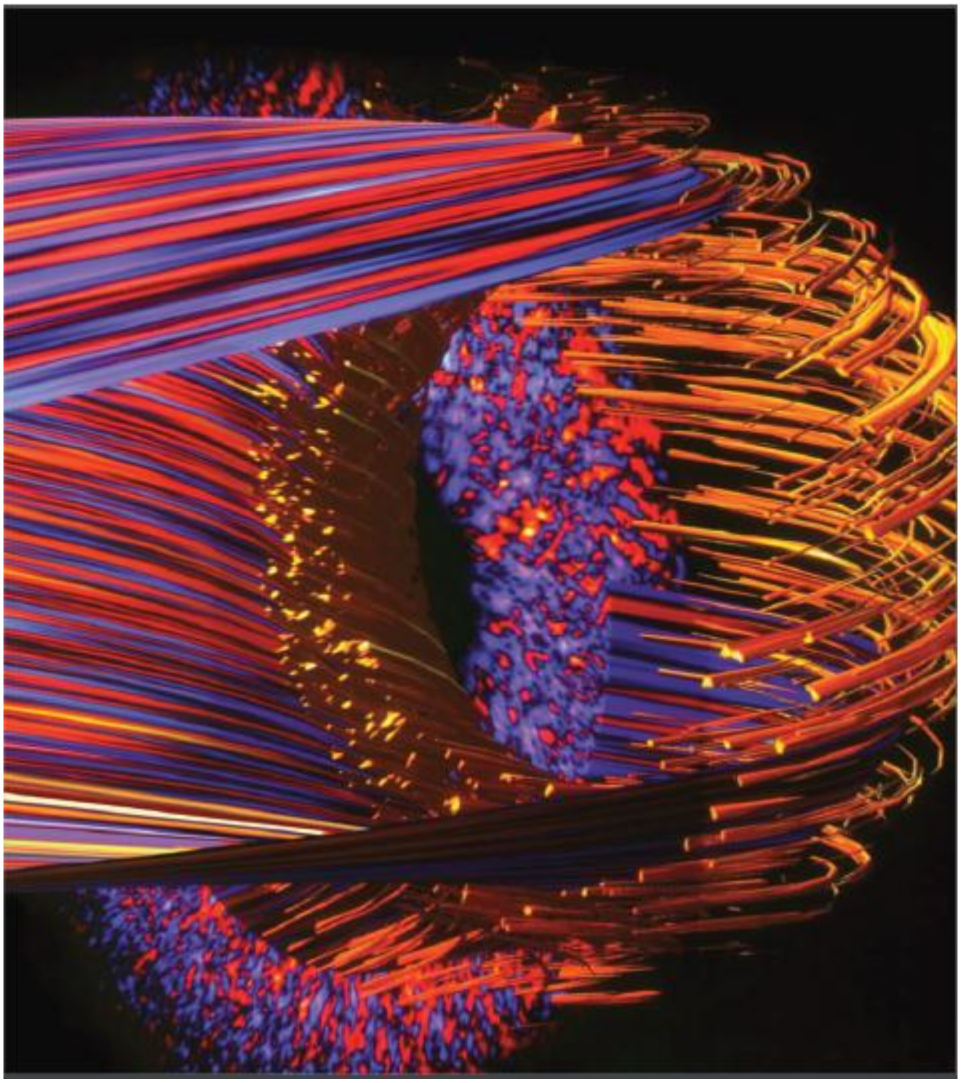

September 30, 2011

The Fusion Simulation Program Execution Plan was prepared for the U.S. Department of Energy Office of Science by the following institutions under these respective contracts:

Princeton Plasma Laboratory (W. Tang)	DE-AC02-09CH11466.
Massachusetts Institute of Technology (M. Greenwald)	DE-SC000-1682
Tech-X, Inc. (J. Cary)	DE-SC0001290
Los Alamos National Laboratory (X. Tang)	DE-AC52-06NA2-5396
Lehigh University (a. Kritz)	DE-SC0001268
General Atomics (V. Chan))	DE-SC0001266
Oak Ridge National Laboratory (D. Hudson, L. Chacon)	DE-AC05-00OR22725
Lawrence Livermore National Laboratory (W. Nevins, L. Daichin)	DE-AC52-07NA27344
Argonne National Laboratory (A. Siegel)	DE-AC02-06CH11357
Lawrence Berkeley National Laboratory (P. Colella)	DE-AC02-05CH11231
New York University (C.S. Chang)	DE-FG02-06ER 54845
University of California San Diego (G. Tynan, P. Diamond)	DE-FG02-04ER 54738
Purdue University (D. Brooks)	DE-FG02-08ER54990
University of Texas (H. Berk)	DE-FG02-04ER-54742

Many others in the fusion energy and advanced scientific computing communities participated in the development of this plan. The core planning team is grateful for their important contributions.

An original version of this plan was submitted on July 31, 2011, in satisfaction of the requirements of the primary contract awarded to the Princeton Plasma Physics Laboratory. This version contains modest updates to the July 31st submission resulting from comments by DOE program offices, advisory committee members, and other interested stakeholders.

The plan is the governing document that establishes the means to execute, monitor, and control the proposed Fusion Simulation Program (FSP).

There are several additional documents referenced within this one and all are supplemental or flow down from this Program Plan. They are not included within this document because of their purpose, format, or size:

- FSP Plan Executive Summary – *a quick overview the FSP's purpose and intentions (MS Word, approximately 2 pages)*
- FSP Plan Summary – *a concise interpretation of the full Plan (MS Word, approximately 30 pages)*
- WBS Cost and Schedule Worksheet – *a multipage spreadsheet containing the proposed FSP work breakdown structure (WBS) and its dictionary, work schedule, required resources, and funding timelines (MS Excel workbook)*
- Detailed Science Driver Planning Reports – *a report containing bottoms-up definitions of plans for addressing critical programmatic issues that require integrated simulations. It was prepared by small groups of technical experts from the fusion energy and advanced scientific computing communities at the request of the FSP planning team (MS Word, approximately 182 pages)*

Table of Contents

Table of Figures

Table of Tables

Fusion Simulation Program Plan

The basic structure of this document, after the introductory sections (Section 0), is as follows:

- the scientific opportunities available to the FSP and the fusion research community (Section 2);
- the organization, policies, and procedures of a program that might best realize those opportunities (Section 3);
- relatively specific plans to realize as many of those opportunities as possible given limited resources (Section 4);
- a proposed management structure (Section 5);
- proposed management policies, procedures, and processes (Section 6); and
- various appendices.

1 INTRODUCTION

1.1 Overview

The overall science goal of the FSP is to develop predictive simulation capability for magnetically confined fusion plasmas at an unprecedented level of integration and fidelity. This will directly support and enable effective U.S. participation in research related to the International Thermonuclear Experimental Reactor (ITER) and the overall mission of delivering practical fusion energy. The FSP will address a rich set of scientific issues together with experimental programs, producing validated integrated physics results. This is very well aligned with the mission of the ITER Organization to coordinate with its members the integrated modeling and control of fusion plasmas, including benchmarking and validation activities [1]. Initial FSP research will focus on two critical areas: 1) the plasma edge and 2) whole device modeling including disruption avoidance. The first of these problems involves the narrow plasma boundary layer and its complex interactions with the plasma core and the surrounding material wall. The second requires development of a computationally tractable, but comprehensive model that describes all equilibrium and dynamic processes at a sufficient level of detail to provide useful prediction of the temporal evolution of fusion plasma experiments. The initial driver for the whole device model (WDM) will be prediction and avoidance of discharge-terminating disruptions, especially at high performance, which are a critical impediment to successful operation of machines like ITER. If disruptions prove unable to be avoided, their associated dynamics and effects will be addressed in the next phase of the FSP.

The FSP plan targets the needed modeling capabilities by developing Integrated Science Applications (ISAs) specific to their needs. The Pedestal-Boundary model will include boundary magnetic topology, cross-field transport of multi-species plasmas, parallel plasma transport, neutral transport, atomic physics and interactions with the plasma wall. It will address the origins and structure of the plasma electric field, rotation, the L-H transition, and the wide variety of pedestal relaxation mechanisms. The Whole Device Model will predict the entire discharge evolution given external actuators (i.e., magnets, power supplies, heating, current drive and fueling systems) and control strategies. Based on components operating over a range of physics fidelity, the WDM will model the plasma equilibrium, plasma sources, profile evolution, linear stability and nonlinear evolution toward a disruption (but not the full disruption dynamics). The plan assumes that, as the FSP matures and demonstrates success, the program will evolve and grow, enabling additional science problems to be addressed. The next set of integration opportunities could include: 1) Simulation of disruption dynamics and their effects; 2) Prediction of core profile including 3D effects, mesoscale dynamics and integration with the edge plasma; 3) Computation of non-thermal particle distributions, self-consistent with fusion, radio frequency (RF) and neutral beam injection (NBI) sources, magnetohydrodynamics (MHD) and short-wavelength turbulence.

The identification of the need for the FSP has been presented in a number of prominent past studies and reports over the past decade [2-6], and the importance of validated predictive simulation capability affirmed prominently in the Department of Energy (DOE) Office of Fusion Energy Simulation (FES) community-wide Research Needs Workshop (ReNeW) in 2009 [7]. Most recently, the Program Advisory Committee (PAC) for the FSP[1] has strongly endorsed both the concept and potential of the FSP. In the Executive Summary of their Report of May 8, 2011, these distinguished scientists have stated that after closely following the development of the FSP Plan over the preceding 18 months, they believe that the FSP will:

- enable significant advances in fusion science,
- substantially increase the value of ITER to the U.S.,
- make major contributions to build the knowledge base required for the Demonstration Power Plant (DEMO) project, and
- provide one of the few opportunities available for the U.S. to provide recognized leadership in the international fusion science community.

They conclude with the statement: "A Fusion Simulation Program <u>of the type proposed</u> provides the most credible path forward for the integrated whole device model that will be <u>highly important for the realization of fusion energy. [8]</u>

1.2 FES and ASCR Mission Needs

Advancing the fundamental science of magnetically confined plasmas to develop the predictive capability needed for a sustainable fusion energy source is a top priority of the FES program, deriving directly from its mission. A validated predictive simulation capability is critical for addressing this priority. The FES program is now moving into the burning (or self-heated) plasma regime through its participation in ITER, an international fusion research facility under construction in Cadarache, France, which will be the world's first facility large enough to achieve burning plasma and investigate its characteristics. ITER is expected to start operations in 2018 and enter its burning plasma phase around 2028. Based on the experience of other scientific communities, including National Nuclear Security Administration (NNSA) and the climate community, it takes more than 10 years to develop, verify, validate, and deploy complex multi-physics codes such as those needed to address the FSP mission. It is therefore important that key activities start as soon as possible, so the U.S. fusion community will be ready to contribute to this international effort and maximize the benefits of its participation. In addition, since the FSP science goals will significantly benefit from the availability of high performance computing (HPC) resources beyond the petascale, an early start will provide timely information to the DOE Office of Advanced Scientific Computing Research (ASCR), whose mission includes assessing the opportunities and challenges of exascale computing in support of the broader DOE Office of Science (SC) mission.

The mission of the FES program is to expand the fundamental understanding of matter at very high temperatures and densities and to build the scientific foundations required to develop a fusion energy source. The associated need to develop a properly verified and validated predictive integrated simulation capability for magnetically confined fusion plasmas is driven by its considerable potential to address a leading challenge for the fusion program – establishing the scientific credibility of magnetic fusion energy (MFE). The mission need is also driven by the U.S. participation in ITER in that experimental proposals requesting run time on ITER are expected to be accompanied by modeling justification to ensure that the targeted scientific results are within ITER's operating capability as well as to ensure that operational limits threatening the integrity of the device will be avoided. In addition, extensive modeling will be necessary to evaluate the data from each ITER discharge and extract scientific conclusions. Addressing this mission need will enhance the credibility of proposed U.S. experimental campaigns on ITER, maximizing the return of our investment in this international effort. By doing so, it will also address a key recommendation of the National Research Council's committee tasked to review the

---

[1] *Douglass Post, <u>Chai</u>r (DoD), Allen Boozer (Columbia U), Leslie Greengard (NYU), Brian Gross (GFDL), Greg Hammett (PPPL), Wayne Houlberg (ITER), Earl Marmar (MIT), Dan Meiron (CalTech), Jon Menard (PPPL), Mike Norman (UCSD), Rick Stevens (ANL).Carl Sovinec (U Wisc), Tony Taylor (GA), Jim Van Dam (U Texas)*

DOE plan for the U.S. Plasma Science Community Participation in ITER. Specifically, the committee suggested that "enabling U.S. ability to contribute substantially to ITER, and maximizing U.S. ability to act on the results produced by ITER, in order to fully reap the enormous scientific and technological reward possible as a result of U.S. involvement in the project" should be a key goal underpinning the U.S. participation.

The FSP directly supports the efforts of the FES to develop integrated simulation capabilities to further its mission and support its strategic goal to "advance the fundamental science of magnetically confined plasmas to develop the predictive capability needed for a sustainable fusion energy source."

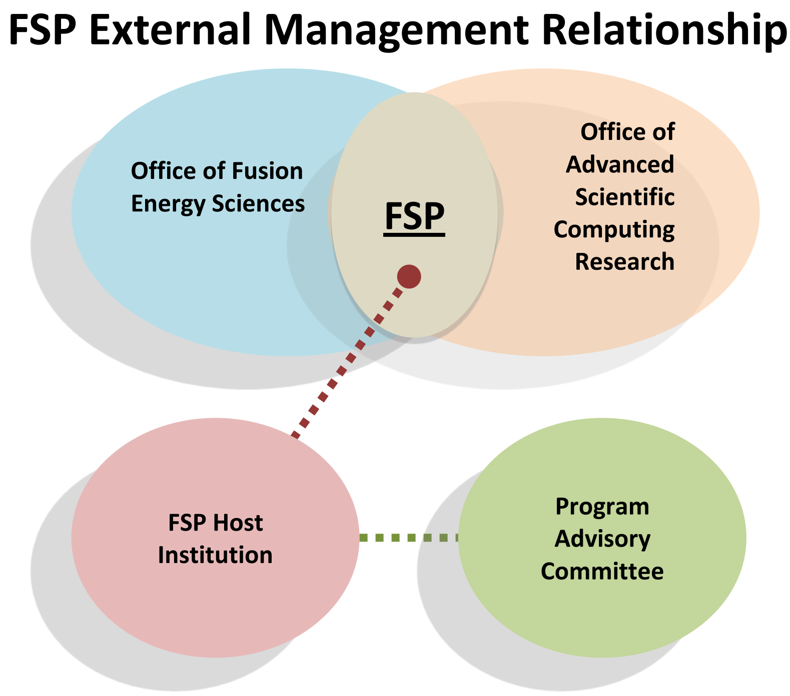

Figure 1: FSP External Management Relationships*

** The DOE Office of Advanced Scientific Computing Research recognizes the synergistic advantages available to it in achieving its stated goals by joining FES in supporting the efforts of the FSP.*

The work of the FSP does not solely reside within the purview of either program office, but at the intersection of interests of both organizations.

Success of the FSP depends critically on leveraging expertise from within the broader advanced simulation community. Building on the foundations established at the FSP community planning workshop in San Diego (Feb. 2011), the FSP Plan features ASCR research contributions that can both accelerate progress in FES as well as advance the ASCR research mission (see Figure 4). Research from within the ASCR community will be required in seven general areas:

(i) **Scalable Solvers** – solver techniques, especially for highly parallel or multithreaded hardware;

(ii) **Time Integration** – improved time integration techniques, especially for coupled partial differential equations (PDEs);

(iii) **Formulation** – Innovative formulation of continuous and discrete models

(iv) **Multi-scale/physics** – advanced methods for multi-scale and multi-physics coupling;

(v) **Data/Meshing** – advanced methods for more efficient data management/analysis, including visualization and meshing;

(vi) **Frameworks** – framework design, including the software challenges of componentization and coupling on HPC systems; and

(vii) **Verification & Uncertainty Quantification** – application and further development of methods in verification and uncertainty quantification.

Although specific to the targeting the FES application domain in the context of the FSP, these categories span in general a wide range of current ASCR activities. These categories are referenced in Table 11: FSP Level 1 Deliverables and Milestones (Quick Reference) in Section 4.8.

1.3 FSP Mission Statement

The goal of the Fusion Simulation Program (FSP) is to enable scientific discovery of important new plasma phenomena with associated understanding that emerges only upon integration. This requires developing a predictive integrated simulation capability for magnetically-confined fusion plasmas that are properly validated against experiments in regimes relevant for producing practical fusion energy.

1.4 FSP Vision Statement

The Fusion Simulation Program will provide the capability to confidently predict toroidal magnetic confinement fusion device behavior with comprehensive and targeted science-based simulations of nonlinearly-coupled phenomena in the core plasma, edge plasma, and wall region on time and space scales required for fusion energy production.

- Integrate the knowledge from key multi-scale physical processes to continually improve fidelity for extending whole-device modeling capabilities beyond current applicability domains.
- Produce a framework in which physics component-codes interact efficiently to enable unprecedented capabilities to compute experimental observables, interpret experimental data, and explore the consequences of theoretical models.
- Incorporate modern software engineering and software quality assurance to ensure the reliability, robustness, and ease-of-use of the tools that are developed.
- Create the most advanced suite of predictive codes under a unified framework and distribute and support it within the fusion community. This will maximize U.S. investments in experimental facilities (especially, ITER) and in HPC resources (especially, the Leadership Class Facilities) to produce the scientific basis for an economically and environmentally attractive source of energy.

2 SCIENCE OPPORTUNITIES

After a competitive peer-review in 2009, the FES selected a national multi-institutional and multidisciplinary team to carry out a two-year detailed planning study for the FSP. The planning team (individuals listed on the FSP web-site (http//www.pppl.gov/fsp) was composed of scientists with a broad range of expertise in fusion plasma science (including theorists, computationalists, experimentalists, and material scientists) and also in applied mathematics, computer and computational science (AM/CS), and software engineering.

The project team was organized into five groups, covering planning areas for the FSP itself: 1) science applications including experimental validation, 2) physics components, 3) software integration and support, 4) software quality including verification, validation, and uncertainty quantification, and 5) management. Overall management for the project definition phase was led by a director (W. Tang) and deputies for plasma science (M. Greenwald) and for advanced scientific computing research (A. Siegel). The management team was rounded out by the heads of the planning groups including V. Chan, X. Tang, J. Cary, J. Hittinger, D. McCune, and A. Kritz. D. Hudson provided program/project management advice. Altogether, personnel from six national labs, two private companies and nine universities were directly funded for the planning activity. In addition, strong input and contributions from the community was actively solicited and incorporated into this effort.

2.1 Motivation and Objectives

The FSP should lead to deeper understanding and improved predictive models by integrating phenomena which are now treated in isolation. These will advance the fusion program's overall scientific mission, improving interpretation of experimental results and embodying our state of knowledge. This should also result in more reliable scenario modeling for existing and future machines – especially ITER.

There will be a number of basic FSP “products," including:

- An FSP simulation suite, namely a set of new scientific software tools and environments that embodies the community’s latest understanding of relevant fusion science and plasma physics in tokamak plasmas;
- New computational tools within a FSP simulation suite to help enable increased understanding and needed advances in the field of plasma physics and fusion energy sciences;
- New capabilities needed to support current and planned fusion tokamak experiments, with ITER being the highest priority focus;
- AM/CS research efforts achieved through a strong partnership on problems of mutual interest;
- Education, training, and team building of the current and next generation of fusion scientists and computer and computational scientists needed to support, develop, and use the FSP simulation suite of software tools.

2.2 Program Development/Prioritization Criteria

As noted, the scale of effort entailed in *toto* about 100 full-time equivalent staff (FTE), beginning in the first year for the scientific program outlined above. This number represents an unrestricted assessment of what level of manpower that these programs could use, but far outweighs the anticipated funding available to the FSP, especially in its early years. Given the likely constraints, it is clear that the scope needs to be more limited and that all elements of the program cannot begin at once. A phased program is also consistent with an orderly buildup of personnel and infrastructure.

To guide the planning, a set of criteria were derived to use in sequencing or prioritizing research. These are:

- A clear need for multiscale, multi-physics integration: The proposed topic should require the capabilities envisioned for the FSP – that is, it should be outside the scope of the current modeling programs. Solving the problem would demonstrate that the FSP “is more than the sum of its parts”.
- Importance and urgency: Importance would be measured against the FES mission to create “the knowledge base needed for an economically and environmentally attractive fusion energy source”. That is, it should answer questions, or solve problems integral to that knowledge base. Urgency should take into consideration schedules, dependencies and critical paths for program elements that the FSP would support – for example ITER.
- Readiness and Tractability: The underlying physics along with the required AM/CS research and computing platforms required to address the proposed topic, should be sufficient to begin work on at the outset of the FSP. The need for the FSP to impact ongoing research at an early date should also be considered. At the same time, some reasonably clear path toward solving the research problem, posed by the topic, should be envisioned.
- Opportunity to open up new lines of research: Attacking the problem should offer the possibility of new insights or potential breakthroughs, particularly those not accessible by other means.

An overarching prioritization criterion is the "buy-in" from the "customer-base" for FSP products with respect to what software capabilities are in greatest demand and urgency from the user communities. The final choices made reflect a realistic level of “market analysis” with linkages to the ReNeW document to the priorities of the Fusion Facilities Coordinating Committee, to ITER, and also other international facilities, e.g., in Asia with experimental capabilities not available in U.S. facilities.

2.3 Initial FSP Science Program Analyses – Results and Justification

The next challenge was to rationalize the bottoms-up science roadmaps and their estimated resource requirements, with a top-down plan that accounted for all needed program elements and that fit within an assumed budget. To carry out this exercise we assumed a funding profile that began at $12M per year and grew to $24M by year 6. Resource requirements were all provided as FTE/year, so to convert these budget assumptions into available manpower, we assumed an average of $400K per FTE year. This was a very rough average of costs for researchers at different levels of seniority and at institutions with widely varying overhead structures. (Clearly a different set of assumptions would lead to an FSP program with different scope.) Then, based on estimates of all program requirements, we made a split between workforce devoted to technical leadership (5-10%), science driver-focused application development (ISA) teams (50-60%) and the cross-cutting enabling technology (ET) teams (40-50%)[2]. We note that the ISA teams would also contain embedded software specialists to provide targeted expertise for components, frameworks, verification and validation). The resource allocation reflects a relative emphasis on infrastructure development in the early years. The ratio of ISA to software team manpower begins at about 50:50 and evolves to roughly 60:40 as the program matures. The result is summarized in the following table.

Year	Funding	Total FTE/year	ISA FTE	ET FTE	Technical Leadership
1	$12M	30	14	13	3
2	$15M	38	20	15	3
3	$18M	45	24	18	3
4	$20M	50	26	20	4
5	$22M	55	30	21	4
6	$24M	60	33	23	4

Table 1: Top-level allocation of resources based on assumed funding profile and manpower costs.

With these levels of resources, the planning team decided to begin with two ISA teams that would grow to full size (10-12 FTE) over three years (allowing time for recruitment and training of new members). In year four, a third ISA would be launched, likely followed by a fourth in year 6 or 7.

Choosing which ISAs to start with was difficult as all of the science drivers addressed critical physics problems. However based on the criteria above especially current readiness, the need for integration and programmatic urgency, it was felt that the strongest cases could be made for the Boundary, Pedestal and Whole Device Modeling. In these areas, two “Killer Apps” were noted to have extreme program importance, particularly for ITER:

- Heat and particle loads with associated impact on PMI, including the nature and impact of edge localized modes (ELM).
- Discharge optimization and prediction, especially the avoidance of disruptions

The initial FSP research will focus on two critical integrated science application areas: ISA1, the plasma edge and ISA2), whole device modeling, including disruption avoidance. The first of these problems involves the narrow plasma boundary layer and its complex interactions with the plasma core and the surrounding material wall. The second requires development of a computationally tractable, but comprehensive model that describes all equilibrium and dynamic processes at a sufficient level of detail to provide useful prediction of the temporal evolution of fusion plasma experiments. The initial driver for the whole device model will be prediction and avoidance of discharge-terminating disruptions, especially at high performance, which are a critical impediment to successful operation of machines like ITER. If disruptions prove unable to be avoided, their associated dynamics and effects will be addressed in the next phase of the FSP.

---

[2] ISA and ET teams are major operational units of the FSP. They are describe fully in Section 3 below.

The Pedestal-Boundary model developed in ISA1 will include boundary magnetic topology, cross-field transport of multi-species plasmas, parallel plasma transport, neutral transport, atomic physics and interactions with the plasma wall. It will address the origins and structure of the plasma electric field, rotation, the L-H transition, and the wide variety of pedestal relaxation mechanisms. This ISA is a key partial integration project – along with future ISAs – that will ultimately be integrated into the FSP WDM which will predict the entire discharge evolution given external actuators (i.e., magnets, power supplies, heating, current drive and fueling systems) and control strategies. Based on advanced components with improved physics fidelity operating within an appropriate integration framework, ISA2 will focus on modeling the plasma equilibrium, plasma sources, profile evolution, linear stability and nonlinear evolution toward a disruption (but not the full disruption dynamics).

The plan assumes that, as the FSP matures and demonstrates success, the program will evolve and grow, enabling additional science problems to be addressed. The next set of integration opportunities could include the following topics:

Disruption Mitigation. If disruptions – the large-scale macroscopic events leading to rapid termination of plasma discharges, including severe impulsive heat loads damaging material components – cannot be completely avoided, mitigating the associated dynamics is critical because ITER can sustain at most a very small number of such full current events. The associated science goal is to minimize the impact of disruptions, including dealing with transient heat and mechanical loads and generation of run-away electrons. This will involve dealing with strongly nonlinear MHD phenomena in large Lundquist number plasmas, addressing coupling to plasma pressure & current and also to atomic physics, neutral and impurity transport, radiation transport, & relativistic electron transport. It will also require assessment of the relationship to an electromagnetic model of the fusion device, including complex wall geometry, power supplies, coils, and control systems. If successful, the expected benefits include: (i) survivability of first wall tokamak components; and (ii) viable steady-state operation of a fusion device.

Core Profiles. The science goal here is improved predictive capability for the temperature, density, current, and rotation profiles in the plasma core, including the internal transport barrier region. This task includes dealing with 3D effects, mesoscale physics, & integration with the plasma edge dynamics. It involves producing self-consistent, global solutions of micro-and macro-nonlinear dynamics on transport time scales. Since mesoscale phenomena (between gyro-radius and device size), overlap with MHD scale, there is no justifiable strong scale separation that can be invoked to simplify this challenging problem. If successful, expected benefits include a predictive capability for plasma profiles that would enable providing profile information needed to determine operational limits (e.g., sustainable plasma pressure) and plasma performance (e.g., fusion yield, bootstrap current fraction), and also to provide confidence in extrapolating core confinement predictions to future devices.

Energetic Particles/Wave Physics. These are dynamical interactions between energetic particles and electromagnetic waves in an MFE plasma that impact the efficacy of auxiliary heating and the fast-particle confinement of fusion products (alpha particles at 3.5 MeV) and of supra-thermal particles from RF and energetic NBI heating. Energetic particles represent potent sources of free energy available to drive instabilities, and their thermalization without loss is critically important. The associated science goals in this area include: (i) a self-consistent description of phase space distribution on long time scales (energy confinement or slowing-down) that are orders of magnitude longer than the Alfvenic time scales for underlying wave-particle interactions; (ii) dealing with strong nonlinearities and mutual coupling to transport through pressure, velocity and current profiles, and fluctuation spectra; and (iii) ultimately delivering reliable predictive capability for fast particle distributions that are self-consistent with fusion reactions, RF and NBI sources, MHD activity and short-wavelength turbulence. If successful, improved predictive capability of fusion yield and key aspects of steady-state operation would be achievable, thereby enabling information essential for ensuring steady-state (long-pulse) performance in burning plasmas such as ITER.

Basic research in these additional ISA areas would also be carried out by the five existing Scientific Discovery through Advanced Computing (SciDAC) projects in MHD, micro-turbulence, wave-particles and RF, all of which were recently renewed for five more years. During this interval, they would carry the principle load in building the foundation for integration activities. The choice for the next ISA would be made during year 3 of the FSP and would depend on progress in these areas.

Year	Total ISA FTE/year	Edge/Pedestal	WDM	ISA 3	ISA 4
1	14	7	7	0	0
2	20	10	10	0	0
3	24	12	12	0	0
4	26	11	11	4	0
5	30	11	11	8	0
6	33	11	10	8	4

Table 2: Proposed level and distribution of effort for ISA teams

3 PROGRAM OPERATION

3.1 The Matrix Approach

The FSP goal is to develop an integrated community computational toolkit that is capable both of carrying out innovative simulations for selected science drivers and that forms the basis for extension to a much broader range of problems. This section outlines the technical organization for the FSP. That is, it describes a structure for carrying out the technical and scientific work of the FSP – specifically, how the science application efforts interact with component development, code infrastructure, quality assurance, production computing and AM/CS cross-cutting efforts.

The mission of the FSP is not limited to a single science problem but rather aims to develop and integrate broader solutions applicable to a grand challenge class of science questions and able to move toward a comprehensive whole device model. In doing so, the plan adopted is structured to achieve economy of scale as it grows in an accountable way to a program capable of producing increasingly greater scientific productivity. The associated goal is to deliver for the first time a suite of true community codes for production computing with user support that is coordinated by means of a national program dedicated for:

- Incorporating modern software infrastructure & developer support – instead of the current practice of diffuse development and support of individual applications;
- enabling modern AM/CS technologies be deployed for rapid sharing of new tools and approaches between applications;
- efficiently integrating physics components with common interfaces & data structures guided by appropriate standards;
- producing FSP-standardized, well-documented tool sets for data preparation, code input validation, data analysis, and visualization FSP standards – instead of the current inefficient customization approach; and
- achieving a world standard for FES with application of modern verification, validation, and uncertainty quantification (V&SQ) methods to yield more rigor & efficiency via strong coordination with experimental facilities (national & international).

Thus, it is appropriate to initiate the FSP with two Integrated Science Applications (ISAs) under simultaneous development. However, this will not by itself lead to an integrated project capable of satisfying the full FSP mission. There would be a strong natural tendency for each ISA team to work independently, resulting in, for example, multiple similar codes for different application areas that are a wasteful duplication of effort. A collection of largely independent research activities of this kind has no real hope of cost-effectively establishing the necessary coordinated infrastructure essential for delivering the software products that would enable ITER to successfully harvest in a timely way the targeted scientific knowledge base needed to deliver on its mission goals, i.e., laying the foundations for moving on to DEMO. The FSP will accordingly employ a matrixed approach, drawing staff from both Integrated Science Application (ISA) Teams and cross-cutting Enabling Technology (ET) Teams.

3.1.1 Integrated Science Application (ISA) Teams

The FSP will be primarily organized around multiple integrated science application teams each with approximately seven to ten members. Each ISA team is multi-disciplinary and charged with executing the simulation development plan for a particular science problem – specifically, those summarized within the science driver reports in Appendix B. Each ISA team will have a single technical lead for day-to-day direction. All ISA teams will be overseen by the Deputy Director for Science who will drive each team toward their technical goals, manage resource distribution, ensure proper interaction with the broader scientific community, and recommend re-prioritization of allocations based on success/failure, new ideas from the community, and new directions from the Program Office.

Through its research and development, the ISA teams will define strategies, address critical problems, exercise the required range of code capabilities, and provide useful tools for the broader fusion community. Working with the Software Integration and Support (SIS) team, they will help define ISA requirements and architecture for common components, infrastructure, and any other enabling science or technology needed to implement the physics integration scheme. For physics components specific to the ISA, the team will work with the component team to adapt or build, ensuring that development work conforms to FSP standards for coding, data exchange and documentation. The ISA teams will ensure regular testing of software components and integrated codes and will document and repair anomalies found in testing using provided infrastructure. Working with the software quality group, the ISAs will carry out a program of verification and uncertainty quantification. They will coordinate partnerships with experimental facilities to plan, execute and analyze validation experiments. They will ensure proper documentation of V&V studies using FSP standard methodologies. At some predetermined interval or when work has progressed to a satisfactory point, the ISA team will authorize the official “release” of FSP code. Finally, the ISA will work with the production support group to prepare code documentation and users’ guides and to provide ISA-specific expertise for user support.

3.1.2 Enabling Technology (ET) Teams

The division into ISA teams will not by itself produce an integrated project. In the FSP, the application teams are matrixed with personnel from various Enabling Technology (ET) teams that focus on the more global aspects of integrated code development. Each ET team will have a single technical lead and will be overseen by one of the two Deputy Directors.

To ensure that each ISA team’s effort is coordinated and integrated as much as possible with other ISA teams and the occasional auxiliary efforts of the FSP, staff from one or more Enabling Technology Groups will be embedded within each ISA. There will be four ET groups, each composed of 3-6 specialists in the following areas:

- Advanced Physics Components (PHYS)
- Software Integration and Support (SIS)
- Software Quality (SQ)
- Production Support (PROD)

Each ET group will have a lead whose responsibility will be to assure that, within the group’s jurisdiction, the ISA teams are taking maximum advantage of commonalities in physics components, infrastructure, quality assurance procedures, etc.

ET staff will receive day-to-day direction from and be directly responsible to their respective group lead. The ISA and ET teams are expected to complement each other’s work and to cooperate fully with one another. The FSP Directorate (see FSP Management Organization, (Section 5.6.1 below) will resolve any ongoing contention among the lines of responsibility.

ET Groups are expected to be permanent functions of the FSP. Their efforts will also be run as projects, either as part of an ISA project or some ET special effort not currently related to an ISA project.

The overall organization scheme is thus matrixed, with primary responsibility for delivering science solutions given to integrated application teams and a set of crosscutting teams with sufficient leadership and coherence to produce widely useful components, tools and infrastructure. Further details of the role of each are provided in the subsequent sections.

3.1.2.1 *Physics Components Group (PHYS)*

The role of the component team (PHYS) is defined in relation to other parts of FSP, particularly the integrated science application effort. The component team reports to the Deputy Director for Science and is responsible for well-defined, reusable physics modules that service more than one application. Each identified physics component will have at least one embedded component team member. The role of the PHYS technical lead is to ensure that the group members working within an application area are developing from a common code base, that code improvements targeting a given application driver are being built into that code base, and that the methods used in the physics components are verified. There are three specific roles for the component team as a whole.

First*,* the component team is a capability organization. It holds the technical capability in developing advanced physics components to be integrated into one or multiple science applications to address one or more science drivers. The component team has both regular members, who are appointed for the entire FSP execution phase, and collaborative members, who participate on the basis of individual component projects.

Second, the component team provides stewardship of the FSP component library by continuously standardizing and maintaining the suite of physics component codes for FSP science applications. It is through the stewardship of the component library that a common set of standards and best practices are introduced and applied to the FSP physics component development. The three primary activities are (i) publishing the component standard on data interfaces and documentation for component developers; (ii) performing acceptance test and review of a newly developed component; (iii) carrying out further improvement, maintenance, and regular testing of the component in the component library.

Third, the component team plans and executes, or manages the execution of, new component adaptation/development projects and the related enabling exploratory research and prototyping.

As an illustrative example, the FSP component team recruits leading subject experts in computational MHD and maintains a suite of standardized MHD component codes. It initiates and carries out development projects of adapting existing fusion MHD code into FSP and of prototyping new physics models and/or numerical algorithms in response to the evolving science driver needs.

3.1.2.2 *Software Integration and Support Group (SIS)*

SIS has the mission of providing the composition software for integrated computation, providing and/or supporting the software for job setup, data analytics, visualization and data management, as well as providing support for software development throughout the project. Software integration has been separated into two areas: 1) On-HPC (HPC = High-Performance Computer) integration, which at present targets the integration of physics components to run together on a High-Performance Computer; and 2) Task Composition (Off-HPC integration), which is concerned with all other integration and development and/or support of the associated modules that are needed to go from initial concept to final research result. For example, included in task composition is the development of specialized fusion plug-ins for reading data into visualization tools. Developer support includes providing and maintaining the collaboration systems, such as software repositories and communication lists, and implementing or assisting with software engineering issues, such as build and test systems and performance measurement.

The SIS group is responsible for the software deliverables and manages the repository. It works under the guidance of a Deputy Director for Code Architecture, who ensures that the overall vision and end goals of cohesion, testing, and release are met in a timely manner. The SIS group also includes cross-cutting teams that supply enabling computational technologies to the application teams. This is in strong analogy with the component team (PHYS) – the SIS role is to ensure minimal duplication of enabling computational technologies across application projects. This includes any and all aspects that make components conceptually related: similarity with respect to: (i) documentation approach; (ii) deployment approach; (iii) ability to read/write; (iv) output formats, use of mesh data types; and (v) visualization. Inter-component coupling is also an option but not necessary. The SIS group also manages versioning and release of the community FSP software, user contributions, and documentation.

3.1.2.3 *Software Quality Team (SQ)*

The Software Quality team (S) has responsibility for ensuring the reliability of the FSP software targeted for release to the community. The technical team leader receives oversight from the Deputy Director for Code Architecture and chairs the Software Quality Board, which is composed of members designated from the other areas. The SQ team leader coordinates the teams tasked with software quality management, which involves software quality assurance (SQA), verification, validation, and uncertainty quantification (UQ). All members of the team undertake the implementation of these activities. Associated tasks of the SQ team include: (i) developing standards for software development and testing; (ii) reviewing plans and progress on software quality activities across the entire FSP program; and (iii) organizing software reviews, prior to release. The SQ technical staff provides crosscutting tools and technologies such as testing systems. In addition, research into new techniques for verification and UQ falls under the auspices of the SQ team.

Each ISA team will have an identified contact from Software Quality. That individual will be responsible for coordinating the relevant activities involving SQA, verification, validation, and UQ. This coordinator will also serve on a Software Quality Review Board, chaired by a FSP-wide software quality manager. The SQ team reports to the Deputy Director for Code Architecture.

3.1.2.4 *Production Support Team (PROD)*

The Production Support Team (PROD) will support production versions of FSP software for research applications by end users. Such end users can come from within the FSP project (e.g., FSP SQ team members performing uncertainty quantification or FSP analysts performing physics validation studies) or from outside the FSP, e.g., experimental facilities or theory and computation base programs. End users will be expected to understand physics issues involved in running the FSP software but will not be required to have FSP code developer skills.

PROD will deploy production versions of FSP software as identified by the code development teams, on specific computational platforms supported by the Software Integration and Support Group (SIS). The team will make available the means for end users to prepare input, submit runs, monitor runs, and examine output. The team will make sure that run output are transferred to FSP data management facilities, with a record of the production code version and copies of all input data preserved.

PROD will provide user documentation for FSP code, provide user support, and, with backing of the Physics Components Group (PHYS) and other FSP code development teams, provide trouble shooting for failed runs and for problems in the supporting software for data preparation, job submission, monitoring, etc. The Production Computing Team will work with users as a key part of the process for continuous improvement of FSP software.

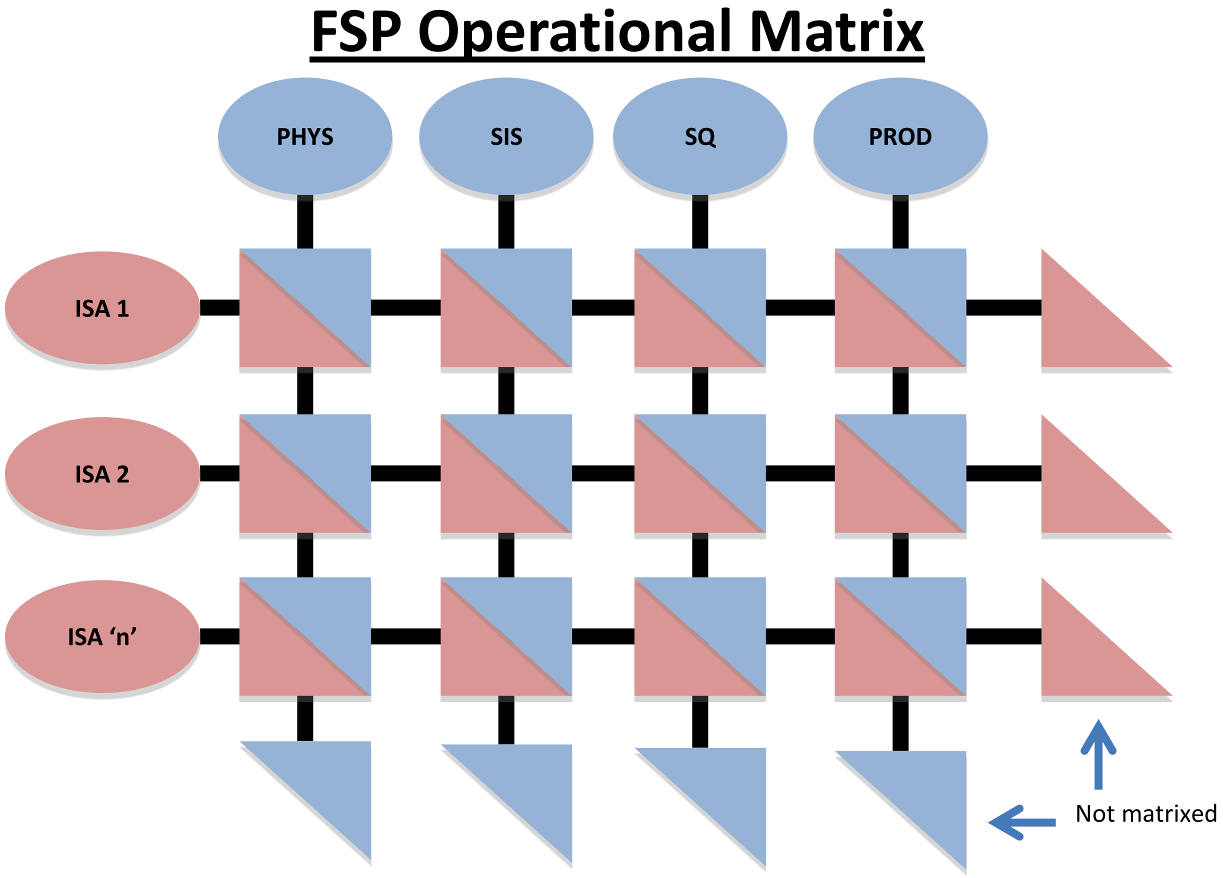

Figure 2: FSP Operational Matrix

This figure graphically represents the FSP operational matrix and the ways in which contributors to the FSP mission relate to one another. It is not an organization chart describing a management structure. In the diagram, each horizontal ISA row of various colored triangles represents a composition of members from both the ISA and ET teams who will work together to produce a particular simulation code release. (The oval shapes are simply row or column identifiers and do not represent people.) A vertical column represents the coordinated effort to leverage, as much as possible, common code solutions, infrastructure, quality control processes, etc. The unpaired triangles represent un-matrixed staff within a respective ISA or ET team. The blue triangles represent ET staff who might, for example, be assigned to work on a physics component that is deemed to be important for future but not currently active ISA efforts. A red triangle at the right end of an ISA might, for example, represent an analyst who is working with users to advance scientific discovery using codes developed in that ISA.

3.2 Application Teams (ISA)

3.2.1 Criteria, policies and procedures

In 2.3 above, each science driver is carefully documented in terms of how and why it was selected as well as a roadmap for developing the scientific capabilities with near, medium, and long-term deliverables. In addition to technical ISA leads, all of the ISA teams will be overseen by a Chief Scientist, whose responsibility is to drive them toward their technical goals and ensure proper impact on the broader scientific community, as well as recommend re-prioritization of allocations based on success/failure, new ideas from the community, and new directions from the Program Office.

3.2.2 Process for developing applications

ISA teams are responsible for following mature software engineering practices appropriate for the release of their ISA to the broader scientific community. The most critical aspects of this process are the management of code releases, documentation, automated testing, and repository management. From an organizational perspective, the key requirement for each ISA is their integration into and adoption of the larger FSP software infrastructure via 1) use of common modules developed/maintained within the physics components group, 2) use of FSP-specific framework standards and tools developed/maintained within the software integration group, and 3) adoption of common software practices via software integration. ISA teams are likely to begin their projects with existing codes based on existing technologies based on custom framework choices, data structures, and enabling computational tools (visualization, meshing, etc.). The applications will then evolve incrementally to conform to common standards defined by the FSP and enforced at the Deputy Director level. These "common standards" include both the use of common framework abstractions, common software process (repository, testing, release, etc.) common enabling tools, and the use of common physics components where appropriate. This process will follow the matrixed organizational approach as described above.

3.2.3 Interaction with other elements in FES program

The ISA teams form semi-autonomous sub-projects within the FSP. At the same time, project leadership must ensure the coherence and "economy of scale" of each ISA into the larger FSP software toolkit. This requires the appropriate balance between customization within each ISA and adoption of overall FSP framework components. To re-emphasize, the FSP structures this process as an incremental evolution – ISAs begin with existing technologies, and continuously evolve those technologies toward common FSP standards while simultaneously exercising and developing the codes. The ultimate goal is for all ISAs to be built off of a common FSP framework.

Each ISA team is responsible for end-to-end code development, testing, and release for its specific area. This process is carried out by using the services and benefiting from the economy of scale of the other crosscutting provider groups that make up the FSP structure. The Physics Component group provides componentized solvers for generic physics that is likely common to two or more ISAs; SIS provides, maintains, and develops enabling computational tools common to ISAs, and the user support group defines a common software process. Interaction of ISAs with each of these crosscutting groups is critical to the success of the FSP as an integrated project.

3.3 Physics Component Team (PHYS)

3.3.1 Organization

3.3.1.1 *Role of the physics component team*

The role of the component team is defined in relation to other parts of FSP, particularly the integrated science application effort. First, the component team is a capability organization. It holds the technical capability in developing advanced physics components to be integrated into one or multiple science applications to address one or more science drivers. The component team has both regular members, who are appointed for the entire FSP execution phase, and collaborative members, who participate on the basis of individual component projects.

Second, the component team provides stewardship of the FSP component library by continuously standardizing and maintaining the suite of physics component codes for FSP science applications.

Third, the component team plans and executes, or manages the execution of new component adaptation and/or development projects and the related enabling exploratory research and prototyping.

As an illustrative example, the FSP component team recruits leading subject experts in computational MHD, and maintains a suite of standardized MHD component codes. It initiates and carries out development projects of adapting existing fusion MHD code into FSP and prototyping new physics models and/or numerical algorithms in response to the evolving science driver needs.

3.3.1.2 *Working with the Integrated Science Application teams*

From the science drivers, the ISA lead consults with the component team to articulate component functionality requirements, as done in current FSP planning for the initial set of science drivers. During the FSP execution phase, the community is engaged to initiate science application integration effort in response to the annual FSP call for white papers. The component team is obliged to interact with the science application proposers on the

requirement/feasibility discussion, but only enters collaborative component code development for the winning science application proposers. In advance of an actual launched science application, the science application proposers can suggest new component development/adaptation. This is especially important for those key physics components where a significant gap between requirements and current capability exists. Advanced development of these components is a part of the risk mitigation strategy of FSP in the execution phase.
The component team gathers component requirements from all ongoing science application integration effort. These provide the basis for prioritization and scheduling on a team level. It is a case of economy of scale for the FSP to coordinate and consolidate the component codes development through a cross-cutting component team.

The component team gathers component requirements from all ongoing science application integration efforts. These provide the basis for prioritization and scheduling on a team level. It is a case of economy of scale for the FSP to coordinate and consolidate the development of components through a cross-cutting component team.

The component team enters an agreement on deliverable schedules with the science application integration effort, for individual science component project. A component point of contact (POC) is formally assigned for individual pieces of component deliverables. The POC may or may not be the primary developer of the aforementioned component code. The primary responsibilities of the POC is to coordinate the changing requirements and development schedule, and in the latter stage, serve as the primary support person to ensure that the component is properly integrated into the science applications.

3.3.1.3 *Working with the Software Integration and Support (SIS) team*

The SIS is responsible for the framework/infrastructure/integration aspects of development. It is a parallel capability organization that provides the stewardship of the FSP framework, namely the software infrastructure for component integration. FSP science drivers demand integrated physics models that are multi-physics and multi-scale. Component factorization must be performed in the context of a specific coupling scheme. The component team works with the framework/integration team to articulate the component functionality requirements in facilitating physics coupling of the components. In terms of software engineering, the component team works with the framework/integration team to specify component data communication and naming convention. It is a key FSP objective to achieve standardization toward common FSP data structures. The external data representation, through the collaborative work with the FSP data management effort, will follow an FSP common standard at the inception of the FSP execution phase. The internal data structures, due to the large amount of diverse FSP component adapted in the early state, will evolve significantly over time toward increased standardization.

Collaboration between the component and framework teams also aims at developing and deploying a common code development technology and standard, applying best practices in quality assurance, and including a verification testing suite.

3.3.2 Criteria, policies and procedures

3.3.2.1 *Component selection criteria and process*

3.3.2.1.1 Component Strategy

The physics component needs in FSP are primarily met through (1) adaptation/upgrade of existing simulation capabilities and (2) development of new capabilities based on existing exploratory research. There can also be (3) an exploratory research arm of FSP component development, which is likely a highly leveraged activity requiring close coordination with the base and SciDAC programs. Its scope and schedule are subject to the funding ramp up profile.

The first part of the FSP component strategy is to establish the FSP component specification. This is a joint exercise between component and science driver efforts, and with extensive collaboration with the framework/integration effort. The science drivers guide the overall process, which themselves are determined by the program direction of the fusion energy program. For any chosen science driver, a technical analysis is performed to specify the required computational physics capabilities. An essential criterion for choosing FSP science drivers is the requirement that the underlying complex physics demands multi-physics and multiscale coupling. The integrated science application software must then require integration of multiple physics components under an FSP framework. The component factorization is driven by both the underlying physics and the anticipated framework/integration needs. It is within this context that the technical specification of a FSP physics component is articulated.

Individual FSP physics component specification provides the technical basis for a gaps analysis which leads to a tailored FSP component development plan. The gaps are measured with respect to existing simulation capabilities in the base and SciDAC programs. There are two courses of action depending on the degrees of the perceived gaps. For small gaps, the most capable FSP candidate component is selected to be adapted and upgraded into FSP deployment. The required improvements and resources are evaluated and become an action item of the component execution plan. If significant gaps exist beyond current capabilities, the FSP will look into existing exploratory research. If an innovative approach is assessed to have passed the proof of principle stage, the FSP will take advantage of this opportunity to launch a new component R&D initiative. In the case that neither the existing code base nor the exploratory research pool produces an credible path for meeting a FSP component need, FSP will contemplate launching its only exploratory research. These constitute the three primary paths of component R&D for the FSP.

It must be emphasized that the component execution plan thus developed in the FSP planning stage will undergo continuous update and revision during the actual execution phase.

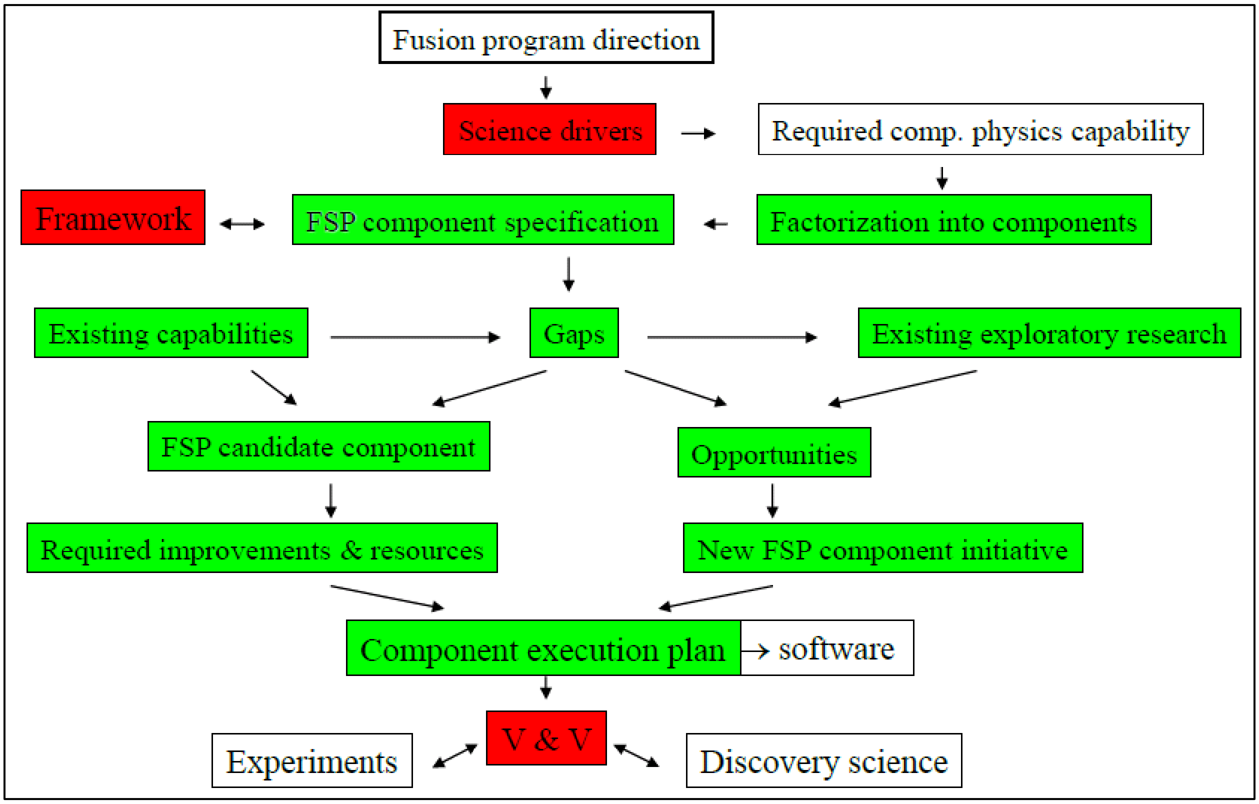

Figure 3: Program interaction graph

3.3.2.1.2 Community Engagement Process

Community engagement plays a prominent and essential role in establishing component needs and readiness. It provides the basis for the FSP prioritization on component development. In the FSP component planning process, two independent, community-based planning activities were carried out. The first was a broad solicitation on the currently available simulation codes in all topical fusion science areas that can potentially serve as a FSP physics component candidate. This was done through the FSP Physics Component Questionnaire, which was distributed to the entire fusion science community via direct email and web publication. The questions were organized into two parts. The first set applied to every component candidate concerning the general physics formulation, numerical methods, computing performance, and software practices. The second set of questions were targeted toward different topical fusion science areas. The response to the component questionnaire from the fusion science community was overwhelming. The detailed response s were collected, analyzed, and posted on the component wiki site.

The second community planning exercise was a targeted solicitation of specific designs on components and coupling scheme to address individual or multiple science drivers. This was done in two stages. The first stage was an open solicitation of design proposals and ideas, with little restriction other than a focus on the science drivers, to the entire fusion SciDAC and proto-FSP community, plus three large code project teams funded by the base programs. The community input provided the basis for a focused design study group with the task of reaching a consensus and/or prioritized component/coupling design to cover each science driver. This paralleled an FSP effort on gathering the overarching framework/integration requirements and software infrastructure design.

With concrete design proposals documented on the FSP wiki, the different parts of the FSP planning were consolidated into the integrated planning teams (IPT) which were made of fusion domain experts on both theory/simulation and experiments, applied mathematician, and computer scientists, and target each of the six science drivers. The IPTs were tasked with not only consolidating and refining the technical requirements, but also specifying the resource requirement and the schedule of delivery. The final IPT reports formed the basis for an FSP program-wide prioritization and scheduling, which results in this FSP execution plan. In the execution phase, the component plan, as other parts of FSP, will undergo continuous revision taking into account the community input, for example, by an annual FSP research forum.

3.3.2.1.3 Establishing component R&D plan across science drivers

The necessity of a unqualified FSP component effort is primarily driven by the fact that there are substantial overlap between the component requirements arising from different integrated science applications targeted at individual science drivers. To realize the economy of scale, component development is planned and carried out across the entire range of the FSP science drivers. There is also the practical consideration that although not all FSP science drivers will have integrated science application effort launched at the early stage of the FSP execution phase, advanced development of key physics component can be initiated by the FSP component team. This will substantially lower the risk of a future ISA for a particularly challenging fusion science objective.

The process for integrating an FSP program-wide component development plan is as follows. The first step is to establish component specification for individual science drivers. This is carried out by first identifying the integrated physics models which are required to resolve the targeted science driver. A component factorization, along with the coupling with the coupling scheme, can then be performed on the integrated physics model. This will lead to a set of component requirements in physics models, computational algorithms, software implementation, all in the context of a specific coupling/integration scheme which facilitates data exchange and time synchronization between different physics components. These concrete requirements form the basis for an analysis on component readiness and gaps.

The second step projects an FSP program-wide perspective on physics component requirements and establishes an integrated research and development plan across the FSP science drivers. The required components are grouped in three categories. The first group is the FSP common components, required by all science drivers. The second group consists of physics components that are shared by multiple science drivers. The third group is the component uniquely required by one science driver. The entire first group and some of those from second group which satisfy the readiness requirement, directly feed into the FSP software infrastructure. They are among the first to be adapted from the base/SciDAC programs. The resource requirements and the deliverable schedule form the core of the initial FSP component development plan. For those components in the second the third group in which a significant gap exists between current capability and the FSP requirement, a set of FSP component common challenges are identified and then fed into the near, mid, and long term research and development strategy. With the current FSP execution strategy, the science driver-unique components, if they satisfy the readiness requirement, their adaptation into the FSP will be carried out by the integrated science application team.

3.3.2.2 *Process for developing components*

Component development is expected to permeate through the entire FSP execution phase. This is the result of science application needs in response to the evolving science drivers.

3.3.2.2.1 Supplying components to selected science application projects

For a funded science application project, the first step after establishing the component requirements is to check if a suitable component is available in the FSP component library. If the answer is yes, a component customization project is initiated. Otherwise, a credible prototype is search in the existing base/SciDAC programs. Once it is identified, a component adaptation/development project is initiative. A POC is assigned in both cases and supports the component integration into the science application. In the unlikelihood that no credible prototype exists, it is fair to assess that there was an error in judgment that the science application project under consideration is prematurely stated.

3.3.2.2.2 Developing components for prospective science application projects

A critical aspect in FSP risk mitigation is the research and development of key physics component well in advance of the integrated science application that requires them. This process of component development is annually visited by the community FSP solicitation to address the science drivers. Members of the fusion research community, either inside or outside the formal organization of the FSP, can propose components to address the science drivers. In the event that a required key component does not exist as a mature code or prototype in the base/SciDAC/FSP programs, the component team will assess if the proposer or anyone else has an innovative and promising approach for such a component. If the answer is yes, the FSP component team can enter an agreement with the proposer on collaborative new component prototyping and development project. Once the new component is developed, it will be committed to the component library, which makes it available for supporting relevant science application projects and project proposals.

3.3.2.3 *Interaction with other elements in FES program*

3.3.2.3.1 Base theory program:

The FSP component development relies on the base theory program for developing the theoretical basis of the physics models to be implemented in the component code. These include the improved closure for the fluid moment equations, a more complete set of gyrokinetic equations that apply to the tokamak edge, and a mathematical rigorous formulation that couples neoclassical and turbulent transport to the quasi-static evolution of the three dimensional magnetic field, etc. The research needs identified in the FSP program will feed into the base program, motivating new solicitation to address them. The base theory program also provides the exploratory research, which upon a proof of principle demonstration, can lead to a more complete component development under FSP.

3.3.2.3.2 SciDAC program

The FSP component development relies on the SciDAC program for component candidates to be adapted into the FSP component library for use in one or multiple science applications.

3.4 Software Integration & Support (SIS)

3.4.1 Organization

3.4.1.1 *Role*

The software integration and support team is tasked with providing common software tools and composition software for both the On-HPC and Off-HPC situations as well providing support for software development throughout the Fusion Simulation Program. On-HPC integration includes the definition of the interfaces through which components can be incorporated in a multi-physics simulation, the development of the reference implementations that illustrate the composition of physics suitable for work on a High-Performance Computer, the development of On-HPC capable software for use throughout the FSP for general purposes, such as I/O. Off-HPC integration includes the development of software for job setup, job submission, data analysis, visualization, and the composition of these tasks, thus going from investigation concept to the recording of generated knowledge. Developer support includes setup and administration of collaboration mechanisms, maintaining regression and unit tests, and setup and maintenance of code repositories, application, creation, and maintenance of build and package management systems, and so forth. More details in all of the will be presented in Section 4.3.

3.4.1.2 *Working with the Integration Science Application teams*

A member of SIS will be embedded in each of the ISAs. That member will be tasked with (1) summarizing the development difficulties faced by that ISA, (2) participating in SIS meetings to look for issues that could be resolved by providing a common software infrastructure, (3) participating in requirements meetings to determine precisely what the common software should be capable of, (4) ensuring that any SIS developed software is meeting the needs of the ISAs, and (5) assist in refactoring the ISAs to use any newly developed common software. For an example, we have already determined that (1) and (2) in the definition phase have identified the need for common metadata structures in files and a common application programming interface (API) for writing those structures. Thus, in the early parts of the project, a sub-team with representation from the ISAs, SIS, and the components teams will set the requirements for any software. Then, the embedded SIS team members will ensure that the newly developed I/O metadata libraries can be used in the ISAs by trying them out within the ISA code base. Any problems will some back to the SIS team for resolution. Ultimately, this will lead to a refactoring of the ISA to use the new metadata libraries.

3.4.1.3 *Working with the Component Team*

The process for working with the components team will mirror that of working with the ISAs, but with some similar and different foci. For developing common enabling libraries, the procedure will be as above, but for developing the integration software, there will need to be APIs developed by both groups, as the component factorization and the component APIs are interdependent. Depending on the component factorization, different API methods will be needed.

For working with the ISA teams and the Components teams, SIS will also have to have continuous interaction for Developer Support issues. The SIS team will adapt build systems to the ISA software, will maintain the software repositories, and will provide the collaboration mechanisms. Continuous interaction will be needed to ensure that SIS is meeting the needs of those groups.

3.4.1.4 *Working with the SQ Team*

In working with the SQ team, the emphasis will be on Off-HPC integration, but the procedures will largely be the same. Joint committees will establish the requirements of the analysts for carrying out computations and analyzing the results. This includes aspects of data management, such as down-selection of data for archival storage in the data management system. It also includes the development of software for obtaining data in the data management system. SIS developed software will be continuously reviewed by the SQ team to ensure that it is meeting their requirements, which one expects to be refined throughout the process.

3.4.1.5 *Working with the Production Computing Team*

The SIS team will need to work with the Production Computing (PROD) in several ways. One joint effort will be in establishing the procedures for software releases. Initially, there will be a study of the release mechanisms in other areas that have similarly complex projects, such as the climate community. Production Computing input will be critical in establishing a release process as PROD will be doing the actual deployments. Similarly, SIS and PROD will have to establish the border between those user issues handled by PROD and those handled by SIS. There is an expectation that PROD will be the recipient of user issues, acting as a triage agent. Those issues that end up being SIS issues solvable with only deeper knowledge of the integration software will be handed off to SIS. (PROD will have similar relationships with the ISAs and the Components.)

3.4.2 Policies and procedures

Software Integration and Support will have to make a number of decisions in terms of priorities and activities. It is expected that demands on this team will quickly outstrip the resources. Clearly, a balance will need to be maintained. If the usability of the software (as required by the SQ team) is not there, science cannot be done. If it cannot be deployed, science cannot be done. At the same time, the basic capabilities (developed in partnership with the ISAs and the Component teams) must exist, or, regardless of deployment, the software will not be used. Finally, there needs to be some focus on research topics to prepare for future needs.

SIS will provide regular reports to its management and advisory committees on how it is maintaining a proper balance among its various constituencies. This will also be reported to any advisory committees. Recommendations form those entities will be requested in writing, and SIS will modify its priorities and plans accordingly.

3.4.3 Integration software development

The development plans for SIS are discussed in more detail on Section 4.3. In summary, we intend to start with the lowest hanging fruit, such as metadata and API definitions, then move to reference framework implementations, API-adherent libraries for, e.g., I/O, and so forth. There will be a continuous interaction with the ISAs and Component Team to assist in refactoring to use the newly developing common software base so as to gain the economies of scale that are possible within an FSP.

The SIS team will adhere to two important methodologies of software development: test driven development and continuous release. Test driven development means that the test is written before the software is written. There are multiple ways to do this. For regression tests, this means writing the input file that would create the desired simulation. Initially, the application (whether the On-HPC application or some executable corresponding to an element of the workflow) fails on the input. As the capability is developed, the application eventually executes and provides a result. Then comes the process of verifying that result, after which the test is incorporated into the regression testing system. With this process, the software can be released at nearly any time with knowledge about its capability limits. This assists in providing users nearly continuous releases.

3.4.4 Interaction with FES program elements

The SIS team will be responsive to the advisory committees of the FSP. In addition, it will be working with the FES program elements primarily through Off-HPC integration, as this is the user-facing part of SIS. For this, SIS will have to be incorporating feedback from users (1) prior to the development of enabling software to get the first draft of requirements, (2) during the development of enabling software through continuous release and comment, and (3) after the release of software for refinements and additions to be planned for the next release.

3.5 Software Quality (SQ)

For the purposes of the FSP, the concept of software quality management is extended beyond the typical software engineering activities of software quality assurance (SQA) to also include simulation software verification, experimental validation, and uncertainty quantification. Thus, the concept of software quality is a holistic concept, including the software implementation, the mathematical approximations, and the physical modeling.

The mission of the FSP verification activity is to produce correct codes with quantified error estimates in simulation results. The associated goals are (i) to find and eliminate programming errors in scientific software based on the discretization of mathematical models; (ii) to demonstrate and document convergence to reference solutions for these mathematical models at the expected rates; and (iii) to develop and to provide the necessary tools and techniques required to provide reliable *a posterior* error estimates for simulation components and integrated simulations.

The mission of the FSP experimental validation activity is to assess and improve physical and computational models by systematic, quantitative comparisons with experimental measurements. The associated goals are (i) to develop and provide the necessary tools and documentation to allow FSP users/customers/stakeholders to determine what level of confidence they will give to predictions made by an FSP simulation and (ii) to provide clear assessments of model and component physical fidelity to help guide their refinement and improvement.

The mission of the FSP UQ activity is to improve the use of simulation results for scientific inference through the quantification of uncertainties. The associated goals are (i) to identify and reduce sources of uncertainty in simulation results; (ii) to develop practical procedures for the routine quantification on uncertainty and integrate UQ into the standard practice of FSP simulation; and (iii) to improve the rigor of validation activities and facilitate more productive collaborations between theory, computation, and experiment.

An overview of these activities and their relationships are provided in Appendix D:.

3.5.1 Organization

3.5.1.1 *Role*

The Software Quality Team has responsibility for ensuring the reliability of the FSP software targeted for release to the community. The technical team leader of the SQ Team receives oversight from the Deputy Director for Code Architecture and chairs the Software Quality Board, which is composed of members designated from the ISAs and other ET Teams. The SQ Team Leader coordinates software quality management activities, in particular, the assignment of members of the SQ Team according to the needs of ISA and Physics Component Teams.

The SQ Team is responsible for the definition and periodic review of standards for software development and testing throughout the FSP. Members of the SQ Team will work in an integrated, collaborative fashion with ISA and Physics Components teams to develop SQ plans (software testing, verification, and validation plans); will participate in the execution of these plans, especially in the analysis of results; and will monitor progress and work with ISA and Physics Components teams to revise plans as necessary. Similarly, SQ Team members will be involved in the development, execution, and review of UQ plans with the ISA Teams. Prior to any significant release, the SQ Team will organize software reviews. The SQ Team is also tasked with the identification, development (as needed), and deployment of crosscutting tools and technologies for SQ, such as testing harnesses, UQ workflow software, and synthetic diagnostics. Finally, research into new techniques for verification and UQ falls under the auspices of the SQ team.

3.5.1.2 *Working with the Integration Science Application Teams*

The type of interactions between the SQ Team and ISAs will differ slightly depending on the specific SQ sub-area.

Members of the SQ team will collaborate with members of each ISA to develop the plans in the testing and verification, will advise in the execution, and will actively participate in the analysis and review of results. Because testing (e.g., unit testing, regression testing, etc.) and verification will be ubiquitous throughout the daily activities of the ISAs, the ISA Team members will have primary responsibility for the standard execution of testing and verification.

Validation analysts from the SQ Team will also be embedded into each ISA Team but, while working closely with other team members, they will safeguard the independence of the code and experimental results as well as the objectivity of the validation metric findings. The validation analysts will coordinate partnerships between experimentalists, modelers, and theorists, both within the ISA and external to the FSP, in order to identify, conduct, document, and refine validation test case studies. Such studies may include the comparison of new and improved models/theories against current FSP capabilities. As part of an ISA, SQ Analysts will partner with scientists to pursue scientific discoveries enabled by new FSP capabilities. Finally, while the responsibility for synthetic diagnostic development falls to the SQ Team, validation analysts will work closely with theorists and modelers in the ISAs on the definition and design of these diagnostics.

The development of effective UQ methodologies for ISA problems will require a close collaboration between physicists, applied mathematicians, statisticians, and computer scientists. Progress will be made by iterative application and evaluation of UQ techniques. Thus, for UQ, experts from the SQ team will be embedded into the ISA teams in order to help design, execute, and analyze UQ studies.

Each ISA team will identify a contact for Software Quality. This individual will be responsible for coordinating the relevant activities involving software QA, testing, verification and validation. This coordinator will also serve on a Software Quality Review Board in order to provide feedback to the SQ Team.

3.5.1.3 *Working with the Physics Components Team*

The type of interactions between the SQ Team and the Physics Components team mirror those between the SQ Team and ISAs, with some modifications. While the testing and verification interactions will be the same, UQ analysis will not be applied to components in isolation; rather UQ analysis will be done only within the context of ISAs. Validation processes will occur as for the ISAs to the extent that experimental data can be obtained to test components in isolation.

The Physics Components Team will identify a contact for Software Quality. This individual will be responsible for coordinating the relevant activities involving software QA, testing, verification and validation. This coordinator will also serve on a Software Quality Review Board in order to provide feedback to the SQ Team.

3.5.1.4 *Working with the Software Integration & Support Team*

The Software Integration & Support Team will work with the SQ Team to adapt and develop program-wide infrastructure to support testing, verification, and UQ. These technologies include, but are not limited to testing harnesses, workflow tools that enable scripted study generation and execution, and libraries that provide verification and statistical analysis tools. Several of such workflow tools already exist, such as the DAKOTA/UQ [3], PSUADE [4], LLNL's UQ Pipeline, and the CalTech PSAAP Center's UQ Pipeline [5] built on Pyre [6]. The FSP should leverage these existing technologies and adapt them for FSP use.

In addition, the development of generic synthetic diagnostics will require standardization of an API and/or data structures, a task identified by the SIS Team. The SQ Team will be an important customer of these data interfaces and will work with the SIS Team in their definition and refinement.

3.5.2 Policies and Procedures

3.5.2.1 *Policies for Verification*

There are no verification procedures that will work for all applications. Verification procedures will in general be easier to apply to individual components than to integrated, multi-physics applications. The FSP will thus follow a tiered approach to the deployment of verification techniques and differentiate between policies and procedures for component development and for integrated application development. Definitions of verification terminology and a best-practices workflow for verification are provided in Appendix D:.

3.5.2.1.1 Component development

The focus of verification activities in component development will be on documentation, code verification through order verification, and calculation verification. It will be the policy of the FSP that component development projects will have the freedom to define and select appropriate test problems and methodologies, but component projects should attempt to apply the most rigorous techniques available.

Documentation must include the governing mathematical model; the discrete approximate model; a priori error estimates; reference solutions; quantities of interest relevant to major use cases; all numerical and model parameters; and all test results. Documentation also includes any scripts and input files used to generate the simulation results. We anticipate that the documentation will be saved and distributed within the framework established by the integrated data management policy outlined in Section 4.3.6.

It is recognized that rigorous code verification will be a challenge due to the complexity of the problem and that benchmarking will continue to be a useful confidence-building tool. Nevertheless, the FSP must be diligent in its

---

[3] http://dakota.sandia.gov

[4] https://computation.llnl.gov/casc/uncertainty_quantification

[5] http://www.psaap.caltech.edu/meetings/sitevisitoct10/posters/McKerns_PSAAP_poster.pdf

[6] http://danse.us/trac/pyre

efforts to move to more rigorous code verification practices. It will be preferred that code verification will be done rigorously through order verification. A suite of reference problems and their solutions will be created that can include known analytic, asymptotic, or manufactured solutions. The latter type of solutions will be preferred, since these can be designed to exercise more complicated couplings in the code. In addition, surrogate error tests based on intrinsic properties (invariants) of the discrete solution, e.g., conservation, where appropriate, should also be included in the component verification test suite. Well-defined code verification tests will be integrated into automated regression testing suites.

In the first few years of the FSP, calculation verification is expected to rely heavily on grid refinement and three-point Richardson extrapolation applied to sequences of four or more grids for *a posteriori* error estimation. A suite of reference problems (without solutions) that demonstrate important behaviors of the component will be created; important quantities of interest for these reference problems will also be identified and defined. Component teams will be encouraged to incorporate more other, potentially more robust *a posteriori* error estimators, such as adjoint or error transport methods. Benchmarking (code-to-code comparison) and self-convergence (Richardson extrapolation using a fine-grid solution in place of the exact solution) will be used to build confidence, even if they are inadequate for proper calculation verification because they do not provide error estimates. In later years, the inclusion of *a posteriori* error estimators into new component development will be a goal.

3.5.2.1.2 Integrated application development

For integrated applications, rigorous verification becomes more difficult for several reasons. Fewer exact or approximate solutions are known, and manufactured solutions are more difficult to devise and implement. Robust *a posteriori* error estimators are more difficult to develop for integrated models. In addition, there are open research issues on how to develop composite error estimates from a configuration of components, which could possibly be dynamic.

As with component development efforts, model documentation (for the integrated application) and code and calculation verification, to the extent that these are possible, will be emphasized. Components to be coupled in an integrated application must have already undergone thorough code and calculation verification as described above. Benchmarking and self-convergence will, in the early years of the FSP, play a key role in the absence of more rigorous code and calculation verification capabilities. In later years, it will be a goal of the FSP to include more robust *a posteriori* error estimators in integrated components for more rigorous code and calculation verification.

3.5.2.2 *Policies for Validation*

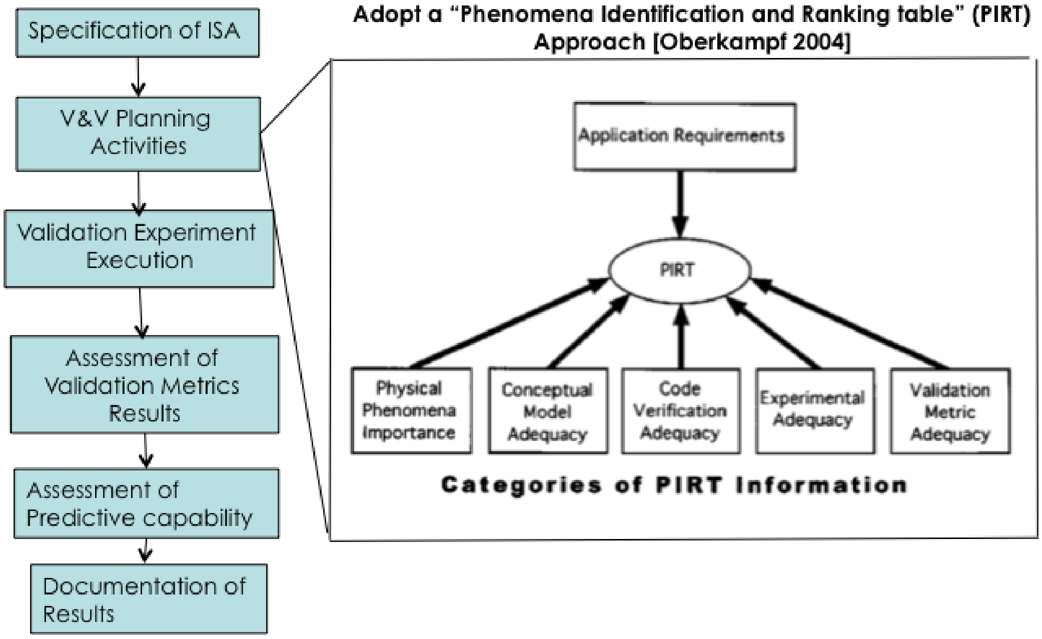

Figure 4: Proposed workflow best practices for the FSP experimental validation

The FSP validation efforts will follow the functional workflow for the "verification-validation-prediction" shown in Figure 4. The Verification and Validation (V&V) Planning Activities box is a crucial step in the validation workflow. The FSP will adopt the Phenomena Identification and Ranking Table (PIRT) approach identified in Oberkampf *et al* [7] to define these V&V Planning Activities.

The Physical Phenomena of Importance are the essential phenomena that must be input to, and described by, the application code. Each phenomenon should be prioritized relative to others, and an explanation of this prioritization included in the PIRT document. In the course of answering these questions, a validation hierarchy will be constructed to help identify a range of experiments, the possible separation of coupled physics, and the levels of complexity relevant to the modeling of the ISA in question. The critical physics in this hierarchy must be extended to the finest granularity required for validation and understanding. An example of such a validation hierarchy for pedestal physics can be found in Appendix D:.

Conceptual model and code adequacy refer to the process of determining if the codes and the physical models implemented therein are adequate and ready for validation. This question entails several aspects. The first is conceptual model adequacy: Are the underlying conceptual models (e.g., ideal MHD, Branginskii equations, δf gyrokinetics etc.) adequate for describing the physical phenomena of importance? Ideally, as part of this process, the most uncertain part of a model should be identified. The second is code verification adequacy: Have the components and frameworks, which represent the numerical implementation of the conceptual models of interest, undergone adequate verification testing? Before validation, each application must complete a sufficiently thorough process of code verification based on experimentally meaningful metrics.

Existing experimental data are unlikely to be sufficient to validate FSP components; new experiments will be required. The validation hierarchy formed in the consideration of critical physics will be used to identify a range of new experiments, based on possible separation of coupled physics and level of complexity. When designing these experiments, the following recommendations should be followed:

- The experiments should be designed to test the most critical or uncertain parts of the physical models.

---

[7] W.L. Oberkampf, T.G. Trucano, C. Hirsch, "Verification, validation, and predictive capability in computational engineering and physics," *Applied Mechanics Reviews* **57** 345 (2004)

- The experiments should help quantify regions of parameter space for which conceptual models are acceptable for the application purposes.
- The experiments should allow precise and conclusive comparisons of calculations with experimental data.
- Required measurements (and diagnostics) must be determined, including the required accuracy and resolution.
- Simulations and experiments should be carried out as independently as possible to minimize prejudicing the outcome.
- Synthetic diagnostics must be developed to allow for direct comparisons between code results and measurements.
- Numerical errors should be estimated through calculation verification and uncertainties should be identified if not quantified.
- The entire process and results, including physics assumption, data reduction techniques, and error analysis must be documented.

An essential component of the validation process is the development and application of metrics for quantifying model fidelity. Metrics are fundamentally algorithms for quantifying agreement between a model and observations. They can be as simple as a chi-squared test of a prediction to a set of measurements, or highly complex and nonlinear functions that incorporate many model predictions and experimental data points. Validation metrics should include the following properties:

- The metric should include an explicit estimate of numerical error in the model calculation or exclusion of such error if it can reasonably be shown to be small.
- The metric should be a quantitative evaluation of the predictive accuracy of the system response quantity (SQR) of interest.
- The metric should incorporate, or at least explicitly include, an estimate of the errors in the measurement(s) the model is to be compared against.
- The metric itself should exclude any indications, either explicit or implicit, of the level of adequacy in agreement between computational and experimental results.

In many cases, the quantities that are of interest and output by a computational or theoretical model are related, but not directly equivalent to what is measured experimentally. Accounting for these differences in model validation is done via diagnostic and comparison-specific models termed synthetic diagnostics[8]. The complexity of a synthetic diagnostic model can be such that it requires verification and validation in and of itself.

The breadth of validated physics that the FSP aims to deliver will therefore require a substantial suite of synthetic diagnostics. These should be implemented as standalone components, independent of specific physics components (e.g., a given microturbulence or MHD code) and possibly shared across ISAs, and subject to the same SQA, verification and validation processes as other components. Fortunately, a number of existing synthetic diagnostics (see Appendix D:) have already been developed that will be used as starting points for the FSP synthetic diagnostic components. The key initial challenge for the FSP will be to translate these various existing models that have generally been implemented by individual researchers for use with single codes into robust, standalone, code-independent components. Prior to any validation activity, required synthetic diagnostics must be determined, developed, and verified.

3.5.2.3 *Policies for Uncertainty Quantification*

There is no single agreed upon way to apply UQ analysis to multi-physics simulations. Resource constraints dictate that the FSP will focus on the forward sensitivity problem that produces uncertainty estimates for

---

[8] R. V. Bravenec and W. M. Nevins, Rev. Sci. Instrum **77** 015101 (2006)

computed results due to uncertain inputs because that is the primary question for the validation use case. However, there other applications that will be of use in the FSP, such as the selection of the most likely model from a set of experimental data and the calibration of models given experimental data with known uncertainty. More information on UQ analysis and a best-practices workflow for UQ are provided in Appendix D:.

The SQ team will continually evaluate existing techniques and UQ infrastructures and periodically recommend best practices and tools for the FSP. The FSP, at least for the first few years of the program, will not restrict consideration to a single UQ approach, but will instead rely on UQ experience developed through iterative collaborative application of techniques to problems of increasing difficulty within the ISAs. Infrastructure to support the large number of simulations common in UQ analysis and the subsequent analysis will be adapted from any one of the number of existing tools (see Appendix D:).

We note that validation has and can be done in the absence of UQ analysis for the simulation output; the addition of UQ makes the validation stronger, more quantitative, and provides direction for addition improvements. Therefore, while it would be preferable to have UQ analysis involved in every validation study, given the state of UQ in fusion simulation and the challenges of the fusion problem, it is unrealistic to expect to have UQ fully integrated in the short term. In this sense, the UQ effort is in the near term more exploratory research than application. As successful UQ techniques are identified and developed, previous validation studies can be re-analyzed and re-interpreted.

Documentation is very important in UQ. A careful accounting of sources of uncertainty must be identified and documented so that appropriate models of the uncertainty can be constructed. There are many decision points and assumptions in the design of a UQ study that must likewise be recorded. All of the relevant parameters and inputs need to be identified, and this information should be recorded. Code development activities must also support code documentation activities, so that this information should be readily available. Documentation also includes any scripts and input files used to generate the simulation results as well as information about the simulations (provenance). We anticipate that the documentation will be saved and distributed within the framework established by the integrated data management policy outlined in Section 4.3.6.

3.5.3 Interaction with other elements in FES program

The success of FSP validation hinges on building a true partnership between FSP and the experimental facilities. Both have to recognize there are significant benefits towards advancing the fusion energy goals by sharing resources and making a commitment to put experimental validation at a high priority programmatically. At https://ice.txcorp.com/trac/2011_FspDefinitionWorkshop can be found a document ("Principles for Collaboration on Major Experiments" describing how the partnership should be implemented which has been prepared based on ideas drawn from existing collaboration agreements used by the three major fusion experimental facilities, and their governing and planning processes. The facility management has reviewed this and provided constructive feedback. The document outlines:

- General principles for collaboration and IP sharing with major facilities
- Interactions with facilities on planning
- Roles for the FSP and the experimental team
- Lessons learned from experimental facilities for FSP in terms of organizing its own research efforts.

It is envisioned that the FSP will provide code suites and computer time subject to allocation process; aid in understanding code capabilities and limitations; dedicated analysts; and consideration of code capability improvements, based on needs of the experiments. Experimental teams will provide facility run time, subject to local planning processes; access to data; support for diagnostic data analysis; and consideration of facility upgrades, based on the needs of the simulation program. The FSP and experimental teams will collaborate on setting priorities; run planning; experiment analysis and interpretation; development of synthetic diagnostics; physics interpretation; and the preparation, presentation and publication of results. These interactions will be facilitated by the validation analysts within the SQ Team.

3.5.4 Research Opportunities

There are many open issues in the application of calculation verification and UQ techniques to multi-scale, multi-physics simulation codes. Possible research areas include:

- Robust *a posteriori* error estimators
- Combination of error estimates between coupled components
- Determination of errors due to coupling
- Error estimation in the presence of models for unresolved physics
- Error estimation for multi-scale problems
- Error estimation for solution-driven model changes
- Methods to ameliorate the curse of dimensionality
- Methods to propagate uncertainties through coupled components
- Dealing with instability, chaotic behavior, or lack of smoothness of response functions
- Efficient incorporation of deterministic error estimates with stochastic uncertainties
- UQ for multi-scale problems

Resources will be limited within the FSP, so the FSP will need to engage external research and development efforts (ASRC, ASC, NSF, etc.) for new verification and UQ tools and techniques.

3.6 Production Computing (PROD)

The Fusion Simulation Program production computing team will be responsible for delivery of production computing and user support services based on software developed by FSP and FSP allied efforts. The set of services provided will be as determined by FSP project leadership based on input from research user groups both within FSP (e.g., the FSP funded analysts and V&V activities) and beyond FSP in the broader Fusion research community: ITER, experimental projects, SciDAC, theory and computation.

The focus of activity of the production computing teams is to support services for end users: researchers who run the software in pursuit of program research goals. In general, end users are not expected to be code developers who port and build their own software from source; the production teams will provide environments with software already built, tested, and ready to use. These teams will provide user documentation – instructions for use, set up of input, interpretation of output, etc.

The production teams will provide trouble-shooting support for when computational services encounter problems. As these services involve deployment of interacting research codes with complicated data and software dependencies, it is expected that the trouble shooting activity will be very important. It is well known from experience: no matter how robust and well tested the individual software components of FSP, the integrated workflows can be expected to encounter problems in a fair fraction of research applications. Therefore, a significant trouble shooting activity is to be anticipated.

3.6.1 Organization

A member of the FSP management team will be responsible for production computing operations. This leader will supervise separate deployment and support activities set up for each computational service as directed by FSP project leadership. Efforts will be organized around deployment of services at a range of computational facilities ranging in power from cluster level computing (e.g., at one or more major fusion labs) to capacity supercomputing (e.g., at NERSC) all the way up to exploitation of leadership class facilities where feasible.

3.6.2 Criteria, policies and procedures

FSP leadership will select and prioritize computational services for production deployment based on the following criteria:

- User demand: determination that a sufficient user base exists to warrant production deployment with reasonable expectation of wide scale use of service;
- Readiness of software: determination that underlying scientific software is sufficiently stable and mature to sustain wide scale use.
- Readiness of hardware: determination that appropriate platforms and data and network resources are available and sufficiently stable to support reliable service execution at scale.

As a decision of deployment involves commitment of significant manpower and computational resources, each decision should be based on a consensus of user groups, software development teams, and computational resource managers. User demand is the most important criterion; in cases where user demand is not met by readiness of software and/or computational resources, other FSP activities may be required to bridge the gap.

3.6.3 Orientation to User Support

The primary mission of the production team is to maximize user productivity by making tested, production versions of the FSP software available to research users in as efficient and effective a manner as possible. Research users are of course expected to have a broad understanding of their scientific applications and research goals, but, they are not expected to have to function as code developers or code installers in order to have access to the software. They should not (for example) be required to read Fortran or C++ source code to determine the meanings of code inputs and outputs.

Instead, the FSP production computing team will deploy tested software on specific production platforms which will be available to users as a service. The team will provide documented, tested methods for preparation of input datasets with working examples. There will be a clear procedure for submission of jobs. The team will provide tools for monitoring the progress of active simulations, and, with the help of software developers, tools for visualization and analysis of the output of completed runs. User level documentation will be provided to clearly explain, in detail, every step of each of these phases of run production.

The team will trouble shoot failed runs and communicate clearly with users on the status of such jobs. If cause of failure is traceable to input data, production team members will give advice to users on input modifications. If the cause is due to a system or hardware failure, the team will restart the job when the problem is resolved. If the cause appears to be in an FSP component and requires further investigation by domain experts, the production team will contact the appropriate FSP software developers.

The production computing team is oriented to serve user needs, in order to maximize the value of FSP software for research applications.

3.6.4 Interaction with other elements in FES program

The production computing teams will interact with FSP software developers—both component developers and the general software infrastructure and support team. Production computing necessarily has a very strong concern for software quality, and FSP is investing significant resources to insure quality. Nevertheless, it is inevitable that production deployment will reveal issues with FSP components that will require attention of component experts going beyond what is available within the teams themselves. In such situations, operations teams will rely on FSP software developers with appropriate component expertise. The role of the operations teams will be to provide the necessary data and information to enable FSP component specialists to precisely reproduce problematic code behaviors, as will be required to enable isolation and repair of trouble spots. Where this is not possible, it will likely be related to software quality issues or conformance to standards, and, the production team will likely need the help of the software infrastructure and support organization.

The production teams will rely on the software integration and support team for robust and reliable build procedures – as will be required to reliably roll out new versions of production software. The production team will need at times to be able to function as code developers and have write access to source code repositories, to commit bug fixes as required.

The production teams will interact with computational resource providers (laboratory clusters, NERSC, LCF sites) and data management facilities. It is bound to happen that certain production deployments will impose hardware requirements that will not have been fully anticipated, and which will require changes in use policies or acquisition of additional computation or data or network capabilities, involving costs that will be handled in discussions between FSP, computational/data resource provider sites, and research user groups.

The production teams will interact with FSP management. A fair amount of project flexibility will be required, as precise anticipation of all costs associated with computational service deployments is not possible.

Above all, the production teams will interact with the research user groups and individual users, from both FSP and the base fusion program. These researchers will indicate priorities for improvements in production computing and in FSP capabilities generally.

4 PROGRAM EXECUTION

This section describes the initial work plan of the FSP in consideration of the limited resources expected to be available.

4.1 Integrated Science Applications

As discussed in Section 2.3 above, the Science Driver plans were developed without reference to particular resource limits. To meet an assumed (and more realistic) resource profile, the FSP will begin with two ISAs, one focused on edge physics, which combines the boundary and pedestal science drivers and the other on whole device modeling, with an emphasis on disruption prediction and avoidance. The reduction assumed funding available naturally results in reduced scope over any fixed period of time. However, in both cases, the FSP effort would represent a dramatic increase in resources applied to these important problems over current levels and should lead to a corresponding increase in scientific progress.

4.1.1 Edge Physics (Boundary & Pedestal) ISA

4.1.1.1 *Introduction*

This document represents a plan for developing powerful simulation tools for magnetic fusion energy devices, merging two of the original Science Drivers owing to the close proximity of the two regions that they model and realistic limits on resources: (1) the warm plasma region known as the scrape-off layer (SOL) where magnetic field lines directly contact material structures together with the associate plasma-wall interactions and (2) the adjacent hotter plasma region known as the pedestal, which is the beginning of the confining closed magnetic field line core. This region is of vital importance for practical fusion energy for reasons explained below and in the Science Driver plans, which are summarized above and included, in full, as appendices to this report.

4.1.1.2 *Overview and Motivation*

4.1.1.2.1 Pedestal physics and simulation

High performance (“H Mode”) operation in tokamaks is achieved via the spontaneous formation of a transport barrier (or “pedestal”) in the outer few percent of the confined plasma. This edge transport barrier strongly improves global energy confinement, and also generally improves global stability, resulting in dramatically enhanced fusion performance and the potential for more cost effective fusion reactors. However, the free energy in the large pressure gradient and the resulting bootstrap current in the pedestal can drive instabilities called Edge Localized Modes (ELMs), which periodically deposit impulsive heat and particle loads on plasma facing surfaces, and may reduce component lifetimes in reactor scale devices. A predictive understanding of pedestal formation and structure, as well as the physics of ELMs, is essential for prediction and optimization of the fusion performance of ITER and future reactors.

The plasma pressure typically increases by 1-2 orders of magnitude from the bottom of the pedestal (very near the magnetic separatrix) to the top, and increases by less than an order of magnitude from the pedestal top to the magnetic axis. Hence, while the pedestal occupies a relatively narrow radial region, it contains far more pressure gradient scale lengths than the core plasma. The impact on global confinement is amplified via coupling to the core plasma where transport is fairly stiff, meaning that the core profiles are closely correlated to critical gradient scale lengths. As a result, the core pressure increases roughly linearly with the pedestal pressure (or "pedestal height"), and the fusion power output scales roughly as the square of the pedestal height, providing a powerful lever for performance optimization of fusion systems. While the performance benefits of H-mode operation are dramatic, there is a potential drawback. The large pressure gradients in the edge barrier lead to large localized currents, via the bootstrap effect, and the substantial free energy present in both the pressure and current gradients drives the ELMs While ELMs are largely benign in existing devices, and can aid in density and impurity control, in future higher power devices, highly impulsive ELM heat and particle loads to plasma facing surfaces, which may constrain material lifetimes.

The pedestal presents a daunting set of challenges to traditional theoretical and computational methods. Because the pressure varies by 1-2 orders of magnitude across the pedestal, and the density, temperature, flow velocity, radial electric field and current also vary substantially, a very wide range of key dimensionless parameters is encompassed in this region. For example, the pedestal plasma often transitions from being nearly collisionless near the top of the pedestal, to strongly collisional at the bottom, requiring methods appropriate for both regimes. More fundamentally, the broad range and overlap of spatiotemporal scales across the pedestal deeply challenges the assumed separation of equilibrium ("macro") and turbulence ("micro") scales upon which most existing theory and computation relies, and thus extensions of basic theory and massive computational resources are expected to be needed. For example, across a single pedestal, the timescales associated with electron drift waves span a wide range (due to the wide variation of equilibrium quantities) which overlaps with the wide range of temporal scales associated with Alfven waves, which in turn overlaps ion drift wave and ion transit temporal scales, which in turn can overlap the fast timescales on which the equilibrium itself is observed to evolve, for example during an ELM. The range of overlapping temporal scales often exceeds six orders of magnitude. A similar overlap is found in physically relevant spatial scales, where the gyroradius and ion drift wave scales can overlap the short gradient scale lengths.

Despite these challenges, there has been substantial recent progress in understanding key pedestal physics issues, and in developing computational tools suitable for pedestal studies. The onset of ("Type I") ELMs, and a crucial constraint on the pedestal height, has been found to be due to the onset of intermediate wavelength MHD modes, known as "peeling-ballooning modes" because they are driven by a combination of the pressure gradient (ballooning) and edge current (peeling or kink) drives. Efficient linear codes have been developed for calculating the peeling-ballooning mode onset condition. Nonlinear simulations using Braginskii, extended MHD, and gyrofluid codes have explored ELM dynamics with increasing physical realism. Static models of the pedestal height and width have been developed by combining the peeling-ballooning constraint with another linear constraint, such as that for stiff onset of kinetic ballooning modes. These models, without any fit parameters, have proved to be reasonably accurate in predicting the pedestal height in the high performance H-mode regime on a number of devices, though a number of extensions can be considered. A set of computational tools have been developed to begin the study of dynamic evolution of the pedestal. Neoclassical transport codes, including fast steady-state solvers, and large-scale initial-value simulations have been developed to treat the pedestal region, and tested, identifying significant ion thermal transport and potential effects due to ion orbit losses. Closed field line gyrokinetic solvers initially developed for the core region inside the pedestal have been extended to include fully electromagnetic perturbations and more realistic collision operators, potentially enabling their use in pedestal studies, both linear and nonlinear. Gyrokinetic codes incorporating both the closed field line (pedestal) region, and the open field line SOL region are under development by a pair of U.S.-DOE projects: the Center for Plasma Edge Simulation (CPES) and Edge Simulation Laboratory (ESL).

The practical goal for pedestal research is to achieve operation with a high pressure pedestal with a profile relaxation mechanism which does not present the material interface with unacceptable transient heat loads – that is to operate with small or no ELMs. For modeling, the goal is to develop the capabilities to understand and predict:

- the onset of edge barriers (or “L-H transition”) as well as the transition from low to high performance H-mode,
- the structure of the barrier in all profiles (with particular initial emphasis on the pressure at the top of the pedestal), and
- the nature of the pedestal relaxation, particularly ELMs, and to identify and optimize methods for reducing transient heat deposition on material surfaces (including ELM-free and small ELM regimes, as well as suppressing or mitigating ELMs via external control techniques, including magnetic perturbations or pellets).

Successful achievement of these goals will require modeling that not only addresses the substantial challenges of the pedestal region itself, but which also couples closely to the open field line region, including the scrape-off-layer, divertor and material surfaces, as well as to the deeper core plasma.

4.1.1.2.2 Boundary/wall physics and simulation

Plasma, neutral gas, and wall processes in the scrape-off layer (SOL) region just outside the magnetic separatrix dividing closed and open magnetic fields line regions play a key role in determining the heat and particle fluxes to material surfaces, both from steady-state or between-ELM periods and from ELMs themselves. While the neutron flux to surrounding walls is broadly distributed, the exhaust plasma fluxes are typically very concentrated owing to anisotropic transport properties of the strong magnetic field even on open field lines. A central issue for future magnetic fusion devices is operating them such that the steady-state peak heat flux to materials does not exceed ~10 MW/m2, which is believed challenging for ITER and an unsolved problem for higher power future devices. For transient heat loads such as ELMs there is a fundamental material melting or vaporization limit of $\Delta S_p \tau_L^{-1/2}$ ~40 MJ m^{-2} $s^{-1/2}$, where ΔS_p is the energy released by the ELM divided by the area affected on the divertor surface, and τ_L is the time for the energy to be lost to the material surface. Among additional major issues are removal of helium ash and tritium, impurity production and transport to the core region, material lifetime, and impact of intense events that periodically eject large energies into the SOL over a short time. Considerably greater detail on these processes and issue associated with them is given in the original Science Driver report mentioned in the introduction. These reports are included with the plan as an separate auxiliary document.

The general focus of the boundary task area is to produce an integrated model of that region that accounts for plasma collisional and turbulent transport, neutral/plasma interactions, and wall interactions, much as discussed in the original Science Driver. However, owing to the reduced scope of the present ISA, a number of the components will need to come from simplified existing models. In particular, neutral models and plasma wall interactions will rely largely on present models, while the coupling of plasma collisional and fundamental turbulent transport will be more completely developed as a fully functional coupled transport/turbulence SOL model does not exist.

The initial focus is on fluid models because of their lower dimensionality compared to kinetic models and because some present-day devices operate in strongly collisional regime. This SOL simulation model must be able to simulate long timescales, ~10-1 seconds for present-day devices using a fixed-temperature wall model owing to wall recycling. The timescale will be much longer when the wall temperature is allowed to evolve. Some resources will be expended to provide coupling between the plasma model and neutrals plus wall interactions (recycling and sputtering) and some cross-cutting resources will be used to improve the implicitness of the numerical algorithms for these models. It is hoped that improvements can be made to plasma/wall interaction physics models through new funding sources such as a possible SciDAC project in FY12.

Three reasons for this prioritization are that

- SOL turbulence and resulting transport across the magnetic field is believed to strongly effect the peak heat flux to divertor surfaces, a major issue for successful operation of ITER and other devices,
- the readiness of 3D fluid turbulence codes to simulate the turbulence with intermittent plasma “blob” transport observed in the SOL, and

- an initial focus on plasma turbulence and transport that includes neutrals is of great relevance to the pedestal region and should provide a direct avenue for coupling or integrating the two regions in this ISA.

In addition, both the SOL and pedestal regions have important kinetic plasma effects that can span long to short Coulomb mean-free paths and thus require an accurate Fokker-Planck operator. Furthermore, the generation of blobs that transport plasma into the SOL likely takes place near the magnetic separatrix, so a portion of the pedestal region should be included for SOL simulations.

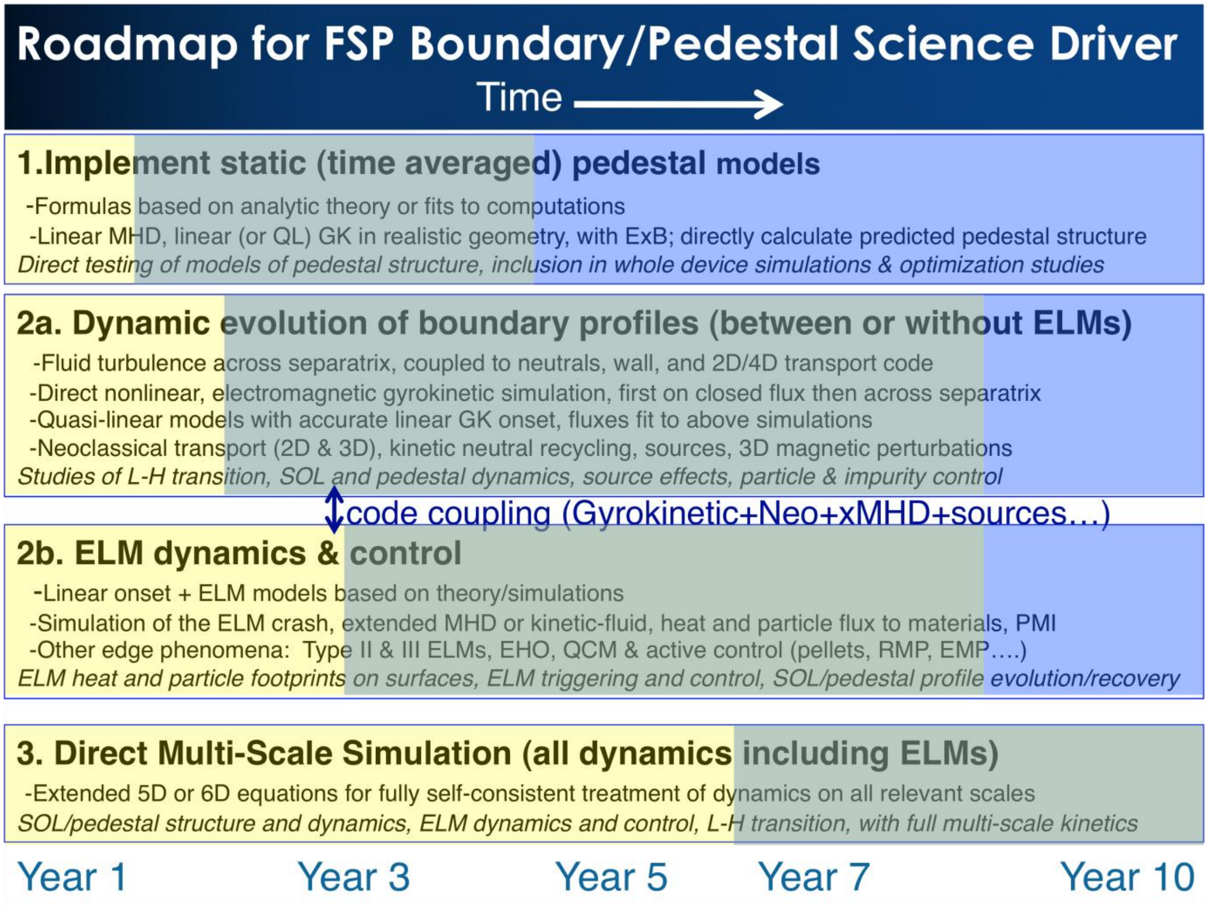

Roadmap for FSP Boundary/Pedestal Science Driver

Time ⟶

1.Implement static (time averaged) pedestal models

-Formulas based on analytic theory or fits to computations

-Linear MHD, linear (or QL) GK in realistic geometry, with ExB; directly calculate predicted pedestal structure

Direct testing of models of pedestal structure, inclusion in whole device simulations & optimization studies

2a. Dynamic evolution of boundary profiles (between or without ELMs)

-Fluid turbulence across separatrix, coupled to neutrals, wall, and 2D/4D transport code

-Direct nonlinear, electromagnetic gyrokinetic simulation, first on closed flux then across separatrix

-Quasi-linear models with accurate linear GK onset, fluxes fit to above simulations

-Neoclassical transport (2D & 3D), kinetic neutral recycling, sources, 3D magnetic perturbations

Studies of L-H transition, SOL and pedestal dynamics, source effects, particle & impurity control

↕ code coupling (Gyrokinetic+Neo+xMHD+sources…)

2b. ELM dynamics & control

-Linear onset + ELM models based on theory/simulations

-Simulation of the ELM crash, extended MHD or kinetic-fluid, heat and particle flux to materials, PMI

-Other edge phenomena: Type II & III ELMs, EHO, QCM & active control (pellets, RMP, EMP….)

ELM heat and particle footprints on surfaces, ELM triggering and control, SOL/pedestal profile evolution/recovery

3. Direct Multi-Scale Simulation (all dynamics including ELMs)

-Extended 5D or 6D equations for fully self-consistent treatment of dynamics on all relevant scales

SOL/pedestal structure and dynamics, ELM dynamics and control, L-H transition, with full multi-scale kinetics

Year 1 Year 3 Year 5 Year 7 Year 10

Figure 5: The three level roadmap for the combined boundary/pedestal ISA*

**The three level roadmap for the combined boundary/pedestal ISA, indicates sets of major tasks (1, 2a&b, 3) that are planned. Each level of the roadmap (1, 2 &3) will begin with a development, implementation and verification stage (shaded yellow), followed by a validation and ongoing development stage (shaded green), and finally a stage of routine application, with minor ongoing development (shaded blue). The major emphasis in the reduced scope plan will be on Level 2, particularly 2a, for both the boundary and pedestal regions.*

4.1.1.3 *Roadmap for the Development of Boundary/Pedestal Simulations*

The goals, challenges and progress described above lend themselves to a three-level plan for the FSP boundary/pedestal ISA effort. This plan, illustrated in Figure 5, addresses both the need to deliver world-leading capability on a relatively short timescale, and the need to address the deeper fundamental challenges associated with pedestal/SOL/wall dynamics, taking advantage of peta- and exascale computing capability as it becomes available.

There are a number of computational approaches that can be applied with increasing physics fidelity but also with increasing challenge to theory and computation. At the first level, the physics of the static (i.e., time-averaged) pedestal can be addressed via linear physics models, based on existing models and their extensions.

At the 2nd level, dynamics of the boundary and pedestal are considered, but a separation is initially maintained between the physics models for the ELM event itself, and the dynamics between, or in the absence of, ELMs. The full dynamics of the SOL/wall will be the initial focus for the boundary area. A wide variety of available and developing tools can be used to treat neoclassical and turbulent transport between ELMs, including 3D fluid and gyrofluid codes and 5D electromagnetic gyrokinetic simulation codes. Full f codes can potentially be used to treat larger perturbations, but will require further development. At this level, the ELM event itself will be treated separately, via calculations of its onset and dynamics with extended MHD or gyrofluid codes. Finally, at Level 3, dynamics across all relevant scales, including ELMs, will be treated self-consistently with a single simulation code. Additional advancements in theoretical gyrokinetic algorithms, and possibly formulations, to allow fully electromagnetic simulations of arbitrary scale electromagnetic modes in cross-separatrix geometry may be required. The most complete models would be 6D full kinetic simulations using the full collision operator. The computational challenge that this would present suggests that its use, at least initially, would be for assessment of the less complete models, though in the longer term, with sufficient computational power becomes available, more extensive use could become practical.

This general outline leads to a corresponding development roadmap with three levels and four major elements, illustrated with a timeline in Figure 5. Note that due to reduced resources, several aspects of the plan will have to rely heavily on theory and code development efforts outside of FSP. In particular, Level 1 will consist largely of implementation of existing codes and models for the pedestal, Level 2 will be the primary area of focus for this ISA, and work in Level 3 will be largely at an exploratory level of effort. As discussed more fully below, the 2-year milestone for the pedestal region lies in Level 1, whereas for the boundary area, the 2-year goal is in Level 2.

- Level 1. Linear models for pedestal structure

  This step would begin with componentization of existing models that solve for static (time averaged) pedestal structure via linear stability analysis, for example, that of peeling-ballooning and kinetic ballooning modes. Improvements can come through use of linear or quasi-linear gyrokinetic calculations, more realistic geometry and inclusion of ExB stabilization. This analysis typically requires hundreds or thousands of independent MHD and/or gyrokinetic stability calculations with trial equilibria. Key issues are robustness, error checking, automation, and, particularly in the case of gyrokinetic calculations, efficiency. Extensive comparison with experimental data sets will be carried out. It is expected that this capability can be made available relatively quickly, allowing a world-leading capability for coupled pedestal-core optimization of fusion systems. (Task A)

- Level 2. Dynamic evolution of the boundary and pedestal via separate inter-ELM and ELM components

  2a. Dynamic evolution of boundary and pedestal profiles between ELMs

  In the near term, dynamics in the boundary region are expected to be addressed with 3D fluid simulations codes coupled to 2D transport codes, and models for neutral and materials physics (Task B). In the longer term, the fundamental tool for calculating boundary and pedestal transport between ELMs is expected to be electromagnetic gyrokinetic simulations of turbulent transport including a realistic collision operator and to separate calculations of neoclassical transport, sources and material interaction (Task C, D). It is envisioned that nonlinear simulations will be employed both for development of simplified transport models, as well as for direct calculations of particle, momentum and heat transport (Task E). Neoclassical calculations will eventually include 3D equilibrium effects, such as neoclassical toroidal viscosity. All of these models would need to be appropriately verified, including extensive verification of reduced dynamic models against direct nonlinear simulations, and validated against experimental measurements.

  2b. ELM dynamics & control with fluid or kinetic-fluid hybrid models

  The models described above would be extended by simulation of phenomena that limit or control the pedestal/SOL pressure gradients. These would include spontaneous plasma behavior [ELMs of various types, Edge Harmonic Oscillation (EHO), Quasi-Coherent Mode (QCM), etc.] and active control through pellets, resonant magnetic perturbations (RMP), electromagnetic perturbations, etc. The work could begin with linear onset from peeling-ballooning calculations, coupled to simple ELM crash models. The next step would be direct simulation of ELM dynamics using extended MHD or two-fluid and/or kinetic-fluid codes (Task E). These codes would need to include realistic calculations of parallel transport and transient heat and particle loads onto material surfaces. Validation experiments could compare ELM (or other mode) structure, dynamic modification of pedestal/SOL profiles, heat and particle footprints and ELM control mechanisms.

- Level 3: Direct Multiscale Simulation

  The prior computational stages use gyrokinetic calculations for modeling the microscale and extended MHD for the macroscale. However, as noted above, these overlap strongly in the edge barrier. Some systematic study will be required to test the assumption of spatiotemporal scale separation, to determine when and how it breaks down and to assess the consequences. Numerical and theoretical progress will be required to develop and implement verified formulations and codes that can simulate multiscale electromagnetic modes and turbulence in separatrix geometry. Several approaches are possible including gyrokinetic treatments without the high-n approximation, kinetic-fluid methods and 6D Vlasov treatments including the full collision operator. The last of these, in particular, will require substantial progress in numerics to be practical. These models would support the most fundamental studies of boundary and pedestal physics including L-H threshold, coupling of turbulence and equilibrium scales, ELMs and ELM control.

4.1.1.4 *Tasks and Milestones*

<u>Tasks for Years 1-2</u> (see Table 4 for effort levels)

A. <u>Static (Linear) Models for Pedestal structure (2-year milestone)</u>

This task will consist primarily of the implementation and testing of existing models of the pedestal structure, based on theory and linear MHD and gyrokinetic calculations; componentization and verification of existing linear MHD and gyrokinetic codes, validation and development of extensions to models

B. <u>Coupled fluid turbulence/transport/wall models (2-year milestone)</u>

The largest gap that will be addressed in this 2-year period is coupling SOL turbulence to long-time plasma/neutral transport using fluid models. In addition, there will be coupling to a wall model, and a near-sheath plasma model. Examples of simulation codes exist for all of the individual processes and some also integrate multiple processes, but a routine coupled transport/turbulence model does not exist. In addition, a smaller amount of work will begin on kinetic models in this period, but full implementation of those will be directed at the 5-year milestone.

- The turbulence in a small region about the separatrix and into the SOL is typically more intermittent and larger amplitude than in the core region. Thus, two strategies will be considered to profile long-time coupling between plasma transport and turbulence. The first is to embed a dynamic neutral model including material recycling within a 3D turbulence code for observed drift-type modes, thus allowing the turbulence code to evolve its own axisymmetric plasma profiles. The second approach is to couple the 3D turbulence code with a 2D transport code (plasma and neutrals) using, for example, the relaxed iteration coupling (RIC) algorithm [Shestakov 03]; some preliminary development has already been done for application of this method to SOL turbulence and transport [Rognlien 05]. These two approaches will be evaluated in the first six months, followed by a focused effort on the most promising. Central questions to be resolve are practicality of very long simulations runs while maintaining particle and energy conservation and the applicability of the RIC method to moderately strong, intermittent transport events.
- Simplified models of plasma recycling at material surfaces are present in existing plasma transport codes. However dynamic wall processes, such as hydrogen accumulation in new conditioned walls (a standard procedure in many tokamak before each discharge) and ejection of hydrogen (out-gassing) in response to wall temperature increases, are not taken into account in a self-consistent manner. Wall codes have now been developed that can describe these time-dependent processes [Hassanein 02,

Pigarov 09]. The task here is to couple an existing model to both transport and turbulent plasma/neutral models, but not to further develop the models unless incremental funds are available. Some initial work has been done in the FACETS SciDAC in this direction that can likely be utilized. Important developments that will be needed are to make the coupling implicit in time as well as the wall code itself to allow appropriate long-time simulations.

C. Preparation of kinetic models (toward 5-year milestone)

- As particles are recycling or sputtered from material surfaces, they penetrate some distance into the plasma before being ionized. It the ionization rate is sufficiently large, the ionization takes place very close to the material and their ion gyro-radii may allow prompt re-deposition to the wall [Brooks 02]. Such a process gives a net sputtering of impurities and is important in determining the evolution of the surface material, especially as it relates to sputtering impurities and separate deuterium and tritium transport during the many particle recycling/re-deposition events. The task is to develop an implicit solver for the shear model and begin work on implicit coupling strategies that minimize the impact of particle noise.
- Coupling fluid and kinetic neutrals is important, especially in the low-density periphery of the SOL. Here the issue of particle noise on the coupling needs to be addressed if the kinetic model is particle-based Monte Carlo [Stotler 01].
- Prompt drift-orbit loss of energetic ions near the separatrix may produce an important heat-flux component to the divertor plate [Chang 04]. Consequently, it is important to eventually include a kinetic ion transport model in the SOL. Likewise, parallel electron transport in the SOL can have energetic tail electrons owing to parallel kinetic effects [Batishchev 97]. Both of these ion and electron kinetic effects will require an accurate Coulomb collision operator, and cross-cutting work will begin on the task of finding a method for efficient calculation of Rosenbluth potentials.

Tasks for Years 3-5 (see Table 4 for effort levels)

The first 5-year boundary milestone is to generalize the basic fluid 2-year model to include kinetic effects for transport across the magnetic field as well as along it. The initial turbulence code that provides the turbulent fluxes will still be an electromagnetic fluid model. Work will be done to develop an electromagnetic kinetic SOL turbulence code, but at the constrained budget level, its completion in the 5-year timeframe is not proposed.

The second 5-year milestone involves dynamic modeling of the pedestal, based on nonlinear, dynamic kinetic descriptions, initially on closed field lines, and then extending across the separatrix and combining with the SOL/divertor/wall simulation efforts.

D. Coupled kinetic-transport/fluid-turbulence; improved wall/sheath models (5-year milestone)

- A 4D kinetic transport model for ions and electrons will be coupled to a 3D turbulence model for long-time transport simulations. As with the fluid model, particle recycling produces a long timescale of ~0.1 s that must be accommodated; the kinetic transport model will thus need to use an implicit time-advance method. The kinetic collision operator will include charge-exchange and a source term for ionization/recombination.
- A kinetic neutral model will be coupled to the plasma model or a sufficiently parameterized, verified reduced fluid model will be used. Implicit coupling will be developed.
- As for the 2-year milestone, a dynamic wall model will be coupled to the plasma/neutral system. Here the generalization to non-Maxwellian particle and energy fluxes will be included in the wall model.
- Impurities will be included in the fluid transport and turbulence models. These in turn will be couple to a near-sheath impurity model for re-deposition of sputtered material. This work will set the stage for adding impurities in the kinetic plasma/neutrals models beyond the 5-year timeframe.

E. Dynamic evolution of pedestal profiles (5 year milestone)

Existing substantial efforts in edge and core gyrokinetics and extended MHD provide a good starting point. Thus, initial efforts will involve adapting existing components to requirements for the FSP. This is a large, broad task and substantial resources will be required. Bulk of effort will initially be towards development, with emphasis switching to new science and V&V in out years

- Componentization and verification of existing nonlinear MHD and electromagnetic gyrokinetic codes
- Design and development of new capabilities (e.g., free boundary equilibrium solver accurate to SOL; ion-electron gyrokinetics (GK) with magnetic perturbations, etc.)
- Experimental validation and new science investigations

Milestone (time)	**Application work**	**Supporting work**
A) Static Pedestal: Linear MHD/kinetic-microturbulence stability boundaries (2 year)	Perform multi-parameter stability studies with gyro-kinetic/ MHD/ fluid codes	Kinetic collision operator
B) Fluid SOL/wall: Coupled plasma transport/ turbulence/ gas/ wall (2 year)	Establish 2-way couplings between components; then couple all components	Implicit solver for fluid turbulence & wall; implicit coupling; framework?
D) Kinetic SOL/wall: Coupled SOL kinetic transport/plasma turbulence/gas/wall (5 year)	Add kinetic transport models for plasma/ gas/; enhanced wall model	Implicit solver for kinetic transport code; implicit kinetic coupling; particle noise; Coulomb collision operator; framework
E) Dynamic Pedestal: Nonlinear MHD/kinetic microturbulence transport (5 year)	Pedestal profile evolution with kinetic code; add neutrals/ begin ELM loss	
Long term) Coupled kinetic pedestal/SOL/wall: Nonlinear evolution of pedestal/ SOL/ wall with self-consistent turbulence, and some multiscale capability (8-10 year)	Couple kinetic pedestal and SOL/wall components; consistent ELM coupling; kinetic SOL turbulence	Implicit kinetic coupling; particle noise; Coulomb collision operator; framework

Table 3: Milestones and Effort Required for Boundary/Pedestal ISA

Year/Task	**Yr 1**	**Yr 2**	**Yr 3**	**Yr 4**	**Yr 5**
Static Pedestal Model	1.5	1.5	1	0.5	0.5
Longscale fluid turbulence. with transport	2	3			
Iterative fluid turbulence / transport	1				
Add impurities to SOL turbulence.			0.5		
Couple wall uptake/ release/ temp	0.5	0.5	0.5	0.25	0.25
Kinetic sheath model	0.25	0.25	0.5	0.5	0.25

Couple fluid/ kinetic neutr.	0.5	0.5	0.5	0.75	0.5
Kinetic collisions with CX, ionization	0.25	0.25	0.25	0.5	0.5
Couple fluid turbulence / kinetic transport			0.75	1	1
ELM simulation	0.5	0.5	0.5	0.5	0.5
Kinetic turbulence	1.5	1.5	2	2	2.5
kinetic turbulence / kinetic transport including. neoclassical	0.5	0.5	1.5	1.5	1.5
TOTAL	**7**	**7**	**7**	**7**	**7**

Table 4: Effort estimates for Boundary/Pedestal ISA*

*Table 4: Effort estimates for Boundary/Pedestal ISA* gives the resources allocated within the ISA group itself for the first five years of the project. A similar level of supporting resources across the organization are expected for successful task completion.

4.1.2 WDM & Disruption Avoidance ISA

4.1.2.1 *Introduction*

The whole device modeling of tokamak discharges involves the integration of different spatial and time scales for the modeling of all discharge phases starting from startup to shut down. High fidelity predictive whole device modeling should accurately account for scrape-off layer physics, plasma wall interactions, core transport, heating and current drive, fast particles, pedestal physics, ELMs and impact on the divertor, 3D MHD modes, as well as other physics issues. The success of a new WDM tool is strongly dependent on careful coupling of diverse physics components. Due to interrelation among physical effects, strong coupling of the physics components becomes essential. The interplay between different physics components introduces a new level of physics fidelity and leads to discovery of new effects that are not available when physics components are considered in isolation.

During tokamak experimental operation, events that rapidly terminate the plasma discharge occasionally occur. The complete and rapid loss of thermal and magnetic energy in these disruptions results in large thermal and magnetic loads on the material wall. For proposed next step experiments such as the International Thermonuclear Experimental Reactor (ITER), the stored energy will be approximately 100 times greater than present day devices, greatly increasing the potential damage of these events. As the computational model for simulating the plasma evolution on long time scales, whole device model is the dominant tool for forecasting the onset of disruption events and ways in which to avoid them.

4.1.2.2 *Overview and Motivation*

There are four overall thrusts that will contribute to the development of a successful WDM capability. These thrusts are:

- High fidelity science components. Each component describes an individual physical effect. Individual components, by themselves, cannot be used to answer questions about ITER performance and used to optimize the ITER scenarios.
- Reliable and flexible framework that set standards for coupling of science components in the WDM suite of codes. The framework should be flexible enough to allow the coupling with 1D, 2D, or 3D components, explicit and implicit coupling, dynamic parallelism, flexible data exchange and storage;

- Verification and validation (V&V) of individual physics components and WDM tool in general. The V&V activity will include the establishment of V&V metrics, set of synthetic diagnostic tools, development of interfaces to experimental data, and legacy transport codes;
- Data visualization, analysis, transport, and storage.

The kinds of physics problems that will be addressed with FSP WDM suite of codes will include the following:

- Predict the plasma confinement, transport and plasma profile evolution in tokamak discharges. To enable simulations on the lifetime of the plasma device, reduced models of the first-principle simulations of plasma turbulence have been remarkably successful in predicting the profiles in many regimes. In the future, convergence in the transport predictions in all tokamak operating regimes is needed.
- Implement and validate linear stability calculations within WDM in order to predict the onset of the macroscopic instabilities that lead to plasma disruptions. The various instabilities disrupt the plasma in different ways leading to a variety of actuator controls for avoiding disruptions. A strong V&V program is needed for modeling these different paths towards disruptions.
- Implement input from the ISA Boundary/Pedestal effort to predict the plasma boundary conditions and the H-mode pedestal. All of the plasma profiles are strongly influenced by the evolution of the plasma boundary, including the capability to compute interactions between magnetic coil currents and plasma currents.
- Predict the sources and sinks that drive all of the profiles in plasma discharges. Sources such as neutral beam injection, fusion reaction products, and radio frequency heating and current drive all involve the computation of fast particle distributions and their interaction with the thermal plasma profiles and plasma edge. These energetic particle distributions also strongly influence macroscopic instabilities so accurate predictions of these distributions are needed to understand how to avoid disruptions.

Some critical physical issues relevant to WDM as well as important gaps and experimental measurements that are needed to address these physical issues are given in the table below.

Issue	**Critical Physics**	**Measurements Needed**	**Important Gaps**
Evolution of plasma profiles from boundary to core	• Coupling of validated models for microturbulence, EP modes, MHD activity and their effects on transport	• Turbulence in density, T_e, T_i, B • Profiles and gradients EP sources and dist. Functions • Precise measurements of current density profile	• Validated component models • Model for interaction between physics models Synthetic diagnostics Nonlinear saturation models for EP and MHD modes • Model linking MHD activity to flux evolution

Prediction, control, and mitigation of instabilities	• Onset, growth rate, and nonlinear saturation for sawteeth, ELMs, RWMs, TMs, NTMs • How these modes affect plasma evolution, • transport and poloidal flux	• Internal mag field fluctuations and structures • 3D arrangement of mag probes	• Validated component models • Model for interactions between models • Effect on profiles and equilibrium • Effect on sources
Interaction of boundary with plasma core	• Effect of heat flux on the boundary and of the boundary on the heat flux	• Profiles and distribution functions in the pedestal and SOL • Impurity generation at wall • 2D radiation profile • Radial electric field and bootstrap current in boundary region	• Validated reliable component models for SOL and PSI • Validated model for density transport in boundary and core • Effects on discharge and PSI of ELM control techniques • Impurity transport
Interaction of boundary wit I tried to get the tables into some consistent style plasma core	• Effect of heat flux on the boundary and of the boundary on the heat flux	• Profiles and distribution functions in the pedestal and SOL • Impurity generation at wall • 2D radiation profile • Radial electric field and bootstrap current in boundary region	• Validated reliable component models for SOL and PSI • Validated model for density transport in boundary and core • Effects on discharge and PSI of ELM control techniques • Impurity transport

Table 5: WDM physical critical issues

It is expected that different elements of the WDM code suite (different combinations of components) will progress at different rates, delivering different levels of physics fidelity. Although the "full" fidelity WDM code desired for modeling ITER discharges may not be available for several years, in the near term codes will be assembled using high fidelity components that are verified and validated/compared to experiments, and able to be placed within a larger simulation. This assembly will be able to answer critical questions in the near term.

4.1.2.3 *GOALS FOR WHOLE DEVICE MODELING/DISRUPTIONS*

The goal of whole device modeling is to provide a comprehensive predictive simulation capability for magnetically-confined plasmas, a capability that integrates the knowledge from key multi-scale physical processes in order to continually improve physics fidelity. This capability is needed to maximize exploitation of fusion experiments, especially ITER, and to establish the scientific basis for an economically and environmentally attractive source of energy. In particular, FSP WDM software must be designed to meet the following needs:

- Scenario modeling to plan new experimental campaigns in existing tokamaks or to extrapolate to planned future devices. Scenario modeling is used to optimize discharge parameters, such as maximizing fusion power production in burning plasmas, and to maximize the effectiveness of planning new experiments.

- Analysis of experimental data by computing quantities that are not measured and to resolve discrepancies between different ways of measuring experimental data.
- Development of discharge profile and shape feedback control techniques
- Prediction of the onset of disruptions and the enabling ways of avoiding them.
- Validation or calibration of theoretical models by comparing simulation results with experimental data.
- Production of self-consistent simulation results that are passed on to other more specialized computer codes.

Fundamentally, three kinds of challenges are faced in WDM modeling:

1. Coupling between different regions of the plasma — such as the coupling between core and edge plasma regions.
2. Coupling between different physical phenomena — such as the coupling between transport and large-scale instabilities.
3. Bridging the gap between short and long time scales, or between microscopic and macroscopic space scales. An example of this last kind of integration would be the simulation of turbulence, which grows on microsecond time scales and sub-millimeter space scales, resulting in transport across the plasma and the evolution of plasma profiles over tens of seconds in a tokamak with dimensions of several meters.

The required whole device modeling capabilities involve self-consistent simulations of the entire plasma discharge that include all of the relevant physical phenomena. Depending upon the requirements of each simulation, the user should be able to choose from a spectrum of models for each physical process. These models should include high physics fidelity models based on first-principles computations as well as reduced models for more rapid computations and validation or empirical models. Full-featured WDM simulations should be capable of simulating the entire time-span of the discharge, from startup to shut down, and the entire spatial scale from the magnetic axis to the interaction between the plasma and the first wall and magnetic coils. There should be seamless access to experimental data or the results from previously-run simulations.

Most of the existing WDM codes are limited to axisymmetric plasmas with simply nested closed magnetic surfaces. There are needs to further develop WDM codes for the open magnetic field region at the plasma edge, including plasma-wall interactions, and to couple the closed magnetic surface regions of the plasma with the open magnetic surface regions. There is also the need to include the three-dimensional effects that result from the formation of magnetic islands, magnetic ripple, resonant magnetic perturbations and macroscopic instabilities in tokamak discharges. . This includes the ability to model the onset and subsequent behavior of the macroscopic instabilities such as sawtooth oscillations, neoclassical tearing modes, resistive wall modes, edge localized modes which can disrupt the plasma. Future WDM codes must include more kinetic modeling — as opposed to fluid approximations — and there must be a closer coupling between fast particle distributions and the more thermal part of the distribution function. A WDM suite of codes must be developed to bridge the gap between high-fidelity turbulence simulations on microsecond time scales and the resulting transport on multiple-second time scales. Integration between fine-scaled kinetic and large-scale macroscopic physical phenomena is needed in order to produce a fully self-consistent simulation capability.

4.1.2.3.1 Goals Associated with Disruption Detection and Mitigation

The proposed science development roadmap was planned to enable the accurate prediction of 1) the onset of disruptions and how to avoid them, 2) the consequence of disruptions and how to mitigate those consequences. The specific questions that we need to answer are:

- How well can we predict the onset of a disruption and what strategies are available to avoid their occurrence?
- How can we eliminate the instabilities that lead to the disruptions?

- What are the effects of runaway electrons and what is the impact of operating regimes on their generation?
- What is the impact of disruptions on the material wall, and how can we better design the first wall to handle the thermal loads?
- What are the forces on the vacuum vessel and support forces during a disruption, and how do we improve their design?
- How can we better design disruption mitigation systems?

Because the mechanism by which the plasma loses its energy to the wall involves long-wavelength instabilities and their nonlinear interaction, extended magnetohydrodynamics (MHD) is perhaps the dominant paradigm for answering many of the above questions. Because the extended MHD codes are less well-suited for exploring large areas of the vast parameter space and long time scales, “reduced models” are valuable to help answer these questions. For this reason, we plan on having two near-to-long term development campaigns oriented around integration efforts with whole device modeling (WDM) codes and with extended MHD codes. Because there are many unresolved physics issues in the WDM, extended MHD, material wall modeling, and models for impurity delivery systems, we will have a parallel development effort in the development of advanced components. The relationship of the physics campaigns with the development campaigns, and the needed integration that these development campaigns entail is shown schematically in the figure below.

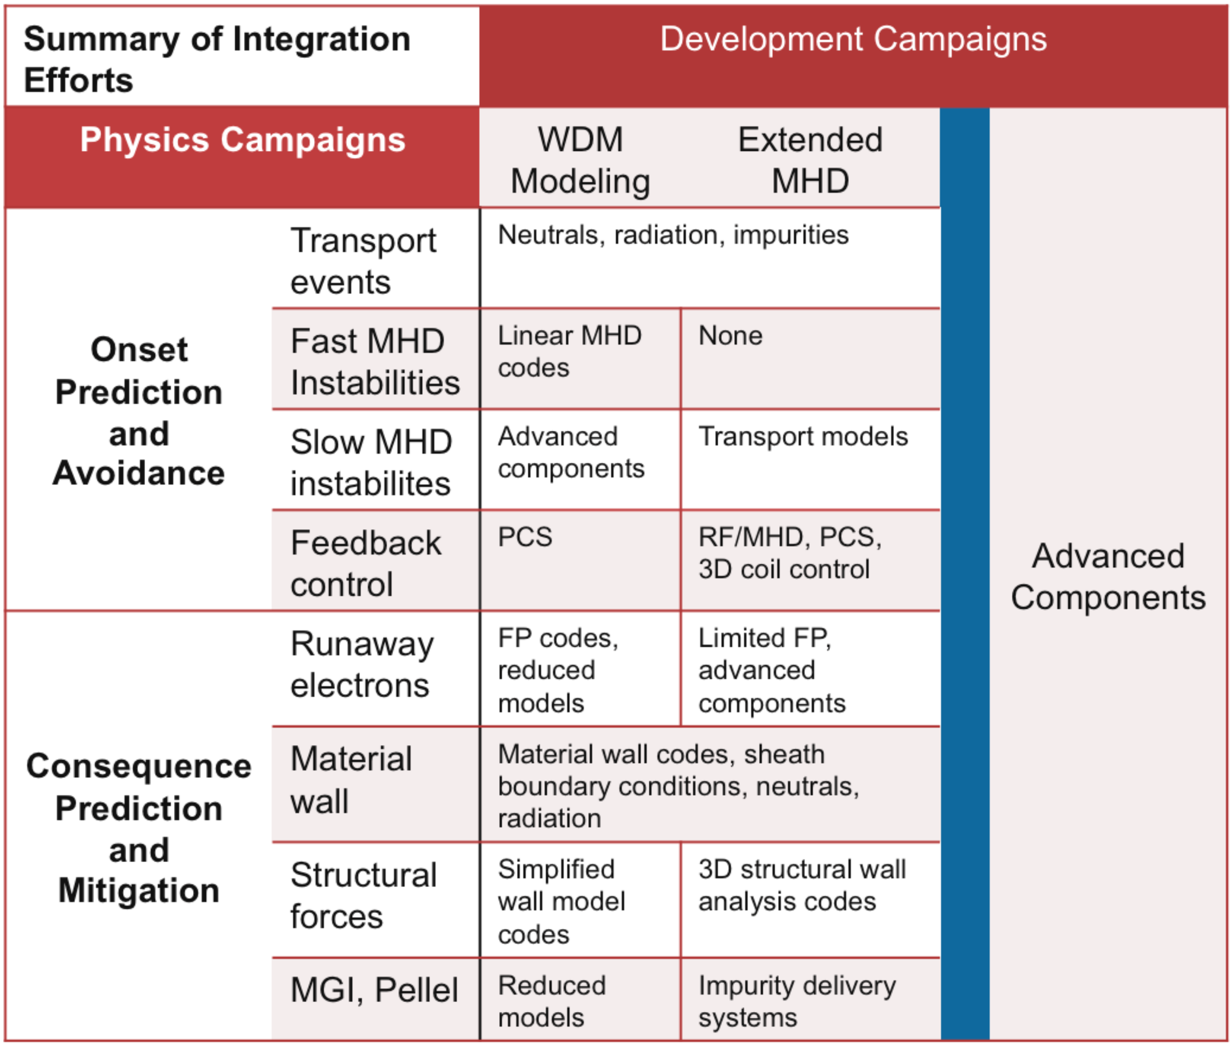

Summary of Integration Efforts		Development Campaigns		
Physics Campaigns		WDM Modeling	Extended MHD	
Onset Prediction and Avoidance	Transport events	Neutrals, radiation, impurities		Advanced Components
	Fast MHD Instabilities	Linear MHD codes	None	
	Slow MHD instabilites	Advanced components	Transport models	
	Feedback control	PCS	RF/MHD, PCS, 3D coil control	
Consequence Prediction and Mitigation	Runaway electrons	FP codes, reduced models	Limited FP, advanced components	
	Material wall	Material wall codes, sheath boundary conditions, neutrals, radiation		
	Structural forces	Simplified wall model codes	3D structural wall analysis codes	
	MGI, Pellel	Reduced models	Impurity delivery systems	

Figure 6: The relationship of the broad physics areas to the development campaign*

**The relationship of the broad physics areas to the development campaign shows a multi-faceted approach for dealing with the problem of disruptions. Not shown are the many reduced models that are likely to be needed for WDM-based development.*

4.1.2.4 *Roadmap for the Development Whole Device Modeling/Disruption Simulations*

The WDM group has constructed four high priority research areas along these lines:

- 2.5D equilibrium and transport solver;
- Self-consistent fast particle treatment for neutral beam, ion cyclotron, and alpha heating and current drive sources;
- Incorporation of turbulence simulation into transport timescale simulation;
- Self-consistent, coupled core-edge dynamics.

These four research areas are being developed from the point of view of deliverables in each given year, incremental progress in physics/code capability, demonstration of specific physics. While these research topics do not cover the whole spectrum of research problems associated with WDM, progress in four identified areas will significantly enhance the predictability of whole-device modeling of tokamak discharges.

A series of observations within the group by members provides some perspective that may not be fully developed yet.

- More than one approach to the framework may be required for the FSP to accommodate phases in its development.
- WDM will likely be faced with making components work in time-dependent simulations after they are validated in individual time-slices by the responsible group. This task includes recognizing all integration timescale issues.
- Fleshing out a group of whole device modeling users should help to direct the needed structure of the WDM tool.
- The experimental data connection will be critical for the WDM tool. The goal is to make this connection as uniform as possible allowing multi-machine comparisons of models.
- Legacy codes are needed to 1) verify new tools developed, 2) facilitate access to interpretation of experimental data.

4.1.2.5 *Tasks and Milestones*

The goals for WDM are to provide comprehensive predictive simulation capabilities for magnetically confined plasmas that integrate the knowledge from key multiscale physical processes across the whole device with progressing levels of physics fidelity. Four high priority physics topics have been identified to support WDM research toward these goals: 2.5D equilibrium and transport solver, self-consistent fast-particle treatment of heating and current-drive sources, incorporation of first-principle gyro-kinetic turbulent simulations into transport timescale simulations, and modeling of ELMs with pedestal, SOL, divertor, and first-wall interactions. Additionally, to support WDM verification and to make efficient use of resource as well as to engage the fusion community, legacy WDM tools and related SciDAC FSP prototype projects useful for WDM will need to be identified and integrated into the WDM framework early in the project phase and used as starting points for development.

A set of high-level milestones and deliverables based on the planned research activities in these areas to meet the WDM goals is given below. These are organized into 2, 5, 10, and 15-year marks and separated into two groups. Below, the overall milestones and deliverables to provide comprehensive predictive simulation capabilities across the whole device with progressing levels of physics fidelity and their validation and applications toward a full ITER discharge simulation is given. The milestones start from legacy WDM transport codes and SciDAC WDM prototypes and move forward toward a set of comprehensive predictive FSP WDM tools with increasing physics fidelity and parallel architectures. In parallel with the development effort, device description aspects and experimental interpretation functions will be established. This is then followed by validation of the physics components against experiments with increasing levels of interactions among crucial physical processes, and demonstration of integration of these components toward ITER applications will be performed.

The FSP Whole Device Model (WDM)/DISRUPTION schedule is broken down into five thrusts plus the central team and production system:

- 2.5D WDM & transport solver (3D equilibrium, 1D transport).
- Disruption Mitigation and Avoidance Tasks
- WDM components for fast particle evolution and sources which take into account RF coupling to fast ions created by neutral beam injection and/or fusion reactions.
- WDM components for evaluation of plasma turbulent transport on transport time scales.
- WDM that couples in edge and wall models (with successive fidelity).
- WDM/Disruption “central team” and production system.

For each of these, estimates are made of:

- Projected schedule of work to be carried out over a 15 year time period.
- Realistic estimate of resources required.

4.1.2.5.1 Central Team WDM/Disruption Milestones

Near-Term 2-3 Years

- Identification and establishment of candidate FSP WDM frameworks
- Identification of candidate legacy WDM 1.5D transport codes and related SciDAC WDM prototype codes for FSP WDM applications
- Componentization of physics modules from legacy 1.5D transport codes and SciDAC WDM prototype codes. Integration of the components to the FSP WDM framework.
- Establishment of device description and experimental interpretation function under WDM framework
- Establishment of validation metrics
- Identification of candidate test cases for verification and validation
- Validation of component models from legacy WDM 1.5D transport codes and SciDAC WDM prototype codes
- Validation of simplified combination of component models in legacy WDM 1.5D transport codes and related SciDAC FSP prototype WDM codes
- Demonstration of selected ITER applications under WDM framework with legacy WDM 1.5D transport codes and SciDAC WDM prototype codes
- Identification of gaps in component models to meet WDM goals

5-10 Years

- Installation of selected FSP components from WDM thrusts and other FSP areas into legacy WDM transport codes and SciDAC WDM prototypes as they become available
- Establishment of FSP WDM prototypes with parallel architectures
- Demonstration of high performance FSP WDM prototypes under WDM framework
- Verification and validation of component models in high performance FSP WDM prototypes
- Establishment of plausibility of validation in complex conditions

- Verification and validation of combination of component models with progressing levels of complexity in high performance FSP WDM prototypes
- Optimization for WDM parallel architectures
- Demonstration of selected ITER applications under WDM framework with high performance FSP WDM prototypes
- Establishment of production system on high performance computing system and documentation
- Deployment of production system on high performance computing system
- Identification of gaps in component models to meet WDM goals
- Establishment of development plan to meet gaps in component models

10-15 Years

- Installation of selected FSP components with more physics fidelity into FSP WDM codes as they become available
- Validation of combination of component models under more complex conditions over the entire discharge evolution in high performance FSP WDM codes under WDM framework
- Demonstration of selected ITER applications under complex conditions over the entire discharge evolution with high performance FSP WDM codes under WDM framework

4.1.2.5.2 Disruption Mitigation And Avoidance Tasks

All Topics

Near-Term 2-3 Year

- Have experimental database of disruptions cases along with the data available to FSP analysts using FSP-developed standards for data storage and access. The experimental database will include control or comparison cases.
- An analysis of the effects of disruptions given the different causes of disruptions will be provided to the computational and theory community to aid in assessing

Disruption onset and avoidance

Near-Term 2-3 Year

- Provide validated WDM capability for enabling predictions of VDE onsets with the uncertainties in the modeling quantified through validation.
- Quantify the limitations of extended MHD for calculating the forces due to the nonlinear evolution of a VDE disruption.

5-10 Years

- Provide capability for WDM and extended MHD components to model the effects of feedback control through integration with PF components.
- Provide numerical analysis of role of impurities in the nonlinear evaluation of VDEs.

10-15 Years

Fast MHD induced disruptions

Near-Term 2-3 Year

- Provide validated capability for using linear ideal MHD codes to predict the onset of fast MHD induced disruptions with the uncertainties in the modeling quantified through validation.

- Quantify the limitations of extended MHD for calculating the forces and fluxes due for fast MHD induced disruptions.

5-10 Years

- Investigate the extent to which impurities affect the nonlinear dynamics of the disruption.

10-15 Years

Slow MHD (Tearing modes and resistive wall modes)

Near-Term 2-3 Year

- Couple free boundary, 3D equilibrium code with islands to neoclassical gyrokinetic code. Gyrokinetic code provides: self-consistent calculation of bootstrap currents in NTM; j from cross-field drifts for purpose of prorating kinetic and flow effects in the equilibrium.
- Begin validation of previous development against saturated NTMs.
- Replace VMEC equilibrium solver with a 3D equilibrium solver that can handle islands and stochastic regions in the Strand-Houlberg 2.5D code. Initially use constant-ψ approximation in islands. (In general geometry, this corresponds to the approximation that $B\oplus\nabla\psi/B\oplus\nabla\phi$ is constant across the island. This is generally a reasonable approximation, even for relatively large islands, as long as the equilibrium is not close to marginal stability for an ideal mode.)
- Provide summary of the ability of the extent to which WDM can predict TM/RWM onset using linear codes combined with reduced models.
- Using NTV rotation models with other momentum transport models, simulate rotation dynamics of RWM/TM including resonant, non-resonant, and wall torques on mode. (Y2 – something into 1.5WDM code)
- Provide quantified analysis of extended MHD ability to predict ability to model rotation evolution of bulk plasma, and also propagation frequency and dynamics of observed in experiment in the presence of external perturbations and ECCD feedback stabilization simulations.

5-10 Year

- Use 3D equilibrium codes (perturbed ideal MHD, stellarator codes) + transport codes to model thermal and particle transport in presence of island(s) and compare to non-linear kinetic MHD simulations
- Sources and transport model incorporated in 2.5D code.
- Validate 2.5D code for full evolution of NTM, including trigger threshold and momentum transport (including NTV terms).
- Begin validation of 2.5D code against RWM time evolution.
- Improve validation of TM/RWM simulations by using energetic particle and CEL-DKE closures. Include interaction with other modes such as Alfven eigenmodes.

Disruption mitigation

Near-Term 2-3 Year

- Provide validated capability to model gas-jet penetration and pellet ablation in 3D in the pre-TQ phase, including radiation and parallel h eat transport and validate TQ onset time against experiment
- Provide WDM capability for modeling gas jet injection.

5-10 Year

- Provide quantified analysis of the ability to model the injection of jet and pellets, especially to describe the mixing/assimilation fraction.
- Provide quantified analysis of WDM capabilities for modeling disruption mitigation experiments.

4.1.2.5.3 2.5D WDM & Transport Solver (3D Equilibrium, 1D Transport).

Near-Term 2-3 Years

- Identification and establishment of 3D equilibrium solvers for WDM applications
- Integration of 3D equilibrium solvers into WDM framework
- Verification and validation of 3D equilibrium solvers
- Adaptation of 1D transport equations in 3D magnetic geometry

5-8 Years

- Adaptation and assessment of reduced transport models and simplified source and loss models in 3D magnetic geometry
- Identification and development of 2.5D transport solvers for WDM applications
- Integration of transport and 3D equilibrium components into 2.5D transport solver under WDM framework
- Verification and validation of 2.5D transport solver against legacy WDM codes and experiments under WDM framework with and without 3D magnetic effects
- Demonstration of selected ITER applications with 2.5D equilibrium and transport solver under WDM framework with and without 3D magnetic effects
- Optimization of 2.5D WDM transport tool for WDM parallel architectures

10-15 Years

- Installation and integration with 3D core and pedestal transport models when they become available
- Installation and integration with 3D SOL, divertor, and wall interaction models when they become available

4.1.2.5.4 Self-Consistent Fast-Particle Treatment of Heating and Current Drive Sources

Near-Term 2-3 Years

- Identification and establishment of essential beam, RF, and alpha-particle physics components for WDM applications
- Evaluation and development of algorithms for integration with WDM tools

5-8 Years

- Installation and Integration with WDM simulation tools
- Verification of validation under WDM framework
- Optimization for WDM parallel architectures
- Demonstration of selected ITER applications

4.1.2.5.5 Incorporation of Gyro-Kinetic Turbulent Simulations into Transport Timescale Simulations

Near-Term 2-3 Years

- Identification and evaluation of first-principle gyro-kinetic turbulent transport simulation tools for WDM applications
- Evaluation and development of algorithms to integrate first-principle turbulent transport tools with WDM simulation tools

5-8 Years

- Installation and Integration of first-principle turbulent transport tools with WDM simulation tools
- Verification of validation of first–principle turbulent transport models under WDM framework
- Optimization for WDM parallel architectures
- Demonstration of selected ITER applications with first-principle turbulent transport models under WDM framework

4.1.2.5.6 Self-Consistent, Coupled, Core-Edge Dynamics

Near-Term 2-3 Years

- Identification of reduced pedestal, SOL, divertor, and first-wall interaction physics component(s) for WDM applications
- Installation and integration of these components into WDM
- Verification and validation of these components under WDM framework
- Demonstration of selected ITER applications under WDM framework

5-8 Years

- Installation and integration of components with more physics fidelity
- Optimization for WDM parallel architectures
- Verification and validation of these components under WDM framework
- Demonstration of selected ITER applications under WDM framework

4.2 Components

4.2.1 Summary of component planning and findings

The FSP component planning accomplished two tasks in the program definition phase. The first is to establish the FSP component candidate database. The second is to establish FSP component designs.

The objective of the first task is to survey the available simulation capabilities by fusion science topical areas. The process is community input via component questionnaires. The outcome is a list of component candidates with detailed description of their capabilities in terms of physics models, computational algorithms and performance. This information is posted on the FSP component wiki.

The objectives of the second task are to (1) transform the physics challenges into computable components which are coupled to resolve the integrated physics models; and (2) establish the component specification for individual science driver in the context of its likely coupling scheme. There were two sequential processes involved. First, design proposals or ideas were solicited and received from the three proto-FSP centers, the five fusion SciDACs, and the three community code projects funded by the base program. The outcome was documented in the Boulder Workshop presentation and the final report. Second, six integrated planning teams (IPT) were commissioned for the six science drivers to consolidate the design proposals into coherent FSP designs for each science driver. The six IPT reports are completed and available for the entire community on the FSP wiki site.

The component candidate database, the Boulder workshop report, and six IPT reports form the basis for the FSP program-wide component prioritization and planning. The details are given the summary findings of the three community planning exercises are highlighted next from both a near term and a long term perspective.

The near-term perspective on FSP addresses the issue of FSP readiness. From the findings of the three community planning exercises, it is evident that there is solid base of existing (component) capabilities and credible integration schemes to produce meaningful integrated software to tackle every science drivers with the first five years. Most if not all expect significant improvement in fidelity beyond current integrated modeling capability in every science driver area. At the same time, limitations are clearly identifiable and identified. The early deliverables of FSP would provide an excellent platform for verification and (in)validation.

It is also evident that the diversity of potential components and integration schemes or approaches for the same science driver reflects the reality that significant gaps exist between current capability and (1) a truly first-principle-based predictive capability and (2) the need to predicting range of current experimental observations.

From a long term perspective, one sees a converging path for FSP research and development. First, all of the provisional science drivers converge to common component R&D needs in key areas. Two cases in point are the (1) two dimensional equilibrium and transport solvers from magnetic axis to chamber wall; and (2) self-consistent three dimensional plasma equilibrium and transport solvers from magnetic axis to chamber wall. Second, Imbedded calculation, which is a technique for multiscale and multiphysics coupling, and called out in a number of science applications, poses similar challenges to coupling/framework. Examples include kinetic closure of XMHD, turbulence transport on transport time scale, RF heating and current drive in an MHD active plasma. Third, the research and development thrusts in the physics integration area are on converging paths. For example, the core transport modeling moves to include the tokamak edge, while the edge transport modeling moves to include the core plasma. Similarly the XMHD modeling moves to include the effect of gyrokinetic turbulent transport, while the gyrokinetic transport modeling moves to include the impact of low-n to medium-n magnetic activities.

While the challenges to achieve predicting simulation for FSP are daunting, the long term perspective on FSP suggests a number of opportunities for risk mitigation. First, because of the converging paths, success of any one of them makes a successful FSP. It is essential to start early in the FSP execution phase and sustain the long term R&D while meeting short term deliverable schedule. Second, different science drivers are increasing being tackled by fewer but highly integrated (physics-wise) components. The long terms prospect of FSO is thus intellectually appealing and clearer in terms of the numbers of components and coupling schemes. Third, the simulation capabilities for the five science drivers converge nicely into the whole device modeling science driver. This makes a natural progression in terms of physics fidelity of FSP predictive modeling of a tokamak discharge.

4.2.2 Targeted FSP components

4.2.2.1 *Prioritization and findings*

One profound finding of the component planning is that there are massive amount of overlap in terms of component needs/specification among the six science drivers. The consolidated FSP program-wide component needs results in much reduced developmental and maintenance cost. Furthermore, the component needs of the two prioritized ISAs essentially cover the needs of the original six science drivers in scope. This finding supports a centralized component library for development and maintenance, and the prioritization becomes an issue of scheduling. Next we present the consolidated FSP physics component list and discuss their readiness, gaps, and developmental strategy.

4.2.2.2 *Profile evolution: quasi-static equilibrium*

Three classes of free boundary equilibrium solvers are required to determine the magnetic field configuration for given plasma pressure and current profiles and the external coil and wall currents. The first class is the free boundary Grad-Shafranov solver for an axisymmetric magnetic field. This is a capability which base program already has and is ready for immediate FSP component adaptation. The second class is the free boundary 3D equilibrium solver for non-axisymmetric magnetic field with closed flux surfaces. In the base program, these have been highly developed as stellerator equilibrium solvers and are ready for immediate FSP component adaptation. The third class is the free boundary 3D equilibrium solver for non-axisymmetric magnetic fields with magnetic islands and stochastic field lines. There are at least two code efforts in the base program in this area. The challenges are in numerical algorithms for inverting the large matrix, and the physics model formulation for specifying the plasma pressure and current in the regions of magnetic islands and stochasticity. Further development with FSP support is desired and anticipated for a FSP component deliverable in the 3-5 year time frame .

4.2.2.3 *Profile evolution: core transport*

The modeling of core transport, which governs the evolution of density, temperature, and current profiles in the center region of a tokamak, strongly depends on whether the magnetic field has closed flux surface or not, and whether the plasma density and temperature are flux function only. The baseline model for core transport is one dimensional model in which density and temperature is a flux function only. There are commonly three descriptions for such 1D transport calculation. The neoclassical transport calculation is mature and ready for FSP component adaptation, except for the determination of the radial electric field. The reduced anomalous transport model has been developed in the base program and is highly successful in modeling actual experiments. It is also ready for immediate FSP component adaptation. The third approach exploits the scale separation between micro-turbulence and the global profile evolution and deploys embedded turbulence transport calculation to evaluate transport flux from gyrokinetic simulations and feed it into the 1D transport equation for long time profile evolution. Much progress have been made in the base program, particularly using the embedded local turbulence calculation. This approach is ready for FSP component adaptation. Comparing the three approaches in a validation campaign is an activity of high value to the fusion program.

Even with closed flux surface, there has been renewed interest in determining the neoclassical radial electric field. Two dimensional transport models, which treat the variation of density and temperature both radially and poloidally, are active research issues. Significant theoretical and computational development within the base program are required before an FSP component can be constructed.

Both ISAs call for transport calculation in a 3D magnetic field which has islands and stochastic field lines. There are two approaches depending on whether the drift-kinetic or gyrokinetic models is employed. The drift kinetic approach has been substantially developed in the base program. Further development is desired for FSP componentization. The gyrokinetic approach is less developed compared with the drift kinetic approach. This is a research issue requiring continuing base and/or SciDAC support for development.

4.2.2.4 *Profile evolution: pedestal transport*

Four classes of FSP components are expected to address the pedestal profile evolution. The simplest, which is also an early deliverable, is a 1D static profile model. This is an ongoing effort in the base program and ready for FSP component adaptation. Improvement of the model is required and hence further development is desired. The pedestal region has strong sheared rotation and pressure gradient, so the issue of radial electric field and bootstrap current evaluation gains even greater importance than the core. 2D static profile model for the pedestal is a research issue and likely requires further development with support of the base program.

The third class of pedestal transport calculation, which still assumes an axisymmetric equilibrium magnetic field, employs embedded electromagnetic turbulence calculation to resolve the particle and heat flux and hence follow the dynamic profile evolution in the pedestal. Either a fluid or a kinetic model can be used for the underlying turbulence calculation. The fluid approach is mature with collisional closure and ready for FSP component adaptation. The actual pedestal plasma is of long mean-free-path, for which the fluid moment closure is a research issue and required base program support for development. There have been substantial efforts applying kinetic model for pedestal transport modeling, using both the continuum and PIC methods. It is recognized in the base program that proper gyrokinetic equation for the pedestal region is a research issue. There are opportunities for both FSP component adaptation and carrying out further development under FSP support. The collision model and implementation are current research issues in the base program. They require significant development with base/SciDAC support.

The fourth class of pedestal transport calculation is performed with a 3D quasi-static magnetic field that no long has closed magnetic surfaces. The transport model can be either drift-kinetic or gyrokinetic equations. There are ongoing effort using the drift-kinetic model. Immediate FSP component adaptation is planned, while the need for further development is clearly identified. The gyrokinetic approach is an active research topic and continuing support from the base and SciDAC program is expected for pre-FSP development.

4.2.2.5 *Profile evolution: scrape-off layer transport*

The tokamak scrape-off layer (SOL) plasmas reside on open magnetic field lines and are subject to rapid parallel transport to the divertor plates. The baseline model for a quasi-static SOL equilibrium must be 2D, i.e. strong radial and poloidal variation. The parallel transport can be modeled either with a fluid model or a kinetic model. For the fluid approach, the collisonal closure for parallel heat flux is mature and ready for immediate FSP adaptation. The kinetic model for heat flux, which would be more accurate, has ongoing effort in the base program, which can be adapted into FSP component. The need for further development is desired. The perpendicular transport in the SOL can be highly anomalous. The reduced model for diffusive and blobby cross-field transport can be straightforwardly adapted into FSP component. Of higher fidelity is the embedded SOL turbulence calculation. This also has a fluid and a kinetic approach. The fluid model with collisional closure is ready for FSP component adaptation, while the further development on closure is desired. The kinetic approach is undergoing active development in the base program. Additional development is required before an FSP component can be adapted. The neutral transport is an indispensible part of pedestal transport. The Monte-Carlo neutral transport is well-developed in the base program and ready for FSP component adaptation. Coupling MC and fluid neutral transport models in different regions of disparate collisionality is a research issue and requires continuing support.

The need for modeling pedestal profile and transport in a non-axisymmetric magnetic field is also established by the two ISA's. The drift-kinetic approach is further along in development. Both immediate FSP component adaptation and further development are planned. The gyrokinetic approach is an research issue requiring further development with base program support.

4.2.2.6 *Profile evolution: sheath transport*

One dimensional sheath model is the de-facto work horse for all existing SOL plasma models. The fluid model, especially the analytical one, is widely used, and can be immediately adapted into the FSP component. The case with oblique magnetic field, which is often the case for tokamak fusion reactor, is a more difficult problem. Further development, whether is supported by base program or FSP, is clearly needed. To that end, kinetic modeling of 1D sheath with oblique magnetic field has ongoing effort in the base program, which can be adapted and further developed under FSP.

A two-dimensional sheath model is needed if the divertor surface is rough or the equilibrium variation has to be taken into account. There were analytical efforts in this area and can be adapted into FSP component. Computational models, based on either fluid or kinetic models, are clearly in need of development.

There is also the concern that the sheath can be quite dynamic. This could be the result of upstream turbulence being convected into the sheath, or locally driven by sheath instabilities. Embedded sheath turbulence calculation can be based on either a fluid or a kinetic model. The usefulness of a fluid model is a research issue. There are ongoing efforts using the kinetic model. Both adaptation and further development are identified such component.

4.2.2.7 *Profile evolution: energetic particle transport*

Energetic particles are produced in a tokamak plasma by fusion reactor, neutral beam injection, RF heating. The energetic ion slowing down is usually modeled by a Fokker-Planck solver or a quasi-linear model. Both are highly developed in the base program and ready for FSP component adaptation and further improvement. The more difficult physics of the nonlinear saturation of energetic particle driven modes and their effect on energetic ion slowing down and radial transport required more powerful tools. This requires a self-consistent calculation of how the modes evolve and how the energetic particle population responds. The kinetic-MHD approach has an o effort in the base/SciDAC program. So is the approach based on gyrokinetic equations. Further development is clearly needs with both approaches.

The runaway electron transport is a key issue in WDM/disruption ISA. The collisional slowing down is modeled by Fokker-Planck solver, which is mature from base program support and ready for FSP component adaptation. The collective effect on runaway electron slowing down is a research issue. Although the theoretical model is actively pursued and amiable for FSP adaptation. Computational development is lacking and significant new development is required of base program and possibly FSP support, due to the importance and urgency of the issue to ITER.

4.2.2.8 *Profile evolution: sources and sinks*

Modeling of neutral beam injection is mature and ready for immediate FSP component adaptation. RF heating and current drive has been modeled with different level of fidelity. The ray tracing approach is mature and ready for immediate FSP component adaptation. The full wave solver has ongoing effort using spectral and finite element approaches, which are ready for adaptation. Modeling the RF sheath and the RF penetration through the SOL plasmas, however, is perceived to have a substantial gap, and significant investment from the base/SciDAC/FSP would be required.

Fueling by pellet injection and gas puffing, and inversely gas pumping as a sink to control boundary plasma, are essential elements in WDM and Boundary integrated modeling. Rudimentary capabilities exist in the base program, which can be immediately adapted into FSP components. Developing a more sophisticated capability, both in terms of mathematical models and computational implementation, has substantial challenges, which requires significant base/FSP support.

Fusion reaction and alpha slowing down are the sources for plasma self-heating in a tokamak reactor. The alpha and neutron birth profile are routinely calculated in existing WDM codes, and can be readily adapted into FSP components. The alpha slowing down calculation shares the same component as described earlier for energetic ion transport. The neutron transport is only modeled for diagnostics purpose, which is also readily adapted into FSP component.

The radiation generation and transport provide an important energy loss channel and can have an appreciable effect on boundary plasma profile and evolution. The immediate needs of FSP can be met by adaptation of existing capabilities in tokamak edge and WDM modeling codes. Further development is desired to treat more challenging problem such as the density-limit disruption in a tokamak.

4.2.2.9 *Profile evolution: wall boundary condition*

The divertor and first wall provide the boundary condition for a whole device plasma simulation in a tokamak. From the plasma perspective, important physics issues include particle recycling, impurity generation and transport, wall electron emission, and dust generation and transport. Current modeling of recycling is rudimentary and empirical. It can be quickly adapted into FSP components. A more desirable capability would be based on coupled molecular dynamics and plasma kinetic simulation. This is a research issue, requiring substantial investment from the base/SciDAC/FSP programs. Impurity generation faces a similar issue. The commonly used data is based on SRIM calculation. More sophisticated MD simulation, on a non-ideal material surface, can substantially improve the fidelity. Again, substantial investment is required here, although it is heavily shared with the recycling capability development. The impurity transport is in better shape by comparison. For example, test particle approach is well development and can be readily adapted into FSP components. Wall electron emission is largely modeled by reduced or empirical model, which can be easily adapted into FSP components. It remains to be seen, through experimental validation, how adequate these models are. Dust generation is more or less a mystery in a tokamak. In the near term, the best hope is a heuristic or empirical model. The longer term might see a MD based capability that could be exercises to produce an improved reduced model. The dust transport, especially away from the wall, is more developed and ready to be adapted into FSP components. Near-wall dust transport requires more work.

From the materials perspective, the change of surface morphology, and the mechanical, thermal, and electrical properties, is of great importance to the viability of a fusion reactor. These are all research issues where little organized research has been carried under the fusion program. Addressing these would require a hierarchy of physics models from density functional theory, to MD, to phase field, to continuum models. Similarly the tritium retention is an issue of paramount importance to tokamak fusion. There is ongoing research in this area, but substantial development is required before a credible model is ready for FSP adaptation.

4.2.2.10 *Off-normal event detection*

Detection of off-normal event can be facilitated by linear and nonlinear stability analysis of a dynamically-evolving quasi-steady-state plasma. For macro-stability, a whole array of linear stability codes is ready for FSP component adaptation. These include ideal and resistive MHD, with or without rotation, and kinetic MHD including the energetic particles and kinetic plasma dissipation physics. Similarly for micro-stability, there is an array of linear stability codes for FSP adaptation to address both electrostatic and electromagnetic gyrokinetic modes. The energetic particle modes are also calculated with great confidence and ready for FSP adaptation. The areas requiring substantial investment are linear macro-stability with 3D equilibrium with magnetic islands and stochasticity, and nonlinear stability such as the NTM onset threshold.

4.2.2.11 *Off-normal event mitigation*

Three classes of off-normal events stand out in FSP planning and they require improved understanding for avoidance by active control and mitigation. The first is disruption. There are three approaches. For fast disruption, initial value extended MHD simulation is the standard tool, and the SciDAC program provides capabilities that are ready for immediate FSP component adaptation. The quasi-static 3D equilibria approach is considered a superior tool for tracking the slow disruption. There are two parallel efforts in the base program. With some additional development, they will be ready for FSP adaptation in the near future. The runaway electron generation and transport are essential element in disruption modeling. This was covered earlier under energetic particle transport in the profile evolution section.

The ELM and divertor/wall loads are key issues in the boundary ISA. There are three approaches to model this problem. The initial value two-fluid simulation is the most developed and ready for FSP component adaptation. A more accurate model would be a drift kinetic closure for the MHD equations. The drift-kinetic MHD is a research issue, requiring substantial development under base FSP programs.

The most ambitious approach is based on initial value electromagnetic gyrokinetic model, which in principle, could also tackle the L-H transition problem. There are ongoing efforts in this area with base/SciDAC program support. Substantial development is required toward an FSP component.

The nonlinear evolution of energetic particle driven modes and the fast particle transport can be addressed by at least two different models. One is based on initial value kinetic-fluid hybrid model, the other is based on electromagnetic gyrokinetic simulation. This area overlapped with the profile evolution needs in energetic particle transport, which was discussed there.

4.2.2.12 *Plasma control*

Simulating plasma control requires the modeling of the actuators and the plasma response. The plasma response could be done by the quasi-static equilibrium model or the fully dynamic model, both of which are shared with those component requirements specified in the off-normal event mitigation section.

The modeling of the actuators includes the image currents in the wall, the external current in the control coils, and the control circuit itself. There are ongoing efforts in all these areas, which lend to an immediate adaptation into FSP components. It is recognized, especially in the control circuit area, that rapid progress in the base program is being made and there is excellent opportunity for substantial upgrade and new development. Additional control of tokamak plasmas include instability suppression by localized current drive and shear flow generation, density control by fueling and gas pumping. Modeling of these actuators has been described in the sources and sinks, and wall boundary condition sections.

4.2.2.13 *In-situ synthetic diagnostics*

An array of in-situ synthetic diagnostic capabilities will be implemented as FSP components or component templates by adapting existing and developing new capabilities. These include diagnostics which characterize plasma profile (density, temperature, rotation with or without 3D magnetic perturbations) and fluctuations (magnetic, density, temperature and its anisotropy, and fast particle population). The goal is to provide synthetic diagnostics capabilities for every known experimental diagnostic technique. It is planned that component templates will be prepared for facilitating rapid component deployment for actual experiment (machine) which tends to have machine specification that significant impact the actual synthetic diagnostic component code.

4.2.3 Prioritized near-term component development plan

At the inception of FSP, two integrated science applications (ISA) will be launched. The first ISA is on whole device modeling, which covers both scenario modeling of a ITER discharge and the issue of disruption avoidance and mitigation. The second ISA is on edge and pedestal plasmas, which addresses ELM control and mitigation and explores a PMI solution to a steady state reactor. Two groups of component projects will be initiated. The first focuses on adaptation of existing codes, while the second involves new development or a significant scale up of existing capabilities.

The component adaptation projects target three classes of capabilities. The first is on profile evolution. These include:

- free boundary Grad-Shafarnov equilibrium solver;
- 1D neoclassical and reduced transport model;
- local turbulence-based transport model;
- pedestal and SOL equilibrium model;
- drift-kinetic pedestal and SOL transport model;
- RF and NBI sources in a axisymmetric plasma;
- gas puffing and pellet fueling model;
- 1D PMI model;
- impurity transport model;
- neutral transport model.

The second is on instability detection. This includes:

- a suite of ideal and resistive MHD stability codes;
- pedestal stability codes; and
- energetic particle stability codes.

The third class of components addresses the nonlinear evolution of a tokamak plasma. Specifically, we will adapt

- initial value 3D MHD codes for core MHD and ELM dynamics;
- initial value gyrokinetic codes for core transport and the energetic particle evolution; and
- initial value Braginskii codes for ELM and SOL turbulence.

The new component development projects cover the component needs from both the two early ISA’s and those to be launched after a full FSP ramp-up. The plasma physics code projects are:

- quasi-static evolution of coupled 3D fields and neoclassical and turbulent transport;
- initial value gyrokinetic solver in 3D fields;
- initial value kinetic MHD solver; and
- initial value gyrokinetic solver for the pedestal/SOL plasmas.

The PMI component project focuses on resolving the plasma recycling at the first wall/divertor. It involves coupling kinetic plasma models for boundary plasma with molecular dynamic models for materials response to plasma irradiation, neutral transport and atomic physics model for ionization and radiation.

4.3 Software Integration & Support

4.3.1 Scope and organization of this section

As a reminder, the Integrated Science Applications (ISAs) will be launched at the start of the FSP and so will be developing capability to address the early targets identified by the Science Driver teams in advance of any integration efforts. It is expected that they will make use of integration software that exists at the time of the start of the FSP, e.g., that from the proto-FSPs or from outside of the fusion community. Hence, the Integration effort associated with the Fusion Simulation Project will be oriented towards developing common software and tools that can be used for later developments of the ISAs, either as the first-launched ISAs add capability or refactor, or for the development of ISAs that are launched later. As well, this section will discuss the developer support that is needed for the undertaking of the Fusion Simulation Project.

In order to break down the problem of integration (sometimes called "composition", which we use synonymously), we have separated it into two areas. On-HPC Integration is defined to be that which is done to couple software on a single high-performance computer, which of course is constructed from many CPUs, but for which upon job launch, a global MPI communicator is defined. Of course, the same integration can occur on a single-CPU workstation. Early in the Definition Project we defined this as Physics Composition, which was a bit of a misnomer, as physics software is also integrated through some method of coupling executables that run on different HPCs or workstations. Another way of defining On-HPC integration is that it is the integration that occurs after release from a job queue and prior to job completion on an HPC.

The term, Task Composition or Off-HPC integration is the coupling outside the HPC environment. This is also called Workflow. Task composition consists of the composition or integration that is outside of the execution on an HPC, from concept to research result. It can include tools for input file preparation, submission to a job queue, job monitoring, data visualization and analysis including comparing with experimental data, data archiving and so forth.

These two integration areas are separated because it is possible that one will use different integration methods in these two areas. Specifically, if one does not have a global MPI communicator, one cannot integrate software using MPI communication. On the other hand, in On-HPC Integration, one generally need not worry about authentication and authorization in the communication.

Use of the terms, framework or workflow, are specifically avoided. These two terms are heavily overloaded and, as discussions during the FSP have shown, have different meanings for different members of the community. Nevertheless, in either On-HPC Integration or Task Composition, we will incorporate the concepts used to describe frameworks [1]. In particular, the composition software should provide enabling tools for composing physics components and analysis components.

We also include in this section some discussion of data management and developer support. Data management covers the down-selection of data for archiving, the actual archiving, and the later retrieval of that data of post analysis. Developer Support includes the installation, adaptation, and maintenance of the tools that allow efficient and effective software development. These include items such as revision control, build and package management systems, regression testing systems, and so forth.

4.3.2 Requirements

The science-driver teams identified requirements that the software integration and support (SIS or "frameworks") effort should satisfy in order to facilitate the achievement of the ISA goals. Some of the requirements were identified by more than one of the ISA's, although the time at which a given requirement is anticipated to be first needed varies between the ISA's. The requirements can be grouped into the various functional software-integration/support categories (On-HPC coupling, task composition/workflows, data management, developer support). The guidance to the ISA teams, and the ISA reports, are most specific with respect to the On-HPC coupling and task-composition/workflow requirements for the SIS activity, although there were also some comments on the other categories that are more general in nature and less specific with respect to the time frame on which they are relevant to the achievement of particular ISA milestones.

In addition to the requirements articulated in the ISA reports, some possibly relevant requirements and considerations are articulated in other reports requested by the FSP planning teams. For example, the proto-FSP assessment report states "We expect that advanced time-advance algorithms will be required for different aspects of the multi-physics integration; It is important that the FSP framework be flexible enough to support any or all of these."

The following requirements were identified in the ISA reports.

High-Performance Composition: The composition software should support running in parallel on large computer systems. It is expected that some of the components used in many of the FSP code simulations will need to be run on either "capability" or larger "capacity" computational platforms. This is identified as a requirement for the WDM year 5 milestones to couple RF and fast ion components, and for core turbulence-core transport coupling. Also, the WDM ISA has as year 5-8 tasks the "demonstration of high performance FSP WDM prototypes under WDM framework" and "Optimization for WDM parallel architectures."

Communication of large data sets between components: For interoperation of the current suite of computational components, it will be necessary for them to be able to access the same data structures of multiple dimensionalities. For 1D and 2D couplings (such as transport fluxes or equilibrium data), the exchange of data is not expected to be a problem, but for 3D (and even higher dimensionalities, such as occur in plasma-neutral coupling), new methods may need to be developed, as the strategies appropriate and applicable to both present day and future platforms will be needed. For 3D there exist multiple strategies. Traditionally on HPC platforms, it was important to have data locality to minimize communication to only the surface data at computational domain boundaries. This will likely get only worse as one moves to higher-end platforms. This indicates a strategy of using extended data structures and domain decomposition. On the other hand, coupling to diagnostic components can be done less frequently, and given that the associated analysis is often done Off-HPC, coupling by files or memory data space may be sufficient.

This requirement was identified as a requirement by the Control of Disruptions ISA, for the coupling of a global gyrokinetic code with a 3D equilibrium code (with magnetic islands). In years 1 and 2, the gyrokinetic code to be coupled will be a neoclassical gyrokinetic code with self-consistent calculation of the bootstrap current while in year 3 it will be a gyrokinetic turbulence code. Volumetric or larger data sets are needed for IB year 1-3 milestones involving coupling kinetic transport and neutral models.

Rapid sharing of the data involved in coupling among components: Here, "rapid" means on a time scale short compared with data manipulation. The coupling may be required every time step or less frequently, for example a turnover time or (a set fraction of the) transport time in a case where component is a turbulence simulation code. The frequency with which the coupling is done, the data volume involved per coupling instance, and the effective data communication rate for a given data communication mechanism (e.g., file-based writes and reads, on-node shared memory access, or inter-node message passing) determine an effective coupling data communication time per time step or coupling iteration. This time should be much less than the time needed by the components to do calculations based on that data. This, in turn, sets a requirement on the particular data communication mechanism needed based on the data communication rate that it can support. Thus, the ISA reports have either identified "tight" in-memory coupling as a requirement or have stated that the slowest of the communication mechanisms, file writes and reads ("loose coupling"), is adequate.

"Tight coupling" (likely some form of in-memory coupling) is identified as a requirement by the following ISA's: The WDM report states generally that the framework should "Provide the infrastructure to enable various types of code coupling satisfying computational considerations, for example "tight in-memory coupling...." Further, this is identified as a requirement for a year 2-3 milestone to couple core and edge dynamics components, and for a year 5-8 milestone to couple core turbulence and core transport. The Pedestal ISA models, which are planned for development; year 3-5 tasks involving coupling transport and sources in evolving magnetic geometry in ELM evolution simulations. For the coupling scheme envisioned for the Wave-Particle Resonances (WP) ISA over the 10–-15 year time frame: "Since the energetic particle (EP) component now includes RF effects due to an ICRF or ECRF induced flux in the MHD equations" (for modeling e.g., sawteeth), "coupling between the EP and RF sources is now tightly coupled." The Boundary ISA identifies this as a requirement for year 1-5 tasks for boundary turbulence-transport coupling and plasma-neutral coupling. The Disruption ISA requires support for in-memory coupling for a year 5 milestone to study the effect of detailed versus simple structural deformation models coupled to MHD dynamics. The Core Profiles ISA has as a year 2 milestone: "Deliver a prototype framework for a time-dependent 1.5D transport solver built from legacy components…" As per the discussion of the WDM requirements, this will likely involve tight implicit coupling. File-based coupling is identified as being sufficient for some ISA coupling tasks, and several of the ISA's propose using this in early versions of their composed software. Furthermore, file-based coupling has the potential advantage that file data structures are more readily permit storage of metadata for self-describing data formats. The Pedestal ISA report, for example, argues that implementation of self-describing data formats for data involved in couplings and that can automate the conversion between different representations would facilitate physics composition. The WDM report gives as a desirable attribute "Efficient and flexible I/O libraries with rich metadata to support large-scale physics components."

Support for implicit coupling: Implicit coupling may be needed, for example when a stiff macroscopic model is involved. The ISA reports are not specific on their particular software needs in support of implicit coupling, but mention differencing, the ability of components to revert to a previous state (which may be more a requirement on the components rather than on the integration software), and nonlinear solvers such as Newton-Krylov (including Jacobian-free Newton-Krylov) solvers.

Implicit coupling was identified as important in the following science drivers: The WDM report states generally that the framework should "Provide the infrastructure to enable various types of code coupling satisfying …algorithmic considerations, for example implicit coupling…" and further, "The framework will need to provide a capability to combine different components under explicit multi-rate, implicit-explicit, and fully implicit time advancement techniques." Implicit coupling is identified by the WDM ISA as a requirement for its activity on 2.5D equilibrium-transport coupling (i.e., coupling a 3D equilibrium component to a 1D transport solver), a year 2-3 milestone on core edge dynamics coupling, and a year 5-8 milestone to couple core turbulence and core transport. The Integrated Boundary (IB) ISA report notes: "The importance of time-implicit coupling depends on the shortest timescale in each component. If a time-implicit component is explicitly coupled to a second component that has a fast timescale, the timestep required for the first component will likely degrade to the explicit timescale of the second." Support for implicit coupling will be needed for IB year 2, -5 and -10 milestones involving boundary turbulence-transport coupling. Requirements related to coupling WDM and material-wall components, with related milestones loosely at years 2 and 10: "Because the transport modeling is over a long time scale, implicit coupling is needed to obtain accurate, self-consistent fluxes between the edge and wall components." The Core Profiles ISA has as a year 2 milestone: "Deliver a prototype framework for a time-dependent 1.5D transport solver built from legacy components…" As per the discussion of the WDM requirements, this will likely involve tight implicit coupling.

Data transformation software as part of integration software: The next three requirements discussed below (support for constrained interpolation, smoothing, and refinement) can be viewed as support for efficient transformation of instantaneous data between representations. Although it is not explicitly mentioned, the use of particle-to-mesh "deposition" is implied for some of the ISA tasks and milestones. This falls into the same category, and could be provided by the integration software.

Support for constrained interpolation, e.g., conservative (of mass, momentum, and energy) interpolation or limited interpolation (i.e., using a limiter to prevent the creation of spurious extrema): The WDM report states generally that the framework should "Provide the infrastructure to enable various types of code coupling satisfying… algorithmic considerations, for example… conservative and limited (to eliminate spurious extrema) interpolation between different grid representations. This is identified as a requirement for a WDM year 2-3 milestone on core-edge dynamics coupling. The Integrated Boundary report identifies this as a requirement, for example for a year 2-3 milestone involving core-edge dynamics coupling, a year 3 milestone involving coupling of kinetic plasma and neutral components, year 5-8 year milestones involving core turbulence-transport coupling. The Pedestal report identifies this as a requirement for year 3-5 milestones involving coupling material wall models to scrapeoff-layer transport.

Support for filtering/smoothing ("data reduction") of data from particle based models coupled to continuum models: The Integrated Boundary report identifies this as a requirement for: years - 1-2 milestones to couple of particle-based sputter erosion/redeposition code for 2D impurities and SOL 2D fluid plasma model; years - 3-5 milestone to couple develop and improve coupling of fluid transport model with particle-based kinetic neutrals model. The Pedestal ISA also gives this as a requirement for a similar physics coupling; year 3-5 milestones for its "level 2 and 3" models, involving coupling material wall models to scrapeoff-layer transport.

Support for stochastic algorithms: This is to providing reproducibility for testing and verification when stochastic (e.g., Monte-Carlo) algorithms are involved, and entails the ability to efficiently produce and use sequences of pseudo-random numbers independent of the computational platform or parallel decomposition. This is given as a requirement for the immediately preceding Pedestal milestones.

Support for refinement of data used in coupling: These are identified as a requirement for the Disruptions ISA for a year 1 milestone involving coupling WDM and linear-MHD components. The scheme envisioned will generate a "cloud" of nearby fine interpolations of a more coarsely represented MHD equilibrium for a statistical quantification of the uncertainties in linear stability boundaries. This is needed because the linear ideal MHD operator is stiff, so that small errors in the MHD equilibrium solution and mapping codes can cause substantial error in the linear MHD solution. An additional requirement raised by this milestone is the ability to initiate and run a large set of small trivially decomposed computational tasks.

Framework flexibility: A requirement for the life of the Core Profiles ISA effort, in order to facilitate development of mesoscale models, is to "guarantee flexibility of the FSP framework to allow researchers to propose, implement and test independent ad-hoc models."

Access to documented standardized tables and efficient interpolation algorithms: This is raised as being important, for example for the treatment of ionization, recombination, sputtering, etc., and is identified as a requirement by the Boundary ISA.

Standardization of database storage and access: The Disruptions ISA identifies this as a requirement to be fulfilled as early as possible for storage and access for analysis to experimental disruption data. Several task-composition/workflow related requirements are discussed in the ISA reports, without an assignment of particular task schedules or milestones for these. The IB, and WDM reports give as a general requirement a *(cross-platform) build system*. The Pedestal report has as an ongoing task, throughout the life of the project, *unifying build systems*.

Several of the ISA reports refer to a need for workflow software level support for visualization. The most demanding of these is in the Pedestal report: "there will be needs for data analysis (e.g., synthetic diagnostics) and visualization of 1D, 2D, and 3D scalar and vector data. As a post processing activity, these may require both interactive and batch processing, both of which should be enabled from within the framework."

Tools for input file preparation and validation are given as requirements for the integration software in the Pedestal and WDM reports. The most detailed specification is, again, in the Pedestal report: "Ideally, the framework would provide tools, perhaps graphical in nature, that would allow for easier configuration and that would automatically generate valid input files for each of the components."

The following are listed and discussed in the WDM and Pedestal ISA reports.

Universal workflow software, in the form of either a graphical system such as Kepler or a script-based solution (such as shell scripts or Python), or perhaps some combination of these.

Ability to record provenance data. For example, from the Pedestal report: "The metadata required to uniquely identify reproduce any simulation (date, platform, component versions, source file identifications, physics composition, parallel decomposition, etc.) must be acquired and preserved. The framework should provide tools to automate the acquisition and storage of provenance information.... In addition, the FSP should maintain a repository of input data files, and the framework should provide a capability to easily obtain data from this repository."

Interactive simulation monitoring: The WDM report explicitly advocates a "web portal or dashboard system that can drive graphical analysis tools and display results in real time" and which would satisfy both interactive simulation monitoring and some integrated data analysis needs.

Restart capability "...to initiate and coordinate the check-pointing and restart for all components combined in a physics simulation..."

4.3.3 Past approaches to development of Integrated Science Applications

As mentioned above, the Integration effort is intended to assist the development of ISAs in the out years, as the ISAs, in order to provide early deliverables, will have to rely on integration software existing prior to the FSP. In this section we describe some of that pre-existing software and some positions that have been stated in the community. In particular, Section 4.3.3.1 discusses a vision for including the legacy frameworks in the FSP, Sections 4.3.3.3, 4.3.3.4, 4.3.3.5, and 4.3.3.6 are descriptions of the integration methods produced by the "proto-FSPs" (the fusion integration efforts that were started 4-5 years ago), and descriptions of the approaches of communities (Climate and European Union [E.U.] fusion) that are outside the U.S. community. In each of these sections we also provide the answers to clarifying questions to the various position providers. This subsection concludes with a discussion of the approaches.

4.3.3.1 *Use of legacy serial frameworks*

FSP can take advantage of legacy frameworks, particularly for whole device modeling (WDM) applications. These frameworks have advantages, because they are familiar to an existing, sizable user community. For example, the TRANSP/PTRANSP WDM production system was used for over 25,000 time dependent tokamak simulations, and over 1200 ITER PTRANSP predictive simulations, in Fiscal Years 2005 – 2010. Examples of resulting publications are referenced [[9]], [[10]].

In the international fusion community in recent years, several dozen users have prepared input data for TRANSP/PTRANSP simulations, and, hundreds have used the output data. The associated data management systems are mature from a fusion research community perspective. This means that they are well integrated with existing experimental data archives (MDSPlus) and there are a rich set of tools for visualization, analysis and post-processing. This level of integration represents many tens of man-years of work, financed over a period of decades by the MFE base program, mainly the experimental projects. Although TRANSP/PTRANSP has never been adapted for supercomputing, it runs well in a serial and low-Np cluster MPI parallel mode. There is an active, collaborative user community that shares template namelists and scripts and other information useful for preparation of input and interpretation of output. TRANSP/PTRANSP also has mature mechanisms for change management, i.e. adding inputs and outputs to adapt to the evolving requirements of research applications. It will advantageous to FSP, especially with respect to prospects for deliverables early in the project, if this work and knowledge can be leveraged rather than reinvented from scratch.

The SWIM proto-FSP framework [[11]] permits the loose coupling of independently built components running as separate processes in a supercomputing environment. It can accommodate serial and low-Np parallel components working in cooperation with high-Np supercomputing models. The TSC WDM has been integrated usually as a serial implementation of the “equilibrium and profile advance” component in SWIM.

It would be useful, and it is proposed, to integrate TRANSP/PTRANSP in the SWIM framework, as an early activity for FSP. It would be useful to configure SWIM so that its task composition and component options are controllable from the user-familiar TRANSP namelist, that access to SWIM services are provided as an extension to the existing TRANSP/PTRANSP run production system, and that the MDSPlus integrated TRANSP outputs are produced as part of the SWIM run execution (in addition to normal SWIM output). This is all achievable with a relatively modest engineering effort (about 1 man-year).

Early benefits to users would include access to high resolution parallel RF supercomputing components such as AORSA and TORIC. TRANSP/PTRANSP already makes extensive use of Plasma States for internal communications and these can be provided for interface to sophisticated MHD components e.g. for evaluation of MHD stability

---

[9] R. V. Budny et al ., Predictions of H-mode Performance in ITER, Nuclear Fusion 48, No. 7 (July 2008)

[10] R. V. Budny et al ., Comparisons of Predicted Performance in ITER H-mode Plasmas with Various Mixes of External Heating, Nuclear Fusion 49, No. 8 (Aug. 2009)

[11] http://www.cswim.org

of an evolving ITER model. Also, on the time scale of FSP project initiation, TRANSP/PTRANSP is expected to include a free boundary MHD equilibrium predictive capability, which opens the door for coupling to scrape off layer plasma models.

The entire TRANSP/PTRANSP WDM control is through a single code-generated namelist structure. This single point of control is a benefit for user “sanity”, but likely won’t scale to all future FSP applications.

4.3.3.2 *Service oriented frameworks*

4.3.3.2.1 Description and/or position

Service oriented architecture (SOA) separates science functions into distinct executables (called extended components here), which developers make accessible to the multiple integration frameworks. The extended components/services can be simple or complex. They are comprised of physics, applied mathematics or computer science executable routines. Each services use APIs to communicate with other services. In this approach, all the possible integration components do not need to be compiled together using one compiler in order to be executed by the integration framework. Each extended component can be compiled in the framework using its own compiler and library options, which allows developers of one service to be completely shielded from the developers of other services. The APIs needs to be simple, user friendly, portable, efficient, and scalable to next generation computer platforms. After the integration APIs are installed in the I/O layer, each extended component can be independently developed and executed, while the same version is used by the integration frameworks. In this way, the service oriented architecture can have various types of extended components: single processor to extreme scale parallel processors, and shared to distributed memories. In this way, the code debugging can be localized to each extended component while coupled in the integration framework, along with the ability to be called by different integration frameworks without additional labor by the code developer. This allows the integration frameworks to use automatically the most recent version extended components, without worrying about the complexities of the `final integration’ of all of the components into a single executable. Since the computer science tools are independent services from the physics components, they are separately developed as flexible and “living” components in rapidly changing software/hardware environments, hence the lifecycle of the SOA framework will be much longer than a rigid integration framework. SOA approaches are developed to handle the complexity which comes from the coordination across the code developers, deployment teams, and research teams. Thus, SOA approaches are commonplace in the enterprise because of the cost savings.

The extended components do not have to reside on one computer. If convenient, some extended components can reside on local data analysis computers, while others reside on remote HPC, during the integrated simulation. In CPES, the integration protocols within an HPC are defined by the ADIOS advanced adaptive I/O protocols. The interfaces between extended components are implement through the I/O layers. Coupling is achieved through files, sockets, memory from another node, or direct memory references with zero copies. The inter platform integration is orchestrated by Kepler workflow. Coupling workflow within an HPC is achieved by Adios, which includes staging memory operation and data management, and run-time simulation control. The mathematical coupling algorithm can be strong or weak. The functionality of ADIOS and Kepler SOA has been demonstrated to petascale computing.

4.3.3.2.2 Answers to follow-up questions

1. The requirements team, from a reading of the ISA requirements document, has identified the following requirements (documents attached). Please provide one sentence addressing each of these, such as "possible by further development", "prototyped", "implemented in special cases", "part of production software installed at the LCFs at <hostname:directory> for general use".

- High-performance framework (runs on an HPC): Production ADIOS/Kepler are installed at LCFs (Jaguar and Franklin) for general use.
- Rapid sharing of the data involved in coupling among components: DataSpaces is available at Oak Ridge National Lab's National Center for Computational Sciences, at least for rapid data sharing.

- Support for implicit coupling: A mathematically tight in-memory coupling case will soon be released.
- Self-describing data formats for data involved in couplings: Available via ADIOS.
- Support for filtering/smoothing ("data reduction") of data from particle-based models coupled to continuum models: Particle data smoothing methods are already used for coupling to continuum models.
- Support for refinement of data used in coupling: Available via ADIOS.
- Ability to initiate and run a large set of small trivially decomposed computational tasks: Available via Kepler workflow engine and ADIOS.
- Universal workflow software: Available via Kepler. ADIOS is being developed for universal intra machine coupling.
- Ability to record provenance data on physics models, computer system, compilers/libraries, etc.: Available via Kepler.
- Input file preparation, staging, and validation: Possible by further development of the eSimMon dashboard.
- File migration between parallel computing facilities and integrated data management systems: Available via DataMover system and Kepler workflow engine.
- Non-interactive data analysis and visualizations via scripted tools or services (such as IDL or Visit): Available via eSiMon Dashboard.
- Interactive simulation monitoring and data analysis through web portal or dashboard systems that can drive graphical analysis tools and display results in real time: Available via eSiMon dashboard.

2. Please provide a listing of other projects that use your approach for:

A. Task composition (aka workflow, inter-computer coupling)

- Kepler is used in Euforia (E.U. Fusion for ITER Application), pPod(phylogenetic analyses), REAP (Real-time Environment for Analytical Processing, SPA (SANParks: Managing Wildlife Populations, Scientific Process Automation at SDM, COMET (COast-to-Mountain Environmental Transect Project), Clotho (Biological Integration)

B. On-HPC coupling (formerly known as physics composition)

- No response.

3. What needs further development that would require FSP level resources?

- To begin with, standard descriptors for data items that are to be commonly exchanged between extended components at all hierarchical levels. Also, intelligent GUIs that would facilitate the construction and launching of complex coupled simulations by physicists who are not necessarily computer science experts.

4. If there were to be a software integration effort aimed in part at developing software for use by any integrating software application (more than one is likely to be the FSP case), what software that you have proposed could become part of that effort?

- Certainly the ADIOS library can be made use of, both as a means of achieving flexible and efficient I/O on petascale and exascale systems and as a way to approach On-HPC application coupling. Kepler is a viable option as a workflow system that provides built-in mechanisms for provenance tracking and supports pipeline parallelism (key to efficient management of workflows with many semi-independent tasks). The DataSpaces tool provides a simple means of conducting memory-based coupling. The eSimMon dashboard is a relatively sophisticated web portal that combines resource and job monitoring, data analysis and visualization, and easy collaboration and data sharing in a single location.

5. In particular, for higher-dimensional couplings, have you developed conservative or limiting (in the sense of removing spurious extrema) transformation software for sending data between components that could move forward into FSP.

- This could be done in the context of the DataSpaces tool, for example.

6. Please provide a list of OS's (specific hostnames for LCFs would be best) to which your approach has been ported.

- No response.

4.3.3.3 *Component composition*

4.3.3.3.1 Description and/or position

FACETS uses a component based architecture [1] for On-HPC multi-physics computing. The FACETS executable is created with implementations for all components linked in, but at run time only the needed components are instantiated. For concurrent execution, the FACETS executable starts with MPI_COMM_WORLD and recursively splits communicators, allocating processors to components, with the splitting associated with any particular composition (i.e. core-edge, core-turbulent transport, etc.) determined by the input file. In the current mode of operation, all of the On-HPC communication is through the communicators, none through files, though neither file communication nor use of separate executables is precluded. Notably, the single-MPI-executable model is compatible with all existing supercomputers.

A component composition approach requires careful design of interface methods to allow easy incorporation of new components. FACETS defines component interface methods for (a) startup/complete phase, (b) data access, (c) update phase and (d) dump/restore phase. The update phase was specifically designed for implementing implicit coupled advances. (FACETS includes both explicit and quasi-Newton based implicit coupling implementations.) Standard names determine data exchanged between the components. These names allow runtime selection of coupling variables.

FACETS allows both internal components, newly developed components that use the FACETS distributed data structures, and external components, which have to now been wrapped legacy components. An example of an internal component is the new core solver [2] using FACETS. This component uses a nested-iteration based implicit solver to speed up the calculation of stiff transport in the tokamak core. The core solver is coupled to an edge component for use in core-edge simulations [3]. This approach has also been used outside of FACETS to create non-fusion components for the solution of plasma fluid equations [4] on structured and unstructured grids. The component composition approach does require modifications to external components to allow incorporation into FACETS, but we have found that these changes are generally small compared with those required to improve the software engineering. Lessons learned: software engineering and degree of validation should be a criterion for component selection, standards are increasingly important in a large, distributed development team.

For task-composition, FACETS uses a scripting approach, as it allows users complete flexibility in problem set up. This scripting is used for loose, file-based coupling of serial executables for problem setup prior to queue submission and for data analysis upon job conclusion.

FACETS addresses several of the whole device modeling (WDM) needs identified in the science driver analyses. In particular it provides a WDM high-performance framework, components for core-edge dynamics, core-turbulence-transport calculations and coupled core-edge-material-wall interaction. Each of these areas requires both explicit and implicit algorithms and leadership class computing facilities (LCF). For example, implicit algorithms are required for core-edge-equilibrium coupling, while explicit algorithms are typically sufficient for core-source coupling; LCFs are required for core-turbulence-transport calculations using gyrokinetics.

Future development will involve new high fidelity physics models and improvements to the coupling infrastructure. For example, we will bring in edge turbulence components starting with fluid turbulence models, eventually incorporating gyrokinetic edge turbulence models. The introduction of 2D surfacial coupling will require implementation of conservative grid-to-grid mapping schemes as well as development of efficient transfer mechanisms of large datasets across different processors.

4.3.3.3.2 Answers to follow-up questions

1. The requirements team, from a reading of the ISA requirements document, has identified the following requirements (documents attached). Please provide one sentence addressing each of these, such as "possible by further development", "prototyped", "implemented in special cases", "part of production software installed at the LCFs at <hostname:directory> for general use".

- High-performance framework (runs on an HPC): FACETS runs on all HPC machines.
- Rapid sharing of the data involved in coupling among components: Yes, via in-memory coupling.
- Support for implicit coupling: Yes, via quasi-Newton based coupling.
- Self-describing data formats for data involved in couplings: Possible by further development.
- Support for filtering/smoothing ("data reduction") of data from particle based models coupled to continuum models: One simply needs to write a FACETS updater to do this, hence possible with very little development.
- Support for refinement of data used in coupling: Possible by further specific development for general case, done in specific cases (e.g., equilibrium refinement using fluxgrid).
- Ability to initiate and run a large set of small trivially decomposed computational tasks: Yes, via task-farming design pattern in FACETS.
- Universal workflow software: No, FACETS uses shell scripts
- Ability to record provenance data on physics models, computer system, compilers/libraries, etc.: Yes, via compile-time provenance collection as well as run-time recording of exchanged data.
- Input file preparation, staging, and validation: Preparation and validation method using FACETS Composer application is being developed.
- File migration between parallel computing facilities and integrated data management systems: No.
- Non-interactive data analysis and visualizations via scripted tools or services (such as IDL or Visit): Yes, via VISIT and python scripts.
- Interactive simulation monitoring and data analysis through web portal or dashboard systems that can drive graphical analysis tools and display results in real time: Prototype FACETS Composer being developed for this.

2. Please provide a listing of other projects that use your approach for

A. On-HPC coupling (formerly known as physics composition)

- All production systems use the single-MPI-executable model: CESM, CHOMBO, VORPAL.

B. Task composition (aka workflow, inter-computer coupling):

- Nearly all projects (CESM) use some type of scripting approach for task composition.

3. What needs further development that would require FSP level resources? (Some of this may come out of (1):

- Development of conservative interpolation schemes for data exchange at 2D/3D interfaces.
- Load balancing of component to achieve optimal performance of components.
- Extension of solvers to allow for embedded turbulence for both core and edge physics.

4. If there were to be a software integration effort aimed in part at developing software for use by any integrating software application (more than one is likely to be the FSP case), what software that you have proposed could become part of that effort?

- The FACETS framework components could become a part of the integration effort: core component, FMCFM, edge component and wall component.
- Vizschema for Hdf5 markup could also be used for output file metadata.

5. In particular, for higher-dimensional couplings, have you developed conservative or limiting (in the sense of removing spurious extrema) transformation software for sending data between components that could move forward into FSP:

- Not yet.

6. Please provide a list of OS's (specific hostnames for LCFs would be best) to which your approach has been ported.

- Mac OS X
- Various flavors of Linux (Fedora, Ubuntu)
- Franklin/Freedom with PGI and Pathscale compilers.
- Intrepid with IBM xlc family of compilers.
- Jaguar.
- Core framework ported to Windows using Visual Studio compilers.

4.3.3.4 *SWIM approach*

4.3.3.4.1 Description and/or position

The Integrated Plasma Simulator (IPS) [1, 2, 3], is a lightweight component framework based on the concepts of the Common Component Architecture (CCA) [4], implemented in Python. IPS components satisfy a simple interface with init(), step(), and finalize() as the primary methods. IPS components are unmodified executables of the underlying codes, wrapped with Python. The wrapper, along with small "helper" executables, adapts the application's native inputs and outputs to the interface that the IPS expects. Components exchange shared simulation data via a "plasma state" file [5], a NetCDF file containing plasma specific data and a Fortran library to access them. Quantities that may not be relevant to all components in the simulation are exchanged via additional files that are agreed upon among relevant IPS components.

The IPS framework provides a vehicle for meeting many of the requirements outlined by the ISAs. The IPS is designed to execute in self-contained fashion as a batch job or interactively as an "On-HPC" high performance coupling framework. The core of the IPS framework runs as a single Python process which may be run on a head node, service node, or in the compute partition depending on the system architecture and the capabilities or limitations of each class of nodes. Components are spawned as separate processes, which in turn invoke the underlying physics codes. The framework supports four levels of concurrency [6], allowing for a highly flexible coupled task execution environment. While the framework itself does not provide for expedited transfer of large data sets among constituent components, the IPS execution environment allows for a straightforward incorporation of an external tool (e.g. ADIOS) into an IPS simulation. Tight, in-memory coupling is not directly addressed by the IPS framework. However the framework can incorporate a tightly coupled component into a larger composite simulation component (e.g. equilibrium and profile advance component incorporating tightly coupled solution of transport, anomalous transport coefficients and MHD equilibrium). Work is currently under way to incorporate FACETS as a single component in an IPS simulation.

The use of standard NetCDF plasma state files in the IPS addresses the requirements for using self-describing data formats for coupling data. The framework does not internally provide tools for data interpolation, such operations are expected to be implemented in the helper codes that adapt codes to the standard plasma state data formats. The use of Python in the IPS framework provide for maximum flexibility at the component and simulation levels. The asynchronous event service of the IPS was recently used to implement a parallel in time (parareal) set of components that was applied to solve a turbulence problem [7], and is currently being applied to problems in gyrokinetics, MHD, and possible applications in climate modeling. The IPS framework is also currently being used to implement coupled modeling for advanced battery design under a DOE Office of Energy Efficiency and Renewable Energy (EERE) program. The IPS simulation environment provides for online monitoring and summary analysis of IPS runs through a "monitor component" and a standard NetCDF monitor file. The monitor file summarizes an IPS simulation in a format that is amenable to comparison with experimental data stored in MDSPlus servers. The ability to map SWIM run summaries to MDSPlus data has just been completed and is currently being tested and rolled out to IPS users.

4.3.3.4.2 Answers to follow-up questions

1. The requirements team, from a reading of the ISA requirements document, has identified the following requirements (documents attached). Please provide one sentence addressing each of these, such as "possible by further development", "prototyped", "implemented in special cases", "part of production software installed at the LCFs at <hostname:directory> for general use".

- High-performance framework (runs on an HPC): Part of production software . In current production use on Franklin (Hopper). In production use on various other clusters at PPPL (stix) and other locations. Not available in a publicly accessible central location on deployment platforms.
- Rapid sharing of the data involved in coupling among components: File based data sharing, no framework support for in-memory data sharing.
- Support for implicit coupling: Prototyping currently in progress for file-based implicit coupling using the IPS framework.
- Self-describing data formats for data involved in couplings: Yes, using NetCDF (via the Plasma State from PPPL)
- Support for filtering/smoothing ("data reduction") of data from particle-based models coupled to continuum models: Standalone prototypes developed and tested – implementation depends on near term priorities/resources.
- Support for refinement of data used in coupling: No direct framework support. Possible with further development.
- Ability to initiate and run a large set of small trivially decomposed computational tasks: Part of production software. Many-task computing support using four-levels of concurrency in the IPS framework.
- Universal workflow software: Not used for Off-HPC aspects. Could be implemented with further development.
- Ability to record provenance data on physics models, computer system, compilers/libraries, etc.: Limited support for simulation data provenance. No direct Framework support for models or compilers/libraries. Possible with further development.
- Input file preparation, staging, and validation: Prototyped.
- File migration between parallel computing facilities and integrated data management systems: Prototyped,
- Non-interactive data analysis and visualizations via scripted tools or services (such as IDL or Visit): Prototyped. moving into production this year.
- Interactive simulation monitoring and data analysis through web portal or dashboard systems that can drive graphical analysis tools and display results in real time.....: In production.

2. Please provide a listing of other projects that use your approach for

A. On-HPC coupling (formerly known as physics composition)

- CAEBAT/VIBE: Computer Aided Engineering of Batteries/ Virtual Integrated Battery Environment, A DOE Energy Efficiency and Renewable Energy (EERE) Funded project uses the a modified version of the IPS to build an integrate battery modeling environment.
- Parareal: Parallel In Time Algorithms for plasma turbulence.

B. Task composition (aka workflow, inter-computer coupling)

- No response.

3. What needs further development that would require FSP level resources? (Some of this may come out of (1).

- More portable and robust configuration management and build system.
- Faster inter-component data exchange (e.g. integrating ADIOS into IPS components).
- Improved data management, archival, and analysis.
- Streamlined simulation setup and execution.
- Improved provenance management.

4. If there were to be a software integration effort aimed in part at developing software for use by any integrating software application (more than one is likely to be the FSP case), what software that you have proposed could become part of that effort?

- The Integrate Plasma Simulator (IPS) can be used to rapidly couple codes that currently exist to provide immediate science results. Capabilities not currently supported by the IPS can be developed in parallel and/or integrated into the IPS environment as they mature.

5. In particular, for higher-dimensional couplings, have you developed conservative or limiting (in the sense of removing spurious extrema) transformation software for sending data between components that could move forward into FSP.

- No. Higher dimensional couplings were difficult to specify interfaces for (e.g. distribution functions), and were also application dependent.

6. Please provide a list of OS's (specific hostnames for LCFs would be best) to which your approach has been ported.

- Jaguar, Franklin, Hopper, and standard Linux clusters, In addition to MacOs, and standard Linux workstations.

4.3.3.5 *CESM approach*

4.3.3.5.1 Description and/or position

The Community Earth System Model, CESM1[12],[13] is an IPCC class community global climate model consisting of geophysical model components, each running on potentially different grids, which periodically exchange boundary data with a mixture of high and low frequency coupling. The CESM management is comprised of Working Groups, which are relatively small teams of scientists that work on individual component models and decide their own development priorities but which are subject to oversight by the CESM Scientific Steering Committee (SSC). In effect the Working Groups are analogous to the Science Driver Teams.

---

[12] http://www.cesm.ucar.edu/

[13] Gent P. R., G. Danabasoglu, L. Donner, M. Holland, E. Hunke, S. Jayne, D. Lawrence, R. Neale, P. Rasch, M. Vertenstein, P. Worley, Z-L. Yang, and M. Zhang, 2010: The Community Climate System Model version 4. J. Climate; submitted

CESM1 coupling architecture is required to satisfy a broad class of requirements. The model must provide a single code base that targets both the needs of model developers and of the external university user community. This same single code base must be able to run low resolution cases on a single processor laptop and also must provide the capability to carry out multi-century ultra high-resolution simulations. Components are not permitted to talk directly to each other, but must communicate via a coupling component in a hub and spoke fashion. The model must demonstrate scalability, both from a coupling perspective, as well as within each model component. The coupling architecture allows for the addition of new model components as well as provides "plug and play" capability of data and active components. This latter functionality is critical for model parameterization development. Finally, the coupling architecture should provide the flexibility of laying out components across processors either sequentially, or concurrently, or hybrid sequential/concurrent, and without any answer changes to model output.

To satisfy these requirements, CESM1 (released in June of 2010) was accompanied by a coupling architecture, CPL7[14], that took a completely new approach with respect to the high-level design of the system. Prior to CESM1, the model (CCSM3) operated as a multiple executable system where components ran as separate binaries over disjoint sets of hardware processors[15]. There were several drawbacks to this approach; the coupler sequencing was difficult to understand, porting and debugging were challenging because of the multiple executable job launch and hardware systems generally do not permit multiple executables to share the same processors. CESM1 is a single executable system where system control is achieved by the introduction of a top-level driver that runs on all processors and calls components via standard interfaces. The model components now have the flexibility to run on all or on an arbitrary subset of the processors. It is important to point out that in both CCSM3 and CESM1, coupling occurs via coupling data structures that are independent of the model data structure and that require copies.

The creation ten years ago of the CESM software engineering group (CSEG) has been critical to the above efforts. CSEG is composed of software liaisons to the CESM Working Groups along with several senior developers that are responsible for coupler development, parallel I/O development and performance optimization. CSEG also is responsible for releasing the model, for user support, and for coordinating the software development with external collaborators such as DOE SciDAC and ESMF.

4.3.3.5.2 Answers to follow-up question

1. Do you use any particular task composition or workflow software?

- Not currently. This is one area that we are actively targeting. Our current workflow involves running the model, archiving the raw output, running various diagnostic packages (written primarily in NCL) on the raw output and then archiving the diagnostic output. The diagnostic packages are currently being run only serially. Furthermore, when running on a non-NCAR platform (e.g. jaguar), we often migrate the raw output back to NCAR (using GridFTP) and run various component diagnostic packages locally. Our collaborators in NCAR's CISL have prototyped using Swift from Argonne (http://www.mcs.anl.gov/research/project_detail.php?id=56) to parallelize many aspects of the current diagnostic workflow. Since our raw output is all in NetCDF, we are also looking at the Pagoda package from PNNL (https://svn.pnl.gov/gcrm/wiki/pagoda) that consists of a set of data-parallel processing tools.

---

[14] Craig A.P., Vertenstein M., and Jacob R., 2010: A New Flexible Coupler for Earth System Modeling Developed for CCSM4 and CESM1, International Journal Of High Performance Computing; submitted

[15] Craig, A.P, R. Jacob, B. Kauffman, T. Bettge, J. Larson, E. Ong, C. Ding, and Y. He (2005). CPL6: The New Extensible High Performance Parallel Coupler for the Community Climate System Model. International Journal of High Performance Computing Applications. 19(3):309-328

4.3.3.6 *EUFORIA approach*

4.3.3.6.1 Description and/or position

The ITM-TF effort aims at developing a modular, versatile and flexible tokamak simulation platform. This platform and its workflow principles must be relevant for any problem of physics, technology, diagnostics of the tokamak. Therefore it targets all applications of Integrated Modeling, from support to the experimental program, tokamak sub-system design/optimization, to first principle calculations.

A key paradigm is that the workflows must be used and designed by physicists. Therefore the structure of the workflows should reflect the physics, while all technical items are hidden from the user. Another key point is the modularity. The workflow should be a suite of physics (or technology) components, each of them solving a given type of elementary physics problem (e.g. equilibrium, wave propagation, synthesized diagnostic, …). Components solving the same type of problem should be exchangeable with no change of the workflow structure. This requires the definition of standardized interfaces (I/O) based on the essence of the elementary physics problem solved.

Therefore the ITM-TF started the process of integrating physics modules by identifying the data units (Consistent Physical Objects, CPOs) that would need to be transferred between physics modules, and then agreeing a standard for each of these CPOs [Imbeaux F. *et al* , Comp. Phys. Commun. 181 (2010) 987]. From a high level ("XML" schema) description of each CPO, Fortran derived types and C++ classes are derived automatically, together with the communication layer that will store and transfer the CPOs between modules. Individual code modules have been adapted to receive their physics inputs as CPOs, and to return physics outputs as CPOs. To build a workflow involving modules, the code modules are wrapped with a tool specifying the input and output CPOs, and then converted into Kepler (the Ptolemy based workflow orchestration engine) actors. Any individual code can be a direct part of the Kepler memory space and directly called, packaged as a web service, a program run locally in the batch queue, or submitted as a GRID or HPC job for remote execution with the data transfers handled transparently as far as the code is concerned. With respect to the developed workflows, any physics component can be replaced by another of the same class without affecting the rest of the workflow; the choice of a particular component being determined by the trade-off between fidelity and speed.

The methodology and approach adopted by the ITM-TF are aimed at rendering the software infrastructure flexible enough so that it can support most of the relevant computer language coding (e.g. Fortran, C++, Python, Java) in an integrated modeling environment with a user friendly interface to build complex scientific workflows using the state-of-the-art competencies developed in EFDA. Modularity of all workflow components and communication between modules via consistent physical objects (CPOs) instead of files makes it possible for future adaptations of the concept and tools to other software communication layers and/or workflow engines.

4.3.3.6.2 Answers to follow-up questions

1. Have you also defined APIs for these for calling components as libraries in C++ and HDF5.

- No response.

2. Are you currently using your technology for coupling components on a single multi-CPU computer (one such that jobs have a defined MPI_COMM_WORLD)?

- No response.

4.3.3.7 *Discussion of the various approaches*

The position paper on integration of legacy frameworks is in part advocating that integration of TRANSP/PTRANSP into the SWIM framework be an early deliverable of FSP, and the adoption of TRANSP input format (namelists and variable names) as a standard for FSP. Since this is an early deliverable, this position must then be taken up by the ISA team, which will be formed by an as yet undefined process. So the first step would have to be the adoption of the SWIM framework for moving forward, and a second step would selecting TRANSP for inclusion. One question that arises is whether the FSP should then consider ensuring that ISAs be able to accept the input formats of other legacy integration frameworks, such as OneTwo and CORSICA. These questions should be addressed in the workshop in February.

The remaining 5 position papers come each from each of the other existing computational integration project, one of FACETS, CPES, SWIM, CESM or EUFORIA. Our classification of these is:

FACETS: Focused on On-HPC coupling through the parallel component, single MPI executable approach that recursively splits communicators and assigns them to components for concurrent coupling. Also developing infrastructure for developing new parallel components. All task composition data exchange is through files, but no adoption of any particular software.

CPES: A service oriented architecture. For On-HPC would seem to require doing multiple-MPI-executable launch with some methodology for component location. For Task Composition, there appear to be two approaches: a Kepler controlled composition, and the service oriented architecture, where the various computational software executables would need to locate the services they need and interact with them through some messaging API. We hope to achieve greater clarity at the February workshop.

SWIM: Focused on On-HPC coupling through a serial Python executable that is launched on the back-end nodes of an HPC and orchestrates the execution of executables that are then coupled through file reads and writes. Makes use of the Plasma State for uniformizing output. Assigns processor sets to the different components for concurrent coupling.

CESM: Focused on On-HPC coupling through the parallel component, single MPI executable approach. Interestingly, this project made a major transition in its most recent release away from a multiple-MPI-executable launch methodology, as discussed in the position paper. All task composition data exchange is through files, but no particular task composition/workflow software was adopted.

EUFORIA: Focuses on Task Composition. Defines standard schema for the data units that must be transferred between executables. The data is stored in HDF5 files with well-defined schemas.

We see that there is a wide range of coupling methodologies for both On-HPC coupling and Task Composition:

- On-HPC coupling by component architecture in a single MPI executable with communication through MPI:
    - FACETS, CESM.
- On-HPC coupling by component architecture in a single MPI executable with communication through files:
    - SWIM

It seems from our discussions with CPES personnel that the On-HPC coupling (e.g., XGC0 using DEGAS) is through the use of DEGAS as a parallel library. It is believed that CPES is using a services approach for inter-HPC coupling, e.g., analysis of computed edge profiles for ELM instability is done by transmitting the data to another cluster where the M3D computation is done.

Task composition through exchange of data in files: ALL.

Task composition using Kepler: CPES.

Task composition through persistently running services: CPES.

Task composition software under investigation? CESM looking at Swift from ANL and Pagoda from PNNL.

4.3.4 On-HPC composition

The goal of this section is to define first processes and targets for the Integration Effort relevant to On-HPC Integration.

4.3.4.1 *On-HPC composition requirements*

The requirement relevant to On-HPC Integration are grouped here as:

- Persistent data and coupling through files
- Self-describing data formats for data involved in couplings
- Ability to record provenance data on physics models, computer system, compilers/libraries, etc.
- Coupling
- High-performance framework (runs on an HPC)
- Rapid sharing of the data involved in coupling among components
- Support for implicit coupling
- Data modification
- Support for filtering/smoothing ("data reduction") of data from particle based models coupled to continuum models
- Support for refinement of data used in coupling; Ability to initiate and run a large set of small trivially decomposed computational tasks.

4.3.4.2 *On-HPC software integration tasks*

The Fusion Simulation Program Integration effort will have sets of activities: (1) short-term activities that will take place in the development of the ISAs, (2) development of standards and enabling activities that will be useful to the ISAs in the short term, and (3) research on methodologies that could become of use in the longer term.

For the ISAs to produce the early deliverables, it is expected that they will build on a base of pre-existing software. For this purpose, the ISAs will need team members from outside of the fusion energy sciences, so called enabling technology team members, under their control. At the same time, the integration teams should be fully aware of these activities in order to discover opportunities for developing common tools. Thus, the Integration Effort will have liaisons on each of the ISA teams to attend their regular developer meeting to obtain an understanding of the types of integration that are being undertaken. This information will feed back into the Integration Effort.

4.3.4.2.1 Development of standards

To develop general coupling software at all, there must be some standards about how such software can be used in a coupled situation. In particular, APIs must be defined for the main types of coupling (files, function calls, messaging, whether MPI or other).

A universal aspect of all of the methodologies for coupling is the use of files. File-based coupling is obviously present in task composition, and it is particularly used by the SWIM project for On-HPC Integration. In addition, all On-HPC executables will have to adhere to some file formats in order for their files to be consumable by any Task Composition software. Thus a first order of business seems to be determining the schemas and formats of files.

Ultimately, one needs to be able to communicate information about data structures (as needed, e.g., for visualization), provenance (how the data was generated, from experimental data sources through software versions), and physics semantics (the physical quantities corresponding to different sets of data). Part of the above effort will be to settle on the formats. Given the wide variety of formats and methodologies in use, we view it as unlikely that the community will be able to settle on a single file format. Regardless, any developed schemas should be translatable into any of the accepted file formats (likely NetCDF, HDF5, and the more recently developed ADIOS-bp), and a message format that could be used in a service oriented or other architecture relying on messaging. These different types of data will be treated in the just named sequence, with reference implementations developed for the agreed upon file and messaging formats.

In addition, the Integration Effort should define the component API for using components through function interfaces. Current proto-FSPs have identified sets of methods including init, advance, revert, finalize, dump, restore. Such an effort would standardize how components are incorporated into single-binary MPI executables and allow for communication with other binaries, whether in running on in the same MPI space through multi-launch or on other machines (especially as needed for analytics). Such communications should be transport layer agnostic.

For each of the above two efforts, the process will consist of first creating committees involving representatives of each of the stakeholder groups to ensure that the developments will meet the community needs. The Integration effort will then be charged with developing a reference implementation that meets those requirements. These reference implementations will be reviewed by the standards committees, with possible results including a request for further implementation or a revision of the standards as needed. This process will iterated to convergence.

4.3.4.2.2 Development of reference implementations for the On-HPC composition software

Upon completion of the standards for coupling, SIS will undertake the development of a reference implementation of the on-HPC framework using those standards. A reference implementation is one that shows that the basic interfaces work, even though what is being coupled might be simpler versions. E.g., NetLib provides a reference implementation of BLAS (Basic Linear Algebra Subroutines), which can then be used to for nearly all computational linear algebra. This allows vendors to provide high-performance version of those libraries (Intel provides MKL, AMD provides ACML) that developers can link to.

In the same way, SIS will provide a reference framework and components and show how they can be coupled in the on-HPC environment. ISA developers can then use this framework as a starting point for developing their own applications. By following this development path, newly developed components will fit into the ISA frameworks automatically.

4.3.4.2.3 Support of software for general use

In addition to the above, the Integration effort will identify candidates for support for general use by the ISAs for On-HPC coupling. An example of a candidate is the ADIOS I/O library. Preliminary work by some of us has shown the need for extensions. The Software Integration effort would work with the developers of ADIOS to define the needs and possibly provide resources to get the work done. Thus, the Integration Effort should act as a reuse team, identifying software being developed by the ISAs for extraction, extension, making more robust as needed, and then helping ISAs make use of it.

Another example is a universal I/O access layer that can address FSP needs for not only storing data, but also adding the metadata as needed for visualization, provenance, and semantics. Currently, there exists this to some extent, e.g., NetCDF can write HDF5 files, and ADIOS-bp files can be converted to NetCDF or HDF5 files. However, it should not be necessary to adopt any of these to have access to FSP data in the other formats – hence the need for a universal access layer.

4.3.4.2.4 Research Tasks

The Integration Effort needs to include research into methods for future computing as well as developing standards and software for present-day computing.

The Integration Team will be interacting closely with the ISA teams. This will help identify additional software that can be extracted for reuse, but it will also help identify patterns in coupling. As these patterns are identified, the Integration Team will work towards providing or providing software that is applicable to coupling scenarios beyond the original use cases. One example might be software for coupling data (such as turbulence fields) through a 2D surface or the data corresponding to, e.g., magnetic equilibrium. The problem is complicated, as a simple broadcast will not scale to a large number of processors. Communication through files will have a similar problem. The issue is to get the correct data to the components that need it.

In the above particular example, there has been significant work by the climate community, which has extensive experience in 2D surface coupling as in, e.g., atmosphere-ocean/land coupling. The process for moving forward will then involve stating the requirements, searching for packages that can (perhaps with modification) do what is needed, and then either adapt those packages or write new software. The Integration Team would then take the responsibility of retrofitting any ISA code to use the common software, so that the FSP can move forward with a reduced code base having less maintenance.

Yet another approach is the emerging concept of DataSpaces, which applications can write into, while others retrieve data from this. This is another way to solve the "MxN" coupling problem that occurs especially for 2D and 3D coupling.

An important question facing the FSP Integration Effort will be whether to build or adopt distributed data structures that can be reused by component developers. At present there exist many sets of distributed data structures, including generally available data structures like Global Arrays and particular distributed data structures. [Nearly every HPC application has its own data structures: Chombo (adaptive mesh refinement [AMR]), MOOSE (unstructured mesh).] This will be a difficult process that will involve prototyping components on different distributed data structures. This process may not succeed, as history shows that HPC application developers are reluctant to adopt and, therefore, lose control over their data structures, perhaps the most fundamental part of an HPC application. However, it is also true that new data structures are created when (as in the example of AMR), old data structures do not meet the need. This may apply to fusion, as for whole-device modeling it may be necessary to marry logically different meshes in different regions. E.g., in the XGC0 axisymmetric edge modeling code, an unstructured 2D mesh is used. In the core in other modeling codes, a structured, field-aligned mesh might be used. Coupling may involve having these two grids overlap with data made consistent on the overlap region through an iterative process. Regardless, it is likely that FSP will need to develop the concept of modeling on multiple, logically different grids, and it may need to understand the distribution of data for parallel computation, perhaps ultimately leading to a distributed data model for this situation.

With regard to service-oriented architectures, the FSP will at this point in time have extensive experience on many fronts. Reference implementation will have been created. User feedback will have been received. Hence at this point the Integration Team can decide on whether further research is needed or whether to being providing generally usable software supporting the service oriented approach.

Initiation of a research program into the development of On-HPC service oriented approaches. Such a research program should address the issues identified by the climate community as well as develop the appropriate APIs and develop an understanding of how one can launch multiple implementations of different components and discover them at runtime. Additional issues will be identified by an issue community committee. A research team would then be developed to address these issues. The issue identifying committee and the research team include a broad range of members, both computer scientists and experienced HPC application developers. The research team would consist primarily of computer scientists but would have an advisory team consisting of both computer scientists and experience HPC application developers.

4.3.5 Task composition

4.3.5.1 *Scope of this section*

Task composition (aka Scientific Workflow) is a formal description of a process for accomplishing a scientific objective, usually expressed in terms of tasks and their dependencies [8]. In HPC simulation this is often used interchangeably with the collection of executable programs, the compute resources they are executed on, the files that are transferred between them, and a description of what processing or transformation occurs within. For small projects many of these steps are carried out by a researcher at a command prompt. In larger projects many of these steps are automated with various tools. FSP will be a large and long HPC development effort and we expect a high level of automation. To enable this transition in the FSP there will be several cross-cutting technologies deployed and supported under this program.

4.3.5.2 *Schema and Engine*

The Schema is the language in which you describe task composition. The Engine is the program that parses the schema and executes the plan. For a task composition engine to be completely general, it ends up being a complete language. Rapid development of different compositions has led to the use of scripting languages as the most common method of describing task compositions. There is some desire for use of Kepler, but this is not observed as being used by working scientists. Developing a sufficiently abstract approach to task composition software that provides sufficient value to computational scientists to justify their learning it over scripting therefore remains a research area.

4.3.5.3 *Provenance*

Requirement: Recording the history of generating a particular element of data. This is the equivalent to the experimentalist's log book information. Each ProtoFSP handles this at some level. FACETS records their compilation and composition information as attributes in their HDF5 files. The SWIM team has each python wrapper code enter progress and generates standard logging files. CPES has their workflow automatically capture progress logging and monitoring information into a MySQL database.

Approach: FSP will integrate provenance with the overall integrated data management. There would be a Catalog in the internet download manager, with a defined schema for describing general provenance. There would also need to be an FSP provenance API that FSP tools would use to register provenance information. This API will need a binding to the most common workflow languages, like XML, python, bash. Thus, all build processing, staging, input generation, batch submissions, architectures, etc. are recorded through a common Provenance API. To as large a degree as possible this information is gathered automatically by the workflow engine. FSP will define this common API and update it as the projects move forward.

4.3.5.4 *Input file preparation, staging and checking*

Requirement: Data for computations involve large numbers legacy input files and formats and input parameters: Ufiles [ASCII TRANSP (wide usage)], EQDSK (ASCII equilibrium file format from EFIT), Plasma State (TRANSP/SWIM NetCDF format), “Fit files” (files that describe profiles used in EFIT fits needed to get profiles inside separatrix. This will remain true for some time into the future. Input data needs to be placed in suitable places for codes to access this information at run time, this is staging. FacetsComposer is an example of a tool that helps facilitate this process today, and the SWIM team is moving to adopt this same technology. CPES has also expressed interest in moving towards a common tool for input validation.

Approach. FSP would be FSP taking over support and distribution of the Input Validator and Assembler. The Assembler would utilize other components in the FSP workflow toolkit, like FSP Provenance Interface. A Validator parses input files for all components that will participate in a co-executing batch job submission. Without some form of Validator a multi-component simulation will have an increasing risk of incompatible input data as simulations get more complex. Furthermore, when the community of users is expanded beyond the inner expert set to a general simulation and verification and validation context then productivity could be hampered to the point of making advanced codes impractical.

4.3.5.5 *Monitoring*

Requirement: During the entire task composition, the elements of this workflow are reporting on their progress to some kind of monitoring tool that lets the users know about the progress of their computation. Elvis is a tool used by the SWIM team to monitor the code progress. FACETS uses FacetsComposer for this work, and CPES uses the eSimMon DashBoard. The DashBoard, a more generally useful product out of the SDM SciDAC Center, makes use of the MySQL database maintained by the CPES development team.

Approach: FSP would look to unify the role of a monitoring tool within the FSP efforts and utilize the Unified Access Layer for data query and the Data Dictionary and Catalog that the Integrated Data Management effort are going to specify and deploy.

4.3.5.6 *Visualization and Analytics*

Requirements: A broad class of batch and interactive data analysis and graphics generation are used throughout the MFE community. For interactive data exploration and analysis people use a preferred tool (VisIt, IDL, Matlab) and have access to specific data in existing data formats (NetCDF, HDF5, ADIOS bp). Non-interactive analyses are expressed as scripts in the existing tool scripting languages (matplotlib, VisIt scripts, IDL scripts) which represent components in a workflow. These workflow elements can then send graphics or statistics into a location for easy use of a monitoring tool. FACETS uses a VisIt back end compute core to generate graphics to send to the FacetsComposer monitoring tool. CPES uses a collection of vector graphics routines to forward image data to a Flash- embedded graphics viewport in the eSimMon monitoring service.

Approach: FSP-supported tools for data movement and conversion will be needed to use existing visualization and analytics tools and scripts. This will give way to IDL/VisIt/matplotlib bindings to the Unified Access Layer API and direct into memory data movement. The products of these tools would then also become data objects within the Integrated Data Management system and available for query and monitoring.

4.3.5.7 *Steering*

For most workshop users simulation steering means a combination of simulation monitoring and the ability to terminate errant simulations. The goal here is to make efficient utilization of limited computing resources while maximizing productive simulation. For most fusion community users this capability already exists. What would be desired is to unify the interface used for this capability. If FSP normalizes on one or two monitoring tools then the addition of a kill capability into this monitoring tool should be sufficient.

For a minority of users steering entails an external simulation control capability. i.e., the ability to change codes and models based on some automatic detection capability. One example would be altering an edge simulation in the presence of some form of instability. Switch to an MHD model for the instability until a recovery occurs, then restart a kinetic model. CPES achieves this with their Kepler workflow engine. However it is a semantic point to call this steering. This is more elegantly handled under the classification of dynamic modeling and not covered under task composition.

There are suggestions that eventually an FSP code should be controllable from an external control design software package like Matlab’s Control System Toolbox. The consequences of this kind of capability were not explored in detail, except that it would require a form of tight coupling between a large coupled parallel simulation and a control logic program. This does not seem to fall under the classification of Task Composition.

Approach: FSP will support a very limited capability to steer computations. Go or Kill. With adequate job monitoring the workflow engine, or a user, can decide to abort a simulation to conserve resources. It is not a rich set of capabilities, but it does address the needs expressed by a large majority of FSP stakeholders.

4.3.5.8 *Verification and Validation*

Requirements: While specifically not a part of the scope of Workflow, effective Workflow design should enable productive and reproducible verification and validation activities. Logging these procedures in a universal Catalog should provide the means to reproduce such activities, and document their validity. It should also facilitate the process of turning one-time verification activities into ongoing regression testing. Once a domain expert has assembled the components and inputs of an effective and correct verification workflow, the FSP team should be capable of making this verification an automatic part of ongoing FSP testing.

In general the UQ process is itself a workflow. The UQ Pipeline out of LLNL, or any ensemble simulation process marshals collections of HPC runs and attendant data, and the provenance of this data, to derive secondary uncertainty data, which are products for the Integrated Data Management archive.

4.3.6 FSP Integrated Data Management

4.3.6.1 *Introduction and Background*

Careful management of data and associated metadata is an important part of any scientific enterprise. This is particularly true for a project of the size, scope and longevity proposed by the FSP. At the same time, it's clear that most current fusion simulation projects lack systematic, project-wide organization of their data. The requirements of the FSP and the opportunity afforded by building a large, coherent program from the ground up suggest a new look at this problem.

An integrated data management system for the FSP would have three main goals:

1. Improve productivity of FSP efforts through better organization and systemization of data and data access.
2. Foster collaboration through more transparent and efficient methods of data exchange.
3. Simplify use of FSP codes.

The integrated data management topic includes all data stored or used as part of the fusion simulation program excepting data transiently produced and exchanged in the midst of high-performance computing. That is, it covers data prepared as input, controls or workflow for simulations along with the final output of codes. It also covers data used in or resulting from verification and validation activities including imported experimental data.

4.3.6.2 *Requirements*

4.3.6.2.1 Basic Requirements

The basic functional requirement for a data management system is an ability to store, locate and retrieve all data in its domain. (In this context, data includes "bulk" data and sufficient descriptive and metadata needed to provide understanding of the origins of the data, to give the data enduring meaning and to allow for efficient searching and browsing.) This rather general requirement is not terribly useful in defining the desired system and needs to be augmented with a set of "non-functional" requirements that specify how the system should behave in order to optimize the utility of the scientific data generated by FSP activities. Among the non-functional requirements we recognize are:

- A unified view of all FSP data.
- Creation of logical data collections.
- Remote access to all data without requirements for multiple remote logins.
- Mechanisms for ensuring data integrity and consistency.
- Ease of use for data providers and consumers aimed at supporting a large community of non-specialists - including capable tools for browsing and searching
- Controlled sharing, that is some specified granularity of access control which would conform to FSP policies and which would use an FSP wide authentication and authorization system.
- Policies and methods for data publication and dissemination
- Persistence - policies and methods for maintaining data integrity over time (especially as technology evolves) and for retiring old data

4.3.6.2.2 Requirements driven by particular aspects of FSP

The non-functional requirements, along with the approach described below, are driven by discussion in the science driver reports along with general considerations for the use of scientific data and the particular nature of the FSP enterprise. The broad scope of the FSP means that it will be supporting a large number and wide variety of data customers spread over a wide geographic domain, accessing a large heterogeneous set of data over a long period of time. The data stored may contain very large arrays, requiring specialized treatment for storage, retrieval and replication. The range of customers includes code developers, practitioners of verification and validation and a substantial end-user community. These groups have very different use cases for data and very different knowledge of specific code implementations. It is worth noting that a central feature of the FSP is the integration of physics and components, requiring most users, even if they are experts on one particular topic, to work in areas in which they are novices. The heterogeneous developer/user base and the long duration envisioned for the FSP, drives requirements for more transparency in data naming and in representation of physical and geometric quantities. It further suggests a deeper need to present a consistent, coherent view of all data and to avoid the "n^2" problem where a large number of application and groups of data items must be customized to each other. It emphasizes the importance of designing software systems that can evolve and adapt with new technologies and new use requirements. The need to import data from experiments or other collaborators brings with it the need to protect their data in a manner consistent with the collaboration policy. These usually specify "no distribution to 3rd parties without approval" which implies a level of differentiated access that must be supported.

4.3.6.2.2.1 Sample use cases driving requirements

4.3.6.2.2.1.1 Use Case: Common data preparation and analysis tools

A group of researchers begin a complex simulation by importing MHD equilibrium and profile data from a number of experiments. They use a common set of tools to transform this data into suitable forms for input into several simulation codes. Visualization tools, which can understand FSP data and metadata formats, help guide this work. Additional common tools check for consistency and validate code parameters and controls against a set of rules defined by developers. The codes are run are several different platforms around the country and the data is stored in nearby archives. Default analysis programs run to post-process output. These use the well-defined API and globally defined data identifiers to run identical analysis on all data sets.

4.3.6.2.2.1.2 Use Case: Documentation

A question arises about the validity of an important result that has been published and presented at several scientific conferences. Researchers use the metadata catalog to determine precisely which simulations were used to produce the figures in question and examine input and output data. A problem with input processing is discovered and the code re-run using the original version and also the current version. Both give results essentially identical to the published work.

4.3.6.2.2.1.3 Use Case: Flexible tools to control data sharing

Unpublished data is imported from an international tokamak for use in validating an FSP code. The experimentalists are collaborating in this effort, but do not want the experimental data distributed beyond the validation team for some period of time. Authorizations are set up, giving only that team read access to this new data set.

4.3.6.2.2.2 Categories of data potentially acquired and stored

The categories listed here need not correspond to any implementation – they are grouped here in a way that corresponds mostly to where the information comes from and what it might be used for. It should also be clear that this list is preliminary and the categories arbitrary to some extent. The point here is show the broad reach of the system we need to build.

- People and their contact information, userid(s) and hooks to authentication
- Authorization information including roles and privileges
- Descriptions of codes and code components including any descriptive or reference material, pointers to published descriptions, citations (examples where the code has been used), code documentation, user guides, namespace information, contact information for code development or support and links to information in code repository
- Description of experimental data sets used for code input or validation, including machine names, shots, dates, times, information on data importation (who, when how, why), a description of the data including information on array shape, size, type, labels, units, independent axes, comments, pointers to data and namespace information
- Code inputs including origin of data (who, when how, why), time-stamped description of data preparation steps and data processing, description of data, including information on array shape, size, type, labels, units, independent axes, comments, other annotation, relation between data items, pointers to data and a data digest for searching
- Code run control parameters including who, when, how, why, origin if part or all comes from other runs, rules for validating code controls and inputs, and results of that validation

- Description of code runs including who, when, how, why, information on computational platform(s) used, a description of how complex codes are built up from components, particulars on code version, build, etc. used for each run
- Code outputs including who, when, how, why, a description of data, including information on array shape, size, type, labels, units, spatial grids/independent axes, comments, other annotation, time-stamped post-processing steps and data, visualization steps, relations between data items, data present and data quality flags, data dearth sciences data system digests for searching, pointers to data
- Description of workflows, who, what, when, how, why with sufficient information to fully describe workflow, time stamped during execution, for scripted workflows this could include full scripts or their equivalent
- Description of data namespace including names, labels, definitions, comments, translation, etc., possibly with a controlled vocabulary and ontology with information on data entry (who, when, etc.)
- Descriptions of data collections and/or Groups of related records
- Outputs of verification, UQ, validation activities with data description, pointers, who, what when, why. Thus the components and processes defined by a Scientific Data Management system are:
- Description of data export (that is any data which is moved outside of domain of FSP data system, e.g. papers, presentations, external databases, etc.), including who, when, how, why and any links to these external data sinks

4.3.6.3 *Data Management System approaches in other related fields*

We have begun a survey of scientific data management systems used by related fields. These will help identify areas of consensus on requirements and approach and may be a source of ideas or software. A very brief (and still superficial) summary of a number of these follows:

- Scientific Data Management Center (SDMC)

  This is an ongoing SciDAC project led by LBNL. It has a number of goals including: Optimizing shared access from mass storage systems including storage resource management and data movement; Parallel-I/O methods for various file formats aimed at enhancing I/O efficiency; Feature extraction techniques and High-dimensional analysis. Some of the tools developed by this project may be directly applicable to FSP, particularly those associated with the FSP data archive described below. https://sdm.lbl.gov/sdmcenter/
- Earth Systems Grid (ESG)

  This system supports distributed access to data, codes, computers and analysis tools in support of climate sciences. An array of middleware and client tools has been developed within a federated architecture – that is where resources are distributed around the network at largely independent sites. The system provides a global view into these federated resources including searching, browsing and data access – all within a common environment. http://www.earthsystemgrid.org
- Program for Climate Model Diagnosis and Intercomparison (PCMDI)

  The PCMDI provides a repository for climate model data from a large set of independent research groups. It sets strong standards for data and file formats, metadata, variable names and definitions, units, etc. It also requires standard and extensive documentation on the computational models used to produce the data. Data access is via file download. http://www-pcmdi.llnl.gov/
- Climate Data Library (IRI/LDEO)

  This site provides a common repository for a wide range of observational data. It provides the data in a wide variety of file formats, including images via download or API. Searching and browsing are supported through the web. It also provides a set of visualization and analysis tools. http://iridl.ldeo.columbia.edu/
- NASA Space Sciences Data Center

  This portal provides a rich set of data, mostly from spacecraft. It encompasses three domains: a planetary data system, a space physics data facility and astrophysics data centers. It is primarily a file-oriented, search and browse system with library-like tools aimed a very large, heterogeneous customer base. http://nssdc.gsfc.nasa.gov/

- Life Sciences Examples
  These include a number of repositories for public data including notably GENBANK a genetic sequence database http://www.ncbi.nlm.nih.gov/genbank/ and the RCSB (Research Collaboratory for Structural Bioinformatics) protein data bank which contains information on biological macromolecules http://www.rcsb.org/pdb/home/home.do
- Hyperstudio
  This is a project which is oriented toward humanities data. It provides a model and powerful tools for browsing and searching heterogeneous data catalogs and timeline visualization tools. http://hyperstudio.mit.edu/
- Fusion Experiments
  Ongoing fusion experiments have a strong track record in data management and data access. These include comprehensive systems for experimental run management, electronic notebooks and code data management for software widely used in support of experimental operations like TRANSP and EFIT. http://www.mdsplus.org
- Past SciDAC projects
  A number of SciDAC projects focused on data management topics and provide valuable lessons learned, including the Particle Physics Data Grid http://ppdg.net/ , the DOE Science Grid http://doesciencegrid.org/ and the Fusion Grid Collaboratory http://www.fusiongrid.org/

4.3.6.4 *Elements of FSP Integrated Data Management System*

We envision two main elements that together would comprise the FSP data store. The first is a comprehensive data catalog (alternately called a metadata catalog), which would hold entries for each FSP activity that involves data. These activities would include at least, data preparation and simulations for any use, verification exercises and validation exercises including the importation of experimental data. The data catalog provides the global data view discussed above. The second and far larger element of the data store would be an archive of bulk data. The overall aim is that between the catalog and the archive, the system holds a complete and accessible description of the data and all of its attributes. The overall system also needs to consider physical and logical methods of data storage, methods for entering data into the catalog and archive and data access methods.

4.3.6.4.1 Simulation metadata catalog

Data use requirements and case studies drive the metadata design. It is meant to enable key FSP activities including:

- Data location through searching and browsing
- Global data access
- Data analysis and visualization through uniform provision of "ancillary" data
- Documentation of all past and current simulation activities

Metadata provides essential information on the context and content of data archives. The metadata catalog provides a consistent, coherent and global view for all FSP metadata. Details of its organization and content will be specified as part of the system design, but certain features are clear. Entries in the catalog will need to be created whenever data-significant activities are begun. This step could be integrated into workflow tools, which will create FSP data as a result of their actions. As data is created and stored in the archive, metadata describing the data provenance would be added to the catalog. Digests of data from the workflows themselves would also be stored. Taken together, this metadata is meant to describe, as fully as practical, how the data was created including information on code configuration and versions, compilers, time stamps, user information, log files, etc. As far as possible, once appropriate mechanisms are in place, this data would be captured automatically (for example by extracting information from the code version control system), without need for manual intervention or excessive burden on data providers. Metadata would also be collected which describe data types, sizes, array shapes, formats, units, independent axes, labels, definitions and comments. In addition, it will be extremely valuable to collect information manually from data providers to explain the motivation and purposes for particular activities or choices. The general approach is to provide automatic capture for metadata which answers the questions "who, what, when and how?" and to require manual input to answer questions about "why?".

The catalog is meant to support searching and browsing - it will be necessary to identify and build a set of tools and higher-level applications for this purpose. To this end, the catalog would contain some amount of high-level data - essentially digests of inputs, outputs, processed data or controls along with any other information which would be useful for locating particular data records. The catalog schema and applications built on it should support standard scientific logbook functions. Self-description and metadata should be sufficient to provide all the additional information required to make stored data useful and to maintain that usefulness over time.

The catalog should have the capability to define collections of related data items using some simple and flexible mechanism. In this way, users could group all information associated with a particular simulations, group related code runs or link simulations to experimental data used for validation. Relationships between data objects should be explicit and stored in the catalog as data.

The catalog would also contain the information needed to read data from the archive. Details on how this would be accomplished await the system design, but the goal would be to provide a seamless path from the catalog to the bulk data. To that end, the catalog would provide information for data "naming" services, discussed in somewhat greater detail in Section 4.3.6.4.1.1 below.

4.3.6.4.1.1 Proposed data model for catalog

The metadata catalog would deal with the following general set of object types (abstract data types). 1) Data, 2) Actions, 3) Sequences, 4) Collections, 5) Comments/Annotation

These are explained below.

- Data: Some form of structured data stored inside or outside the database schema. (If outside the schema it would contain descriptions and pointers to the data. That is, the catalog has a domain which extends beyond itself - to a data archive.)
- Actions: Create, move or transmute data from one form to another. (These would include all data importing and exporting, fetching, storing, pre-processing, code running, post-processing, etc. By import and export, I mean actions which move data into or out of the catalog domain)
- Sequences: Particular sets of objects organized (linked) to represent workflow. (An open question is how general we want to be about the topology and connectivity. Can we convince ourselves that we only need relatively simple types? or that tools exist for representing and manipulating more or less arbitrarily linked structures?)
- Collections: Other simple lists of various kinds of related objects
- Comments/Annotation: Unstructured text with additional fixed attributes (who, when…) that can be associated with any other object, including other comments. Within their text, comments should be able to point/link to any object inside the catalog domain or link to URL's outside of it.

4.3.6.4.1.2 Access to data catalog

In addition to APIs for data access, we anticipate creation of several web based interfaces, built to enhance browsing and searching. One important interface paradigm to support would be a scientific logbook. Workflow for manipulating inputs, running codes and processing outputs is a natural application for this sort of interface. The workflow describes a sequence of steps, manual or automatic, carried out over time. Documenting these steps, along with time-stamped comments creates a record of events very much like a traditional lab notebook to which additional links (including images or movies), comments and notes can be added.

4.3.6.4.2 Data Archive

The data archive contains the bulk of the data stored by FSP researchers. It can be viewed as a collection of data objects which can be accessed individually or in groups. The format and organization of the archive will be part of the overall system design. Currently "raw" data is stored in several file formats including HDF5, ADIOS-BP and NetCDF for simulation data and MDSPlus for the experiments. The design will need to consider whether it is best to convert all data into a common format or (perhaps more likely) to support access to a variety of underlying formats through a common access layer. The server structure, that would support the archive, needs to be defined as well, but it seems likely that the archived data will be accessed through more than one server and likely from more than one site. Issues of data replication, data merging and data consistency need to be addressed. Tools developed by the SciDAC SDMC are good candidates for addressing many of the challenges associated with the archive.

Since simulation data has usually not been provided with permanent archives, the question of what data to keep and for how long arises. It seems to us that a reasonable position is to keep entries in the data catalog "forever" but to maintain a flexible position with respect to retention of data in the archive. Obvious "bad" or obsolete runs could be disposed at any time by the data originator. When data is removed from the archive, corresponding entries would be marked as deleted in the catalog but those catalog entries themselves would remain. This process is analogous to the use of a traditional lab notebook where notes may be marked through, but pages from the notebook are never thrown away. As the archive grows in size, policies on data retention will have to balance utility against cost. In the case of very large data arrays, it may be useful to keep, in addition or instead of raw data, data digests or decimated data (a la Google Maps), which represent the data in a more compact form. This could include processed data or graphics of any kind as well.

The FSP will need to take account of data exported outside of the FSP domain in the form of external databases (e.g. ITPA) or in publications. Tracking these data exports through some form of tagging in the catalog enhances the overall traceability of data originating in the program. Just as one wants to understand the provenance of data stored in the archive, one wants to keep track of which data was exported where, when, why and by who. This feature may connect to a policy on data retention, where rules for retaining the original copy of exported data may be appropriate. (One suggested option for linking to data in publications is through Digital Object Identifiers, http://en.wikipedia.org/wiki/Digital_Object_Identifiers , supported now by many publishers.)

Some of the same approaches and mechanisms could support data replication within the FSP domain. In this case, where multiple versions of data are in use, it is essential to protect data consistency and integrity - there should never be more than one writeable version of any data.

4.3.6.4.2.1 Data Access

To simplify data access, the general approach suggested here is to provide a universal access layer (UAL) which carries out all loading and retrieval for FSP data. This layer could get or put a data object or collection of data objects and could be qualified by a combination of conditions (including range in any dimension) on data attributes. Reference to data objects would be through a globally unique, fully qualified name (to be defined) obviating the need of users to know where or how the data is actually stored, hiding details of underlying file formats, indexing or parallel access technologies. Data objects could support multiple methods, for example data_get could return data into memory or into a local file in a selected set of formats. A low-level API will be needed which would support all data with higher level APIs built, as needed, on lower-level services. The API might support the equivalence of "query estimation" - that is, a user could get an estimate of how much data a call would return before executing it. This may be particularly useful for manipulation of the very large data sets envisioned. Support for transparent access to data subsets or decimated data will need to be provided. The metadata catalog should contain sufficient information to access any particular piece of data through either interactive or automatic processes.

4.3.6.4.2.2 Namespace management

We can anticipate a large name space that must be supported in a project of this breadth. (For comparison, we note that a typical major fusion experiment has on the order of 100,000 named data items and that the ITER design plans for about a million.) It would not be practical for one individual to learn or manage all of these names. Some degree of standardization may become FSP policy, but it is prudent to consider approaches that allow a high degree of transparency without complete standardization. What is envisioned here is a structured namespace driving a capable data dictionary application (or applications). The namespace might be stored in tables as part of the data catalog schema. While the view of the database would be global and centralized (with replicas as needed), management would be distributed - data providers would manage the namespace of their own data through a common set of tools. Data object attributes could include data names, definitions, units, labels, etc. Vocabulary would be controlled, that is particular terms like the time unit "seconds" could be represented and spelled in only one way to enable effective querying. (Of course, words would be added to the vocabulary as needed). Population of the data dictionary would be enforced through the UAL, a data object could not be stored until its name was defined. The dictionary could also support name translation from code to code or code to experiment through an ontology of some level of formality. It is noted that a hierarchical name space could most easily support browsing - an important consideration for enabling data location. While the naming system will need to be defined as part of the system design we can assume the existence of unique identifiers for each data item which would include globally named data services for remote access.

4.3.6.5 *Interactions with other FSP tasks*

4.3.6.5.1 Task composition (workflow)

The plan for integrated data management described here is consistent with the needs of computational workflow (task composition) and for production computing. Workflow that heavily involves data preparation, movement and staging is tightly coupled to data management and would be both a customer and a provider for stored data and metadata. Workflow tools would create entries in the data catalog and insert values as part of their operation, associating data in various states of pre- and post-processing. For code run preparation carried out "off-line", that is without a network connection to the catalog server, data would be stored locally and would have to be merged at some later time.

4.3.6.5.2 Production Computing

It is anticipated that the same data management system will be used by developers and users, including verification, validation and UQ activities. Production computing differs from "development" computing mainly in the number of runs performed. The technology underlying the data catalog must be able to support efficient access to tens of thousands of runs - a constraint, but not a severe one. It is possible that production computing would increase the number of sites at which FSP codes are run and where its data is stored. Design for the data catalog must take this into account and provide for mechanisms to unify the view of the data regardless of its location.

4.3.6.5.3 Data Centers (coordinate with user support/operations)

The physical realization of the data systems for the FSP also deserves attention. The volume of data as measured by the total number of bytes or files or named items is likely to be very large. Tools will be needed to support efficient staging, movement and access to "big data" - the large arrays that modern simulations can produce. Data will need to be preserved through traditional back-up strategies and protected against disaster. Some consideration of long-term data preservation and strategies for dealing with changes in storage hardware and software will need to be part of the FSP program.

Maintenance of a system supporting the level of coherence and organization outlined above, strongly suggests the existence of a relatively small number of dedicated data centers. Who would "own" and operate these centers has not been seriously discussed yet - but it seems clear that the FSP staff must have a leading role in administering the data at the centers. Consideration of data transfer rates suggests co-location at major OSC computing centers and/or at centers for FSP computing (if they exist). Since post-processing of simulation output can also require significant data movement, it may also be wise to locate substantial computing power for these operations at the data centers. The nation-wide design of data and computing systems used by the FSP should try to align data production, movement and consumption in a way which optimizes overall throughput and productivity.

4.3.6.6 *Tasks and Technical challenges*

A set of important challenges has been identified. These will need to be addressed at some level of detail in the next six months, as part of the project definition, and during the early phases of the program implementation.

- Catalog database design: The schema (tables, attributes, relations, etc.) which supports the data catalog functions is likely to be complex. Design of this schema will need to take into account use cases and performance requirements.
- Supporting the universal view: Given the distributed nature of FSP activities, mechanisms for synchronizing the single coherent view of all data provided by the catalog, will be required.
- User interfaces/portals: for entering and viewing catalog and archive data
- Name space management: Here there are high-level design decisions which have broad implications across many parts of the FSP. Standards will need to be defined and a set of user tools specified.
- Universal Access Layer: This critical piece of software must be highly capable and easy to use.
- Efficient access to "big data": Approaches for storage and use of the very large data sets that may be generated need to be outlined. Caching, while maintaining data consistency, distributed computing or other strategies for improving overall performance will need to be investigated.
- Query estimation: assessing the size and resources required before large data transfer
- Integration with authentication and authorization systems
- Migrating users from older approaches including adequate training and documentation.
- System evolution: The overall system outlined here is moderately complex with many moving parts. How to begin with modest implementations and build in capabilities as needed over time.
- Design processes: We must define an approach to developing and evaluating conceptual designs starting with the project definition and continuing throughout the life of the FSP program.
- Governance and decision processes: The management processes that oversee the development and maintenance of the data management system and oversee authorization, data quality standards, data retention policies, etc. need to be defined.
- Change management: Some structured approach must be in place to plan, document and control the evolution of the system over time.

4.3.7 Developer support

4.3.7.1 *Scope of this section*

One of the main products of the FSP is software. Proper developer support is critical for producing fusion simulation software that is robust, reliable, portable, scalable, usable, and maintainable. Based on community input on the current state of development support in the fusion community, we have identified a number of development support areas which are described in the remainder of this section, including a discussion of specific requirements, the approach recommended for the FSP and associated costs, and any technical challenges we have identified. Most of the effort estimates are based on data provided by the following projects: Community Earth System Model (CESM) project (10 core developers and 100-200 active contributors); the FLASH Center (9 core developers and many external contributors); and PETSc (20 active developers).

The requirements discussed in this portion of the report may have to be revised and extended based on the approach taken by the ISA, system integration, task management, data management, and user services working groups.

4.3.7.1.1 Requirements

4.3.7.1.1.1 Revision control

Requirements. Revision control requirements include easy remote access from multiple platforms, local control over repositories, the ability to federate (mirror) repositories, and the nesting of multiple separate repositories.

Approach. The main feasible choices are Subversion (http://subversion.tigris.org) and Mercurial (http://mercurial.org). Other options lack some important features, e.g., CVS lacks atomic commits, symbolic links handling, binary differences, and merging). Some advantages of Mercurial over Subversion include better automation of merges and more flexible federation of repositories. Another advantage of Mercurial and other distributed version source control systems in general is that final commits to the main repository (similar to the "trunk" in Subversion conventions) are annotated with all the comments of any intermediate commits and are more robust upon server failure (since any checked out version can become master if necessary). Nesting of repositories is supported through the subrepositories feature which allows a collection of repositories to be treated as a group (similar to the Subversion externals property). A disadvantage is the need to check out the entire history of a repository when it is first cloned, which can be time-consuming, especially through remote connections. This is a one-time cost, however, and would thus not hamper development significantly. In light of these considerations, the FSP recommends that FSP developers use Mercurial for revision control in support of the development of all FSP software and associated documentation. In addition, Subversion mirroring should be provided for ease of access by users who may not be familiar or do not wish to use Mercurial.

To support reliable access and automation of development tasks (e.g., testing), FSP should maintain a central collaboration server that provides hosting for repositories and other services.

Technical challenges. The main challenge to the FSP community is the need to integrate previously unversioned code into a revision control system and for some, the learning curve of a new tool.

4.3.7.1.1.2 Naming conventions and namespaces

Requirements. Prevent name collisions in coupled fusion simulations, enable support for both loosely coupled (not single address space) and tightly coupled (single address space) simulations.

Approach. Whenever multiple codes are linked into a single executable, there is always the potential for symbol collisions (i.e., multiple source codes producing identical, potentially incompatible linker symbols). The FSP will define symbol naming conventions for all languages used in FSP software development and require their use to minimize the chances of symbol collision. The easiest approach is to require that every FSP component add a unique, short string to the beginning of all their module names and free functions in Fortran code. C++ source code should use an unique namespace, and C source code should also prepend a unique, short string to type definitions and function names. This simple discipline helps avoid symbol collisions, and it can also help in debugging because it's easy to go from a symbol name to the associated FSP component.

Technical challenges. The main challenge is the conversion of currently non-namespaced codes to use new naming conventions. Some of this process can be automated with existing refactoring tools (e.g., Eclipse) or custom scripts developed for this purpose. Another challenge is in maintaining these codes as the computer architectures change, and the coding teams change.

FSP support. The SIS group must develop documentation of the naming convention and assist developers in selecting unique namespaces for existing and new components. To support single address-space integration, the callable interfaces of existing codes (public interfaces of components) and any newly developed components and other software infrastructure must be refactored to include unique naming or alternatively, namespaced interface wrappers can be created. For well-structured codes, this would incur minimal developer cost because only small portions of legacy codes would require modifications. The SIS effort required for creating documentation would be approximately 0.25-0.5 FTEs for initial documentation and for ongoing developer support related to naming conventions and namespace-related refactoring support.

4.3.7.1.1.3 Communication services

Requirements. Enable and encourage communication among developers and between developers and users. Because of the large numbers of FSP participants, it will be important to enable developers to access to communications (e.g., teleconferences) in which they did not participate but which may be relevant to their work.

Approach. With respect to developers, the availability of multiple flexible communication channels is essential in enhancing productivity, preventing misunderstandings, and minimizing distractions. To that effect the FSP will provide infrastructure to enable different means of developer communications, including mailing lists (e.g., with Mailman), central web-based resources (wiki, bug trackers, software directory, static web pages), teleconferences with desktop sharing and recording (e.g., with ReadyTalk), creation of web casts and screen casts of presentations (e.g., with Screenflow). FSP will create guidelines on how the developers should use these tools.

Technical challenges. Installation and support of multiple servers and managing upgrades. Security challenges at Labs.

4.3.7.1.1.4 Software access

Requirements. Ability to access, build, and use specific versions of FSP software and up-to-date documentation.

Approach. FSP will provide a central directory of software at the FSP collaboration server that includes references to a canonical location (URL) of each top-level package (component or set of components). Using this directory, FSP can provide access to and archiving of both development versions and public releases. In addition to providing a uniform interface to accessing software, FSP should provide mechanisms for linking related information (e.g., documents, standards) and ensuring that the central directory is up-to-date. To support limited cross-platform builds, FSP will identify or develop a package management tool and require that all required software metadata (e.g., third party package dependencies) be specified using a prescribed format by all component developers. In addition to package dependencies, FSP will define build conventions (e.g., makefile targets) and expectations on the degree of portability and robustness of package builds that enable uniform builds of FSP components. The specific choice of build system technologies will depend on the solutions adopted for system integration.

4.3.7.1.1.5 Software release process

Requirements. Versioning conventions, central directory of FSP software.

Approach. FSP will define common versioning requirements and a common release schedule based on physics milestones, requiring at least annual releases. All software will be accessible through a central directory maintained on an FSP server. Individual developer groups must provide a URL and other required metadata for inclusion in the directory. Software metadata must include dependencies on third-party or other FSP software, including version numbers and/or repository tags.

Technical challenges. To prevent the directory from becoming stale, identification of dead links to software in the central directory must be automated.

FSP support. Definition of software metadata and ensuring that software releases meet requirements requires approximately 5% of FSP development effort (by dedicated support staff) initially. Based on data from CESM, PETSc, and FLASH each release would require about 2.5% of general developer effort for the two weeks preceding the release.

Issue tracking

Requirements. Ability to submit feature requests and bug reports, assign them to developers, and track their resolution. Support both email and web interfaces.

Approach. Provide centralized issue tracking service for all FSP software development as part of the FSP collaboration server. Issue tracking using the Roundup software has been used successfully for a number of years by other projects, including the Common Component Architecture (CCA) Forum. An important feature supported by Roundup and lacking in some centralized project management systems (e.g., Trac) is the support of email interfaces in addition to web-based issue tracking forms. This significantly lowers the barrier to communication between new users and developers because issues can be created and discussed entirely over email without requiring the participants to set up accounts or use a web interface. The experience of existing projects (CCA Forum, PETSc) shows that the community is much more likely to communicate effectively if they can do so over email rather than web-based interfaces. Thus, we recommend that FSP adopt a system that supports email issue management, such as Roundup or the all-in-one project management system RedMine.

Technical challenges. Maintaining a reliable issue tracking service and ease of creation and configuration of new trackers.

4.3.7.1.1.6 Software standards

Requirements. FSP software should follow common versioning conventions, be accessible through a central directory, and contain sufficient up-to-date documentation. ISAs must have debugging strategies for both stand-alone components and coupled applications.

Approach. FSP will create a comprehensive description of the requirements that an FSP-compliant software package. Create and maintain versioning and documentation requirements and assist physics components developers in implementing them. Assist ISAs with defining and implementing software standards that enable code coupling and debugging.

Documentation. All FSP software packages must provide the following types of documentation, which will be defined

Internal (API) code documentation for developers of the package. This type of documentation will necessarily evolve with the code and should be integrated into it as much as possible, with web and other versions generated automatically using tools such as Doxygen.

Tool (e.g., IDEs, scripts, build systems) documentation for developers.

Documentation for users (normally non-developers); created and maintained as part of the user support FSP services.

Metadata. All FSP codes must be accompanied with metadata describing nonfunctional requirements and capabilities (e.g., build-related requirements and options, system requirements).

Testing. All FSP software packages must include a set of regression tests that can be executed automatically on a set of predefined platforms.

Portability. FSP packages, including their prerequisites, must be portable to a set of agreed upon architectures to be determined in the beginning of the FSP project.

Technical challenges. A gap in existing documentation approaches exists in ensuring that software documentation is correct a robust automated system for creating "live" documentation that runs examples and includes output as part of the documentation.

4.3.7.1.1.7 Testing

Developer support for testing is critical to the development of reliable software. Software testing policies and support are a shared responsibility of both the SIS and SQ missions. Details on software testing are provided in Section 4.4 on Software Quality as a component of the overall Software Quality strategy.

4.3.7.1.1.8 Performance

Requirements. Components must be able to measure their own performance at an adequate level (at least CPU seconds) to identify bottlenecks and performance degradation/improvement with new development.

Approach. Provide preinstalled performance measurement tools on FSP hardware resources and documentation on using them to supplement the tools' own documentation. Include performance tests in regular unit and regression tests.

Technical challenges. The main challenge is in establishing the performance monitoring capabilities of physics components and defining a minimally intrusive approach to performance measurements of simulations on a variety of platforms.

4.3.7.1.2 Summary

In this section we briefly summarize some of the key early steps that FSP should take in the area of developer support.

4.3.7.1.2.1 FSP-wide conventions and standards

Concepts Only

FSP must develop guidance documents and up-front understanding of the following.

- Software management and release protocols;
- Testing and documentation expectations;
- How requirements are specified, prioritized, and assigned; how does the project governance structure support this (e.g., through a board of the major customers)?

This document represents the very first steps in the direction of creating such guidance.

4.3.7.1.2.2 FSP collaboration server

To effectively support multiple aspects of development, FSP will maintain a central collaboration server, with services including: revision control hosting, wikis, issue tracking, mailing lists, project web pages, software directory (including software releases and testing results), and documentation.

4.3.7.1.3 Technical Challenges.

- Security
- Project-wide authentication and authorization
- Ensuring adoption of new (to some) technologies

4.3.8 Summary and resource requirements for software integration and support

SIS has been allocated four FTE at the start of FSP increasing to 5-6 towards the end. This level of funding is insufficient to accomplish all of the goals that have been elaborated upon in this section. Hence, we have pushed out many items, and for the near term we concentrate on the following tasks in on-HPC integration.

- Define implementation independent metadata schema and language APIs (6 mo. at .76 FTE = 0.38 FTE-year)
- Common I/O capability available to the fusion community (18 mo. at 1.33 FTE for 2 FTE-year)
- All ISA providing commonly structured output (12 mo. at 1 FTE for 1 FTE-year)
- On-HPC component interface defined (6 mo. at 1.14 FTE for 0.57 FTE-year)
- Reference implementation of on-HPC framework (24 mo. at 1.5 FTE for 3 FTE-year)
- ISAs refactored to use SIS integration framework (18 mo. at 2 FTE for 3 FTE-year)
- Assessment of 3D data couplings (24 mo. at 1 FTE for 2 FTE year)

The layout of these items in time are shown in the WBS. To avoid repetition, we refer the reader to the WBS for the task composition and developer support tasks to be undertaken in our first five years.

4.4 Software Quality

As discussed in Section 3.5 and Appendix D:, the complexity of the FSP could justify verification, validation, and UQ efforts far beyond the available resources. The software quality activities must therefore be thoughtfully limited in scope and directed towards key questions. Presented in this section is a preliminary roadmap within the context of the two initial ISAs of a scientifically defensible Software Quality (SQ) effort.

Most resources will be directed towards verification and experimental validation. In the context of verification, this includes the traditional software verification activities of unit, integration, and regression testing in addition to numerical algorithm (code) verification. The emphasis of verification efforts will be on code correctness, and most of the activities are sufficiently generic so as to apply well to ISAs and component development with little specialization. In contrast, an extensive experimental validation effort is presented that is tailored specifically to the two ISAs.

In support of experimental validation, the primary goal is to develop probability distributions of computed quantities that accurately reflect the numerical errors and uncertainties in other inputs. Less ambitious mileposts on the road to this destination are still of use. The minimal requirement for justifiable comparison with experiment is that systematic code verification is done on all components. Confidence can be increased with reliable calculation verification capabilities, as these can provide crude discretization error bounds on quantities of interest (QoIs). The inclusion of coupling error estimates further refines the knowledge of uncertainty in the calculated QoIs. Ultimately, the ability to include the uncertainty of inputs in computed QoIs not only improves the conclusions of validation activities, but also provides the means to direct efforts to reduce uncertainties and improve models. Thus, UQ in support of validation need not be an all-or-nothing proposition, so long as the limitations of the conclusions are well understood and explicitly documented.

The SQ execution plan therefore initially emphasizes verification. Both calculation verification and uncertainty quantification efforts will commence within the first year of the FSP, but these activities will be necessarily focused and exploratory in nature. Effective UQ methodologies for a limited set of specific questions will be developed. Over time, the UQ scope will grow to encompass a broader array of applications and will more directly impact experimental validation efforts.

4.4.1 Verification

Lifecycle verification and testing are extremely important for the reliability and correctness of any software. For simulation software, there are two aspects to verification. Simulation is the process of constructing a (reduced) mathematical model of a physical system, followed by the definition of a numerical solution approach, and finally implementing the solution in software. As such, simulation software contains both numerical and non-numerical algorithms that must be verified, and the techniques are different for each. Non-numerical algorithms generally have well-defined results or behaviors and are therefore directly testable. Numerical algorithms produce approximations to physical models for which solutions are generally not known. For clarity, we will refer to these two verification activities as *software testing* and *numerical algorithm verification*, respectively, for disambiguation. The risks of not performing sufficient testing or verification are many, including unreliable results, inability to establish the fidelity of computations, difficulties detecting bugs and correcting them, and long lasting unresolved problems.

4.4.1.1 *Software Testing*

We focus on verification as it affects the implementation of FSP component and integration software.

4.4.1.1.1 Approach

In the first year of FSP, we will develop a methodology for test-driven development. Testing requirements will be defined and provide documentation for verifying physics components and integration software within the limits of available developer resources (e.g., community feedback indicates that most developers are willing to spend up to 20% of their time on verification). We outline some of the development processes that support software testing.

4.4.1.1.2 Unit and integration testing

We expect that the integrated simulations implementing solutions to the proposed science drivers will consist of both legacy and newly developed codes. Because it is difficult or impossible to create unit tests for large code bases that do not already have them, FSP will require interface-level unit tests only for new components and integration software infrastructure developed during FSP. Components based on large legacy codes will be treated as black-box units and tested accordingly. All tests should be included in the source code's revision control repository. In addition to unit tests, integrated scenario tests must be developed and included in an FSP-wide regression test suite. FSP must provide software infrastructure to support the creation and execution of unit and integration tests for all FSP implementation languages, as well as include web-based reporting and archiving of results. FSP must also define a standard "diff" interface and implementations of comparison operators for different data types and output formats provided by components. If a component produces nonstandard (custom) output, the component must provide an implementation of the "diff" interface for all relevant data types. FSP must ensure that regular testing occurs on key platforms, including LCFs. FSP web portals should provide access only to components with input checking as a minimal precaution against malicious or improper use. FSP should attempt to leverage the many existing tools available for unit testing (e.g., see http://en.wikipedia.org/wiki/Category:Unit_testing_frameworks).

4.4.1.1.3 Regression testing

Regression testing is the process of partially retesting a modified program (e.g., as new functionality is added or a bug is fixed). Regression testing can also be used to identify performance degradation. Regression tests rely on data management and task composition for creation and execution of reproducible tests. The following regression test practices should be followed by all FSP software packages:

- When a bug fix is implemented, a corresponding test is committed to the code's repository.
- When an integrated test fails, the FSP framework must be able to execute the unit or regression tests corresponding to the components involved in the integrated simulation.
- Continuous integration: test whenever changes are committed to a repository (at most daily).
- Input checking. Each component must check the correctness of the supplied inputs as a standard part of the simulation workflow (see report sections on Task Composition and Data Management).

FSP codes must include metadata describing test requirements, for example the costs of individual tests in terms of core-hours estimates. FSP should consider leveraging existing test harnesses that may be available, but often such test harnesses are written specifically for the application at hand.

4.4.1.1.4 Technical challenges

Providing as much testing automation as possible will be crucial in reducing the long-term verification effort in FSP and enabling more extensive testing. Some regression tests correspond to problems that occur only for larger core counts and would require significant computational resources.

4.4.1.1.5 FSP support

The FSP collaboration server will host or mirror test results for all FSP codes. Based on community feedback and practices in the CESM, PETSc, and FLASH projects, we have determined that over the life of the project, approximately 5-10% of each developer's time will be dedicated to verification: creating and running tests. The maintenance of a (semi-)automated test system would require a 0.25 FTE on an ongoing basis; as a developer support activity, maintenance would be the responsibility of the SIS team. The SIS and SQ teams will collaborate on the initial effort on defining a test-driven development methodology and creating/adapting tools to support the automation of testing on FSP and LCF resources; this task will require approximately 1 FTE for one year.

4.4.1.2 *Numerical Algorithm Verification*

The activities for both ISAs and any additional component development are quite similar, so we do not distinguish. The code and calculation verification activities will be led and primarily conducted within each ISA and component project, but at-large numerical analysts from the SQ team will actively assist in study definition, analysis of results, review of documentation, and development and implementation of new *a posteriori* error estimation techniques.

Each component, whether part of an ISA activity or component development project, will have to undergo a formal verification process. The basic process, to be refined by the SQ team and adapted over time, will be:

- Develop pre-requisite documentation
  - Develop (assemble) documentation of the mathematical models, the discretizations, the parameters, a priori expected rates of convergence, test problems with and without solutions, physically important quantities of interest for the test problems, and prior verification and benchmarking activities. This is necessary not only for good verification practice, but also as part of the componentization process and necessary for the validation process evaluation. It is expected that for existing codes, much of this material exists, but it will still need to be compiled.
  - If only benchmarking has been done previously, define appropriate test problems, either from theory or using the method of manufactured solutions, if possible. Manufactured solutions are preferable, as these can improve verification test code coverage.
- Conduct code verification on the identified problems
  - Design the studies, considering Lp-norms in addition to previously identified metrics. Grid convergence studies using Richardson extrapolation will be the most common method for the demonstration of convergence at the anticipated rate.
  - Develop the necessary discrete problem definition. This development includes the formulation of input files as well as any modifications to the code, for example, the inclusion of prescribed sources to enable manufactured solutions.
  - Automate the study generation, execution, and analysis. Combined with the necessary input files and output files to be used as reference solutions, this provides a set of tests suitable for future regression testing.
  - Identify any convergence problems, correct, and repeat until satisfactory convergence results are obtained.
  - Document the results.
- Explore calculation verification techniques
  - Design the studies, primarily considering relevant QoIs.
  - Perform grid convergence studies at four or more resolutions. Using Richardson extrapolation, determine the rates of convergence for any three successive resolutions. If consistent results are found, estimate the error in the QoIs. If not, investigate the cause of the inconsistency.
  - If Richardson extrapolation fails, other a posterior error estimation techniques can be used as appropriate. Many of these will require additional theoretical and code development, but not necessarily intrusion into the code.
  - Document the results.

Based on the Science Driver reports and the subsequent planning workshop, the degree to which existing codes have undergone a formal verification procedure such as that as outlined above varies greatly. For those that have undergone extensive testing already, most of the effort will be to capture the required information in a single document and to make the tests suitable for regression testing. For new components, the development team should consult the SQ team on the inclusion of native error estimation capabilities, such as adjoint or error transport methods.

A similar process will be applied to integrated applications, although the results are expected to become ambiguous as the problem complexity increases. Special care needs to be taken in the problem definition in order to ensure that a mathematically consistent problem is defined for any grid convergence study. As more advanced *a posteriori* error estimators are developed and made available, it is expected that code and calculation verification on integrated applications will be revisited.

4.4.1.3 *Tasks and Milestones*

Since the ISA and component plans have not yet been formulated at a suitable level of detail, absolute verification tasks and milestones based on particular components cannot be constructed. Nevertheless, relative tasks and milestones can be formulated. We note that, prior to any planned validation activity, at minimum the code verification of each component must be completed. Here we present a summary of the tasks and milestones; a more detailed breakdown is provided in the WBS.

Year 1:

- In consultation with ISA and component projects, the SQ team formulates, refines, and documents project-wide software verification policies and procedures.
- In consultation with ISA and component projects, the SQ team formulates, refines, and documents project-wide numerical algorithm (code and calculation) verification policies and procedures.
- The ISA and component projects, in collaboration with numerical analysts from the SQ team, develop and document verification plans for each code that is componentized / developed.
- Consultation between the SQ team and the ISA verification leads should begin to ensure agreement on the targeted validation problems and metrics.
- In consultation with ISA and component projects, the software testing tool requirements are identified.
- Existing software testing tools for unit, integration, and regression testing are identified and evaluated. Where no suitable tool exists, development is begun on tools that meet the project requirements.

Year 2:

- Ongoing application of the verification process.
- A tool suite of software testing tools is made available to the whole of the project.
- Numerical analysts from the SQ team will begin developing strategies for the application and development of advanced error estimation techniques.
- As soon as the capability exists to simulate for the milestones, calculation verification efforts should begin to determine techniques that can provide reliable error estimates in quantities of interest important to experimental validation.

Year 3:

- Ongoing application of the verification process.
- Numerical analysts from the SQ team will continue to develop strategies for the application and development of advanced error estimation techniques, with special emphasis on coupling errors.
- The SQ team will begin to consolidate the best of the developed grid convergence study tools into a single suite of tools for study creation, execution, and analysis. This preferably should be integrated into a program-wide workflow tool.
- Initial calculation verification techniques will be available for use in more systematic validation studies.

Year 4:

- Ongoing application of the verification process.
- SQ team makes initial release of a tool suite for verification.
- Continued work on advanced error estimation techniques.

Year 5:

- Ongoing application of the verification process.

- Continued work on advanced error estimation techniques.
- Release of verification tool suite including new capabilities.

4.4.1.4 *Resources*

Each ISA should dedicate at least 0.5 FTE towards these activities. The appropriate level for component development will depend on the size and scope of each component development team. Initially, approximately 1 FTE of numerical analysis expertise will be required in the SQ team to support these activities.

4.4.2 Experimental Validation

4.4.2.1 *Plan for ISA 1: Boundary and Pedestal*

4.4.2.1.1 Goal and Focus

This ISA aims to develop the capabilities for modeling the outer region of the tokamak from the top of the edge pedestal to approximately a millimeter into the first wall with the goal of quantitatively predicting the density and temperature of the pedestal and the heat and particle loads leaving the plasma and impacting the plasma first wall and divertors. The model should cover phenomena over a wide range of timescales from the steady-state (time-averaged) heat and particle fluxes to larger transient fluxes induced by off-normal and loss of performance events such as disruptions and edge localized modes (ELM).

While a high pedestal is optimal for overall fusion performance, the free energy in the sharp gradients of the pedestal can also drive intermittent instabilities called Edge Localized Modes (ELMs). Though ELMs are generally benign on existing devices, they deposit heat loads on material surfaces, which could constrain material lifetimes on ITER, and operation with ELM control or in regimes with small or no ELMs is desired. Hence a validated understanding of the L-H transition, pedestal structure and ELM dynamics is crucial to the successful operation of ITER. Furthermore, ELM events, bursty transport and fuelling via neutrals couple the pedestal to the open field line scrape-off-layer (SOL) region and the material surfaces.

Normal operation of ITER and fusion reactors requires successful channeling of plasma heat flux from the core region to the scrape-off-layer (SOL) where the open magnetic field lines guide the heat and particles to the divertor. In the divertor, it is necessary that the heat be conducted away safely over a sufficiently large surface area or radiated in the presence of impurity ions to avoid material heat flux limits of ~10 MW/m^2. In the presence of off-normal events such as disruptions and ELMs, heat will escape across the magnetic field and impinge on highly localized spots, which gives rise to a limitation of approximately $E_L \tau_\Lambda^{1/2} \sim 1$ [MJ s$^{1/2}$], where E_L is the energy ejected into the SOL over a time scale of τ_Λ seconds. The large power deposited on the first wall will rapidly erode material facing the plasma resulting in significant shortening of the wall lifetime and thus potentially reducing availability of a fusion power plant. Avoidance or mitigation of disruptions and ELMs is essential for a fusion power plant. Because of the uncertainties of the physics in the tokamak boundary, it is likely any modeling of the boundary will be made up of a combination of first-principles and reduced (even empirical) models in the foreseeable future. A heavy burden will be put on experimental validation to quantify the fidelity of each component as well as the integrated model of this region. Fortunately a wide range of existing devices with pulse length ranging from a few seconds to hundreds of seconds, and operating with very different boundary conditions are available for this purpose.

4.4.2.1.2 Critical Issues

The critical issues that can impact the heat and particle loads as well the edge transport barrier and the maximum plasma pressure at the top of the pedestal include

- Startup
- L-mode, H-mode, L-H transition
- Pedestal structure
- ELM avoidance and mitigation

- First wall (FW) & divertor PMI, loads on high heat flux PFCs
- Evolution of FW & divertor PFCs (material migration, mixed & redeposited materials, etc.)
- RF antenna/SOL interactions
- **Impurity generation and transport**
- **Steady-state operations with self-consistent plasma & wall modeling**
- Termination & shutdown.

The **green** highlighted issues are discussed in the pedestal science driver report, the **blue** highlighted issues are discussed in the boundary science driver report, and the **bold** highlighted issues straddle both areas. Each issue forms an extensive experimental validation campaign. Clearly, the boundary and the pedestal are closely coupled through many shared physics and code capabilities. They are also coupled to ISA 2: Whole Device Modeling (WDM) with focus on disruption avoidance, and other science drivers (see Table 6).

The way these critical issues are structured, they are amenable to experimental validation in a multi-level approach as suggested in the validation best practices guidance (Section 3.5.4 above). In the pedestal area, the issue of L-H transition focuses on a one-time very rapid transition, the dynamics of which changes the character of the edge plasma completely and in fact establishes the edge pedestal. Key theoretical models to be validated would include the cause of the transition, and the behavior of the neoclassical and turbulent transport before and after. The pedestal structure issue centers on the time-averaged behavior of the pedestal in the H-mode. Prediction of stability boundary and transport evolution in a stable pedestal would be subjects for validation. ELMs are intermittent, performance degrading events. The physics to be validated describes the dynamics leading to the crossing of the pedestal stability boundary, the nonlinear consequence of the instability, and the mechanism resulting in the return to stability followed by repetition of the whole cycle. The time-averaged pedestal behavior is clearly dependent on the ELM dynamics and vice versa, which would be addressed by the next level validation (see Figure 7). This will have to be coupled to the L-H transition at yet another level of complexity to validate the evolution of the pedestal from start-up.

Application Area	Capability Needed From ISA 1	Capability Provided to ISA 1	Capability shared with ISA 1
ISA 1: Boundary & Pedestal	Boundary to Pedestal: Heat, particle, momentum fluxes Neutral and impurity fluxes	Pedestal to Boundary: Heat, particle, momentum fluxes	Boundary/Pedestal: Gyrokinetics Fokker-Planck collisions Kinetic neutral transport
Science Driver: Wave-Particle	Plasma profiles Fluctuation levels	Local heat deposition from fast particles and RF	Parasitic RF losses and impurity sources
ISA 2: Disruptions		Transient local heat and particle loads	Atomics and neutral physics, radiation transport
ISA 2: WDM	Reduced models for boundary, especially fueling, fuel retention, impurity sources		

Table 6: Shared physics and code capabilities of ISA 1 and other science drivers

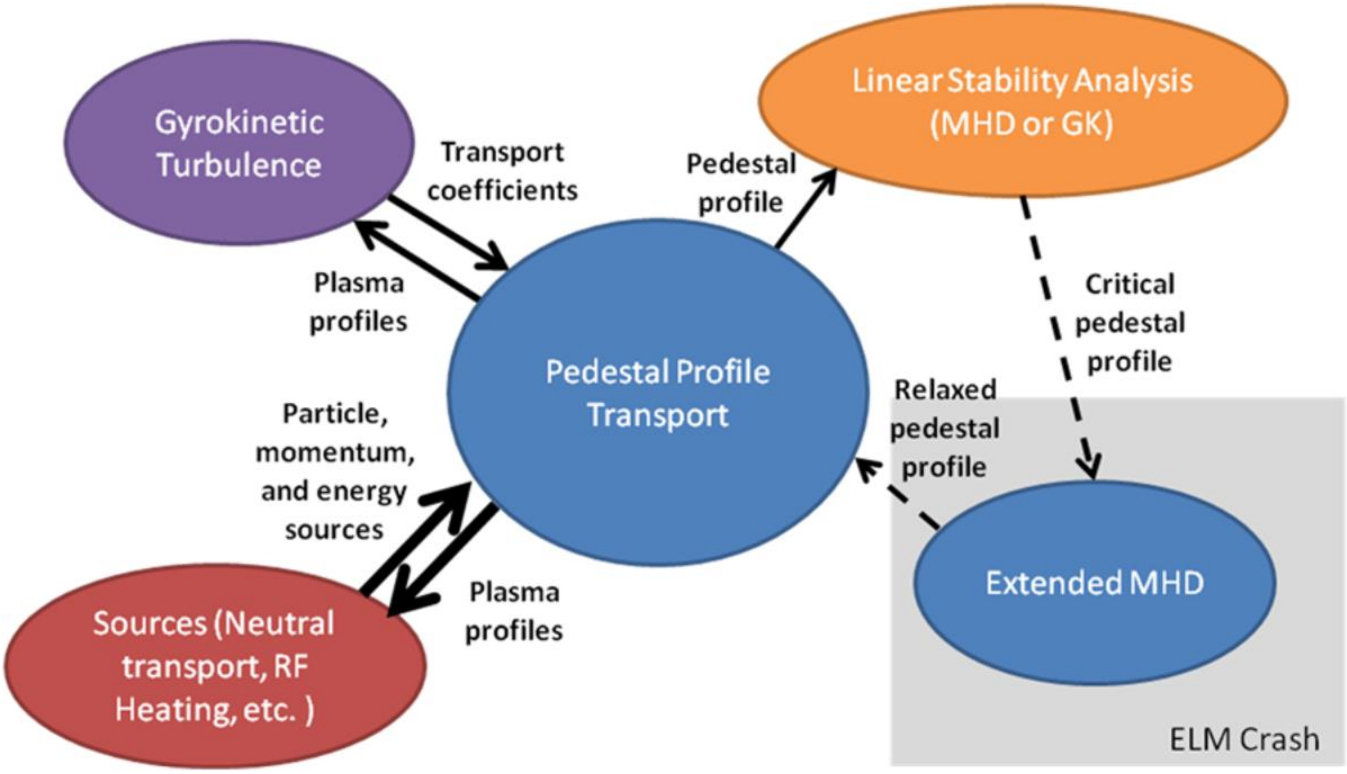

Figure 7: Representative dynamics (Level 2) pedestal component configuration

Analogously, in the boundary area, modeling the loads on high hear flux PFCs suggests validation of a short time scale phenomenon such as how the materials behave under a significant heat pulse from an ELM or a disruption. Conversely, modeling of material migration requires validation of long time cumulative effects. The two require very different designs of experimental campaign. One can also plan validation based on the separation in physics. An example is the RF antenna design that changes the sheath electric field, which impacts impurity production. Separately, the impurity transport can be examined under a fixed background. Keeping in mind the main focus of this ISA is to predict the impact of heat and particle loads on PMI, the code development plan has to eventually couple the heat and particle fluxes from the core and pedestal to the SOL and PFCs. The validation plan for this multi-physics coupling will have to be consistent with the code development timeline.

4.4.2.1.3 Validation Template

Borrowing from past experience in experimental validation, we can construct a template for testing and validating physics processes that reflects the hierarchical strategy. The template has five guiding principles with timescale indicated.

- Most processes have predicted implications for one or more profiles (1-3 years)
  - Make good, time-resolved measurements of profiles to see if predicted limits occur where predicted (e.g., critical gradients)
  - For most processes, we have some capacity to calculate these limits and capabilities are being improved
- Simultaneously, make measurements of phenomena which should appear when predicted limits are reached (e.g., rise in fluctuations with expected characteristics) (1-3 years)
  - For most processes, we can predict qualitative behavior of these phenomena
- Perform steps 1 and 2 over a wide range of plasma conditions, chosen to stress the important parameters of the processes (1-3 years)
  - If a process survives steps 1,2 and 3, we would have good confidence that it is important

- Longer term, need to make quantitative tests of the relevant phenomena (e.g., fluctuation amplitudes) (3-5 years)
  - This will generally require theoretical/modeling advances which are now underway
  - May require diagnostic advances
- For processes that survive steps 1-4, need development of integrated models (transport models or frameworks which incorporate important processes) (3-5 years or longer)
  - Determine how the processes interact
  - Validate integrated predictions against experiment.

4.4.2.1.4 Metrics

In order for this template to provide useful results, quantitative validation metrics will need to be developed for each model application of interest (i.e. different metrics will be needed for studies of L-H transition physics, pedestal structure, ELM dynamics, etc.). These metrics are needed to both establish the fidelity of current models (and thus the confidence that should be assigned to their predictions), and to track improvements in model fidelity as they (and available computing resources) improve. While the requirements for validation metrics is discussed in detail in the best practices section (Section 3.5.4 above), some key features of these metrics are that they should: incorporate an assessment of the numerical error in the model results, as well as both model and experimental uncertainties, and reflect the inherent key sensitivities of the models being considered. In general, a suite of “simple” metrics (which assess model fidelity for a single physical parameter) will be needed, with these simple metrics combined into composite metrics to provide more holistic assessments of model performance.

As an example of a possible metric suite, consider the case of the H-mode pedestal structure. The most basic metrics might be comparisons of the model-predicted pedestal height and width against experimental measurements, using a simple parameterization to characterize both model and experiment results. Here the experimental uncertainties are assessed based upon the fitting of the measured data points to the parameterization, and the model uncertainties via propagation of uncertainties in the experimental input parameters through the model. More advanced metrics might relax the assumption of a single pedestal width or height, and compare the predicted and measured structure of various profiles (e.g., n_e, T_i, T_e, E_r), or replace the use of parameterization comparisons with calculations of chi-squared “goodness of fits” for model predictions to measured data points. Additional constraints, such as predictions for turbulence statistics such as amplitudes, fluxes, and correlation lengths could be incorporated to supplement tests of the predictions of equilibrium pedestal profiles. The additional constraints added as the metrics are refined should be chosen with an aim of identifying key model strengths and weaknesses, thereby providing clear guidance for the theorists and modelers on which aspects of the model need the most improvement.

4.4.2.1.5 Readiness Assessment & Resources

To assist in the development of a validation schedule, we have tabularized (Table 7 and Table 8) under each high priority issue, the critical physics that need to be evaluated, the readiness of the modeling/simulation capability and the experimental readiness in a self-consistent way. The green color in the tables indicates short-term (1-2 years) readiness, the blue color indicates medium-term (3-5 years) and the red color indicates long-term. Using these tables as guides, the next level details can be worked out by the validation team in the execution phase of the FSP.

Issue	Critical Physics	Model/Simulation Readiness*	Experimental Readiness
Cross-field plasma transport	• Micro-turbulence/blobs; transport from strong, intermittent events • Mesoturbulence/ELMs • Coll. & turb. Transport	• Couple SOL fluid plasma transport/turbulence • Couple (2D, 2v) kinetic SOL plasma with nonlinear F-P collision	• Fluctuations: reflectometry, probes, BES, gas-puff imaging • Profiles & flows: Thomson scattering,

	Role of magnetic topology/shear, X-point and wall/divertor contact	model capable of full short-to-long mfp • Extend fluid turbulence to foot of pedestal • Fluid ELM simulation for SOL response • Couple evolving MHD equilibrium to account for shifting separatrix	reflectometry, probes. • Distribution functions: charge-exchange recombination for ions & divertor Thomson for electrons
Heat and particle loads	• Surface fluxes from integrated plasma, atomics phys., neutrals, currents • Fueling, recycling, retention • Shear physics • Radiation transport • Private-flux region transport	• Couple neutral model, initially fluid • Develop and extend kinetic Monte Carlo neutral transport • Couple dynamic wall model for hydrogen wall uptake/ recycling with dynamic 2D SOL plasma model	• Particle fluxes: probes, D-alpha emission profiles • Heat fluxes: IRTV, thermocouples, probes • Near-surface tile analysis of hydrogen depth profiles • Radiation transport: spectroscopy • Private flux transport: probes, divertor Thomson
Material surface evolution	• Plasma surface interaction & resulting evolution • Surface chemistry • Effect of coatings • Dust generation	• Initiate full coupling between near-surface, particle-based sputter erosion/redeposition code for 2D impurities and SOL 2D fluid plasma model • Couple initial surface evolution model and near-surface plasma model	• Surface evolution, surface chemistry, & effect of coatings: DiMES/MiMES-style probes; near-surface tile analysis of element depth profiles; scanning electron microscopy of surfaces; in-situ surface diagnostics (e.g., DIONISOS & MAPP)

Table 7: Boundary physics validation assessment table

Issue	Critical Physics	Model/Simulation Readiness	Experimental Readiness
Pedestal structure and dynamics	• Micro-meso instability • Quasilinear and neoclassical transport • Nonlinear turbulent transport • Particle and energy sources and sinks • Neutral and atomics physics	• Linear Peeling-Ballooning stability analysis for static pedestal structure • Linear electromagnetic gyrokinetics (EM GK) • 2D Neoclassical transport and flows • Static pedestal models based on coupled linear physics • Nonlinear EM GK turbulence simulations • 3D Neoclassical transport including	• Compare gradients within barrier to linear MHD and GK mode onset criteria • Measured edge current comparisons with neoclassical • Dynamic profile evolution • Turbulence comparisons with models •

		stochastic field and orbit loss • Reduced transport models based on nonlinear simulations • Couple to particle and energy sources •	
Relaxation mechanisms	• Nonlinear extended MHD and gyrokinetic models for ELM onset, nonlinear evolution and effects on plasma • Coherent mode stability, nonlinear evolution and effects on plasma • 3D equilibrium effects including non-axisymmetric magnetic fields • Pellet and other ELM triggering sources	• Linear onset from P-B calculations coupled to simple ELM crash models • Direct simulation of ELM dynamics using extended MHD, 2-fluid or kinetic-fluid codes	• Mode structure comparisons with linear calculations • ELM dynamics • Fast profile evolution • 3D equilibria • Multiscale and multi-channel fast dynamics
Transition physics	• L-mode turbulence and transport • Turbulence suppression mechanisms • Feedback loop • Transitions from low to high performance H-mode	• L-mode turbulence simulations with 3D codes • Couple linear or quasilinear gyrokinetic code with realistic geometry and ExB stabilization • Couple transport from core-pedestal • L-H Physics	• L-Mode turbulence characterization • Flow and E_r evolution • GAMs and zonal flow dynamics • Fast dynamics across transition • Fast evolution of flows and E_r across transition

Table 8: Pedestal physics validation assessment table

For analyst manpower estimate, we summarize Table 7 and Table 8 into several validation tasks

- Validation of pedestal structure and dynamics
- Validation of pedestal relaxation and transients
- Validation of plasma-wall interaction
- Validation of coupled pedestal/boundary physics.

The manpower requirement and the validation tasks timeline are given in the WBS. Considerable basic research will be needed to develop more quantitative PMI models although reduced models might be tested early on. For this reason, validation of combined pedestal & wall-divertor interaction will be beyond five years.

4.4.2.2 *Plan for ISA 2: Whole Device Modeling (WDM)*

4.4.2.2.1 Goal and Focus

The goal of this ISA is to build up capabilities for WDM, beginning with existing framework approaches and including components for profile evolution, stability assessment and nonlinear evolution (disruption prediction) including active control. This ISA would tackle the disruption prediction campaign outlined in the disruption science driver report. A key reason for the focus on disruption avoidance in the WDM development is because ITER can only withstand a few unmitigated disruptions a year. It needs WDM for plasma control development to achieve Q=10 while avoiding disruptions. It is envisioned that the WDM capability when developed to maturity will enable the ITER plasma control system (PCS) to meet the challenge of fusion burn control and event handling i.e. keep discharge available for physics exploitation and avoid disruptions and prolong discharges if possible. As an abstraction for the PCS, WDM will integrate all the necessary physics to simulate the plasma response to external influences. Magnetic field coils, heating and current drive sources, and plasma transport properties determine equilibrium shape and profiles. Pedestal/ELMs, fueling, and impurities strongly influence fusion performance. Heating, current drive, fueling, and 3D field actuators strongly influence plasma MHD stability and thus disruption avoidance. Disruption mitigation is required when disruption is unavoidable. Experimental validation will have to be planned to test the fidelity of each physics element, as well as binary and multiply coupled physics.

4.4.2.2.2 Critical Issues

A possible flow-chart for WDM-based stability forecasting is shown in **Error! Reference source not found.** below.

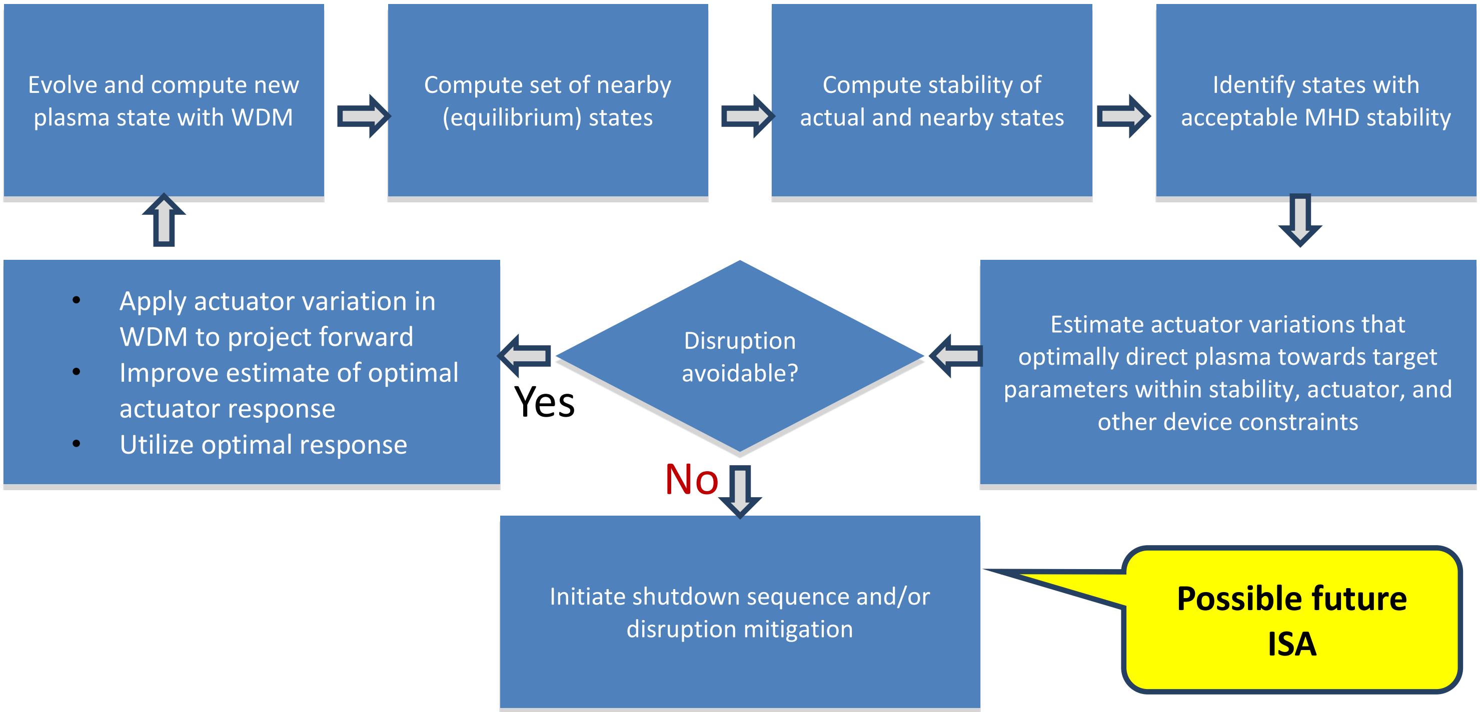

Figure 8:Flow chart for WDM-based stability forecasting

Each box in this figure represents an extensive validation campaign. Following the best practices guidance described in Section 3.5.4 above, a hierarchical series of validation steps should be designed to evaluate the physics. Take for example the box: Compute set of nearby (equilibrium) states. One might start with validating an axisymmetric equilibrium and its sensitivity to measurements of $\langle J \cdot B \rangle$, q_{95}, l_i, etc. 3D effects are often important in tokamak equilibrium solutions. Next hierarchy up in validation will have to include error fields, TF ripples, RMP coils and magnetic islands. Effects of energetic particles on kinetic profiles will have to be accounted for. Further considerations will include the impacts of 3D fields on transport and equilibrium profile modifications. The edge pedestal has a profound contribution on the equilibrium. Both the edge bootstrap current and the pressure gradient can quantitatively alter the equilibrium hence the stability of the tokamak plasma. It is clear that diagnostics for measuring the current profile, the fast ion pressure profile, and the edge current and pressure are critical for the validation campaign. An essential list of diagnostics (including synthetic diagnostics) should be identified for the validation campaign designed for each box in **Error! Reference source not found.**.

At the next level of complexity, validation will address the stability prediction capability, which will fully utilize the validated equilibrium models. The stabilities relevant to tokamak disruptions can be classified into six types (Figure 9): external kink, vertical displacement event (VDE), internal kink (sawteeth), lock mode, tearing (TM) and neo-classical tearing (NTM) mode, and resistive wall mode (RWM).

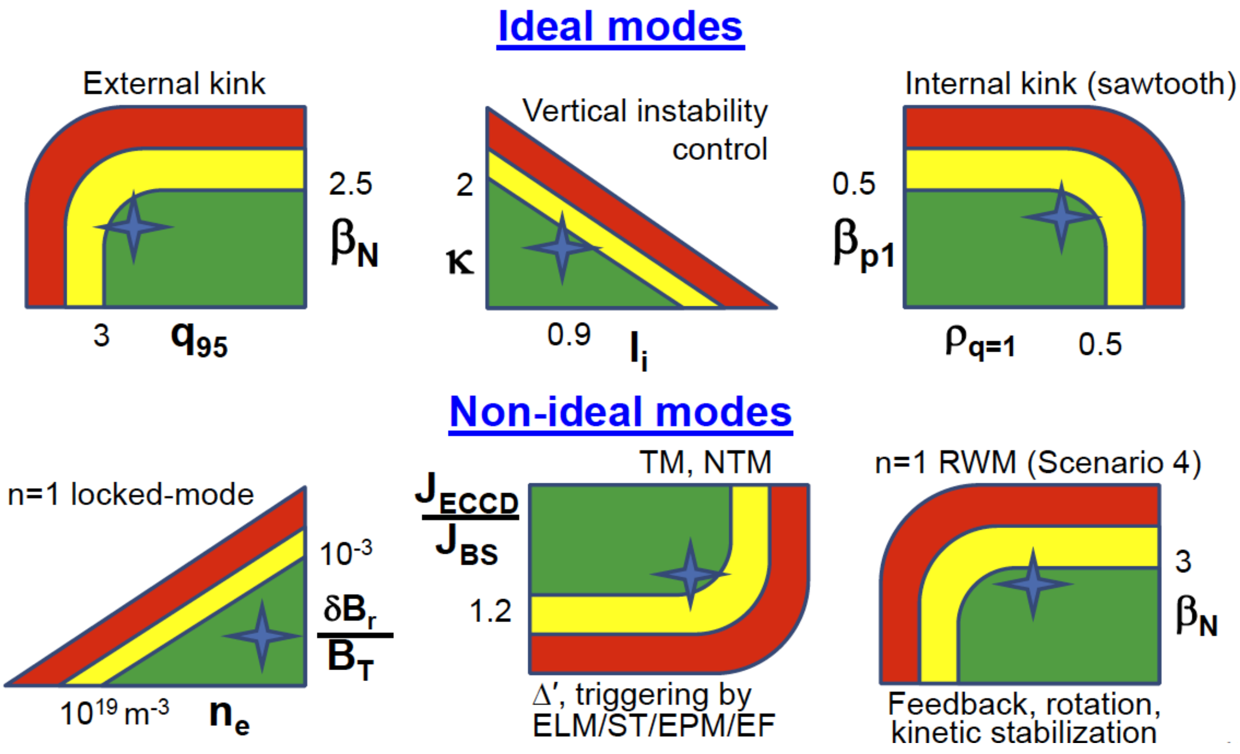

Figure 9: Classification of instabilities responsible for tokamak disruptions

Each of these instabilities requires validation of different critical physics issues. For the pressure and current-driven external kink modes, validation should focus on the WDM-experiment comparisons of disruption probability versus proximity to ideal limits. For VDE, the focus should be WDM-experiment comparisons of plasma response to varied κ, shape and li, also understanding of control noise, impact of disturbance and implications for ITER. For sawtooth instability and control, validation activities will evaluate 2D and 3D equilibrium and transport response to sawtooth, equilibrium and sawtooth control using validated actuators, and NTM triggering by sawteeth. Under locked-modes and error-field correction, the validation activities will include 3D (perturbed) equilibrium calculations needed in WDM for locked mode threshold, and establishing theoretical understanding of locked-mode threshold scaling. For NTM stability and control, the focus will be evaluating the fidelity of WDM combined with nonlinear extended MHD for understanding NTM stability/triggering thresholds, and WDM with 3D equilibrium for understanding transport response to NTM. RWM validation will continue to assess the validity of perturbative versus self-consistent RWM models, WDM-experiment comparisons of RWM stability thresholds, and the ability of actuators to modify equilibrium to optimize RWM stability. Since the proposed validation tasks are extensive, it is recommended that in case further prioritization is needed, the first focus should be on the VDE and NTM induced disruptions. These two are the most frequently observed causes for disruption reported on JET and other tokamaks.

Transport profile evolution governs the dynamics of a discharge leading to the eventuality of a stable plasma or a disruption. WDM is the tool used to simulate the profile evolution dynamics. Extensive experimental testing is needed to quantify the accuracy of the predicted profiles, which are essential for calculating stability thresholds. 3D magnetic fields play an important role in many disruption scenarios. 3D fields from MHD modes can damp rotation and induce disruption. The status of understanding of particle and momentum transport needs significant improvement. Fast ion transport by 3D magnetic perturbations can impact NBI deposition and alter the pressure profile. Applied 3D fields can provide useful control tools to improve stability and transport. Each effect depends on multi-physics, which in the near-term can only be modeled by WDM with reduced models. For this reason, it is important to quantify the impact of loss of fidelity in going from first principles to reduced components in WDM simulations. A plan to experimentally evaluate the fidelity of WDM in simulating multi-physics in the next five years is strategically reasonable and highly useful, for example in application to determine optimal actuator and transport response to avoid disruption.

4.4.2.2.3 Accuracy Requirement

"How accurate does a model have to be?" is a relevant question for FSP to consider. The need to avoid disruptions implies FSP components must be of sufficient accuracy to enable robust control and preserve sufficient distance from controllability boundaries. For example, to ensure reliable positional control in a vertically elongated tokamak plasma, the vertical instability growth rate must be monitored in real-time and kept below a critical value. Assuming Gaussian statistics, in order to achieve an incidence of disruption below one per year in DIII-D (in the absence of hardware or other system faults), this critical growth rate must be ~30% below the "moderate risk" growth rate for which the typical noise and disturbance environment produces ~5-10% disruptivity. An FSP growth rate predictor with 5% accuracy would thus require the predicted growth rate to remain ~35% below the moderate risk growth rate. It is the size of the required margin from the controllability boundary (~30%) that typically determines the required accuracy for real-time monitoring. In contrast, the model accuracy required to ensure robust closed-loop stability characteristics is often much less stringent. For example, design of a robust linear control algorithm for vertical control typically requires no better than 20% error in predicted growth rate. The range of variation in accuracy required by these examples illustrates the importance of specifying a "target" uncertainty (TU) for a given FSP component such as the WDM.

4.4.2.2.4 Readiness Assessment & Resources

To assist in the development of a validation schedule, we have tabularized (Table 9 and Table 10) under each high priority issue, the critical physics that need to be evaluated, the readiness of the modeling/simulation capability and the experimental readiness in a self-consistent way. The green color in the tables indicates short-term (1-2 years) readiness, the blue color indicates medium-term (3-5 years) and the red color indicates long-term. Using these tables as guides, the next level details can be worked out by the validation team in the execution phase of the FSP.

Issue	Critical Physics	Model/Simulation Readiness	Experimental Readiness
Fast MHD-induced disruptions (VDEs, ideal MHD)	• Stability of low-n modes • Nonlinear VDE evolution • Uncertainty quantification of stability boundaries • Control of actuators for stable equilibrium access	• Use WDM to simulate onset of VDE (force balance and control) • Extend MHD component capability to model impurities, radiation and wall (reduced model)	• VDE evolution: useful validation data available on many tokamaks – vertical position, halo/Hiro currents, heat flux patterns. • Low-n modes stability and uncertainly: analysis of closely spaced, high quality EFITs from existing or new experiments to test stability codes against data.

			• Dedicated experiments needed to test real time stability analysis and algorithms to avoid unstable state through profile control actuators in variety of operating scenarios.
Tearing mode-induced disruptions	• Accurate closures for MHD equations including energetic ions • Evolution of tearing modes on transport time scales including rotation dynamics and interaction with external structures • Threshold physics of neoclassical tearing modes	• Couple neoclassical gyrokinetic code to 3D equilibrium with magnetic islands • Develop 3D equilibrium solver that can handle islands and stochastic regions • Couple 3D equilibrium with 2.5D WDM code • Couple gyrokinetic turbulence code with 3D equilibrium	• Variety of existing data on NTM threshold and evolution
Resistive wall modes-induced disruption	• Accurate closures for MHD equations including energetic ions to accurately capture RWM stability • Evolution and control of RFA on transport time scales including rotation dynamics and interaction with external structures	• Incorporate kinetic effects in extended MHD • Couple self-consistent rotation with MHD and transport	• RWM stability: validate with closely spaced high quality kinetic EFIT reconstructions including fast ion pressure from existing or dedicated experiments

Table 9: Disruption modeling assessment table

Issue	Critical Physics	Model/Simulation Readiness	Experimental Resources Needed
2.5 – 3D free boundary equilibrium generation and discharge evolution	• Model field errors, magnetic islands, applied mag. perturbations • Evolution of plasma and machine parameters • Dist. Functions in 3D space • Self-consistent	• Componentize 3D equilibrium with nested flux surfaces and prescribed boundary conditions • Couple flux surface averaged equilibrium quantities with 1D transport • Componentize 3D equilibrium with	• Measurement of magnetic field at multiple toroidal locations external to the plasma

	treatment of EPs from NBI, ICRF, and fusion products	islands and prescribed boundary conditions - 2.5D reduced model transport simulation with island evolution	
Evolution of plasma profiles from boundary to core	- Coupling of validated models for microturbulence and EP modes, and their effects on transport - Onset and evolution of internal transport barriers - Effect of large-scale instabilities on transport	- At least one componentized solver module with access to all reduced transport models embedded - Improve model for poloidal and toroidal momentum transport - Extend solver/transport component to include first principles transport models in both fluid-based and kinetic based reduced WDMs	• Fluctuation measurements of n, T_i, and T_e in kHz to MHz range in core and pedestal • Measurement of magnetic field fluctuations in kHz range in core and especially in pedestal
Prediction, control, and mitigation of instabilities	- Onset, growth rate, and nonlinear saturation for sawteeth, ELMs, RWMs, TMs, NTMs - How these modes affect plasma evolution e.g., transport and poloidal flux	See Disruption Section	• Measurement of magnetic field at multiple toroidal locations external to the plasma • Measurement of n_e, T_e, T_i profiles with high time resolution • Measurement of magnetic field fluctuations in kHz range in core and especially in pedestal
Interaction of boundary with plasma core	- Effect of heat/particle flux on the boundary and of the boundary on the heat/particle flux - Effects of neutrals, large- scale instabilities, particle losses - Onset and dynamics of the H-mode pedestal; L-H transition	- Componentize reduced pedestal and edge models - Couple of 3D equilibrium, kinetic neoclassical and extended MHD codes	

Table 10: WDM validation assessment table

For analyst manpower estimate, we summarize Table 9 and Table 10 into several validation tasks

- Validation of plasma equilibrium states
- Validation of profile evolution from boundary to core
- Validation of fast MHD-induced disruptions
- Validation of slow MHD-induced disruptions
- Validation of transport-MHD coupled disruption simulations.

The manpower requirement and the validation tasks timeline are given in the WBS. The disruption validations will focus on identifying the disruption precursors and parametric disruption boundaries with high accuracy. Disruption dynamics and mitigation techniques will not be validated in the first five years with the prescribed funding constraints.

4.4.3 Uncertainty Quantification

The stated goals of the FSP program are justifiably grand objectives, but the difficulty of the problem and the limited resources mandate a pragmatic and restricted approach to UQ. Based on the Science Driver reports and the subsequent planning workshop, there appears to be little existing effort in the fusion community on the quantification of uncertainty in simulation. There is thus much work to do, and a great deal of education that must occur of both the fusion scientists and the UQ experts. The expectation is that the UQ effort will start out small and highly focused, but grow in scope and program relevance over time.

It is critical to define a small number of concrete problems for UQ analysis in the context of the two initial ISAs, not for the whole range of possible computations. To do this, the FSP UQ effort will need to establish an ongoing dialogue between physicists, numerical analysts, UQ experts, and computer scientists in order to define and to refine the targeted UQ problems. UQ analysis will not initially be attempted on large, integrated application codes. Instead, hierarchical sequences of increasing complexity (and physics fidelity) will be investigated to develop the knowledge and the methodologies necessary to attempt UQ analysis on integrated applications. Indeed, the focus of the FSP UQ activities will be on the application of existing techniques specifically to fusion problems. For the most part, limited resources will require the FSP UQ effort to leverage the ongoing development of new UQ analysis techniques from SciDAC, ASCR base programs, ASC PSAAP centers, and other external researchers in the field.

Potential UQ activities for each of the two initial ISAs will be discussed. Selection of the most beneficial and tractable activities will be left a decision to be made by the SQ team in consultation with the ISAs. A general roadmap for the UQ activities will then be presented.

4.4.3.1 *UQ Plan for ISA 1: Boundary and Pedestal*

There are many target goals for the Boundary/Pedestal area including prediction of the L-H transition, predicting the onset of ELMs, and predicting the heat flux into the diverter plates. In the first five years of the FSP, the ELM instability threshold will be the target best modeled by the available codes.

The Boundary/Pedestal development plan is already organized into a hierarchical sequence of models of increasing complexity suitable for a staged UQ analysis. This organization can be used to define a sequence of UQ studies for the purposes of developing, testing, and improving a UQ methodology. Specifically, a proposed order would be:

- Determine uncertainty in ELM instability threshold and critical profile (height and width) with geometry and pedestal profiles as inputs using:
  - Linear MHD stability analysis
  - Linear gyrokinetic stability analysis
  - Quasi-linear gyrokinetic analysis
- Determine uncertainty in profiles and fluctuations in the plasma boundary layer from input fluxes and geometry using:

  - Fluid turbulence models coupled to neutral transport and atomic physics
  - Kinetic plasma models coupled to neutral transport and atomic physics
  - Hybrid kinetic-fluid models that span the collisionality range coupled to neutral transport and atomic physics

Each stage of these two targeted hierarchies builds on the results of the previous, and each hierarchy can be further extended to include models of increasing physics fidelity that track the planned code developments. Based on the milestones and deliverables of the Boundary/Pedestal ISA, the capabilities required to execute the above plan will be available in time for UQ analysis.

4.4.3.2 *UQ Plan for ISA 2: Whole Device Modeling and Disruptions*

As before, an incremental approach to the application of uncertainty analysis will be taken towards the WDM ISA. In a manner similar to that for the Boundary/Pedestal, a sequence of hierarchical models from the Whole Device Modeling and Disruptions ISA can be identified for a staged UQ analysis. Specifically, the order would be:

- Determine uncertainty in the plasma with geometry and pedestal profiles as inputs using a hierarchy of core plasma transport models:
  - GLF23
  - TGLF
  - GYRO

- Determine uncertainty in equilibrium reconstruction and instability threshold with geometry and pedestal profiles as input using
  - Mapping code coupled to DCON stability threshold code
  - Perturbed Grad-Shafranov equilibria coupled to mapping code coupled to DCON stability threshold code
  - Add initial GLF23 transport step with fixed sources for self-consistent equilibria

The first study would allow for comparison of uncertainties between reduced models and more complete models. Use of reduced models by WDM is an important practical issue, and such a study would provide further insight into the range of applicability of reduced models. The second study builds on the first step and addresses the sensitivity and uncertainties to prescribed and self-consistent equilibria including the sensitivity to errors in the reconstruction of magnetic field geometry.

In addition, the WDM ISA provides an opportunity to consider a holistic UQ target of high value to ITER: the prediction of disruptions. Because of the size of and the planned operating conditions for the ITER tokamak, disruptions can cause severe transient loads on the machine. According to participants in the 2011 workshop, the target for ITER is to minimize the number of disruption events to about 1 in every 100 discharges. In addition, a large amount of data relevant to disruptions exists across a broad set of experiments. The amount of available data and importance of the disruption problem make this problem an ideal first candidate for UQ analysis on a full-physics application. The WDM ISA and SQ team should consult to identify an integrated application instantiation suitable to investigate this problem and to develop a systematic strategy for the ensuing UQ analysis. The problem should be sufficiently restricted for tractability, e.g., limited to disruptions of an H-mode plasma, in order to reduce the problem space and to down-select from the supporting experimental data.

4.4.3.3 *Tasks and Milestones*

UQ analysis will occur in accordance with the policies and procedures described in Section 3.5.2.2, where an iterative, exploratory process is recommended to make the analysis tractable. Documentation of the approach, including all assumptions is key. Here we present a summary of the tasks and milestones; a more detailed breakdown is provided in the WBS.

Year 1:

- In consultation with ISA and component projects, the SQ team formulates, refines, and documents project-wide UQ policies and procedures.

Year 2:

- In consultation with ISAs, ISA scientists, validation analysts, and UQ expert from the SQ team select one or two problems to address hierarchically with UQ analysis and document these. Specification of the UQ problem begins, including:
    - Identification and documentation of the desired results (metrics, QoIs) of the UQ analysis.
    - Identification and documentation if input parameters and data.
    - Quantification and documentation uncertainties in input parameters and data.
    - Identification of strategies for UQ analysis
- ISA scientists and the SQ team begin the iterative process of exploring parameter space for the chosen problem(s).

Year 3:

- ISA scientists and SQ team continue exploration of parameter space, applying a variety of techniques to reduce the dimensionality of parameter space. Results are documented.
- As appropriate, SQ team applies both deterministic and statistical techniques in order to select the best methods for the problem(s).
- Preliminary results of UQ studies are documented and reported.

Year 4:

- Full sensitivity and uncertainty analyses are applied to the initial levels of the problem hierarchies. Results are documented.
- More comprehensive results of disruption UQ study are documented and reported.
- SQ Team begins development of UQ tool suite for project-wide use.

Year 5:

- Based on previous results, full sensitivity and uncertainty analyses are applied to problems of increasing complexity. Results are documented.
- More comprehensive results of disruption UQ study are documented and reported.
- First release of UQ tool suite to FSP project.

4.4.3.4 *Resources*

Each ISA should dedicate at least 0.5 FTE towards these planning and analysis activities. Approximately 1 FTE distributed among applied mathematics, statistics, and computer science expertise will be required in the SQ team to support these activities, but this effort level will grow over time as focus shifts from software testing and verification activities to more intensive UQ studies.

4.5 Production Computing

4.5.1 Operations and User Support Requirements

It is envisioned that the ultimate beneficiaries of FSP production computing will be end users – magnetic fusion energy scientists using FSP tools to plan or analyze experiments, test theoretical ideas, propose new devices, etc. The empowerment of end users is a core motivation of FSP. Specifically, it is expected that FSP will strengthen the position of U.S. scientists as they propose experiments for ITER. It is expected that very detailed simulation will be required in support of every ITER experimental proposal, far more so than is the case for current day tokamaks. The importance of such preparatory simulation is to be expected in view of the cost of running ITER shots and of the strict requirement to protect ITER from potential damage due to plasma instabilities (ELMs, disruptions).

FSP production computing teams will provide deployment of FSP software in production mode, with documentation and user support to cover all aspects of FSP production computing.

FSP expects to place a very strong emphasis on high performance computing (HPC) as a means to access new levels of fidelity in simulation. End users cannot be expected each individually to master the software engineering intricacies involved in delivering a high performance high fidelity multi-physics simulation built up from physics components to form a complete simulation software system comprising 100,000s to 1,000,000s of lines of code. Therefore, production computing support will be needed to allow end users access to proven, productive FSP workflows.

Support of production computing entails the delivery, to end users, of all software and documentation needed to realize production workflows. User support covers every aspect of the run production cycle:

- Assembly and verification of input data:
    - Tools to generate and verify control data for simulations.
    - Tools to access and verify experimental data and other time dependent inputs.
    - Access to supporting engineering data such as tokamak, neutral beam, and RF antenna descriptions.
    - Access to supporting physics data such as atomic and nuclear reaction rate tables.
    - Thorough documentation of input data assembly process and for run control variables.
    - Staff able to respond promptly to user queries.

- Run submission:
    - Authenticated remote access to computational resources including HPC.
    - Access to queue position information on submitted runs pending execution.
    - Thorough documentation of run submission process.
    - Staff able to respond promptly to user queries.

- Monitoring and user intervention during run execution:
    - Web based tracking of status of executing runs.
    - Visualization of intermediate run results.
    - Owner interface to stop/restart/delete executing jobs.
    - Owner ability to modify ("steer") run input data during execution, where appropriate.
    - Thorough documentation of run monitoring and user intervention options.
    - Staff able to respond promptly to user queries.

- Analysis, visualization of output of completed runs, with data management:
    - Appropriate data management (including archiving and/or disposal, and record keeping) with catalog of completed runs.
    - Authenticated remote access to run data by all authorized project team members and collaboration partners.
    - Tools for analysis and visualization, meeting user defined requirements.
    - Thorough documentation of available tools for analysis and visualization of run results.
    - Staff able to respond promptly to user queries.

- Trouble shooting:
    - Prompt, skilled debugging of all aspects of production work flows and associated software.
    - Assembly of datasets to enable precise reproduction of FSP component failures, for forwarding to domain experts (FSP component developers), when required.

It is clear from the interactions with likely FSP users at the 2010 and 2011 FSP workshops, that such levels of user support are requirements for broad usability of FSP research software.

Experience from earlier efforts has shown conclusively that a strong commitment to production computing and user support does in fact enable user productivity. Thus, the TRANSP/PTRANSP "Fusion Grid", which features a support philosophy aligned with the above, has seen sustained growth in worldwide use (see figures, below). However, it is also clear that such support is expensive. The interaction with users frequently reveals new research opportunities, e.g., in diagnostics simulation or other validation related activity – but a full response to such opportunities can be labor intensive and may fall outside the scope of prior planning.

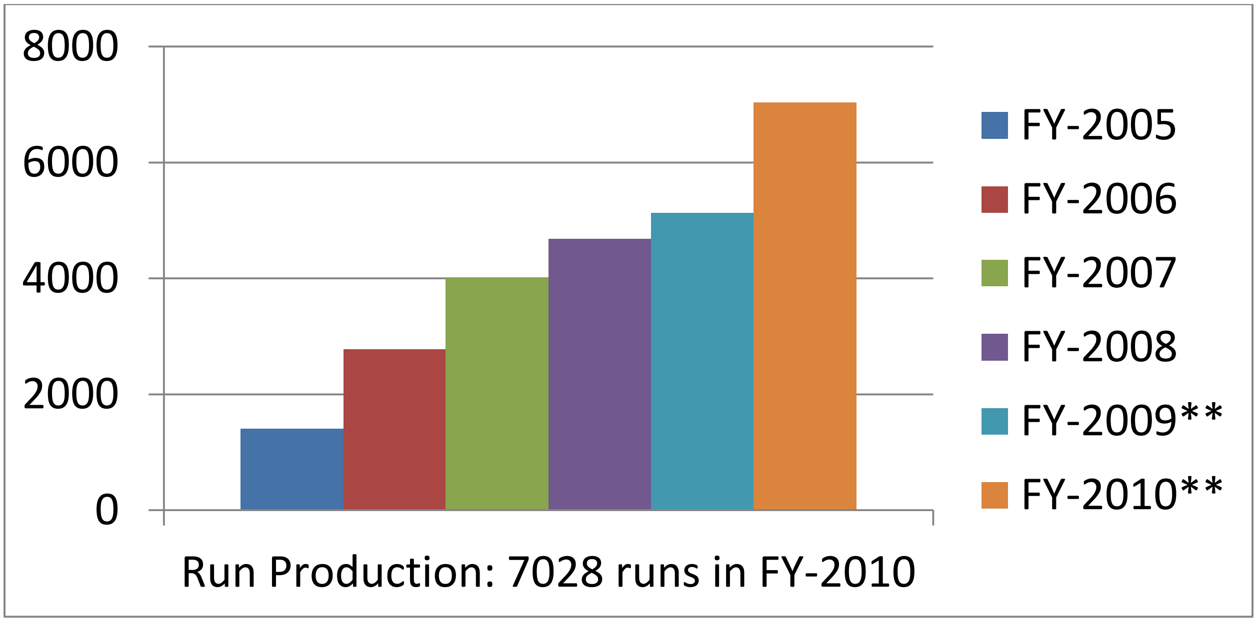

Figure 10: Recent TRANSP/PTRANSP Fusion Grid run production

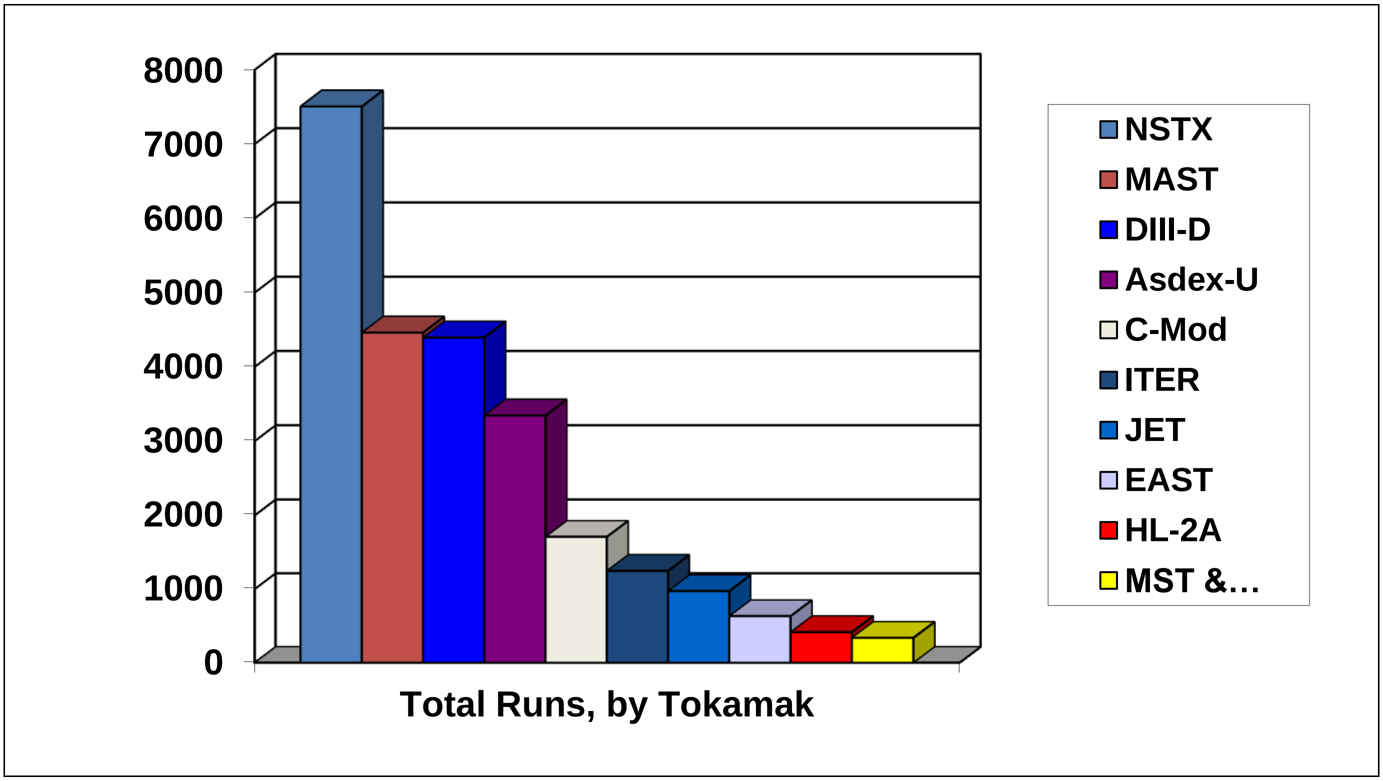

Figure 11: Recent TRANSP/PTRANSP run production, by tokamak*

**About 40% of service utilization, over 10,000 runs in a six year time period, is by international users in the MFE research community*

Preliminary estimates indicate that satisfaction of user requirements, at the level described, can easily entail a cost of 5-10% of FSP project resources.

To control costs, it is important that production operations be carefully planned. This generally entails limiting run production to a small number of sites, at which the expertise (especially for trouble shooting of crashed runs) can be concentrated, with necessary debugging tools, stable compute platforms and systems technical support provided. Such sites should be able to sustain production loads in the range of 1000s of runs. At the start of FSP, the two sites selected for production will be PPPL (where the current “Fusion Grid” operations are centered), and NERSC, where policies have been aligned with capacity HPC and are open to meeting such user requirements as remote job submission. Present day leadership class facilities, with their restrictive access policies and research focus on pushing the limits of HPC, are generally not appropriate (i.e. not sufficiently stable) for production use, but their pioneering efforts do show the way to the production (capacity) HPC of the future (this leads to broad HPC use in which the wider scientific benefit of LCF research will be realized).

4.5.2 Use by Experimental Facilities

Facilities would make use of production versions of FSP ISA applications – whole device modeling as well as detailed edge/pedestal physics modeling – for experimental planning and data analysis. This activity will exercise codes well beyond the range of FSP software development testing, and so, naturally provide a vehicle for much of the required validation effort for FSP software.

The physics requirements for experimental data analysis are only slightly reduced from that of FSP “pure predictive” applications. Even when plasma conditions are largely set by observation, computational simulation still requires detailed modeling of heating, current drive, and aspects of plasma conditions that are not fully specifiable from measurement (such as fast particle distribution functions). In fact, to ensure feasibility of validation against direct observations, constraints on accuracy of source models is general much more severe than is the case for predictive studies of future reactor designs.

Full predictive capability is still needed by existing experiments for planning purposes. Theory-based predictive models are validated against existing shot data. Where these fail, empirical methods can often be used, particularly in the case of planning of experiments on existing devices in which the new shots are modest extensions of shots already carried out on the same facility.

It is expected that Experimental Facilities will work with FSP analysts and collaborate in the validation of physics models as implemented in FSP software.

4.5.3 Use by FSP Analysts for Validation, Verification, and UQ

Although FSP analysts will also have access to development versions of FSP software, it will be advantageous in many cases for analysts to act as users and access the production system, taking advantage of its greater level of automation, support, and integration with data management. Verification and UQ activities often require the ability to carry out 100s or 1000s of runs; often only a well-planned and well supported production computing operation, with full engagement of data management, can sustain such volumes.

The production computing team will work with FSP analysts to assure that their needs are met.

4.5.4 Data Archives

Experimental facilities have generally incorporated the results of analysis and experimental planning simulations as part of their own data archives, using local mechanisms to keep track of runs and support authenticated data access by scientific collaborators.

FSP plans to provide its own data management facility. This will be coordinated with and share much of the technology of data management operations of the experimental facilities. The production computing team will work with FSP data management experts to coordinate operations and assure that production computing works in a manner that is compatible with and takes advantage of data management capabilities.

Where appropriate, the production computing teams may provide software to assist users in exploration of existing data archives.

4.5.5 Documentation and User Support

The operations teams will certainly make available to end users all documentation of components and applications as provided by code developers and associated theory and computation research groups. Beyond this, however, it is more specifically the job of operations teams to provide user level documentation: the detailed, step by step practical instructions for actual execution of simulations and access to output data for interpretation of results. In the long run the operations teams will provide user level documentation for every step needed for production use of FSP software.

4.5.6 Education, Outreach and Training

4.5.6.1 *Education:*

The need for the infusion of young talent into the FSP activities over the course of a program extending beyond ten years will clearly require active engagement with associated education programs. In addition to establishing connections to university programs producing the best young scientists in theoretical, experimental, and computational plasma physics, it will also be important for FSP to attract the attention of graduate students in applied mathematics and computer science with interests in the plasma science applications domain. Establishing FSP postdoctoral positions will be especially important as a foundational component in developing “analysts” whose multi-disciplinary skills are key to a vibrant and productive FSP.

4.5.6.2 *Outreach:*

It is to be expected that outreach activities to large audiences covering broad descriptions of FSP research and capabilities will be carried out on a continuous basis by the FSP. This will involve engagement at high profile meetings and workshops and presentations of seminars/colloquia at universities and laboratories – both national and international. Such presentations will likely lead to inquiries by much smaller groups or individual scientists interested in learning actual use of specific FSP software tools. This would in turn lead to FSP training demands.

4.5.6.3 *Training:*

The training of the requisite talent base for FSP activities will span a multi-disciplinary set of topics. For example, the FSP production computing support teams can be expected to help provide some of the key training support for broader FSP project efforts. Within this context, operations teams will also encourage and assist the formation and training of user groups – with experience often demonstrating that the most efficient education and training often happens when users help each other. Communication among users can often be more effective in training and overcoming user misconceptions than are educational programs conceived from within a development or support group of a software project. For large complicated software systems with correspondingly modest sized user groups, effective training usually amounts to apprenticeship – novice users learning and working with and assisting experienced users. This takes place one user at a time. Therefore, the FSP operations teams will provide materials to aid this process – while recognizing that the main motivation/initiative is expected to come from the user groups themselves.

4.5.7 Deliverables

There is a conceptual challenge in specifying production team deliverables for a software project that has not yet started. The nature of production computing challenges (and associated labor cost) depends sensitively on the operational characteristics of productions systems, as executed on specific computational resources by specific users for specific research applications. This is not accurately knowable prior to user engagement with working systems delivered by other elements of the FSP project.

It is tempting to postpone the formation of FSP production computing teams, under the presumption that they will not be needed until other parts of the project reach maturity. The problem with this is that significant lead time and planning is needed to set up successful operations; if initiative is postponed until needs become obvious to all, it will be too late to respond effectively. Therefore, development of production capability needs to start on “day 1” of FSP. This requires that specific hardware and software systems be identified for “day 1” production deployments.

Fortunately, significantly useful production ready software systems are available from both the MFE base program and SciDAC “proto-FSP” efforts, with capabilities that bear on the FSP ISA applications of Whole Device Modeling (WDM) and edge/pedestal dynamics (EDGE). Initial efforts should focus on deployments based on these existing systems. These systems bring important resources (known to be important from past use experience), such as connections to existing experimental databases (MDSPlus), that will also be needed by natively developed FSP systems as these emerge into production. Where earlier systems overlap with FSP in capability, the co-deployment of FSP and legacy systems will aid greatly in benchmarking and establishment of user confidence in new emerging FSP tools.

Here are the elements of a successful strategy:

- About 5-10% of FSP project effort directed to operations and user support (2 FTE at project initiation).
- Specific identification of hardware resources for production:
  - Computational centers – where the services execute.
  - Data centers – where the long term results of completed runs are stored.
- Specific identification of applications and workflows for production deployment.
  - Initially: existing systems put forward jointly by user groups and development teams, and deployed in collaboration.
  - Later: new FSP developed workflows.
- Proactive support extended to known user groups.

The production computing team would need to expand with the numbers of production applications and runs, the number of users, and the number of computational production platforms. The FSP will start with an initial allocation of effort; subsequent rates of expansion will be set according to user demand. FSP management may need flexibility to accommodate itself to the flexibility of this demand.

4.5.8 Schedule and Resource Requirements

Initial operations team size: 2 FTE. Planning guideline: 5-10% of FSP project labor resources, but actual level to be determined during project according to user demand for services and level of difficulties involved in delivery of services. Indicators of difficulty are the numbers of crashed runs, degree of involvement of operations team personnel in debugging of crashed runs, degree of involvement of operations team personnel in development of user productivity tools (e.g., for monitoring of running jobs, visualization and analysis of simulation outputs, etc.), amount and quality of required user documentation.

- Year 1: identification of initial set of applications, targeted users, data center, and mid-range production facility (e.g., Linux cluster at PPPL with option for clones at other major labs). First production runs on cluster.
- Year 2: identification of first “capacity” HPC production platform: NERSC. First production runs at NERSC. Reliable, well supported production operations on Linux cluster production facilities.
- Year 3: Add applications; improve reliability on all production systems. Establish comprehensive user level documentation (i.e. step by step instructions) for all services.
- Year 5: Develop design proposal for Leadership Class Facility (LCF) production computing.
- Year 7: First documented LCF production application.
- Year 10: Reliable LCF production computing.

4.5.9 Major Milestones

- Year 5: NERSC (capacity HPC) production facility for FSP Whole Device Model and pedestal/edge model.
- Year 5: total production in range of 5000 runs/year on all platforms.
- Year 10: LCF HPC production computing capability.
- Year 10: total production in range of 20000 runs/year on all platforms.

4.6 HPC Resources

The Fusion Simulation Program (FSP) planning activity was engaged in active discussions with NERSC – both with a direct visit and associated discussions with the NERSC leadership at LBNL and also in a major NERSC Resource Requirements Workshop held in Washington, DC on August 3-4, 2010.

The Fusion Simulation Program (FSP) mission is to provide predictive capability for the behavior of magnetic confinement devices via science-based simulations of nonlinear, coupled phenomena on time and space scales required for fusion energy production. This will require multi-scale, multi-physics integration well beyond current capabilities. The mission will be accomplished through improvements and innovation in physics formulation, numerics and algorithms along with the use of increasingly powerful computer architectures. A rigorous verification, validation, and UQ program will be an integral part of the FSP, requiring significant computational resources on its own. Productions services, with a large user base are also planned.

Over a 5 year horizon, notional estimates for FSP HPC needs can be summarized as follows:

- If fully funded in a sustained manner, the FSP is envisioned to rough doubly scale and scope of MFE computation program.
- A rough estimate can be made by extrapolating from related computational programs in MFE (especially the proto-FSPs), leading to the following:

  - Large jobs using in aggregate >1,000 core-hours on 1M cores
  - 10,000s of small runs using 1000s of cores
  - 100s of medium scale runs using 10,000s of cores
  - Memory requirements from 0.1 GB/core for largest jobs to 2 GB/core for small and medium runs

The size and scope of the FSP will drive significant computational and storage requirements. Specifically:

- More CPU hours (In addition to normal growth, the FSP will roughly double the size of MFE computing with a focus on some of the largest, most demanding computational problems.)
- Fast turn-around for smaller jobs, especially in support of code development, verification and validation.
- Support for production computing including a "Simulation as a Service" model. Requires some level of federated authentication and authorization.
- Integrated data management, long term storage and advanced cataloging of modeling and experimental data.

Support for off-line analysis of large data sets – on systems “close” to storage to facilitate data access.

With regard to “capability computing” requirements, the resources reside at the Leadership Computing Facilities (LCF’s) at ORNL and ANL. Since access to such resources can only come at present from participation in the INCITE Program, this will be the approach followed by the FSP unless other arrangements are made available by ASCR.

With regard to expected FSP “capacity computing” requirements, estimates were formulated in the aforementioned interactions at LBNL and also at the NERSC Workshop in Washington, DC.

- FSP Planning estimates for capacity computing requirements and growth can be summarized as follows:
- Tuning of Systems for job mix – find most cost-effective platform for each job with flexibility (e.g., priority sometimes needed for small jobs!)
- Special requirements with respect to memory, storage, etc.
- Availability of needed libraries and other supporting software
- Ability to set priorities within FSP domain
- Adequate CPU hours for software development (advanced components & frameworks), for V&V + UQ testing, and for production services

With regard FSP data storage needs, only rough notional estimates – as in the case with computational needs – can be made and indicate:

- Aggregate archival storage is likely to be in multi-PB range in 1000s to 10,000s of files per year
- Temporary storage needed by jobs during runs are also predicted to go into the PB range
- As noted in the previous slide, we are planning to catalog all FSP runs across all platforms regardless of physical location
- UAL (universal access layer) planned for location independent data access

4.7 Interactions with Advance Computing and Mathematics

4.7.1 Overview

The initial FSP science driver planning exercise yielded potentially significant software engineering, AM/CS research & development activities required to meet the corresponding integrated science application simulation goals. The list of research topics falls into several broad categories: 1) innovative time integration techniques, especially for coupled PDEs; 2) new solver techniques, especially for highly parallel or multithreaded hardware; 3) issues related to data management/analysis, including visualization and meshing; 4) framework design, including the software challenges of componentization and coupling on HPC systems, and 5) issues in software verification and uncertainty quantification, The fusion community has had significant experience and notable success in recent years in the SciDAC and Proto-FSP programs, forming mutually beneficial partnerships with the AM/CS community to accelerate and enable progress in advanced fusion simulation. Through these collaborations a significant portion of the AM/CS community has gained a strong familiarity with fusion applications and the underlying mathematical formulations and computational challenges. A similar type of close collaboration is likely to be successful and will be critical to the success of the FSP.

This section describes and outlines the specific AM/CS research activities that were identified as critical to the goals of the FSP execution plan – specifically, the two initial ISA projects, aspects of the initial advanced component work scope, and the framework and integration activities. From an organizational perspective, FSP manages chose not to isolate these activities in any single management construct, but rather to distribute them naturally throughout the FSP organizational structure to ensure maximum integration with the motivating application. As the goal is to define this research as a partnership between the FES and ASCR community, we do not include in this section software engineering challenges that, however complex, do not contain a significant research component

4.7.2 Process for identification of AM/CS research

Using the detailed science driver reports as a starting point, an initial list of candidate AM/CS research topics were identified and categorized by the FSP planning team. In addition, to clarify and identify in more detail required research areas, the FSP management team made specific requests for clarifications from the authors of the science driver documents. These clarifications provided further insight and detail into the open research issues associated with each candidate topic, for example eliminating some candidate topics while elevating the importance of others. The results of this preliminary set of processes are posted on the community website as input to the February 2011 AM/CS and validation FSP Community Workshop.

4.7.3 FSP AM/CS and Validation Workshop

One major focus of the February workshop was to refine and vet the early candidate list of critical AM/CS research topics previously identified. Thus, the FSP selected approximately 30 representatives from the ASCR community – specifically, whose research involves applied mathematics, computer science, or software engineering within the context of HPC science applications. Another forty or so participants from the fusion community were invited in order to create a fertile context for dialog between the application needs and enabling research. It is important to point out that, among the thirty AM/CS participants; about one-third have had significant previous and/or ongoing collaborations with the fusion community as participants in SciDAC and/or Proto-FSP projects. This balance was intentional and aimed to combine accumulated experience with fresh perspectives and new potential partners.

In addition to plenary talks, the workshop was organized as a sequence of matrixed breakout sessions populated to ensure participation from individuals with expertise in each of the selected topical areas. The selected areas are depicted in Figure 12 and closely mirror the high-level topics identified above.

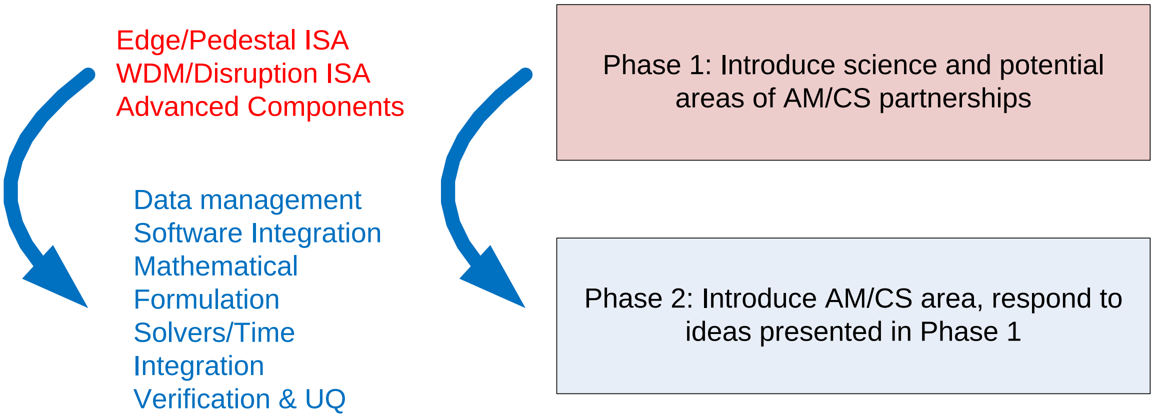

Figure 12: Depiction of application groups (red) and AM/CS research groups (blue) for General Atomics workshop

Leads from the Physics Components and two ISA group initially gave a brief summary of their initial perspective on AM/CS challenges within their work scope (as outlined by the FSP planning process). This was followed by a series of breakouts led by members of the AM/CS community, who processed the initial input and developed ideas/provided feedback to help better match opportunity with existing capabilities, identified where current research trends intersect application needs, etc. This process was then repeated to further refine and tighten the final product.

4.7.4 Overview of working group progress

This section provides a more detailed overview of the activities in each AM/CS breakout group and how they lead to the preliminary set of findings outlined in the subsequent sections.

4.7.4.1 *Frameworks and Integration*

The specific goals for the breakout group were:

- Review requirements in the areas of software integration and support and production computing.
- Identify settled and outstanding user-facing issues.
- Identify settled and outstanding developer-facing issues.
- Settle on near-term plans and needed research.

The majority of the discussions in this group centered on critical complex software engineering activities that however were likely not strong candidates for AM/CS partnerships. There are three possible exception: 1) techniques/standards for design scalable HPC software components, particularly as they relate to the adaption of legacy components; 2) techniques for implementing tight coupling (e.g. JNFK) compatible with software componentization, and 3) the development of an appropriate abstraction/tool for scalable parallel data structures. Each of these concepts is being further developed for potential inclusion in the final version of this report.

4.7.4.2 *Integrated Data Management, Analysis and Visualization*

Discussions focused on a common approach to Off-HPC data management to support a consistent, coherent and unified view of all FSP data with common access methods. The nature and scale of the FSP drives requirements in this area beyond current practice. Workshop participants generally agreed on a set of functional and nonfunctional requirements. Key to meeting the requirements will be a unified simulation catalog that holds all FSP related metadata including data provenance. Its scope will be broad and cover as much FSP activity as practical including standard scientific logbook functions. The catalog will support namespace management via a comprehensive data dictionary and hold pointers to a distributed archive. The bulk data will likely be co-located at computation sites. Common access method for all data would be provided through universal access layer (a common API). Data is accessed via globally unique identifiers, and will not require users to be aware of physical location of data or storage internals. The combination of a common API and location via the metadata catalog effectively federates the distributed archive into single shared FSP data store. Operation of the catalog must be closely coupled to workflows and software integration activities, which are responsible for On-HPC data management. Implementation will be phased in over a number of years with priority going to creating the simulation catalog and the tools needed to populate and query it, including those needed to document simulation workflow. A number of specific needs and challenges for data analysis were identified – particularly those associated with synthetic diagnostics, feature identification and extraction of data characteristics for validation. The consensus was that visualization requirements were currently being met. However, serious issues are present at the exascale for visualization and analysis. More analysis and visualization will have to be carried out concurrent with computation – a stiff challenge to the way science is currently done.

4.7.4.3 *Solvers And Time Integration*

The solvers and time integration breakout group develops recommendations for applied mathematics research activities in the areas of solvers and time discretization that addresses the computational challenges inherent in the scientific goals of the Fusion Simulation Program (FSP). The group makes specific recommendations on the choices of linear and nonlinear solvers, including preconditioners, and the time discretization that can reduce computational expense, enable coupled, multiphysics , multiscale simulation, and/or improve calculation robustness. The breakout group evaluates the numerical challenges identified in the science driver areas and recommends existing approaches or new line of research that could address the challenges. The relative difficulty of finding a solution to each challenge is also estimated.

The group took the input from the pedestal/boundary, whole device modeling (WDM)/disruption, and advanced components breakouts on their prioritization of science and computational challenges, and ranked them into four categories:

A – adequate with current fusion technologies;

N – near-term (1-3 years) challenges which existing applied math technologies should suffice;

M – medium-term (4-6 years) challenges which requires nontrivial improvements to existing technologies;

L – long-term (7-10 years)challenges which requires significant advances in numerics.

In addition to specific findings for each individual challenges identified by the science driver/advanced component breakouts, the solvers/time-integration group identified a set of common themes which are particularly illuminating. On time integration, the group finds that need for implicit, tightly coupled time-stepping is ubiquitous across ISAs and originates in the desire to overcome time-scale disparity. To address this need of FSP, research is needed on alternative time integration (e.g. parallel in time) and approaches to couple multiple physics models through an implicit integration in a generic and adaptable way that promotes both stability and accuracy. The group finds that the need for scalable nonlinear solvers arises as the result of implicit temporal discretization. This becomes a bigger challenge due to the across-the-board emphasis on tight model coupling, both in components and ISAs. For example, the nonlinear solver includes a kinetic closure of thermal plasma and energetic particles for the extended MHD component. The group finds that Jacobian-free Newton-Krylov-type methods offer many advantages, and can be readily adapted into FSP. With a few exceptions, the group finds that the need for scalable linear solvers originates in the context of a nonlinear iteration loop and in the need for preconditioning. While most applications have sparse matrices, several applications (e.g. gyrokinetic continuum approach, RF, and 3D MHD with Fourier representation) require dealing with dense matrices, which presents a challenge at the ultrascale. Opportunities identified for near- and mid-term research and FSP applications by the breakout group include the adoption of hierarchical, multi-level methods as subcomponent multi-level solves for physics-based preconditioners, exploiting data locality, reducing memory traffic and reducing communication, and exploiting hybrid and less synchronous programming models. It is of interest to note that many fundamental research challenges for scalable solvers are identified with specificity in the context of the complex multiphysics scenarios of FSP.

4.7.4.4 *Mathematical Formulation*

The charge to the mathematical formulation (MF) group was to identify research topics within the Integrated Science Applications that involve innovative approaches to the formulation of the key governing equations and/or their discrete representation. This was intended to complement the solvers group, whose focus was centered on techniques to solve systems of simultaneous equations assuming a given mathematical formulation. The MF group also chose to extend their discussions to algorithmic research related to enhanced code performance on innovative HPC architectures/programming models currently coming online (and likely on the path to exascale). A significant portion of the discussion in the MF breakout focused on issues related to extended MHD codes – both kinetic closure and consistent hybrid codes with the goal of self-consistent evolution of energetic particles including integrated effects from MHD, RF, and micro-turbulence. There was also significant discussion of the development of techniques to resolve the range of timescales from instability to confinement, as well as innovative spatial discretizations for large spatial anisotropies. Another specific topic of frequent discussion was the development of extended MHD codes with fixed kinetic closures – both the development of the closure models, issues in discretization (finite elements vs. finite volume) and Eulerian vs. particle approaches. Finally, there was considerable discussion on strategies to significantly increases the efficiency of fluid codes, either via improved scalability, improved use of new architectural features, or innovative meshing strategies (e.g. some form of AMR). For PIC codes the discussion focused on mathematical approaches to sampling noise reduction, as well as number of issues related more to improved performance using GPGUs and other new architecture features (PGAS, etc.).

4.7.4.5 *Verification and Uncertainty Quantification (V&UQ)*

The charge given to the V&UQ team was to assess the proposed FSP V&UQ strategy in order to identify areas of weakness, to suggest improvements, to identify alternative methodologies, and to identify areas that will require substantial advancement in V&UQ techniques. The goal was to define a realistic strategy for implementing a meaningful, integrated V&UQ effort within the FSP that increases in rigor and sophistication over the life of the program. The breakout group was attended by a cross-section of computational and experimental physicists, computer scientists, numerical analysts, and UQ experts from physics, applied mathematics, engineering, and statistics.

A reoccurring topic in the discussions was the importance of documentation – of the models, the codes, the parameters, the sources of uncertainties, the assumptions used in analysis, etc. – for the success of any V&UQ program. A systematic approach to documentation and archiving of the potentially vast quantities of data generated in V&UQ studies is key. A dynamic database for archival was proposed.

There was consensus as to the importance of more rigorous code verification practices for the FSP. Interest was expressed in more use of the Method of Manufactured Solutions (MMS), especially if tools to simplify and automate its application can be made available. Caution was expressed, however, that MMS is difficult to apply in the complicated geometries common to tokomak simulation.

On the topic of calculation verification, i.e., quantifiable, *a posteriori* numerical error estimation in calculations with unknown solutions, it was generally agreed that it is a necessary component of UQ and, while many techniques exist, the application of these methods to coupled, multi-physics codes is still an open research topic. In particular, an appropriate accounting of coupling error will require the development of new techniques, and existing techniques will either be difficult to incorporate into existing codes (e.g., error evolution) or will be difficult to formulate for complicated operators (e.g., adjoints). New component development within the FSP should attempt to include some form of *a posteriori* numerical error estimation. Most of the discussion in the two breakouts focused on UQ. In particular, a dialogue between UQ practitioners and fusion energy scientists was begun. For the former, their inquiries centered around: what is UQ, how is UQ analysis done, what can be obtained from UQ analysis. For the latter, the interest was in establishing the scope of the problem in fusion energy simulation, identifying the characteristics of the problem (e.g., the number of parameters, the nature of the solutions), and identifying the key quantities of interest (QoIs). It was determined that systematic UQ has not really been applied in fusion energy simulation, but this subject area has the advantage of being data-rich. The application of UQ analysis techniques to a problem as difficult as the fusion simulation problem has not been done and will require incremental and iterative steps involving UQ experts, fusion scientists, computer scientists, and numerical analysts. It was recommended that a concrete problem be defined in the context of the two ISAs, and examples of investigation of hierarchies of components of increasing complexity from both the WDM/Disruptions and Edge/Pedestal areas were developed. The subject of disruptions was identified as the critical target moving forward, in the context of ITER, and this should be the initial focus of UQ investigations; as such work must be done to define the size of the parameter space for this problem. Uncertainties in input parameters must be nailed down by the fusion energy community. All existing approaches to UQ were suggested to be used, since there are no clear methods that will succeed, and different approaches can provide complementary information. Such an approach will necessitate an active involvement and dialogue between UQ experts and the fusion simulation community. New technologies will likely need to be developed, including methods of dimensional reduction and techniques to combine uncertainties and numerical errors. Over all, it was agreed an ongoing dialogue must continue to educate the fusion community about UQ and to educate UQ analysts about fusion energy simulation challenges and needs.

Concern was expressed over available resources for the V&UQ efforts. Specifically, based on V&UQ efforts on “simpler” multi-physics problems like climate, it is expected that FSP will require at least as much manpower effort, if not more (nominally 8 FTE). In addition, massive amounts of computer time will be required, and it is unclear how this will be acquired.

4.7.5 Detailed formulation of AM/CS needs

Based on the entire process described above, culminated with the General Atomics workshop, the FSP management team has begun to refine and finalize a final list of embedded AM/CS subprojects on the critical path to the goals of the FSP. This process is not complete, but a strong flavor of our progress to date, particularly in the critical area of scalable solvers and time integration, is summarized in the following sections.

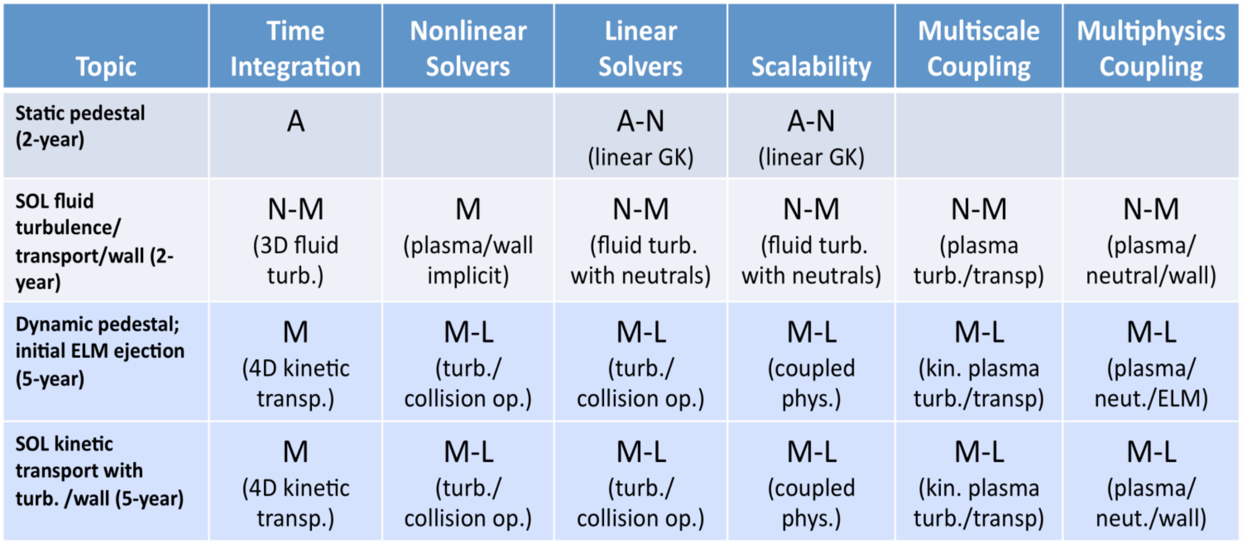

Topic	Time Integration	Nonlinear Solvers	Linear Solvers	Scalability	Multiscale Coupling	Multiphysics Coupling
Static pedestal (2-year)	A		A-N (linear GK)	A-N (linear GK)		
SOL fluid turbulence/ transport/wall (2-year)	N-M (3D fluid turb.)	M (plasma/wall implicit)	N-M (fluid turb. with neutrals)	N-M (fluid turb. with neutrals)	N-M (plasma turb./transp)	N-M (plasma/ neutral/wall)
Dynamic pedestal; initial ELM ejection (5-year)	M (4D kinetic transp.)	M-L (turb./ collision op.)	M-L (turb./ collision op.)	M-L (coupled phys.)	M-L (kin. plasma turb./transp)	M-L (plasma/ neut./ELM)
SOL kinetic transport with turb. /wall (5-year)	M (4D kinetic transp.)	M-L (turb./ collision op.)	M-L (turb./ collision op.)	M-L (coupled phys.)	M-L (kin. plasma turb./transp)	M-L (plasma/ neut./wall)

Topic	Time Integration	Nonlinear Solvers	Linear Solvers	Scalability	Multiscale Coupling	Multiphysics Coupling
XMHD	A-N	N-M	A-N	N-M	M-L	L
MHD-K	N-M	M-L	A-N	N-M	M-L	A-N
GK	A-N	M-L	A-N	A-N	M-L	L
3D MHD EQ		A-N	N-M	N-M	M-L	L
RF		N-M	A-N	A-N	M-L	L
EDGE	N-M	N-M	N-M	N-M	M-L	L
NEUTRAL TRANSPORT	A-N	N-M	N-M	A-N	M-L	L

Topic	Time Integration	Nonlinear Solvers	Linear Solvers	Scalability	Multiscale Coupling	Multiphysics Coupling
WDM – Soln. Methods	N-M-L	N-M-L	A	N	M-L	M-L
WDM – Stiff Turbulent Transport		N-M			N-M	N-M
WDM – Coupled XMHD / Kinetic Hybrid	N-M	M-L	N-M	N-M	M-L	M
Equilibrium Soln. 1.5D & 2.5D		N	N	N-M		
Disruptions – Stability/ Bifurcation with Uncertainty		M	M	N		
Disruptions – Initial value XMHD solvers	A-N	A-N	A-N	N-M	M-L	M-L
Disruptions (non-ideal) - XMHD Eigensystem analysis		M-L	M-L	M-L	M-L	M-L

Figure 13: Assessment of the specific near, medium, and long-term applied math research challenges for the Boundary/Edge ISA, Whole Device Model ISA, and component development milestones.*

** The abbreviations used in the chart are:*

A: Adequate with current fusion technologies
N: Near-term challenge: 1-3 years, existing applied math technologies should suffice
M: Medium-term challenge: 4-6 years, nontrivial improvements to existing technologies
L: Long-term challenge: 7+ years: significant advances in numerics are required

4.7.6 Challenges

- Issue 1: Provide coupled long-time fluid plasma SOL turbulence/profile evolution by either:
  - efficient direct long-time turbulence model evolving its own profile; better preconditioners needed
  - or, periodic coupling of transport/turbulence codes; large-scale, intermittent fluctuation technique

- Issue 2: Coupling plasma efficiently to SOL neutral (volumetric fluid / then kinetic); need scalable preconditioning, Monte Carlo implicitness and noise reduction

- Issue 3: Implicit wall response time integration, coupling to SOL neutrals/plasma (surfacial); little experience, method and impact of coupling on performance of other SOL components needed

4.7.6.1 *General Questions:*

- Concise math description?
  Yes, though some ad hoc physical flux-limiting of plasma/neutral fluxes
- Challenges well understood?
  Significant previous experience except for wall response and coupling; some experience on turbulence/transport coupling; Monte Carlo implicitness still a challenge
- Technology / research solution time horizon?
  Apply existing technology over 2 years to fluid models; as higher dimensionality (fluid -> kinetic) and more components (e.g., wall complexities) extends to 5-year/beyond
- Possibility for concrete math and science partnership?
  very high

4.7.6.2 *Issues for Linear/Nonlinear Solvers/Timestepping:*

- Common linear and nonlinear solvers used?
  preconditioned Newton-Krylov methods (40%-70% total time)
- Obvious new nonlinear / linear solvers / preconditioners to apply?:
  Developing preconditioners for turbulence
- Time integration issues?
  Desire larger timesteps

4.8 Major Deliverables and Milestones

The following table contains a high level estimate of Level 1 deliverables and milestones relative to the start of the Program. Progress toward Level 1 milestones is formally reportable to the DOE Program Office(s). Level 2 and Level 3 milestone are of those tracked by to the FSP Directorate, Executive Committee, Team Leads and other FSP staff and stakeholders. A complete Level 1 through 3 milestone table is listed in Appendix C:. The expected dates for many of the items listed below and in the appendices should be understood as the time when they become viable entities which may require ongoing refinement and extension rather than always being finished products or completed activities. The definitions for the bulleted items in the Applied Math and Computer Science Research column can be found in Section 1.2 above.

Deliverable / Milestone	Applied Math / Computer Science Research	Expected Date from Award
1.4: FSP operational – management and technical teams sufficiently staffed, funding mechanisms in place, FSP policies and procedures in effect.		6 months
5.1.2.4: Common I/O capability with consistent metadata available to fusion community		21 months
5.4.3: Availability of all FSP software on line and continuously updated		21 months
3.2.8: First release of FSP 1.5 WDM code	• Scalable Solvers • Time Integration	22 months
4.3.1-4: Release of library of adapted components including Grad-Shafranov solver and embedded turbulence model	• Multi-scale/physics • Frameworks	23 months
2.2.4: First release of static model within FSP	• Scalable Solvers	32 months

framework		
5.1.5: Reference implementation of On-HPC integration software, Release 1 (concurrent components, low-dimensional couplings)	• Data/Meshing • Frameworks	32 months
3.5.5 Release code with gyro-kinetic turbulent simulations included	• Scalable Solvers • Time Integration • Multi-scale/physics • VUQ	41 months
2.3.3: First release of coupled SOL model	• Scalable Solvers • Time Integration • Multi-scale/physics • Data/Meshing • VUQ	53 months
4.4.1: Complete new component development of profile evolution with 3D equilibrium	• Scalable Solvers • Time Integration • Formulation • Multi-scale/physics • Data/Meshing	59 months
2.5.3: First release of coupled kinetic SOL model	• Scalable Solvers • Time Integration • Formulation • Multi-scale/physics • Data/Meshing • VUQ	68 months
3.4.8: Release WDM code with 3D core and pedestal models	• Scalable Solvers • Time Integration • Multi-scale/physics • Data/Meshing • VUQ	73 months
3.6.5: Release code with combined ISA1 and ISA2 components	• Multi-scale/physics • Data/Meshing	83 months
2.6.2: First release of dynamic pedestal model	• Scalable Solvers • Time Integration • Formulation • Multi-scale/physics • Data/Meshing • VUQ	88 months

Table 11: FSP Level 1 Deliverables and Milestones (Quick Reference)

4.8.1 Schedules and Resource Requirements

A resourced WBS and detailed schedule is in the process of being developed. Effort resource estimates are provided throughout various sections above.

4.9 Program Scope

The eventual description of scope will be more detailed and will list things that are FSP will undertake and, importantly, things that will be explicitly excluded.

The FSP Program's objectives will be achieved through activity in the following major areas:

- R&D to enable scientific discovery of important new plasma phenomena with associated understanding that emerges only upon integration.
- Establishing a productive partnership with ASCR which will provide sustained access to state-of-the-art AM/CS tools and resources, as well as substantial access to the most powerful HPC resources available to the SC community.
- Production and vetting with the FES user community of a suite of software tools requiring development of a predictive integrated simulation capability for magnetically-confined fusion plasmas that are properly validated against experiments in regimes relevant for producing practical fusion energy.
- While complementary R&D will be planned and executed, the following major areas are not within the scope of the funded FSP Program:
- Funded R&D activities in the FES Base Programs for Experiments, Theory, & Modeling
- Funded R&D & LCF activities in the ASCR Base Program

4.10 Performance Metrics

The successful operation of the program, once basic deliverables have been accomplished, will be measured by the following deliverables: (1) delivery of new code; (2) quality of code ensured via QA standards; (3) impact on delivery of new science by users measured by survey & feedback; (4) impact of associated operations/user support measured by user survey; and (5) impact of Education/Outreach/Training (EOT) measured, e.g., by active involvement of graduate students and post-docs in the FSP.

4.11 Program Assumptions and Constraints

The current FSP plan is delivered to DoE after a 2-year dedicated effort. If approved by DoE, subsequent success in the operational/execution phase of the FSP will be constrained by budget and associated federal funding decisions. At present, this plan represents an ongoing process that provides a credible assessment of the resources necessary to address the mission need as well as the timeline for delivering the required capabilities. However, an estimate of the resources and schedule for meeting this need can be guided by the experience of other communities which have developed simulation capabilities of comparable complexity and scope. The Accelerated Strategic Computing Initiative (ASCI) Level 1 University Centers Program by NNSA funded at $25M / year to develop the predictive capability required to certify the U.S. nuclear stockpile in the absence of testing provides a good example. The estimate in the current FSP plan considers the total funding from both program offices – FES and ASCR – for this joint interdisciplinary effort. Regarding schedule, the general consensus opinion (from many applications domains) is that it takes approximately 10-15 years to develop a validated simulation capability of the scale required to address the FSP mission need, including a 4-5 year “ramp-up” period during which the funding is increased to its full level as the staffing of the project team is being built up.

In order to provide an initial schedule, the FSP Execution Plan made the following assumptions:

- That the process for the eventual awarding of a DOE contract to execute the FSP will include decisions on which institution will host the FSP and who the Director and Deputy Directors will be. Therefore, the Directorate will be in place at Program start.
- That the funding for the Directorate will be in place at Program start and that management and oversight mechanisms for the entire FSP effort will be in place and fully functional within two months of Program start.
- That the initial ISA teams will be selected by some, as yet unspecified, process that will identify partner institutions and the ISA technical team lead within three months of start and that the ISA teams will be fully staffed within six months of start. That the ET teams (lead and other staff) will be selected by the FSP Directorate and be fully in place within 3 months of Program start.
- That the Project Management Office will be staffed and fully functional within one month of Program start.

- That the effort levels start at $12M for the initial year and increase by $3M per year until $24M when the FSP is maximally configured and operational.

The most important assumption in the planning process is that DOE will delegate sufficient authority to the FSP such that its directors can manage the Program in an effective manner. That requires establishing a process whereby FSP can control or otherwise effectively influence the redirection of program resources when changes to the plan are required either to compensate for emerging deficiencies or to take advantage of opportunities.

5 MANAGEMENT ORGANIZATIONS AND RESPONSIBILITIES

The FSP management and communications hierarchy will be comprised of personnel from DOE-SC, FES, ASCR, a lead/home institution, the FSP Management Team, and various participating organizations sponsoring members of the Program Implementation Team (PIT), i.e., the people who will be directly involved in day-to-day FSP efforts.

Day-to-day management of the FSP is the responsibility of the FSP Management Team which is a triumvirate (a Director and two Deputy Directors) drawn from the consortium of universities, labs, and the private sector that will be chosen to execute the FSP. FSP Management Team oversight is provided by an Integrated Program Team (IPT – described below) and with additional advice and guidance provided by an external Program Advisory Committee (PAC).

Communications and reporting requirements among the various stakeholders mentioned below and with others will be described in detail in the FSP Communications Plan.

5.1 Department of Energy

Ultimate authority and responsibility for managing Department of Energy programs and facilities resides with the Secretary of Energy. The DOE Office of Science has been delegated responsibility for R&D in the FES (theory, modeling, and experiment) and the ASCR (HPC capabilities, applied math, and computer science) crucial to achieving goals described in the Department's Strategic Plan. The Office of Science provides overall program policy and guidance, technical oversight, and budgets for implementing its assigned role. Specific responsibility for design, implementation, and operation of the FSP is assigned to DOE SC's Office of Fusion Energy Science (SC 24, FES) and the Office of Advanced Scientific Computing Research (SC 21, ASCR).

The mechanisms by which FES and ASCR will provide funding to the FSP have not yet been decided. It is assumed that the DoE expectations for the FSP will be articulated in a Memo of Understanding (MoU)-type statement from FES and ASCR.

5.2 DOE FSP Program Managers

The DOE Program Managers' responsibilities are shared between FES and ASCR with FES having the lead role. The Program Managers' roles and responsibilities are summarized as follows:

- Define program mission requirements and objectives;
- Function as DOE HQ points of contact for program matters;
- Oversee program progress and help organize reviews as necessary;
- Coordinate with other DOE HQ organizations as needed to execute the program;
- Budget for funds to execute the program; and
- Control changes to program baselines in accordance with this execution plan

5.3 Integrated Program Team

The purpose of the FSP Integrated Program Team (IPT) is to provide strategic planning, coordination, and communication for the FSP in order to ensure the program's objectives are achieved on schedule, within budget, and consistent with quality, environment, safety, and health standards. The IPT meets on a bi-weekly basis and is the primary multi-way communication channel for the FSP management to communicate with DOE to discuss progress and current issues of concern. The work of the IPT includes:

- Ensuring that program management is carried out with integrity and in compliance with applicable laws.
- Supporting the FSP Director in the performance of project management responsibilities;
- Developing an appropriate program contracting strategy;
- Assuring all program interfaces are identified, completely described/defined, and managed to completion;
- Identifying appropriate and adequate program performance metrics and ensuring that they are met;
- Support independent periodic reviews and assessments of program performance and status against established performance parameters, baselines, milestones, and deliverables, taking corrective actions as appropriate;
- Planning and participating in, as necessary, ad hoc program reviews, audits, and appraisals;
- Supporting development of all federal Critical Decision packages or their equivalents;
- Reviewing and commenting on program deliverables,
- Reviewing baseline change requests (at the appropriate levels) and supporting change control boards as requested;
- Supporting the preparation, review, and approval of completion and closeout documentation of specific FSP elements; and
- Delivering a quality, cost-effective program.

The IPT is comprised of a core membership plus any additional subject matter experts as required on an ad hoc basis. The core IPT membership includes:

- DOE FSP Program Managers (for FES & for ASCR)
- FSP Director
- FSP Deputy Directors (for Science and for Code Architecture)
- FSP Head of Program Management Office
- Representative of lead institution management

5.4 Program Advisory Committee

The PAC is an external group of experts, reporting to the Director of the FSP home institution and providing advice on a broad range of technical and managerial issues. It would meet approximately once per year and address a charge formulated by the Director of the lead institution and the FSP Director.

5.5 FSP Lead/Home Institution

The lead/home FSP institution will be the ultimate responsible entity to DOE for FSP performance.

FSP Partner Organizations

The FSP partner organizations are those that will participate in the execution phase of the FSP after a selection approach that is yet to be determined. In the 2-year planning phase, the FSP home institution and partner organization resulted from a peer-review process. This multi-institutional, multi-disciplinary team was comprised of 6 national labs (PPPL, ORNL, LANL, LBNL, LLNL, ANL), 2 companies (General Atomics, Tech-X), and 9 universities (MIT, Princeton, Columbia, NYU, UCSD, Chicago, Lehigh, Purdue, Texas)

5.6 Program Implementation Team (PIT)

Program Implementation Team are the people who will be directly involved in day-to-day FSP efforts.

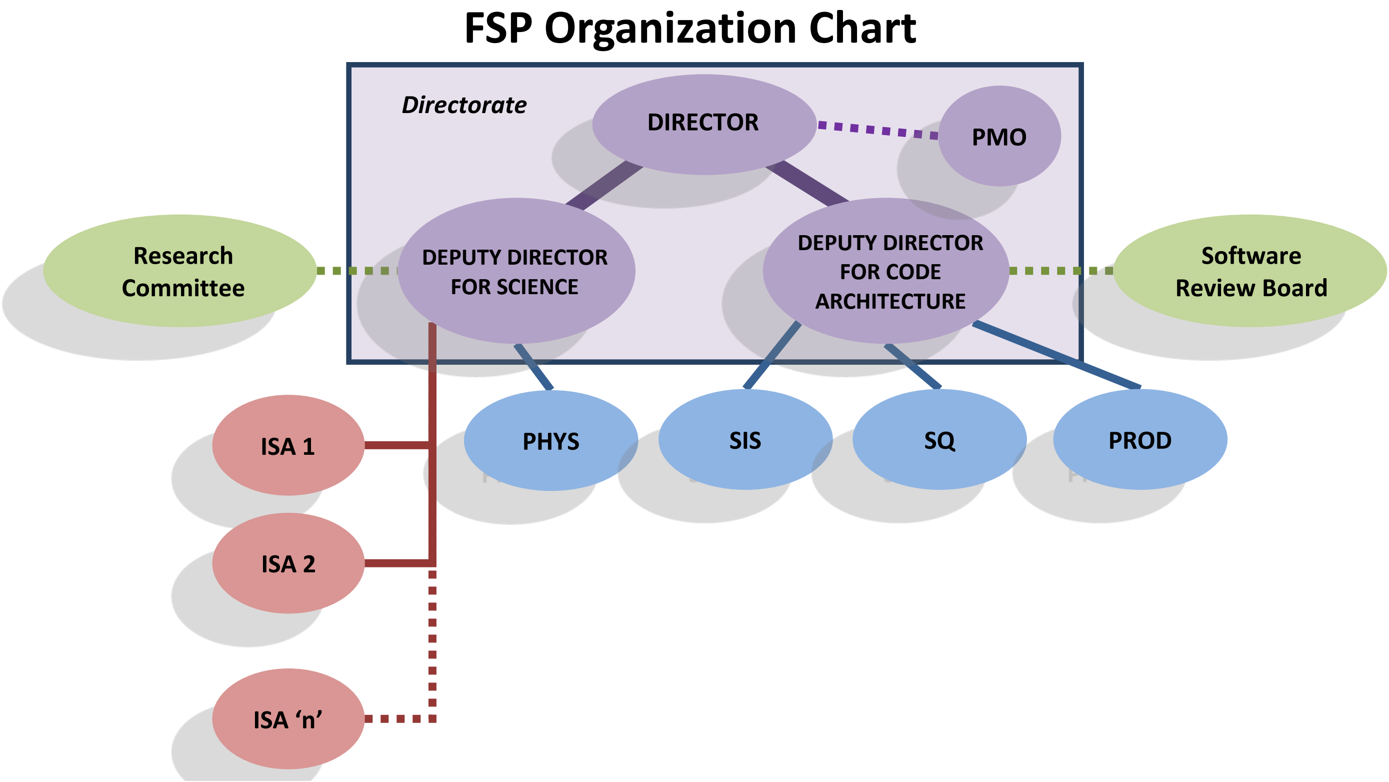

Figure 14: FSP Organization Chart

5.6.1 FSP Director

The FSP Director has responsibility for ensuring overall scientific and software development goals of FSP are properly executed. He/she is the principal contact with DoE and with the management of the lead institution and implements the FSP through the M&O contractor for the lead/home institution, which is responsible for overall program coordination, execution, and facility operation.

The FSP Director serves as the principal contact with DoE and directly oversees the FSP Directorate, which includes the Deputy Director for Science, Deputy Director for Code Architecture, and the Head of the Program Management office. The Deputy Directors for Science & Code Architecture act as a top level communication channel to ensure that the cross-cutting functional groups work together in a seamless way. The Director is ultimately responsible for ensuring that the overall scientific and software development goals of the FSP are being successfully executed. He/she must make final decisions on project prioritization, funding allocations, and personnel using a well-defined process based on input from the core management team, the larger FSP leadership group, and external peer review.

Key decisions on are made by the FSP Director FSP plans and priorities are guided by input from an external Program Advisory Committee (PAC) and by internal Research Committee and Software Review Board; i.e.,

5.6.2 FSP Deputy Director for Science

The FSP Deputy Director for Science reports to the FSP Director and oversees the Integrated Science Applications (ISA's) and Physics Components development teams. This position drives the scientific goals of the project, ensuring that the application projects are well balanced and making adequate progress. He/she will suggest funding allocations (by area) to the Director based on the need to balance short/long-term progress, address the priorities of the Program office, and respond to community feedback. The Deputy Director of Science also serves as chair of the FSP Research Committee that will advise the Director on questions of planning and priorities.

5.6.3 FSP Deputy Director for Code Architecture

The Deputy Director for Code Architecture reports to the Director and directly oversees the Software Integration, Quality Assurance, and Operational/ User Support groups. The Deputy Director for Code Architecture is ultimately responsible for the management of the overall FSP code repository as integrated software and drives both the AM/CS research and applied project software goals of the project. This position also ensures primarily that an integrated “community suite of tools” flows from the science driver projects and lives within a proper software development lifecycle, including documentation, testing, versioning, and repository management.

5.6.4 FSP Head of Program Management Office

The Head of the Program Management office reports to the FSP Director and directly oversees a small group of assistants in an office of Project Management. He/she establishes the standards, policies, & procedures to be followed for FSP project management and coordinates project tracking & reporting for the Directorate. This includes responsibility for formally representing and tracking the work of the FSP as a set of milestones and deliverables with clear deadlines, interdependencies, and levels of effort. It will work with the Deputy Directors and FSP management team leads to help cast scientific goals in terms of a concrete plan, including a mechanism for accommodating appropriate levels of flexibility and accommodating longer-term research where appropriate.

5.6.5 FSP Technical Leads

Each of the six major operational divisions of the FSP, i.e., ISAs, Physics Components, SIS, QA, and Production Computing, has a technical lead with the responsibility to manage the effort within the division and to coordinate its effort with the other divisions.

5.6.5.1 *Integrated Science Applications (ISA)*

The Head of Integrated Science Applications reports to the Deputy Director for Science. He/she has overall responsibility for coordinating the ISA’s, including experimental validation and application of Validation and Uncertainty Quantification principles for these activities. This requires efficient coordination with the heads for Physics Components, Software Integration, and Quality Assurance (QA) to ensure that each of the science goals is accomplished following a “living roadmap” for each ISA with appropriate standards met. Each Application Area Lead of the ISA’s (Core Profiles, etc.) reports to the Head of Integrated Science Applications and oversees a small (approx. 7-10 person) integrated team tasked to execute the WBS (work breakdown structure) for carrying out research and developing code needed for each of the targeted FSP science drivers. This includes developing end-to-end capabilities – from scientific definition and formulation through problem verification, uncertainty quantification, and validation to software release. The ISA teams will include members matrixed from the Physics Component and Software Integration groups to ensure that a common set of components and enabling computational tools are being developed across the FSP application projects. The ISA Area lead will also initiate the release of associated new software capabilities to the Community in coordination with other groups

5.6.5.1.1 ISA Leadership Criteria

There are two key aspects of the ISAs, both critical to the success of the FSP – scientific discovery and software development. Both the execution/analysis of innovative simulation and the development of a usable “community code” are necessary outcomes for the success of the FSP. Thus, selection of ISA leads is critical to the success of the FSP. Ideally, each ISA lead is a domain scientist with a strong background in some aspect of numerical simulation and software design. This is often a difficult set of constraints to fulfill. Typically, the technical lead will want to appoint a Chief Architect responsible for overseeing the ISA as a mature code development project, independent of its scientific goals and achievements.

5.6.5.2 *Advanced Physics Components (PHYS)*

The Head of Physics Components reports to the Deputy Director for Science and has overall responsibility for coordinating the development of re-usable physics components in coordination with the Integrated Science Application teams and in significant collaboration with the Deputy Director for Code Architecture. Each physics component area of need has been identified via a systematic “gaps analysis” of the targeted FSP Science Drivers. The leaders of each of such areas reports to the Head of Physics Components and ensures that new products for a given ISA are being developed from and built into a common code base. Their role is to ensure that the group members working within an application area are developing from a common code base, that code improvements targeting a given application driver are being built into that code base, and that the methods embedded in the physics components are verified (vs. problem-specific verification, which is the domain of the ISA teams). The physics component teams are not simply thought of as providing services for the ISA teams. Since it is unnatural for a simulation team to wait on members of an external group to supply their physics capabilities, these component developers are actually part of the ISA team and largely share the same goal. However, it is important to point out that they additionally bring an eye of generality that will not otherwise drive the ISA leads. Under the guidance of the Head of Physics Components and under the overall purview of the Deputy Director for Code Architecture, the physics component groups can also work on long-term research that does not directly affect the near-term ISA goals – but which can potentially benefit a wide range of future scenarios.

5.6.5.2.1 PHYS Leadership Criteria

Leading the advanced physics component requires domain expertise in both fusion sciences and advanced scientific computing. Experience in managing teams of scientific software developers and interacting with theory community on model formulation is also desired. With the matrix management structure of FSP, it is essential that the component team leader has the vision on working with ISAs and other ETs to carry out physics components development for the entire FSP program with uniform standard.

5.6.5.3 *Software Integration and Support (SIS)*

The Head of Software Integration reports to the Deputy Director for Code Architecture and has overall responsibility for coordinating the integration of FSP physics components – both “legacy” and new codes in coordination with the Integrated Science Application teams and in significant collaboration with the Deputy Director for Science. He/she oversees/coordinates the teams tasked with ensuring that the overall vision and end goals of cohesion, testing, and release are met in a timely manner. This includes responsibility for focusing crosscutting teams on producing the physics integration tools needed by the ISA teams, – while ensuring minimal duplication. These teams include: (i) Physics Coupling and Integration: (ii) Task Composition/Workflow: (iii) Integrated Data Management; (iv) FSP developers support; and (v) Enabling Computational Technologies & Tools (Legacy & New). In strong analogy with the Physics Components group, the role of the Software Integration group is to ensure minimal duplication of enabling computational technologies across the ISA’s. The tasks for the associated team leaders (reporting to the Head of the Software Integration) would include any and all aspects that make the physics components conceptually related with respect to similarity of: (i) approach to documentation; (ii) approach to deployment; (iii) ability to read/write same data formats; and (iv) use of the same mesh data types and visualization approaches. Inter-component coupling is also an option but not necessary. The software integration head also manages version control and coordinates release of FSP codes initiated by Head, Integrated Science Applications (ISA’s).

5.6.5.3.1 SIS Leadership Criteria

Leadership of the Software Integration and Support team will require the usual skills of technical management. However, there will be a much greater need for human interaction and negotiation skills as the SIS team lead will be responsible for communicating with the other team leads to ensure that the SIS takes on development of software that is needed by the ISA teams and the SQ team, working with the Production Computing team on deployment and defect management and correction, and working with the component team on component integration. Thus, it will be important to select someone with management experience in a situation of inter-team dependency (as opposed to having only single line management skills).

5.6.5.4 *Software Quality (SQ)*

The Head of Software Quality reports to the Deputy Director for Code Architecture and has overall responsibility for coordinating the activities of the teams ensuring the reliability of the FSP frameworks and components targeted for release to the community – including standards for verification, validation, and uncertainty quantification (VVUQ). Before release of the FSP software tools developed, the SQ team will also oversee the internal ("alpha" testing) as well as external "beta" testing of these products. While implementation of the activities noted would be the responsibilities of all members of the FSP team, the SQ team would include dedicated technical staff providing crosscutting tools and technologies (e.g., testing systems, etc.). The Head of Software Quality also chairs a Software Quality Board, composed of designated members from other areas. This group would develop standards for software, development and testing; review plans and progress on software quality activities across the entire FSP program. The SQ head would organize software reviews, prior to the release FSP products.

5.6.5.4.1 SQ Leadership Criteria

The responsibilities of the Software Quality Team require a diverse team comprised of software engineers, applied mathematicians, statisticians, uncertainty quantification experts, and physicists. Thus, the leader of this team will require a broad technical background that enables him or her to understand enough of the issues and challenges faced by the team to ensure timely completion of tasks and milestones. Ideally, the SQ Team Leader will have experience in managing multidisciplinary teams and will understand the languages and cultures of the disciplines engaged in the SQ activities. In addition, if the SQ Team grows sufficiently, the SQ Team Leader will likely appoint one person as the leader of each of software testing, verification, validation, and uncertainty quantification subareas.

5.6.5.5 *Production Computing (PROD)*

The Head of Production Computing reports to the Deputy Director for Code Architecture and has overall responsibility for coordinating the group comprised of teams responsible for organizing and coordinating the release of FSP software and an appropriate subsequent level of "customers" support. Overall, this includes: (i) FSP User Support – including Job Monitoring, Bug Tracking, and General Troubleshooting/Triage; (ii) Documentation; and (iii) Creation and Maintenance of FSP web sites and any other tools required to support users both internal and external to the FSP.

5.6.5.5.1 Production Computing Leadership Criteria

Production computing team leadership should reside with an experienced technical person familiar with the exigencies of research driven production computing and comfortable working with a research user community in a role involving significant elements of engineering support as well as collaboration with the FSP code development groups for physics components and the various enabling technologies.

5.6.6 Research Committee

This committee is composed of FSP leadership (Technical Leads along with the FSP Directorate) and includes representatives of major collaborating groups. It is chaired by the Deputy Director for Science to help advise the Director on a broad range of research planning activities including assessment of priorities for R&D, preparing work proposals and organizing publications and presentations. This internal committee will discuss relevant issues and make recommendations to the FSP director – with findings and recommendations that are well documented.

5.6.7 Software Review Board

This board is chaired by the Software Quality Lead and includes designated members from other areas. The function of this board is to: (1) provide standards for FSP software development and testing; (2) review plans and progress on software quality activities across the entire FSP program; and (3) facilitate software reviews prior to release.

Specific responsibilities of the FSP Director with regard to the overall successful execution of the FSP, include:

- Executive level management of the design, acquisition, and transition to operations of the FSP to ensure all program requirements are fulfilled in a safe and cost-efficient manner;
- Financial authority and accountability as delegated by DoE and the lead institution to develop budgets and control FSP work within approved technical, cost, and schedule baselines, and control changes to approved baselines in accordance with established configuration management procedures;
- Management and direction of procurements within the authority delegated by DoE and the lead institution, including the authority to request the execution and delivery of contracts and agreements, and also of purchase orders, assignments, and instruments and documents of any kind relating to the acquisition, sale, or disposition of products, services, materials, supplies, and equipment relating to and necessary for the proper execution of the FSP; and
- Overall responsibility to hire and manage the human resources necessary to execute the FSP and ensure an effective transition to operations, including the overall responsibility for managing the human resources systems within the authority delegated by DoE.

6 MANAGEMENT PROCESSES

The FSP will be managed as a mixed life-cycle effort whereby some program activities will be managed using project management techniques and some managed as ongoing operations. It is assumed that the Program would not be required to follow DOE O 413.3b or similar. Nonetheless, Program activities will follow program and project best practices. Project activities will be defined, documented, and tracked in a commercial project management software application as a resource-loaded, cost-estimated, and scheduled work breakdown structure. To simplify the cost accounting, progress-tracking , and reporting for the entire program, the operational components of the program will be registered in the same project management system file along with budgets and schedules as appropriate; however, the operational components will be flagged as non-project items so they do not skew any earned value or other important project control mechanisms that may be utilized.

6.1 Program Baselines

Program cost and schedule baselines are represented by a WBS and integrated schedule which, for practical publication reasons, is reported in spreadsheet form as auxiliary document “FSP WBS Cost Worksheet.xlsxm”. The spreadsheet was also used for gross calculations of cost and cost profiling estimates. The spreadsheet pages are:

- WBS – an extract from the MS Project application showing Task Name (activities and roll-ups); start, duration, and finish; activity dependencies (Predecessor links); WBS dictionary (Notes column); effort rate or effort amount in FTEs; and some calculated fields showing time from project start in months, quarters, and years.
- FSP – the estimated cost profiles by months, quarters, and fiscal years.
- Milestones – Level 1, 2, and 3 milestones (important checkpoint or delivery dates). Level 1 Milestones are those that are of direct interest to (and will be reported to) the DOE program office(s); Level 2 to the FSP Directorate and Research Committees; Level 3 to team leaders and team members.

- FTE Allot (7 worksheets) – estimated cost profiles for each major WBS level: MGMT, ISA-1, ISA-2, PHYS, SIS, SQ, and PROD.

The scope baseline is to be inferred from the various descriptions in this plan and the WBS detail.

Once the Program begins, the WBS will be reevaluated, fully resourced, and logically scheduled in order to provide FSP management with a current, effective work plan and cost estimate.

6.1.1 Work Breakdown Structure

The FSP is officially a program that will be influenced by Office of Science management guidelines such as DOE Order 413.3A. per instructions from the FES, this not mandatory for the FSP but a model to be tailored in practical application. The program will also follow, as appropriate: (i) the processes specified by the Project Management Institute; and (ii) the incorporation of the guidelines for Earned Value Management as described in ANSI/EIA-748-A-1998. Thus the FSP will reflect the “best practices” of project management. The program’s performance baseline will be determined using an integrated analysis of logic-driven, resourced-loaded activities that follow its work breakdown structure (WBS). The WBS, with budget and schedule, will be formally maintained initially in a Microsoft Project Management System with the option to move to a more sophisticated system is necessary at some later time.

The WBS for the FSP is consistent with the conventional structure for DOE sponsored programs. Under that structure, each major WBS element representing a major program deliverable is further broken down into a hierarchy of lower level elements. The lowest level WBS elements describe the activities that must be performed to provide the deliverables represented at the higher levels. Every element within a work breakdown structure level has an associated cost estimate (possibly zero, e.g., for a simple checkpoint milestone) and schedule. This method provides both a “top-down” and “bottom-up” perspective for the effort and other resources required by the program.

Some scheduled activities for the FSP may be reliant on the development and delivery schedules of commercial vendors from which computing equipment or software are acquired. (if needed). As such, vendor milestones may be incorporated into the FSP WBS in order to document those dependencies.

6.2 Change Control Process

Any stakeholder may submit a request to modify the execution of the program. All change requests must be formally submitted to the FSP Directorate for preliminary evaluation and possible distribution to the appropriate decision authority. All change requests that may impact a program baseline along with any resulting decisions will be properly documented in accordance with the procedures established in this section and in the separate. More detailed FSP Configuration Management Plan. A formal change control request template is available either in hardcopy or as an online MS Word form that may be printed or e-mailed. Requestors will provide personal identification and contact information; reason for the change; expected cost, schedule, and technical impacts; and new or modified risks resulting from the change. More detail on change controls and configuration management in general can be found in Appendix E:.

6.2.1 Approval Levels For Change Control

The FSP Director is responsible for implementing the baselines and is ultimately responsible for ensuring the change processes are successfully followed. Actual decision authority, however, is vested in a hierarchy of roles depending on the impact that the decision may have on the program.

6.2.1.1 *Change Approval at Level 1*

The DOE FSP Manager(s) must approve changes that would significantly modify the scope of the program (i.e., the set of deliverables), delay a Level 1 Milestone by greater than one month, or that would require any change in Total Program Cost (TPC) or an allocation from contingency funds of $1M or more in aggregate or over $500,000 for any single allocation.

6.2.1.2 *Change Approval at Level 2*

Level 3 changes include changes that will materially affect the final product, delays impacting Level 2 or lower milestones, and any change in control account budgets including contingency use.

The FSP Director will convene a Change Control Board (CCB) whenever necessary to discuss, obtain advice, and support requested changes. The CCB includes the FSP Director, who chairs the board, the Deputy Directors, and any other FSP stakeholder deemed appropriate by the FSP Director. The CCB will, most often, be comprised of members of the Research Committee but could be expanded for highly significant change requests.

6.3 Contingency Management

A program contingency fund will be established to cover cost overruns that may occur as known and unknown risk events occur. The amount in the fund is determined after analyzing budget estimates along with all identified program risks. This technique is discussed in the FSP Risk Management Plan. Authority to reallocate some amount of the contingency funds rests with the appropriate person as described by the change control procedures.

Schedule contingency is calculated and controlled in a way similar to calculating cost contingency and is likewise discussed in the FSP Risk Management Plan.

6.4 Risk Management

This topic is covered in detail in the FSP Risk Management Plan which describes how the program team will identify, track, and manage the various risks events that may impact the program. Essentially the plan is to follow project management best practices by being very proactive in conducting frequent risk identification and analysis sessions and by developing risk management tactics appropriate to all known risks. An example set of risks could include those associated with the following FSP targeted areas:

- Science Drivers: (1) underlying physics models not sufficiently complete to adequately resolve scientific issues consistent with experimental reality; and (2) major challenge of reaching agreement on importance of any given science driver due to varying needs in different parts of FES community
- Frameworks: (1) chosen framework technologies may prove incompatible with future computational architectures; and (2) existing components found to be insufficiently engineered and/or robust for use in the more demanding framework environment
- Components: balancing the needs of delivering advanced physics code software products and the exploratory research needs for producing the physics capabilities required to resolve the FSP Integrated Science Application (ISA) challenges.
- Experimental Validation: even with premier plasma diagnostics, there are practical limitations of experimental measurement to comprehensively measure all important parameters with the needed spatial coverage and resolution
- Verification and UQ: dealing with challenges associated with integrated vs. single physics – especially with fidelity assessments deploying as yet untested models for uncertainty quantification.
- Production Computing: lack of sufficient experience with objective software product testing (“alpha,” “beta,” …) and with customer support for a much large user community than encountered today
- General Risk: DoE-SC needs to acquire the experience/knowledge for managing a major software R&D effort of the scale proposed for the FSP.

Additional risk management detail may be found in Appendix F: Risk Management.

6.5 Communications

This topic will be covered in detail in a FSP Communications Plan which will describes how and to whom program status and other information is communicated. This plan will be written once the Program begins, since the formal reporting structures have yet to be decided.

Essentially, regular formal status reporting schedules will be established for the FSP Team as well as with a possible FSP Federal Project Director (t.b.d. by DOE-SC), and to DOE FES/ASCR (monthly). In addition, there will be reports to other stakeholders as requested, and presentations in more public forums such as annual conferences and research symposia.

Complementing reports and presentations on progress will be outreach activities to promote the FSP new capabilities to the scientific community and to solicit new partners in solving scientific challenges.

6.6 Integrated Safety Management

A key component of a successful program is to ensure that safety, heath, and environmental issues are addressed early in a program's life cycle and fully integrated into all program activities. This topic is covered in detail in the FSP ES&H Plan. Essentially, the FSP Team will follow all relevant safety procedures required by participating institutions as well as follow all appropriate institutional health and safety guidelines. Special emphasis is placed on electrical safety.

6.7 Cyber Security

FSP will follow all cyber security regulations and guidelines of the participating institutions. FSP managed computers will be unclassified and open systems that contain only scientific research data and no personally identifiable information (PII). No Privacy Impact Assessment is required for these systems.

The FSP will meet FISMA, OMB and NIST requirements.

7 REFERENCES

Referenced reports used in the FSP planning activity include:

(1) "The ITER Integrated Modeling Programme," W. Houlberg, (private communication) May 3, 2011.

(2) J. Dahlburg, et al ., FESAC Report, [http://www.isofs.info/FSP_Final_Report.pdf (2002)]

(3) D. Post, et al ., FES Report leading to "Proto-FSP Projects [Journal of Fusion Energy 23, 1 (2004)].

(4) A. Kritz, D. Keyes, et al ., FSP Workshop Report (2007),
http://science.energy.gov/~/media/fes/pdf/workshop-reports/Fsp_workshop_report_may_2007.pdf

(5) W. Tang, et al ., FESAC FSP Report (2007)
http://science.energy.gov/~/media/fes/fesac/pdf/2007/Fesac_fsp_report.pdf

(6) F. R. Bailey, et al ., ASCAC FSP Report (2008)
http://science.energy.gov/~/media/ascr/ascac/pdf/reports/Ascac_fsp_report_final.pdf

(7) R. Hazeltine, D. Hill, et al ., Report of the Research Needs Workshop (ReNeW) (2009)
http://burningplasma.org/web/ReNeW/ReNeW.report.web2.pdf

(8) FSP Program Advisory Committee Final Report (May, 2011)
http://www.pppl.gov/fsp/documents/FSPAC%20reportMAY2011.pdf

(9) “Scientific Grand Challenges: Fusion Energy Sciences and the Role of Computing at the ExtremeScale,”PNNL-19404,212pp(2010).
http://www.er.doe.gov/ascr/ProgramDocuments/Docs/FusionReport.pdf

8 TABLE OF ACRONYMS

AM/CS	Applied Math/Computer Science
AMR	Adaptive Mesh Refinement
ANSI/EIA	American National Standards Institute/Electronic Industries Association
API	Application Programming Interface
ASCI	Accelerated Strategic Computing Initiative
ASCII	American Standard Code for Information Interchange
ASCR	Advanced Scientific Computing Research
AT	Advanced Tokamak
base/SciDAC	fusion base programs/Scientific Discovery through Advanced Computing programs
CAM	Control Account Manager
CCA	Common Component Architecture
CCB	Change Control Board
CET	Centers for Enabling Technologies
CI	Configuration Item
CISL	Computational and Information Systems Laboratory (NCAR)
CPES	Center for Plasma Edge Simulation
CPO	Consistent Physical Objects
CPU	Central Processor Unit
CS	Computer Science
CSEG	Community Software Engineering Group
CTF	Component Test Facility
DEMO	Demonstration Fusion Reactor
DOD	Department of Defense
DOE	Department of Energy
DOE-SC	Department of Energy – Office of Science
ECCD	Electron Cyclotron Current Drive
ECRF	Electron Cyclotron Radio Frequency
EERE	Energy Efficiency and Renewable Energy
EFDA	European Fusion Development Agreement
EFIT	Equilibrium Fitting Code
EHO	Edge Harmonic Oscillation
ELM	Edge Localized Mode
EM	Electromagnetic
EMP	Electromagnetic Pulse
EOT	Education/Outreach/Training
EP	Energetic Particle
Er	Radial Electric Field
ES&H	Environmental Safety and Health
ESL	Edge Simulation Laboratory
ESMF	Earth System Modeling Framework
ET	Enabling Technologies Team within the FSP
EU	European Union
ExB	Electric Field - Magnetic Field cross product
FEM	Finite Element Method
FES	DOE-SC Office of Fusion Energy Sciences
FFCC	Fusion Facilities Coordinating Committee
FISMA	Federal Information Security Management Act
FSP	Fusion Simulation Program
FTE	Full Time Equivalent
GAMs	Geodesic Acoustic Modes
GFDL	Geofluids Dynamical Laboratory
GK	Gyro-Kinetic
HPC	High Performance Computing (or Computers)
ICRF	Ion Cyclotron Radio Frequency

IDM	Internet Download Manager
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
IPCC	Intergovernmental Panel on Climate Change
IPS	Integrated Plasma Simulator
IPT	Integrated Program Team
ISA	An Integrated Science Application Team within FSP
ITB	Internal Transport Barrier
ITER	International Tokamak being built in Southern France
ITM-TF	Integrated Tokamak Modeling Task Force (European Union)
ITPA	International Tokamak Physics Activity
JFNK	Jacobian-Free Newton-Krylov
JxB	Cross product of current density and magnetic field vectors
LCF	Leadership Class Facility
LH	Lower Hybrid
LHRF	Lower Hybrid Radio Frequency
MD	Molecular Dynamics
MFE	Magnetic Fusion Energy
MHD	Magneto-Hydro-Dynamics
MMS	Method of Manufactured Solutions
MoU	Memo of Understanding
MPI	Message Passing Interface
NASA	National Aeronautics and Space Administration
NBI	Neutral Beam Injection
NEO	Drift Kinetic Neoclassical Code
NIST	National Institute of Standards and Technology
NNSA	National Nuclear Security Administration
NPA	Neutral Particle Analyzer
NRC	National Research Council
NSF	National Science Foundation
NTM	Neoclassical Tearing Modes
NTV	Neoclassical Toroidal Viscosity
Off-HPC	Does not require a High Performance Computer
OMB	Office of Management and Budget
On-HPC	Requires a High Performance Computer
OS	Operating System
OSC	Office of Science
PAC	Program Advisory Committee
PCR	Project Change Requests
PCS	Plasma Control System
PDE	Partial Differential Equation
PEP	Program Execution Plan
PERT	Project (or Program) Evaluation and Review Technique
PF	Poloidal Field
PFC	Poloidal Field Coil
PHYS	The Advanced Physics Components Team within FSP
PII	Personally Identifiable Information
PIT	Program Implementation Team
PMI	Plasma Material Interface or Plasma Material Interaction
PMO	Project Management Office
POC	Point of Contact
PPPL	Princeton Plasma Physics Laboratory
PROD	Production Support Team within FSP
Proto-FSP	Fusion Simulation Program Prototype
PSAAP	Predictive Science Academic Alliance Program
PSI	Plasma Surface Interactions
QAP	Quality Assurance Plan
QCM	Quasi-Coherent Mode

QoIs	Quantities of Interest
R&D	Research and Development
RF	Radio Frequency
RIC	Relaxed Iteration Coupling
RMP	Risk Management Plan
RWM	Resistive Wall Mode
SciDAC	Scientific Discovery though Advanced Computing
SIS	The Software Integration Support Team within FSP
SOA	Service Oriented Architecture
SOL	Scrape-Off-Layer
SQ	Software Quality or the Software Quality Team within FSP
SQA	Software Quality Assurance
SQR	System Response Quantity
SRIM	Stopping Range of Ions in Matter
SSC	Scientific Steering Committee
Te	Electron Temperature
Ti	Ion Temperature
TM	Tearing Mode
TPC	Total Program/Project Cost
TQ	Thermal Quench
TSC	Tokamak Simulation Code
TU	Target Uncertainty
UQ	Uncertainty Quantification
V&V or VV	Verification and Validation
VDE	Vertical Displacement Event
VMEC	Variational Moments Equilibrium Code
VVUQ	Verification Validation and Uncertainty Quantification
WBS	Work Breakdown Structure
WDM	Whole Device Modeling
XMHD	eXtended MagnetoHydroDynamics

APPENDIX

Appendix A: Addressing the RFP

In the current FSP Plan, the following topics in the original Request for Proposals (RFP) have been addressed.

FSP Deliverables – With the co-leads (Kritz and Keyes) of the 2007 FSP workshop report as part of the current FSP Planning Team, the list of prioritized deliverables outlined in that document have been critically evaluated and modified as appropriate in articulating the science opportunities and goals. More specifically, each of the 6 Science Drivers discussed in detail in Appendix B: of the Plan are directly related to the "Critical Issues for Burning Plasma Experiments" and "Physics Components Essential for Integrated Plasma Simulations" sections within Chapter 2 of the 2007 FSP workshop report. The associated roadmaps were guided by both the near-term and longer term priorities of FES stakeholders with respect to national as well as international needs, including ITER. The planning study has included a systematic assessment of the resources (in terms of Full Time Equivalent [FTE]) and mix of expertise (plasma physics, material science, applied math, and computer science) necessary to successfully accomplish the Integrated Science Application (ISA) goals. This has accordingly entailed detailed descriptions of the method or approach that will be followed for determining the required resources and reassessing the list of deliverables for the FSP, as well as for developing clear and compelling Work Breakdown Structures. More specifically, this has involved:

- Comprehensive assessment of the present computational capabilities of the fusion community in terms of major simulation codes, numerical algorithms, computational science tools (data management, visualization, code performance tools, etc.), computational frameworks, interface standards, code scalability, and other related issues. Detailed information of this kind can be found in the full length versions of Science Drivers reports in the Appendix of this FSP Plan. Identification of major gaps and weaknesses, and suggestions for the path forward are addressed – with respect to scientific opportunities contained in the Science Drivers discussions within this Plan – as well as in the targeted goals of the ISA's.

- Integration and coordination of the FSP with the projects in the FES SciDAC portfolio, including the process for incorporating results from the FES SciDAC Centers into the FSP have been addressed by a detailed assessment of the SciDAC proto-FSP projects and by the articulation of targeted relationships/collaborations in the Program Execution of the current plan. This has also encompassed:
  - Integration and coordination of the FSP with other SciDAC (non-FES) Centers, and in particular with SciDAC Institutes and Centers for Enabling Technologies (CETs), as well as with efforts supported by the ASCR Applied Mathematics program, as described in this Plan;
  - Integration and coordination with the FES analytic theory and modeling program, including the process for incorporating improved theoretical models into the FSP simulation codes and engaging the help of the FES theory community to address gaps in the physics models implemented in the FSP codes; associated examples are provided in the Science Drivers reports in the Appendix; and
  - Integration and coordination with the materials community for the purpose of addressing the plasma-materials interaction challenges – especially with respect to the ISA on Edge Physics and within the Science Drivers reports in the Appendix.

- Details of the FSP vision and approach for developing a successful and credible Verification and Validation plan, including interaction and coordination with the FES experimental and diagnostic communities is addressed in detail in the current plan.

- Interaction and coordination with international integrated modeling efforts-in particular those undertaken by our ITER partners in support of the needs of the international ITER Organization (IO) are specified in various parts of the plan. This has been informed during the planning process by productive interactions in workshops, meetings, etc. involving, for example, the E.U. integrated modeling activities.

- High Performance Computing (HPC) Resource Requirements-as a major computational activity, the success of the FSP will critically depend on the availability of HPC resources. The FSP Plan describes in sufficient detail the current approach in for determining the required HPC resources for carrying out the various FSP tasks, including the appropriate mix of capacity and capability resources. Resources to be considered include the-current and projected- capabilities at the SC leadership computing facilities, as well as other resources (national or local) that can be reasonably expected to be available to the FSP researchers.

Referenced reports used in the FSP planning activity include:

http://www.science.doe.gov/ofes/programdocuments/reports/FSPWorkshopReport.pdf

http://www.ofes.fusion.doe.gov/FESAC/Oct-2007/FESAC_FSP_report.pdf

http://www.sc.doe.gov/ascr/ASCAC/Reports/ASCAC_FSP_REPORT_FINAL.pdf

Appendix B: Science Drivers (abridged version)

The challenges to theory and simulation are well known – solutions must address a wide range of temporal and spatial scales, intrinsic nonlinearities, strong anisotropies and, particularly near the plasma edge, a rich set of non-plasma physics phenomena. Historically, the approach taken has been to divide the problem into separate domains, each with a limited range of scales. Thus we have RF codes which work on time scales comparable to the inverse cyclotron frequency, turbulence codes working on the inverse diamagnetic frequency and so forth. A similar logic divides up the problem spatially between core, pedestal boundary layer and plasma-wall interactions. However, while substantial progress has resulted, this approach is fundamentally inadequate for many problems. Clean scale separation is only an ideal, in reality, strong ordering is often not justified. Further, additional physics enters in important ways – nuclear reactions, atomic physics, neutral transport, radiation transport, plasma-material interactions cannot be ignored or treated as small perturbations.

Thus the motivation for the FSP is to foster scientific discovery that emerges only upon integrated, multi-physics, multiscale simulation of magnetically-confined fusion plasmas. Detailed planning for the FSP began by selection of a set of Science Drivers. These are a set of compelling scientific problems chosen to focus FSP's initial research, to define and exercise the required range of capabilities and to produce a set of useful tools for the broader fusion community that would substantially impact ongoing research. The Science Drivers could also be described as a set of evolving use cases – defining the requirements for physics components, software infrastructure and experimental validation. The planning team identified six of these drivers which spanned a wide range of needed physics capabilities. These six were:

- Boundary Layer: including turbulence, atomic physics and plasma-wall interactions
- Pedestal: formation, structure and relaxation
- Core Profiles: Including nonlinear turbulence and MHD
- Wave Particle Interactions: including fusion products and RF
- Disruption: detection, avoidance, mitigation and effects
- Whole Device Modeling

A multi-step process was undertaken to develop the overall FSP program plans based on these Science Drivers. The first step was to develop a science "roadmap" for each – that is a step by step plan for building scientific capabilities matched to a set of important fusion physics problems. At each step, the requirements for physics components and software infrastructure were defined along with the needs for verification, validation and uncertainty quantification. Based on these roadmaps schedules and milestones were defined and resource requirements were estimated. Within this section is a summary of these plans, the details of which may be found in the attached document "FSP Science Drivers – Detailed Reports from Community Teams". Not surprisingly, the sum of required resources generously exceeds any expected funding level. Section 2.3 will explain the subsequent process, where the schedule and priorities were adjusted to mesh development elements, including infrastructure and other cross-cutting math and software technologies needed and to match the anticipated funding profile. The result is an overall self-consistent scientific program plan for the FSP, outlined principally in Section 4 of this report.

The Science Driver plans contain extensive discussion of the technical challenges, proposed approaches, requirements for component development, software infrastructure, verification and validation and overall estimates of required resources. Each report contains extensive information in the following organization:

- Background and motivation;
- Goals;

- Components: Including requirements for physics codes (components) that need to be integrated in order to achieve the stated goals and plans for developing new components or adapting existing components;
- Framework requirements: Analysis of the requirements for composition of the physics components (including data exchanges and algorithms) and requirements for workflow (task composition);
- Validation requirements: Plans and priorities for validation of critical physics associated with the science driver including measurement requirements;
- Connections to other work: Needs for collaboration with other efforts within the FSP as well as requirements for work to be accomplished outside the FSP;
- Schedule and resources: A projected schedule of the work to be carried over a 15 year time period including an estimate of resources required;
- Milestones: Suggested high-level goals and milestones (at roughly the 2, 5, 10 and 15 year marks).

While these reports, totaling over 185 pages, are too lengthy to be included in full in this document, the complete reports are essential resources for FSP planning and can be found as an attached document and online at:

http://fspscidri.web.lehigh.edu/index.php/Main_Page#Integrated_Science_Application_Plans.

The sections below summarize each of these longer reports.

B.1 Boundary

B.1.1 Boundary Background

The boundary region in a fusion device includes a narrow plasma region and the near surface of adjoining materials. Processes in this region determine the distribution of plasma particle and heat fluxes to surrounding materials and the associated response of the material (*e.g.*, heating, erosion, and tritium trapping). Simultaneously, the eroded material becomes part of the ionized plasma and its intrusion into the hot core region must be understood and controlled. Issues associated with plasma exhaust, material erosion, tritium trapping, dust, impurity intrusion, RF interactions and response to off-normal events are among the most challenging for the successful development of practical fusion energy. A predictive simulation model of this region requires coupling of disparate physics models describing plasmas, neutral gas, radiation, solid and possibly liquid materials operating on a wide range of space and time scales.

B.1.2 Boundary Motivation

Fusion reactors represent a tremendous extrapolation in power loading and pulse length. At the same time, it is impossible to simultaneously match Scrape-off Layer (SOL) and core dimensionless parameters with current experiments, making empirical prediction rather uncertain. Improved understanding and predictability would impact a number of critical programmatic issues including: 1) Implications for the selection and lifetime of plasma-facing material components, 2) Acceptable levels of tritium co-deposition in re-deposited material and tritium trapping in bulk surface material, 3) Impact on the plasma including core plasma contamination by surface emitted material, 4) Safety issues associated with accumulation of dust that can be easily dispersed during an unintended vent, and 5) Coupling with pedestal and core to impact fueling, toroidal rotation, edge transport barrier, and tokamak density limits.

B.1.3 Boundary Goals

The survivability of fusion plasma facing materials places constraints on the impinging plasma fluxes. A boundary plasma model capable of predicting those fluxes will allow future devices to be designed and operated in a manner consistent with those constraints. Such a model should, first, be able to reproduce the parametric scaling of the following quantities in existing experiments, and, second, incorporate a fundamental understanding of the underlying physical processes, allowing the model to be extrapolated to future devices with confidence. Goals for the model include:

1. Heat loads to material surfaces both during steady state operation (in L-mode and H-mode between ELMs) and in transients (ELMs, disruptions)
2. Fluxes of particles to material surfaces, including those of deuterium, tritium, helium, and all impurities.
3. Fluxes of particles back into the boundary plasma due to plasma-material interactions, including:
    - Impurity generation by physical and chemical sputtering,
    - Recycling of deuterium and tritium,
    - Removal of deuterium, tritium, helium and other particles from the system by pumping mechanisms.
4. Transport of those particles through the boundary plasma and the resulting sources of particles, momentum, and energy in the pedestal and core plasma.
5. Tritium recycling, transport, and retention in materials; implicit in the above, but listed separately because of its importance.
6. Particle, momentum, and energy sources in the boundary and core plasma due to external fueling, including gas puffing, pellets, and other techniques.
7. Modification of plasma facing materials due to plasma fluxes and externally applied treatments (*e.g.* boronization), including erosion, re-deposition with mixed materials, dust generation (and transport).

B.1.4 Boundary Roadmap

The boundary simulation was divided into 6 tasks.

Task 1: Couple fluid plasma turbulence, transport, and neutrals in the SOL

Years 1-2: Couple SOL fluid plasma transport to turbulence models with existing micro-turbulence and transport codes using either iterative coupling or long-time turbulence simulation with continuously evolving profiles. Next neutral models would be coupled in. These would initially be fluid calculations, likely embedded in plasma fluid codes for coupling efficiency and verified against Monte Carlo models. The computational framework could be a continuation of some of the development begun in the FACETS proto-FSP.

Years 3-5: Couple impurities and radiation transport models, assessing impact of turbulence on impurity transport. The fluid turbulence models would be extended to the foot of the pedestal region and begin to include long toroidal wavelength electromagnetic modes (ELMs). The model would accommodate an evolving MHD equilibrium to account for motion of the separatrix. The neutral model would be extended to include additional species and equations to solve for neutral temperature. A kinetic neutral model would be coupled in (likely Monte Carlo) with attention to methods for reducing statistical noise.

Years 6-10: Model 3D effects for the MHD equilibrium, plasma transport (including peaking factors for plasma-material interactions) and radiation transport

Task 2: Coupling plasma-material interaction models with plasma transport

Years 1-2: Couple a dynamic wall model for hydrogen wall uptake and recycling with a dynamic 2D SOL plasma model. Then implement full coupling between near-surface, particle-based sputter erosion/re-deposition code for 2D impurities and SOL 2D fluid plasma model and resolve possible particle-noise issues. Provide the interface and a reduced material model that uses as input ELM and disruption characteristics, i.e., frequency, duration, and power, and can output the material response corresponding to non-melting (acceptable) or melting (non-acceptable) condition.

Years 3-5: Couple initial surface evolution model and near-surface plasma model. Couple kinetic SOL to dynamic SOL models. Improve near-surface model coupling to MD model;

Years 6-10: Couple 3D SOL code to 3D near-surface and PMI codes. Include 3D impurity transport, surface evolution, improved plasma/material interaction models

Task 3: Coupling kinetic plasma turbulence and transport in SOL

Years 1-2: Couple (2D, 2V) kinetic SOL plasma with nonlinear Fokker-Planck collision model capable of full short-to-long mean-free path (leverage CPES and ESL developments). Implement initial coupling (perhaps non-conservative) of kinetic plasma code to kinetic neutral model; demonstrate strong recycling and near steady-state. Develop and extend kinetic Monte Carlo neutral transport component:

Years 3-5: Couple kinetic (first electrostatic, then EM) turbulence to kinetic transport from foot of pedestal to wall. Improve (conservative, more efficient) coupling of kinetic plasma code to kinetic neutral model then apply similar technique to nonlinear neutral transport problems in kinetic Monte Carlo code.

Years 6-10: Couple kinetic impurities to main ion transport, extending kinetic domain well into pedestal; (either couple to pedestal model or extend domain of single kinetic model). Couple Kinetic ELM simulations to model ejection and heat footprint. Develop hybrid fluid-kinetic neutral transport component:

Task 4: Coupling SOL and Pedestal plasmas

Years 2-5: Begin extending fluid and kinetic transport well across separatrix (see Tasks 1 and 3)

Task 5: Coupling RF antennas/physics with SOL and PMI models

Details have yet to be worked out, but modeling would address the interaction the boundary plasma and material with radio frequency (RF) antennas, and associated electromagnetic fields, requiring much better integration with boundary models. RF sources inject power into the SOL plasma and potentially drive large RF sheaths, and in turn, the plasma gives rise to antenna sputtering. All of these transient and RF processes produce supra-thermal particles and thus ultimately require kinetic descriptions in 3D.

Task 6: Atomic physics models

Years 1-5: Develop tractable characterization of high-Z atoms. Calculate kinetic details for hydrogen molecular physics and incorporate into kinetic neutral transport model. Identify and obtain data for molecular species pertinent to mixed material environment of ITER. Assemble improved data and simplified models for breakup of hydrocarbon molecules:

B.1.5 Summary of Boundary Resource Estimates and Milestones

Resources (FTE/year) required to fully implement this program have been estimated as follows:

Task	Years 1-2	Years 3-5	Years 6-10
Task 1	**2.5**	**3.5**	?
Task 2	**6**	**4**	**2**
Task 3	**5**	**6**	**4**
Task 6	**7**	**7**	?
Totals	**20.5**	**20.5**	

Table 12: Effort estimates for boundary research

With those level of resources, the following tables summarizes achievable high-level goals and milestones:

Boundary Milestone	**Year from inception**
Self-consistent SOL models with fluid plasma turbulence and transport (heat-flux width)	2
Dynamic coupling between PMI model and SOL plasma (integrated particle inventory)	2
Electrostatic kinetic turbulence and transport in SOL	5
Surface evolution model	5
Extension of kinetic transport and turbulence into pedestal or coupling with pedestal model	10
Tritium transport and retention	10
Electromagnetic kinetic turbulence and transport	10-15
3D kinetic transport – peaking factors	15

Table 13: Milestones for boundary research

B.2 Pedestal

B.2.1 Pedestal Background

High performance ("H Mode") operation in tokamaks is achieved via the spontaneous formation of a transport barrier (or "pedestal") in the outer few percent of the confined plasma. This edge transport barrier strongly improves global energy confinement, and also generally improves global stability, resulting in dramatically enhanced fusion performance. The pedestal presents a daunting set of challenges to traditional theoretical and computational methods. Because the pressure varies by 1-2 orders of magnitude across the pedestal, and the density, temperature, flow velocity, radial electric field and current density also vary substantially, a very wide range of key dimensionless parameters is encompassed in this region. For example, the pedestal often transitions from highly collisionless near the top, to strongly collisional at the bottom, requiring methods appropriate for both regimes. More fundamentally, the broad range and overlap of spatiotemporal scales across the pedestal deeply challenges the assumed separation of equilibrium ("macro") and turbulence ("micro") scales upon which most existing theory and computation relies. Further, perturbations can be large compared to the background equilibrium, for example during repetitive instabilities (ELMs) or so-called "blob" transport, presenting a challenge to perturbative methods. Flows and sources, including impurity radiation and atomic physics, are expected to be important, bifurcations and operation near marginality must be considered, and geometry is complex, particularly in problems for which coupling to the boundary region outside the separatrix is strong. In addition, in plasmas with ELMs, the pedestal region does not generally reach a steady state, but rather continues to evolve throughout the ELM cycle. The ELM itself is a highly complex event, involving both MHD and transport physics, and extending from the pedestal region, where it is primarily driven, out into the open field line region, and finally onto material surfaces, with coupling back to the deeper core via the evolving pedestal profiles.

B.2.2 Pedestal Motivation

The plasma pressure typically increases by 1-2 orders of magnitude from the bottom of the pedestal to the top, and increases by less than an order of magnitude from the pedestal top to the magnetic axis. Hence, while the pedestal occupies a relatively narrow radial region, it contains far more pressure scale lengths than the core plasma. The impact on global confinement is amplified via coupling to the core plasma where transport is fairly stiff, meaning that the core profiles are closely correlated to critical gradient scale lengths. As a result, the core pressure increases roughly linearly with the pedestal pressure (or "pedestal height"), and the fusion power output scales roughly as the square of the pedestal height, providing a powerful lever for performance optimization of fusion systems. While the performance benefits of H-mode operation are dramatic, there is a potential drawback. The large pressure gradients in the edge barrier lead to large localized currents, via the bootstrap effect, and the substantial free energy present in both the pressure and current gradients drives the ELMs While ELMs are largely benign in existing devices, and can aid in density and impurity control, ELMs deposit a highly impulsive heat and particle load on plasma facing surfaces, which may constrain component lifetimes in reactor scale devices. A predictive understanding of pedestal formation and structure, as well as the physics of ELMs and enhanced confinement regimes without ELMs, is essential for prediction and optimization of the fusion performance of ITER and future reactors.

B.2.3 Pedestal Goals

The practical goal for pedestal research is to elucidate a path toward operation with a high pressure pedestal with a profile relaxation mechanism which does not present the material interface with unacceptable transient heat loads – that is to operate with small or no ELMs. For modeling, the goal is to develop the capabilities to understand and predict:

1. The onset of edge barriers (or "L-H transition") as well as the transition from low to high performance H-modes including the prediction of the transition in terms of input power.
2. The structure of the barrier in all profiles (with particular initial emphasis on the pressure at the top of the pedestal), including prediction of large scale radial electric field (Er) and plasma rotation. This also includes better understanding of plasma fueling across the pedestal region.
3. The nature of the pedestal relaxation, particularly the wide variety of ELM types and ELM-free H-modes and to identify and optimize methods for reducing transient heat deposition on material surfaces (including ELM-free and small ELM regimes, as well as suppressing or mitigating ELMs via external control techniques, including magnetic perturbations or pellets).

B.2.4 Pedestal Roadmap

The goals, challenges and progress described above lend themselves to a three level plan for the FSP pedestal effort. This plan addresses both the need to deliver world-leading capability on a relatively short timescale, and the need to address the deeper fundamental challenges associated with pedestal dynamics, taking advantage of peta- and exascale computing capability as it becomes available. All of these models would need to be appropriately verified, including extensive verification of reduced dynamic models against direct nonlinear simulations, and validated against experimental measurements.

Level 1. Linear models for pedestal structure

This step would begin with componentization of existing models that solve for static (time averaged) pedestal structure via linear stability analysis, for example, that of peeling-ballooning and kinetic ballooning modes. Improvements can come through use of linear or quasi-linear gyrokinetic calculations, more realistic geometry and inclusion of ExB stabilization. Extended models could include shorter wavelength driftwave modes (e.g., electron temperature gradient modes) and neoclassical effects. This analysis typically requires hundreds or thousands of independent MHD and/or gyrokinetic stability calculations with trial equilibria. Key issues are robustness, error checking, automation, and, particularly in the case of gyrokinetic calculations, efficiency.

Level 2. Dynamic evolution of the pedestal via separate inter-ELM and ELM components

2a. Dynamic evolution of pedestal profiles between ELMs: The fundamental tool for calculating pedestal transport between ELMs is expected to be electromagnetic gyrokinetic simulations of turbulent transport, coupled to separate calculations of neoclassical transport and sources. In some limits, gyrofluid or Braginskii simulations may also be employed. Neoclassical calculations will eventually include 3D equilibrium effects, such as neoclassical toroidal viscosity. Source models should include neutral transport and pellet fuelling – and eventually to a more complete model of the boundary plasma, recycling, impurity sources, etc.

2b. ELM dynamics & control with fluid or kinetic-fluid hybrid models: The models described above would be extended by simulation of phenomena which limit or control the pedestal pressure gradients. These would include spontaneous plasma behavior including ELMs of various types, Edge Harmonic Oscillation (EHO), Quasi-Coherent Mode (QCM) as well as active control through pellets, resonant magnetic perturbations (RMP), electromagnetic perturbations, etc. The work could begin with linear onset from peeling-ballooning calculations, coupled to simple ELM crash models. The next step would be direct simulation of ELM dynamics using extended MHD or two-fluid and/or kinetic-fluid codes. These codes would need to include realistic calculations of parallel transport and through coupling to boundary models, compute transient heat and particle loads onto material surfaces. Validation experiments could compare ELM (or other mode) structure, dynamic modification of pedestal profiles, heat and particle footprints and ELM control mechanisms.

Level 3: Direct Multi-Scale Simulation

The prior computational steps use gyrokinetic calculations for modeling the micro-scale and extended MHD for the macro-scale. However, as noted above, these overlap strongly in the edge barrier. Some systematic study will be required to test the assumption of spatiotemporal scale separation, to determine when and how it breaks down and to assess the consequences. Numerical and theoretical progress will be required to develop and implement verified formulations and codes which can simulate multi-scale electromagnetic modes and turbulence in separatrix geometry. Several approaches are possible including gyrokinetic treatments without the high-n approximation, kinetic-fluid methods and 6D Vlasov treatments including the full collision operator. The last of these, in particular, will require substantial progress in numerics to be practical. These models would support the most fundamental studies of pedestal physics including threshold, coupling of turbulence and equilibrium scales, ELMs and ELM control.

B.2.5 Summary of Pedestal Resource Estimates and Milestones

Resources (FTE/year) required to fully implement this program have been estimated as follows:

Task	Years 1-2	Years 3-5	Years 6-10	Years 11-15
Level 1	3.75	3	.1	.1
Level 2	5	7	6	4.5
Level 3	1	1	4	4
Totals	10	11	10	8.5

Table 14; Effort estimates for pedestal research

Suggested high level goals and milestones:

Pedestal Milestone	Year from inception
Componentization of Level 1 models	2
Initial validation of Level 1 models against pedestal height and width observations	2
Initial coupled pedestal-core optimization of ITER base case with Level 1 models	2
Componentization of Level 2 models	5
Pedestal turbulence simulations on closed field lines	5
Initial development of reduced dynamic models from nonlinear simulations	5
Coupled simulation of between-ELM transport with reduced models, and ELM events	5
Pedestal turbulence simulation on closed and open field lines	7-10
Validation of calculated ELM heat flux on material surfaces	7-10
Validation of pedestal turbulence simulations	7-10
Componentization of Level 3 models	15
Verification of Level 2 components using Level 3 component	15
Direct multi-scale simulations of transport and ELMs	15
Simulation of the L-H mode transition	15

Table 15: Milestones for pedestal research

B.3 Core Profiles

B.3.1 Core Profiles Background

This task entails the development of validated, predictive simulations of core profiles for temperature, density and momentum for all plasma species and on time scales relevant to plasma evolution. It is not restricted to MHD-quiescent discharges and thus must consider, in addition to collisional and turbulent transport processes, the effects of MHD modes such as sawteeth, ELMs and more importantly, neoclassical tearing modes (NTM) and fast-particle driven modes. A successful model should also take account of any other mesoscale phenomena including turbulence spreading, self-organized criticality, super- or sub-diffusion and so forth. Coupling to the edge plasma are likely more complex than simple boundary conditions on profiles and may entail consideration of fluctuations and flows.

Turbulence, driven by gradients in the plasma temperature and density, is the dominant mechanism for the transport of particles, momentum, and energy. This turbulence is rigorously described by the coupled gyrokinetic and Maxwell equations. In addition, there is a well-developed theory for neoclassical transport, which provides an adequate description of the residual collisional transport present in the absence of plasma microturbulence. Most previous efforts to simulate the transport-time-scale evolution of core plasma profiles are based on the solution of the 1-D transport equations using approximate transport coefficients obtained through some combination of analytic theory, simulations, and experimental results. The advent of terascale computing opens the possibility of directly computing the turbulent fluxes for input into the transport equations. This capability has been successfully demonstrated by both TGYRO and TRINITY. While these models have mainly focused on the local transport model, where turbulence and transport on each flux surface are assumed to be entirely independent and driven solely by local plasma profiles, they may also be applicable to situations when mesoscale phenomena are important. Ideally, magnetic confinement devices employ a magnetic geometry in which the field lines cover a set of nested toroidal surfaces (the magnetic surfaces). Particles, momentum, and energy are transported rapidly within each magnetic surface, and more slowly across magnetic surfaces. This allows us to reduce the problem of computing core profiles of density, plasma flow, and temperature from a problem in three spatial dimensions to a one-dimensional problem (in the flux-surface label, which we take to be r in this discussion). One must also solve a subsidiary equation for the equilibrium magnetic geometry with the given the plasma profiles. However, in real devices this degree of symmetry may not be attained. The flux surfaces may be modified by 3D magnetic perturbations from MHD instabilities, field errors or applied perturbations. It's clear from experimental evidence, that if these perturbations are sufficiently strong, underlying turbulence and transport can be affected. A complete model for core profiles will need to account for these effects in some cases.

A representative and important problem concerning 3D effects is the interaction of plasma turbulence with NeoClassical Tearing Modes (NTMs). NTMs are slowly growing MHD modes which are generally unstable in high-performance plasmas. Macroscopic quantities such as the flows or profiles interact with a magnetic island and the island, in turn, may drastically influence the underlying turbulence, leading to changes in the flow pattern and to core plasma profile evolution. This importance of the interplay between flows and turbulence in the presence of an island is widely appreciated and at the same time has never been addressed self-consistently. Critical questions which should be addressed include: i.) what determines the critical island size? ii.) what are the effects of polarization currents and how do flows evolve when the island is present? iii.) do density and ion temperature gradients flatten, and if so, by what mechanism? iv.) what type of confinement state or regime results? In particular, might internal transport barriers form?

B.3.2 Core Profiles Motivation

Prediction of core plasma profiles addresses what is perhaps the most basic question for magnetic confinement. For a given input power and fueling level, how hot and dense is the resulting plasma? The answer to this question translates directly to the ratio of fusion to input power and thus to the ability to generate net power from a fusion device. The predicted plasma profiles also determine the macroscopic stability of the plasma and the quantity of bootstrap current which is essential for achieving true steady state in the tokamak. In general, solution to this problem requires detailed simulation of plasma turbulence – one of physics' grand challenges.

B.3.3 Core Profiles Goals

The ultimate scientific goal is a validated transport model which reliably predicts, for each plasma species, profiles of density, temperature, rotation and current and their evolution on transport time-scales. That is, it would encompass all the phenomena that set the core profiles including turbulence (in all relevant fields and at all relevant scales) and nonlinear MHD (i.e. soft limits as opposed to collapses or disruptions). It would to need to solve for turbulence in 3D perturbed equilibria and include the physics that controls transport barriers. Ultimately, understanding plasma transport is central to the design of an engineering test reactor and commercial power plants based on magnetic fusion. In the process, this topic addresses scientific grand challenges including nonlinear coupling of dynamics across a broad range of spatial and temporal scales.

This goal will be addressed via three parallel paths:

1. A local transport model based on coupling many gyrokinetic simulations distributed across the radial profile
2. A transport model based on global simulations which is capable of addressing mesoscale phenomena on transport time-scales
3. Developing a plan for incorporating boundary interactions between the core and pedestal including fluctuations and flows.

A key component of this development plan is the continual validation of these physics models against experimental data throughout its entire duration.

B.3.4 Core Profiles Roadmap

The program to meet the goals outline above are broken into 4 distinct but overlapping tasks:

Task 1: Comprehensive validation of local and global models

Year 1-2: Begin work on existing codes within the FSP environment, helping to insure that the FSP mechanisms for data access/storage, user interfaces, etc. properly supports validation Begin development of appropriate synthetic diagnostics. Plan validation campaigns.

Years 3-5: Complete initial validation assessment study of relative and absolute local and global gyrokinetic turbulence models for slowly evolving, MHD-quiescent plasmas spanning low-β_N Ohmic discharges to high-β_N H-modes. Complete initial validation assessment study of self-consistent fast particle profile and Alfven eigenmode saturation in varying plasma conditions. As global gyrokinetic models are available. Analysts need to adapt tools (like synthetic diagnostics) to these codes and plan the year 5 validation campaign. Depending on the success of the validation campaign thru year 3, consider expanding the number of experimental analysis to improve coverage of major U.S. facilities and/or initiate validation campaigns on major international facilities. Complete three validation milestones, (1) comparing local and global models in MHD-quiescent plasmas; (2) in presence of significant Alfven eigenmode activity; and (3) self-consistent islands in the presence of turbulence.

Years 6-10: First iteration of core/edge coupling algorithm is complete. Analysts must develop plans for validation of coupled edge/core model.

Task 2: Local Model Development

Year 1-2: Recruiting personnel, learning code use within FSP environment and helping to insure that the FSP code development effort properly supports Validation (data access/storage, appropriate models, user interface, etc.). Local model becomes operational to meet milestone “demonstration of local transport model”. Incorporate fast particles in anticipation of year-3 milestone.

Year 3-5: Support validation milestone relating to fast particles. Update codes as required to improve both code fidelity and user interface. Prepare for year-5 validation milestones comparing local and global models. Update codes as required to improve both code fidelity and user interface. Support year-5 validation milestones comparing local and global models. Update codes as required to improve both code fidelity and user interface.

Task 3: Global Model

Year 1-2 Develop required components (3-D GK code, 3-D magnetic geometry component, 3-D transport solver), and begin coupling these components.

Year 3-5: Complete coupling of 3-D components to produce functional mesoscale transport model. Prepare for year-5 validation milestones comparing local and global models; computing self-consistent turbulence in presence of 3-D islands: validation of free-boundary equilibrium in presence of islands plus turbulence Update codes as required to improve both code fidelity and user interface. Support year-5 validation milestones comparing local and global models; computing self-consistent turbulence in presence of 3-D islands: validation of free-boundary equilibrium in presence of islands plus turbulence Perform self-consistent calculation for narrow islands to determine NTM threshold width. Update codes as required to improve both code fidelity and user interface.

Task 4: Core/Edge Coupling

Year 3-5 Finalize plan for core/edge coupling. Recruit personnel. Develop required packages, prepare for Year-5 milestones. Complete code development and demonstrate initial core/edge coupling model.

B.3.5 Summary of Core Profiles Resource Estimates and Milestones

Resources (FTE/year) required to fully implement this program have been estimated as follows:

Task	Years 1-2	Years 3-5	Years 6-10
Task 1: Validation	3	3-10	3-10
Task 2: Local Model	2	2	2
Task 3: Global Model	3	3	3
Task 4: Core-Edge Coupling	0	2	2
Totals	8	10 - 15	10 - 15

Table 16: Effort estimates for core profile research

Suggested high level goals and milestones:

Core Profiles Milestone	Year from inception
Identify initial verification and validation test cases (including metrics to be used for each case), and integrate all needed experimental data into a generally accessible database.	2
Complete verification assessment and documentation of existing components and frameworks.	2
Deliver prototype framework for a time-dependent 1.5D transport solver built from legacy components (e.g. "FSP0") and Demonstration of local transport model operation within the FSP code	2
Complete initial validation assessment study of relative and absolute local and global gyrokinetic turbulence models for slowly evolving, MHD-quiescent plasmas spanning low-β_N Ohmic discharges to high-β_N H-modes.	3
Complete initial validation assessment study of self-consistent fast particle profile and Alfven eigenmode saturation in varying plasma conditions.	3
Couple global gyrokinetic codes to a 3D equilibrium code with islands.	3
Incorporate kinetic and flow effects in 3D equilibria through gyrokinetic calculation of $\mathbf{j}_\perp$.	3
Complete initial validation assessment study of relative and absolute local and global gyrokinetic turbulence models for quickly evolving, MHD-quiescent plasmas (internal transport barrier formation, core response to edge BC change via L-H transition).	5
Complete initial validation assessment of relative and absolute local and global gyrokinetic turbulence models in presence of significant Alfven eigenmode activity, with focus on fast particle transport.	5
Complete code development (and demonstrate operation) for 1st cut on core/edge coupling based on local transport model.	5

Calculate self-consistent turbulence in the presence of magnetic islands through the coupling of a 3D equilibrium code to a global gyrokinetic code that can properly handle neoclassical effects as well as turbulence.	5
Validate self-consistent solutions for free-boundary equilibria with magnetic islands in the presence of turbulence against tokamak data for saturated tearing modes.	5
Self-consistent calculations for narrow islands determine the NTM threshold width.	5
Complete 2nd round of validation assessment study of relative and absolute local and global gyrokinetic turbulence models MHD-quiescent plasmas.	10
Complete self-consistent calculation of the evolution of 3D equilibria in the presence of turbulence.	10

Table 17: Milestones for core profile research

B.4 Wave-Particle Interactions

B.4.1 Wave-Particle Background

The realization of fusion energy requires the efficient production of well-confined suprathermal ions (alphas, injected neutral beams, RF heated ions) and the efficient transfer of their energy to the core of the thermonuclear burn. In addition, efficient and reliable methods are needed for heating, current drive and MHD stability control using radio frequency power in the ion cyclotron (ICRF), lower hybrid (LHRF), and electron cyclotron (ECRF) range of frequencies, as well as high-energy neutral beam injection. Suprathermal ion heating is generally efficient and reliable in existing experiments, in part due to the low velocity of the injected neutrals compared to the Alfvén velocity. However, Alfvén wave instabilities are expected in future reactors and in ITER, where the bulk of the suprathermal ions will have velocities exceeding the Alfvén velocity, with consequences that cannot be reliably predicted at present.

Effective RF control of the plasma requires a detailed quantitative understanding of the various mechanisms that can dissipate wave energy as well as the effective control of these mechanisms to efficiently and reliably target the wave energy. Recent advances in RF theory and simulation, coupled with experimental advances enabled by new diagnostics and continuing detailed measurements, have led to an unprecedented understanding of the physics of RF heating and current drive in the plasma core of axisymmetric toroidal magnetic fusion devices. However, the existing models cannot predict the net amount of power that will be coupled into the core of the tokamak plasma instead of dissipated in the launcher, nearby vessel components or plasma edge. Questions remain about the detailed self-consistent interaction of the ICRF and LHRF waves with energetic particles created by the RF waves, by NBI, or by fusion reactions.

B.4.2 Wave-Particle Motivation

Understanding these processes is essential for predicting fusion reactor performance, where almost all power needed to sustain the plasma comes through highly energetic ions. Significant progress has been made in understanding the range of wave-particle phenomena relevant to the production and dissipation of suprathermal ions and electrons. However a predictive capability is far from available at present. In the area of collective Alfvénic instabilities, integrated modeling capability is needed to predict the spectrum of unstable modes and to evolve the modes self-consistently with the particle distribution. In both the ion cyclotron and lower hybrid range of frequencies, predictive capability is needed for edge dissipation mechanisms and for the self-consistent description of waves in the presence of high-energy ions in the plasma core. The nonlinear saturation and dynamics of Alfvénic instabilities depends critically on the fluxes of particles in the phase space of the energetic particles. These fluxes are heavily modified in the presence of RF fields as compared to say coulomb collisions and classical slowing down. In addition, the presence of the Alfvénic instabilities will modify the flux further and this will also have an impact on the particle distribution generated by the RF waves. Hence the physics of collective instabilities in the presence of RF waves and suprathermal particles is a strongly coupled problem requiring advanced simulation capability.

B.4.3 Wave-Particle Goals

Solution to the Wave-Particle problem will require a number of distinct modeling activities that are ultimately integrated into a self-consistent description of RF fields, fast ions, MHD instabilities and turbulent background. Targeted goals would include:

1. Address in present experiments, the role of Alfvénic modes in modifying the fast ion distribution in the plasma core and enhancing particle losses to the first wall.
2. Develop a multimode capability to predict the self-consistent mode amplitudes and fast ion distribution, with prescribed sources and sinks of fast ions, in regimes where multiple Alfvénic instabilities are expected.
3. Extend this capability to address additional interactions of the energetic particles with the background plasma (such as turbulent fluctuations and tearing modes) and integrate into WDM.
4. Predict wall loading and fast ion pressure relaxation in advanced tokamak (AT) regimes with RF heating and Alfvénic instabilities.
5. Predict parasitic power losses for RF heating and current drive in the edge plasma, launcher and vessel walls. Models could be developed in the near term and compared to experimental studies in order to develop a simplified but validated model for RF wave propagation in the edge plasma regions.
6. Predict wave-particle plasma interactions that lead to sawtooth stability under the combined influence of RF tail ions, neutral beam ions, alpha particles and localized current drive schemes such as electron cyclotron current drive (ECCD). In particular, assess sawtooth stability in ITER in the presence of Alfvénic instabilities and RF heating.
7. Develop a self-consistent description of the suprathermal ion population under the combined influence of a turbulent background, tearing modes and the RF field.
8. Integrate wave-particle models with turbulence simulation codes to help address the role of Alfvénic instability induced zonal flows on the background turbulence and on the role of background turbulence in modifying the distribution of the energetic particles.

B.4.4 Wave-Particle Roadmap

Task 1: Understand the physics of multimode induced redistribution of fast ions

Years 1-2: Develop quasi-linear model with multiple modes, realistic sources and resonance overlap using linear eigenmode solutions. Develop PIC code with linear eigenmode solutions and realistic sources. Develop nonperturbative kinetic eigenmode solver including realistic distributions of energetic particles in fully nonlinear codes.

Years 3-5: Integrate reduced models in WDM code and validate against experiment. Extend fully nonlinear simulations codes to a slowing down time scale. Integrate RF source into reduced model

Years 6-10: Integrate RF source into fully nonlinear models.

Years 11-15: Integrate fully nonlinear code with RF source into RWM and sawtooth stability models and turbulence models. Integrate fully nonlinear models as a replacement to the NUBEAM package (or equivalent) in the WDM code

Task 2: Understand the effect of the edge plasma, RF launching structure, and tokamak vessel on the coupling of RF waves in the ion cyclotron range of frequencies (ICRF) and lower hybrid range of frequencies (LHRF):

Years 1-2: Develop finite element method (FEM) description of edge with linear and nonlinear boundary conditions. Develop coupled edge to core description using a core spectral solver coupled to an FEM edge description through an admittance matrix, for example the TORIC + TOPICA codes. Perform preliminary 3D simulations of core to edge ICRF wave dynamics utilizing spectral solvers (AORSA +

TORIC) extended to a cold, linear plasma model in the edge. Complete development of PIC codes (VORPAL) for simulating linear and nonlinear RF wave interactions with the plasma edge

Years 3-5: Use 3-D field reconstructions from spectral codes extended to edge and coupled spectral / FEM models to simulate ICRF fast and slow wave excitation including surface wave excitation and RF sheath formation in present day tokamaks (NSTX, DIII-D, and Alcator C-Mod). Use 3-D field reconstructions from coupled spectral / FEM model to simulate LH wave coupling in present day tokamaks. Formulate conductivity operator in FEM basis in 2-D. Develop coupled core – edge FEM RF solver based on new conductivity representation and verify code against spectral solvers (AORSA and TORIC)

Years 6-10: Use new FEM-based core-to-edge solver to assess surface and slow wave excitation in the ICRF regime using 3-D field reconstructions. Use new FEM solver to assess electric field high points on ICRF launching structures and compare with experimental measurements. Simulate long distance ICRF and LHRF coupling in ITER. Include effects of wave scattering and edge plasma variations on coupling in wave solvers

Years 11-15: Building on existing 1-D full-wave PDI simulation experience, develop 2D / 3D full-wave simulation capability for describing three wave parametric decay instability (PDI), including finite toroidal extent of pump wave and compare simulated decay spectra with measurements. Perform 3-D simulations of parametric decay instability using hybrid codes that employ an electron fluid description and a particle treatment for ions. Perform 3-D simulations of RF sheath formation using hybrid codes that employ an electron fluid description and a particle treatment for ions

Task 3: Understand the role of finite ion orbit width effects and mode conversion to short wavelength modes in ICRF heating schemes:

Years 1-2: Complete integration of full-wave / Monte Carlo description of ICRF – fast wave particle interaction using statistical particle lists and 4-D quasilinear diffusion coefficient. Validate model against experiment with synthetic diagnostics for neutral particle analyzer and fast ion Dα imaging – examine interaction of ICRF fast waves (low and high harmonic) with neutral beam ions. Use synthetic diagnostic codes for reflectometry and PCI to validate simulations of mode converted ICRF waves against experimental measurements

Years 3-5: Employ energetic particle distributions modified by the EP components in wave propagation model. Evolve EP distributions in full-wave / Fokker Planck solvers and pass back to EP component

Years 6-10: Use reduced models (full-wave + continuum Fokker Planck with finite orbit width effects) to study interaction of ICRF fast waves with NBI ions and fast fusion alphas in ITER.

Years 11-15: Use parallel framework to perform time dependent simulations where EP component, EP sources, and WDM models are iterated in time.

Task 4: Understand how LHRF generated nonthermal electron tails can be used for localized control of the current profile:

Years 1-2: Validate nonthermal electron distributions simulated by coupled full-wave / electron Fokker Planck model using synthetic diagnostic codes for hard x-ray emissivity and current density (Motional Stark Effect). Develop theory for fast ion (fusion alpha) – LH wave interaction and implement in full-wave solver

Years 3-5: Compare predictions of ray tracing / Fokker Planck model against more complete full-wave / Fokker Planck treatments to determine conditions under which reduced ray tracing description is adequate. Assess interaction of LH waves with fast alphas for an ITER discharge

Years 6-15: Use parallel framework to perform time dependent simulations of LH current profile control in present day devices and in ITER.

Task 5: Understand how thermal electron distributions and nonthermal ion distributions generated by ICRF and ECRF waves can stabilize of destabilize MHD phenomena in plasma.

Years 1-2: Finish closure theory for including driven currents due to electron cyclotron current drive (ECCD) in the MHD equations – (cases where the electron distribution is minimally distorted and the RF effect can be included through an RF flux term). Using a parallel framework, numerically implement this closure scheme using a ray tracing code to evaluate the ECRF – induced flux in the MHD equations

Years 3-5: Validate simulation capability for NTM and sawtooth control via ECCD against experiments using the parallel framework capability developed in (a).

Years 6-10: Finish kinetic closure theory for including energetic ICRF distributions in the MHD moment hierarchy (case where the ion distribution function is anisotropic).

Years 11-15: Use a parallel framework to numerically implement closure schemes needed to include the effect of energetic ICRF tail in MHD codes.

Task 6: Understand the effect of driven RF waves on plasma rotation, plasma flows, and the scrape-of-layer (edge):

Years 1-2: Continue to validate existing theories for toroidal plasma rotation via ICRF and LHRF waves against experiment using the simulated wave fields from core wave solvers.

Years 3-5: Develop new theory for toroidal rotation drive and plasma flow generation via LHRF and ICRF waves if needed. Use qualitative predictions of edge RF dissipation in coupling scheme shown in Fig. WP-2 and simulate using a parallel framework

Years 6-15: Once RF rotation theory is developed and validated, perform time dependent simulations of existing discharges and ITER using a parallel framework. Use parallel framework to couple the edge ICRF and LHRF wave solutions with gyrokinetic edge particle transport and stability codes (see coupling scheme in Fig. WP-2), in order to understand the interactions of RF with ELMs and to understand impurity generation from sheath interactions with the vessel

B.4.5 Summary of Wave-Particle Resource Estimates and Milestones

Resources (FTE/year) required to fully implement this program have been estimated as follows:

Task	Years 1-2	Years 3-5	Years 6-10	Years 11-15
Task 1	4	4	2	2
Task 2	6	5	5	7
Task 3	6	2	3	4
Task 4	2.5	2	5	5
Task 5	3	2	3	4
Task 6	1	3	8	8
Totals	22.5	18	26	30

Table 18: Effort estimates for wave-particle research

Suggested high level goals and milestones

Wave-Particle Milestone	Year from inception
Demonstrate capability to simulate fast ion transport and redistribution using a reduced model analysis on a slowing down time scale with sources and sinks of fast ions (no RF) and linear eigenmode solutions. Begin validation of reduced models with experiment.	2
Demonstrate capability to simulate linear nonperturbative eigenmodes for inclusion in reduced models and validation of fully nonlinear codes. Begin verification of eigenmode solver with linear mode structures.	2
Begin to simulate multiple Alfvénic instabilities using fully nonlinear codes with realistic sources and sinks of fast ions. Begin validation of nonlinear code solvers.	2
Demonstrate capability to simulate linear 3-D ICRF and LHRF wave fields in the plasma edge.	2
Have coupled full-wave / Fokker Planck simulation capability in place to treat finite ion orbit width effects.	2
Demonstrate capability to simulate fast ion transport and redistribution using a reduced model analysis using nonperturbative kinetic eigenmode solvers. Continue validation effort with updated eigenmode solver. Integrate reduced model into whole device simulation code for assessing effects of fast ion redistribution and loss on discharge evolution and vessel safety, particularly for ITER.	5
Extend simulation of multiple Alfvénic instabilities using fully nonlinear codes for longer durations, approaching the slowing down time of the energetic particles, with accurate description of sources and sinks. Begin validation of fully nonlinear solvers against experiment on slowing down time scale.	5
Validate simulation capability for linear ICRF and LHRF wave coupling against experiment.	5
Validate simulation capability for core ICRF wave physics with finite ion orbit effects against experiment.	5
Demonstrate capability to simulate coupling between EP sources and EP component (reduced models) by passing RF induced non-thermal ion distributions.	5
Validate capability to quantitatively simulate RF sheath effects against experiment.	10
Have closure scheme(s) formulated for including non-thermal ion distributions in MHD equations.	10
Perform simulations using a self-consistent coupling between the EP sources and the EP component based on the closure relations formulated for the MHD equations.	15
Perform simulations using a self-consistent coupling between the RF waves and plasma edge, which includes the effects of non-linear RF edge dissipation mechanisms in edge transport and stability codes.	15

Table 19: Milestones for wave-particle research

B.5 Disruption Detection, Avoidance, Mitigation and Consequences

B.5.1 Disruptions Background

During tokamak experimental operation, disruptive events that rapidly terminate the plasma discharge occasionally occur. The initial triggers for disruptions are varied, but all lead to large scale magnetohydrodynamic (MHD) instabilities that destroy the magnetic geometry and lead to rapid loss of plasma confinement. The key scientific challenges include strongly nonlinear MHD, including kinetic effects, with large Lundquist number coupled to plasma pressure and current profile evolution; relativistic electron transport in stochastic magnetic fields; atomic physics; neutral and impurity transport; radiation transport; plasma wall interactions and an electromagnetic model of machine with its complex wall geometry, power supplies coils, control systems and diagnostics. The effects of disruptions include severe heat loads, JxB forces and run-away electron generation.

B.5.2 Disruptions Motivation

Disruptions pose a serious threat to current-carrying devices like a tokamak. The complete and rapid loss of thermal and magnetic energy in these disruptions results in large thermal and mechanical loads on the material wall. For proposed next step experiments such as the International Thermonuclear Experimental Reactor (ITER), the stored energy will be approximately 100 times greater than present day devices greatly increasing the potential damage of these events. Relativistic runaway electrons generated by disruptions could severely damage internal components. Exacerbating the risk to the machine and increasing the engineering challenges, the disruption phenomena are often highly non-axisymmetric increasing local thermal and mechanical stresses.

B.5.3 Disruptions Goals

This science driver aims to obtain an improved predictive capability for the onset of disruptions to aid in avoidance through plasma control and in the development of algorithms for triggering disruption mitigation actuators, and to model the dynamics of mitigated and unmitigated disruptions in order to understand how to limit their effects. Achieving this goal would improve the viability of the tokamak as a practical energy source and enable the robust operation of tokamaks by allowing more aggressive operating regimes and by enabling faster recovery from off-normal events. The proposed science development roadmap was planned to enable the accurate prediction of 1) the onset of disruptions and how to avoid them, 2) the consequence of disruptions and how to mitigate those consequences. The specific questions that we wish to answer are:

1. How well can we predict the onset of a disruption and what strategies are available to avoid their occurrence?
2. How can we eliminate the instabilities that lead to the disruptions?
3. What are the effects of runaway electrons and what is the impact of operating regimes on their generation?
4. What is the impact of disruptions on the material wall, and how can we better design the first wall to handle the thermal loads?
5. What are the forces on the vacuum vessel and support forces during a disruption, and how do we improve their design?
6. How can we better design disruption mitigation systems?

B.5.4 Disruptions Roadmap

The tasks given in this section includes tasks that for experimentalists, computationalists, analysts, and theorists, and although these roles are denoted, close cooperation, including inclusion into the code design process, will be required for the success of this endeavor. We also make note of the milestone that each task can address.

Task 1: 1.5D Whole Device Model (WDM)

Year 1-2: Use WDM to simulate onset of VDE in extensive scan of experimental database. Benchmark instability and controllability threshold and early growth rate against linear stability calculations and experiment. Enable WDM codes to refine and perturb equilibrium as simulation progresses. Using existing models, investigate the generation of Runaway Electrons.

Year 3-5: Enable WDM code to launch ideal MHD codes for multiple toroidal mode numbers and analyze stability boundaries. Integrate WDM codes and extended MHD codes with Poloidal Field component to study interaction of feedback system on dynamics of disruption.

Task 2: 2.5D Whole Device Model (WDM)

Year 1-2: Couple neoclassical gyrokinetic code to 3D equilibrium solution with magnetic island, incorporating self-consistent bootstrap current from gyrokinetic code. Complete code modifications needed to couple $j_{\perp}$to 3D equilibrium solver to incorporate kinetic and flow effects. Simulate gas-jet penetration and pellet ablation in 3D in the pre-Thermal Quench phase, including radiation and parallel heat transport. Perform WDM simulations integrating in disruption mitigation techniques.

Year 3-5: Perform nonlinear simulations of Thermal Quench with 3D jet/pellet model and validate mixing/assimilation fraction versus species against experiment with improved impurity models. Interface extended MHD codes to more complete wall models, including newly developed components. Compute currents and forces induced in realistic 3D conducting structure, utilize reduced models of plasma material interaction to compute surface damage, impurity generation. Perform WDM simulations integrating in disruption mitigation techniques

Task 3: Extended MHD

Task 4: Component Development

Year 1-2: Develop a PF component capable of reuse by existing codes. Extend MHD component capabilities to include improved ability to model impurities, radiation, and simplified wall models. Improve existing wall component to enable three-dimensional wall effects, gaps, double-walls, and blankets.

Year 3-5: Develop models for runaway electron confinement and evolution during current quench – in particular transport of RE during successive transitions from stochastic to (partially) closed flux surfaces during q-evolution and evolving island overlap. Improve handling of kinetic effects and flow in 3D MHD equilibrium. Incorporate effects of turbulence in 3D equilibrium using gyrokinetic code that can handle turbulence. Replace VMEC equilibrium solver with a 3D equilibrium solver that can handle islands and stochastic regions using transport equations for flux diffusion outside and inside islands. Incorporate the improved 3D equilibrium in 2.5D code. Incorporate sources and transport model, including models for NTM triggers and momentum transport in 2.5D code.

Task 5: Validation

Year 1-2: Develop experimental database of disruptions and **a**nalyze to study consequences and effects. Improve validation of extended MHD codes by comparing with more localized measurements. Begin validation of modified equilibrium solver against saturated NTMs, using pressure and net current profiles consistent with experimental data. Validate thermal quench onset time against experiment.

Year 3-5: Statistically analyze the results of simulations of WDM with linear MHD analysis when applied to the experimental database. Validate initial extended MHD simulations against experiments. Validate extended MHD components against measured rotation evolution of bulk plasma, including propagation frequency and dynamics of mode observed in experiment and ECCD feedback stabilizations. Study rotation of tearing mode/resistive wall mode with external perturbations for cases with large perturbations and fast slowing down times in extended MHD codes. Determine verification and validation metrics for F-P, WDM, and extended MHD simulations of runaway electrons**.** Provide data from WDM/extended MHD disruption simulations to more detailed PMI and structural analysis codes and perform initial assessment and validation of effects of disruption on the wall. Validate 2.5D code for full time-evolution of NTM, including trigger threshold and momentum transport (including NTV terms). Validate 2.5D code against RWM time-evolution.

B.5.5 Summary of Disruptions Resource Estimates and Milestones

Resources (FTE/year) required to fully implement this program have been estimated as follows:

Task	Years 1-2	Years 3-5	Years 6-10	Years 11-15
Task 1: 1.5D WDM	8	6		
Task 2: 2.5D WDM	4	5		
Task 3: Extended MHD	10	10		
Task 4: Component Development	3	4		
Task 5: Validation	5	5		
Totals	30	30		

Table 20: Effort estimates for disruptions research

Suggested high level goals and milestones:

Disruption Detection, Avoidance, Mitigation and Consequences Milestone	Year from inception
Provide validated WDM capability for enabling predictions of VDE onsets with the uncertainties in the modeling quantified through validation.	2
Provide validated capability for using linear ideal MHD codes to predict the onset of fast MHD induced disruptions	2
Couple free-boundary, 3D equilibrium code with islands to neoclassical gyrokinetic code.	2
Provide validated capability to model gas-jet penetration and pellet ablation in 3D	2
Compare computed disruption forces using simplified models axisymmetric models of 2D walls with 3D models	2
Replace VMEC equilibrium solver with a 3D equilibrium solver that can handle islands and stochastic regions in the Strand-Houlberg 2.5D code.	3
Provide capability for WDM and extended MHD components to model the effects of feedback control through integration with PF components.	5
Investigate the extent to which impurities affect the nonlinear dynamics of the disruption.	5
Use 3D equilibrium codes (perturbed ideal MHD, stellarator codes) + transport codes to model thermal and particle transport in presence of island(s) and compare to non-linear kinetic MHD simulations	5
Improve TM/RWM simulations by using energetic particle and CEL-DKE closures. Include interaction with other modes such as Alfven eigenmodes.	5

Provide quantified analysis of WDM capabilities for modeling disruption mitigation experiments.	5
Model RE confinement/transport during current quench	5
Provide report on the quantified differences of extended MHD results using more realistic wall models and the simplified wall models.	5
Provide validated software for modeling disruption onset, mitigation, and consequences using the WDM modeling via integration with linear MHD, plasma feedback, external source, material wall, and structural wall components.	10
Provide new ability for studying the nonlinear, three-dimensional evolution of the instabilities that lead to disruptions, the consequences of the disruptions, and ways to mitigate their consequences.	10

Table 21: Milestones for disruptions research

B.6 Whole Device Modeling

B.6.1 WDM Background

The whole device modeling (WDM) of tokamak discharges involves the integration of different spatial and time scales for modeling of all discharge phases starting from discharge startup to discharge shutdown. High-fidelity predictive whole device modeling should accurately account for scrape-off layer physics, plasma wall interactions, core transport, heating and current drive, fast particles, pedestal physics, ELMs and impact on the divertor, 3D MHD modes, as well as other physics issues. The success of a new WDM tool is strongly dependent on careful coupling of different physics components. Due to interrelation among physical effects, strong coupling of many of the physics components will be essential. The interplay between different physics components introduces a new level of physics fidelity and should lead to discovery of new effects that are not evident when physics components are considered in isolation. WDM will depend critically on:

1. High fidelity science components;
2. Reliable and flexible framework that set standards for coupling of science components in the WDM suite of codes.
3. Verification and validation of individual physics components and WDM tool in general.
4. Data visualization, analysis, transport, and storage.

B.6.2 WDM Motivation

Whole device modeling is critical for development and optimization of scenarios on experimental devices. Especially for large, expensive facilities like ITER, the need to use run time efficiently and safely will demand extensive modeling as part of the experimental proposal process. WDM is also essential for design of new machine especially future burning plasma devices like CTF and DEMO. Important highly-coupled physics problems that will be addressed with FSP WDM suite of codes will include:

- Predicting the sources and sinks that drive all of the profiles in plasma discharges.
- Predicting the plasma confinement, transport and plasma profiles in tokamak discharges.
- Predict the onset, frequency and consequences of macroscopic instabilities including the plasma disruptions.
- Predicting the plasma boundary conditions ranging from plasma-wall interactions through the scrape-off-layer and the H-mode pedestal.

B.6.3 WDM Goals

The goals for WDM are congruent with the goals for the Fusion Simulation Project as a whole — to provide a comprehensive predictive simulation capability for magnetically-confined plasmas that integrates the knowledge from key multi-scale physical processes with continually improved fidelity. FSP WDM software must be designed to meet the following needs:

- Scenario modeling to plan new experimental campaigns in existing tokamaks and planned future devices.
- Analysis of experimental data to compute the time evolution of plasma profiles that are not measured and to resolve discrepancies between different ways of measuring experimental data.
- Validation or calibration of theoretical models by comparing simulation results with experimental data. Development of discharge control techniques
- Production of self-consistent simulation results that are passed on to other more specialized computer codes.

Depending upon the requirements of each simulation, the user should be able to choose from a spectrum of models for each physical process including high physics fidelity modes based on first-principles computations or reduced models for more rapid computations and validation or empirical models. There should be seamless access to experimental data or the results from previously-run simulations.

B.6.4 WDM Roadmap

The FSP Whole Device Model (WDM) schedule is broken down into four thrusts plus the central team and production system :

Task 1: 2.5d WDM & transport solver (3d equilibrium, 1d transport).

Year 1-2: 1D transport equations working with simplified sources and boundary conditions and reduced transport models. All reduced models in place for basic equilibrium and transport evolution. Componentization of VMEC and PIES. Computation in place for flux surface averages needed for 1d transport equations (as derived from the 3d MHD equilibrium).

Provide Flux diffusion calculation in 3D equilibrium. First 2.5d simulations with prescribed boundary 3d equilibrium and closed flux surfaces; benchmark with TASK 2.5d.

Year 3-5: Reduced models in place for MHD stability assessment and MHD island representation. Nonlinear MHD stability assessment working; ability to output 3d initial condition data to non-linear 3d MHD stability code. Source models adapted for 3d. 2.5d prescribed boundary simulation including magnetic island evolution, with adapted reduced transport models. 2.5d free boundary simulation with nested flux surfaces in the confinement region. Coupling to non-turbulent gyrokinetic code, simulations of neo-classical island evolution. Model ready for detailed validation against observed 3d effects in tokamaks and stellarators. Universal applicability of reduced models not expected.

Year 10-15: 3d edge/SOL/wall coupling. Core transport models specifically adapted for 3d. Core model useful for prediction of 3d effects in experiments, using input data for boundary conditions. Core/edge combined model useful for prediction of coupled core edge system behavior.

Task 2: WDM components for fast particle evolution and sources which take into account RF coupling to fast ions created by neutral beam injection and/or fusion reactions.

Year 1-2: Port of selected legacy WDMs. Verify componentization of important codes; improve as required. Install RF code / CQL3D and NUBEAM combination in legacy WDMs. Test using RF-SciDAC identified shot data.

Year 3-5: Production deployment and support in high performance FSP WDM.

Task 3: WDM components for evaluation of plasma turbulent transport on transport time scales.

Year 1-2: Development of at least one componentized solver module with access to all reduced transport models embedded. Componentized solver installed in legacy and/or FSP prototype WDMs, verified and available for production use. Establish kinetic-based reduced WBM with MHD perturbation included.

Year 3-5: Design for extension of solver/transport component to incorporate 1st principles transport models in both fluid-based and kinetic based reduced WBMs. Production deployment for reduced model validation. Transport time scale 1st principles simulation validated when overlapping MHD not present. Design for turbulence time scale 1st principles simulations embedment into WDM. Deploy available edge transport components.

Years 6-10: Verification and validation of transport time WDMs, with space-time embedded turbulence-time-scale first principles gyrokinetic codes. Development and testing of coupling techniques to 1st principles gyrokinetic turbulent transport components, as transport-timescale gyrokinetic formalism becomes available.

Year 10-15: Wide production use of transport-timescale WDM using 1st principles turbulent transport models.

Task 4: WDM that couples in edge and wall models (with successive fidelity).

Year 1-2: Initiate componentization of linear static Pedestal: Work with Boundary Layer effort on wall coupling. Support for production and validation applications.

Years 3-10: Explore reduced model methods for incorporating additional core-edge dynamics.

Incorporate higher fidelity pedestal components, including full turbulence models.

Task 5: WDM "central team" and production system.

Year 1-2 Establishment of software component standards review process. Development of data standards planning group: Adaptation of Plasma State or successor. Data system integration. Development of plan for extraction of components from legacy WDMs: First WDM validation capability on NERSC. First FSP component installation in legacy WDMs: Collaboration with base program funded legacy WDM teams. Production system documentation and user support. Code development to assist users (portal access, dashboards, etc.).

Year 3-5: Testing/development of high performance FSP WDM prototypes. FSP component review, installation, and testing. Production system documentation and user support. First production use of high performance FSP WDM.

B.6.5 Summary of WDM Resource Estimates and Milestones

Resources (FTE/year) required to fully implement this program have been estimated as follows:

Task	Years 1-2	Years 3-5	Years 6-10	Years 11-15
Task 1: WDM & Transport Solver	4	5	7	7
Task 2: Fast Ion Evolution	1.5	1.5	1.5	1.5
Task 3: Turbulence on Transport Time Scales	4	6	8	8
Task 4: Coupled Edge/Wall	1	1	1	1
Task 5: Central Team & Production System	3.5	5	5	5
Totals	14	18.5	22.5	22.5

Table 22: Effort estimates for WDM research

Suggested high level goals and milestones:

Whole Device Modeling Milestone	**Year from inception**
Integration of 3D equilibrium solvers	2
Adaptation of 1D transport equations in 3D magnetic geometry	3
Evaluation and development of algorithms to integrate first-principle turbulent transport tools with WDM simulation tools	3
Integration of reduced pedestal, SOL, divertor, and first-wall interaction physics component(s) for WDM applications	3
Installation and Integration of first-principle turbulent transport tools with WDM simulation tools	6
Integration of transport and 3D equilibrium components into 2.5D transport solver under WDM framework including optimization of 2.5D WDM transport tool for WDM parallel architectures	7
Self-consistent fast-particle treatment of heating and current-drive sources including integration with WDM simulation tools; verification of validation under WDM framework; optimization for WDM parallel architectures and demonstration for selected ITER applications	8
Installation and integration of self-consistent, coupled, core-edge dynamics components with more physics fidelity	8
Installation and integration with 3D core and pedestal transport models when they become available	10
Demonstration of selected ITER applications with first-principle turbulent transport models under WDM framework	
Installation and integration with 3D SOL, divertor, and wall interaction models when they become available	15

Table 23: Milestones for WDM research

Appendix C: Milestone Table

The following table contains a high level estimate of important deliverables and milestones relative to the start of the Program. The expected dates for many of the listed items should be understood as the time when they become viable entities requiring ongoing refinement and extension rather than finished products or completed activities.

Level 1 Milestones	**Date**
MGMT (WBS 1.4) FSP Substantially Operational	0.4 Years
ISA-2 (WBS 3.3.5) Release of improved 1.5 WDM code	1.6 Years
ISA-1 (WBS 2.2.4) First release of static model within FSP framework	2.7 Years
SIS (WBS 5.1.5) Reference implementation of On-HPC integration software, Release 1 (concurrent components, low-dimensional couplings)	2.7 Years
ISA-2 (WBS 3.2.8) First release of FSP 1.5 WDM code	3.4 Years
ISA-2 (WBS 3.5.5) Release code with gyro-kinetic turbulent simulations included	3.4 Years
ISA-1 (WBS 2.3.3) First release of coupled SOL model	4.4 Years
ISA-1 (WBS 2.5.3) First release of coupled kinetic SOL model	5.7 Years
ISA-2 (WBS 3.4.8) Release WDM code with 3D core and pedestal models	6.1 Years
ISA-2 (WBS 3.6.5) Release code with combined ISA1 and ISA2 components	6.9 Years
ISA-1 (WBS 2.6.2) First release of dynamic pedestal model	7.3 Years
Level 2 Milestones	**Date**
ISA-2 (WBS 3.9.2) 1.5-D profile evolution with boundary and recycling/fueling models	0.4 Years
ISA-2 (WBS 3.2.1) Identify legacy WDM 1.5D and WDM prototype codes and determine elements to be used in FSP	0.5 Years
ISA-2 (WBS 3.2.2) Determine initial FSP 1.5WDM applications	0.5 Years
SIS (WBS 5.5.1) Collaboration tools and repositories operational	0.7 Years
ISA-2 (WBS 3.3.1) Determine next level FSP WDM applications	0.7 Years
SQ (WBS 6.1.5) Definition of SQ policies and procedures	1.2 Years
SIS (WBS 5.4.3) Release 1 of software availability and validity monitoring packages	1.7 Years

SIS (WBS 5.1.2.3) I/O Library ready for use by ISAs	2.2 Years
SQ (WBS 6.7.3) Preliminary UQ analysis results	2.7 Years
SIS (WBS 5.4.2) Release 1 of software for standard visualizations	3.2 Years
SIS (WBS 5.1.3) ISAs refactored to use I/O layer	3.2 Years
ISA-2 (WBS 3.6.1) Identify and install reduced pedestal, SOL, divertor, and first-wall interaction physics component for WDM applications	3.8 Years
ISA-2 (WBS 3.6.3) Optimize for WDM Parallel architectures	4.1 Years
SIS (WBS 5.1.6) ISAs refactored to use SIS implementation layer for 1-2D couplings	4.2 Years
SQ (WBS 6.7.6) UQ analysis result Phase II	4.7 Years
SQ (WBS 6.7.9) UQ analysis result Phase III	4.7 Years
ISA-1 (WBS 2.7.4) Nonlinear/coupled micro-mesoscale transport	4.7 Years
SIS (WBS 5.6.3) Automated builds of complete software suite	4.8 Years
ISA-1 (WBS 2.8.3) Nonlinear evolution of ELMs and crash dynamics	4.9 Years
SIS (WBS 5.6.4) Test systems applied to ISA 1,2 and first components	5.7 Years
ISA-2 (WBS 3.4.6) Install and integrate 3D core and pedestal models	5.7 Years
ISA-2 (WBS 3.11.2) Dynamics from stability onset to disruption	5.8 Years
PHYS (WBS 4.3) Component library adaptation	6.3 Years
PHYS (WBS 4) ET - Components	6.3 Years
ISA-1 (WBS 2.10.1) Heat and particle impact on wall and back reaction during ELMs	7.6 Years
ISA-2 (WBS 3.12.2) Control of evolution for disruption avoidance	7.8 Years
ISA-1 (WBS 2.8.5) Externally applied 3-D magnetic field effects	8.9 Years
ISA-2 (WBS 3.9.3) 2.5-D profile evolution with boundary transport, sources and sinks	9.1 Years
ISA-1 (WBS 2.10.2) Heat and particle load control and impact on core performance	9.3 Years
Level 3 Milestones	**Date**
ISA-2 (WBS 3.2.3.2) Determine modules that will be loosely coupled	0.5 Years
ISA-2 (WBS 3.2.3.1) Determine modules that will be tightly coupled	0.5 Years
ISA-2 (WBS 3.2.3.3) Identify gaps in component models to meet WDM goals	0.7 Years
PHYS (WBS 4.3.3) Embedded turbulence model Part I	0.9 Years
PHYS (WBS 4.3.1) Free boundary Grad-Shafranov solver Part I	0.9 Years
SQ (WBS 6.2.4) Definition of verification plans	0.9 Years
ISA-2 (WBS 3.2.6.1) Integration of tightly coupled components	0.9 Years
ISA-1 (WBS 2.9.2) Private flux region heat and particle transport	0.9 Years

SQ (WBS 6.3.3) Release of SQA tool suite	1.2 Years
ISA-2 (WBS 3.2.4.1) 1D core transport model	1.3 Years
ISA-2 (WBS 3.2.4.2) Pedestal equilibrium model	1.3 Years
ISA-2 (WBS 3.2.4.4) NBI model	1.3 Years
ISA-2 (WBS 3.2.4.3) RF heating and current drive model	1.3 Years
ISA-2 (WBS 3.2.4.5) Fueling model	1.3 Years
SQ (WBS 6.2.7) Definition of validation plans	1.4 Years
ISA-2 (WBS 3.10.1) VDE disruption model and control	1.4 Years
ISA-2 (WBS 3.2.6.2) Establish connection with loosely couple components	1.5 Years
SQ (WBS 6.2.10) Definition of UQ plans	1.7 Years
ISA-2 (WBS 3.5.1) Identify and evaluate gyro-kinetic turbulent transport simulation tools for WDM application	1.7 Years
PHYS (WBS 4.3.4) Embedded turbulence model Part II	1.9 Years
PHYS (WBS 4.3.2) Free boundary Grad-Shafranov solver Part II	1.9 Years
ISA-1 (WBS 2.8.1) ELM onset	1.9 Years
ISA-2 (WBS 3.8.2) 3-D free boundary equilibrium	1.9 Years
ISA-2 (WBS 3.3.3) Introduce more advanced and first principal components	2.4 Years
PHYS (WBS 4.3.5) Linear stabilities (MHD)	2.9 Years
SQ (WBS 6.6.2) Release of verification tool suite	3.7 Years
PHYS (WBS 4.3.6) Energetic particle/kinetic stability modules	3.9 Years
SQ (WBS 6.8.2) UQ Tool suite Initial release	4.3 Years
PHYS (WBS 4.2) Team standing-up and continuous functions	4.9 Years
SQ (WBS 6.4.3) Synthetic diagnostic release I	4.9 Years
ISA-2 (WBS 3.4.4) Adopt and assess reduced transport models in 3D magnetic geometry	5.5 Years
ISA-2 (WBS 3.4.5) Adopt and assess simplified source and loss models in 3D magnetic geometry	5.7 Years
SQ (WBS 6.8.4) UQ Tool suite second release	6.2 Years
ISA-2 (WBS 3.6.2) Install and integrate components with higher physics fidelity	6.3 Years
PHYS (WBS 4.3.7) Refactoring pedestal/SOL equilibrium module from ISA1 for WDM	6.3 Years
SQ (WBS 6.6.4) Release of revised verification tool suite	6.6 Years
SQ (WBS 6.4.5) Synthetic diagnostic release II	7.8 Years

Table 24: FSP Milestone Table

Appendix D: Verification, Validation, and Uncertainty Quantification

Verification, Validation, and Uncertainty Quantification are important components of a comprehensive strategy to develop an integrated suite of predictive computer models for complex magnetic fusion plasma behavior. These activities are fundamental to the establishment of confidence in simulation capabilities and ultimately will lead to a more synergistic relationship between computation and experiment. As such, verification, validation, and uncertainty quantification must be an integrated aspect of the FSP.

The quantitative comparison of the physics model to experiment is limited without a quantitative understanding of the accuracy and errors in the computations that approximate the physics model as well as a good understanding of the errors embodied in the experimental measurements. One can validate the computer calculation with respect to experiment but be misled in understanding the quantitative comparison of the underlying physical models, unless one has a good understanding of the uncertainty in the simulation results arising from uncertainty in input data and parameters required by the numerical and physical models.

There are many different kinds of errors and uncertainty that affect the embodiment of the physics model in a computer code and that can potentially affect the comparisons of the simulation results to experimental observations. These include, but are not limited to:

- systematic and stochastic measurement error;
- limitations of theoretical and phenomenological models, i.e., what phenomena are included and what are omitted;
- limitations of the representations of the data and models (how models and data are simplified for use in analysis);
- accuracy of the data used in physical models
- accuracy of computation and approximation, including both deterministic and stochastic approximation errors, iteration errors, and finite-precision arithmetic errors;
- random phenomena (aleatoric uncertainty).

For the FSP to provide a robust predictive capability that can be confidently used in decision making processes, the errors and uncertainties from simulations must be identified, quantified, and, where possible, reduced. Concentrated efforts in verification, experimental validation, and uncertainty quantification are therefore critical for the success of the FSP.

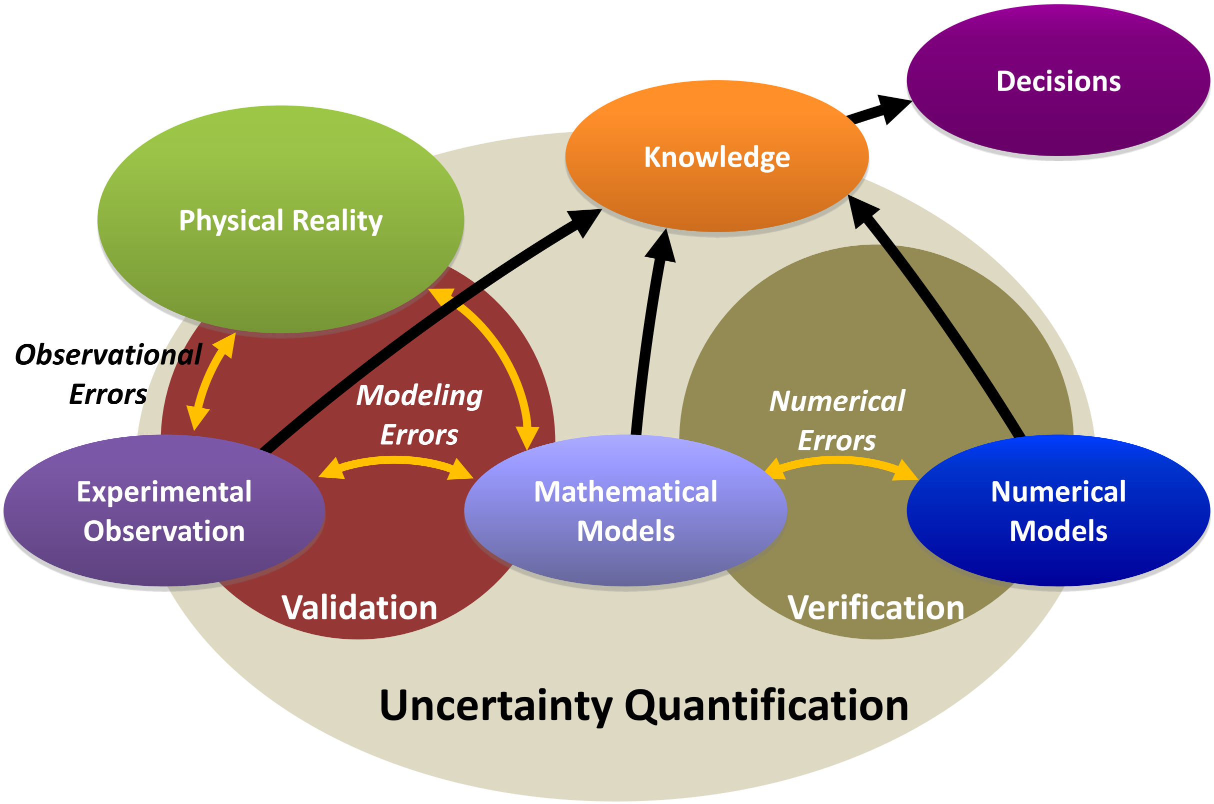

Figure 15: Relationship between experiments, mathematical models, numerical models, errors and processes.[1]

D.1 Overview of Activities

For predictive simulation, model and data uncertainties must be propagated through computational models systematically in order to obtain quantities of interest with quantified uncertainty. Uncertainty Quantification (UQ) is the quantification of all sources of uncertainty that may affect such an inference. Quantities of Interest (QoIs) are those quantities to be obtained from a simulation or an experiment. Sensitivity Analysis is intimately coupled to uncertainty analysis. In deterministic methods, sensitivity analysis is done first and used to conduct an uncertainty analysis; the opposite is the case for statistical approaches. In the forward problem, the sensitivity of a result to specific input data, e.g., initial conditions and model parameters, is quantified. There is also an inverse sensitivity analysis problem in which the sensitivity in input parameters and data from experimentally measured results is inferred. This is the case in Calibration, where experimental data are used to tune model parameters.

In the context of the FSP, we are primarily but not exclusively concerned with the comparison of simulations and physical models to experimental data, that is, Validation. UQ techniques can be applied to establish the usefulness (validity) of a code or model within a particular use or regime based on quantitative comparison with experimental data. Both UQ and sensitivity analysis apply equally well to the scientific results of both experiments and simulations; experiments possess observational errors. QoIs from both types of activities require confidence intervals to evaluate the degree of agreement properly. Sensitivity analysis results can be used to identify those data for which better knowledge would provide the most leverage in affecting the comparison to experiment. This is one use of uncertainty and sensitivity analysis, but by no means the only.

The physical model is typically a set of ordinary or partial differential equations that themselves are an approximation of physical reality. Ideally, physical models are compared directly with experimental results.

---

[1] Adapted from T. Oden *et al* ., ”Computer Predictions with Quantified Uncertainty, Part I”, SIAM News **43**(9) 1 (2010).

When numerical results are compared with experimental results, the code is being used as a surrogate of the physical model. Thus, errors due to this difference must be carefully included in any UQ estimate for simulation data. Verification processes play two roles here: first, to demonstrate that the code approximates the intended physical model and second, to estimate the numerical errors inherent in any computed result.

A depiction of many of these activities, in relation to various sources of error and uncertainty, can be found in Figure 15. In the pursuit of knowledge, we have three ways to investigate reality: experimentally, theoretically, and computationally. Experimental observation is most directly connected to physical reality, but there are observational errors inherent in experiment. Mathematical models developed in theoretical investigation are based off of specific experimental results and/or more general physical principles. The nature of modeling is that some level of physical detail is neglected or approximated, leading to modeling errors. The mathematical models resulting from theoretical development are not in general solvable in closed form, and thus numerical models that generate approximate solutions are used. Of course, the process of discretizing the mathematical models and solving them in discrete arithmetic leads to another level of approximation error. These discrepancies are identified and associated with the roles of validation and verification. Note that UQ can be thought of as a much larger activity that can play a broader role.

A depiction of a rigorous validation process with integrated verification and UQ is provided in Figure 16. From the problem definition, parallel experimental and theoretical/numerical activities are undertaken. At first, these efforts might collaborate directly, for instance, in pre-test studies and for calibration. Once in the validation mode, however, there should be a separation between the validation effort and the simulation effort. Note the role of code verification in the numerical model development, in contrast to the role of calculation verification and UQ in the final analysis of simulation results. Validation does not require the UQ/calculation verification analysis step; indeed, science has successfully made great progress without quantified uncertainty in theoretical/simulation results. The inclusion of these analyses improves not only the quantitative nature of the validation, but also provides additional information on subsequent studies to pursue to improve the experimental data and/or the theoretical/numerical models. Finally, although we only explicitly identify UQ in the final step of experimental and numerical predictions, UQ activities can occur in many other places in the process (e.g., in the calibration, model or experiment revision, etc.)

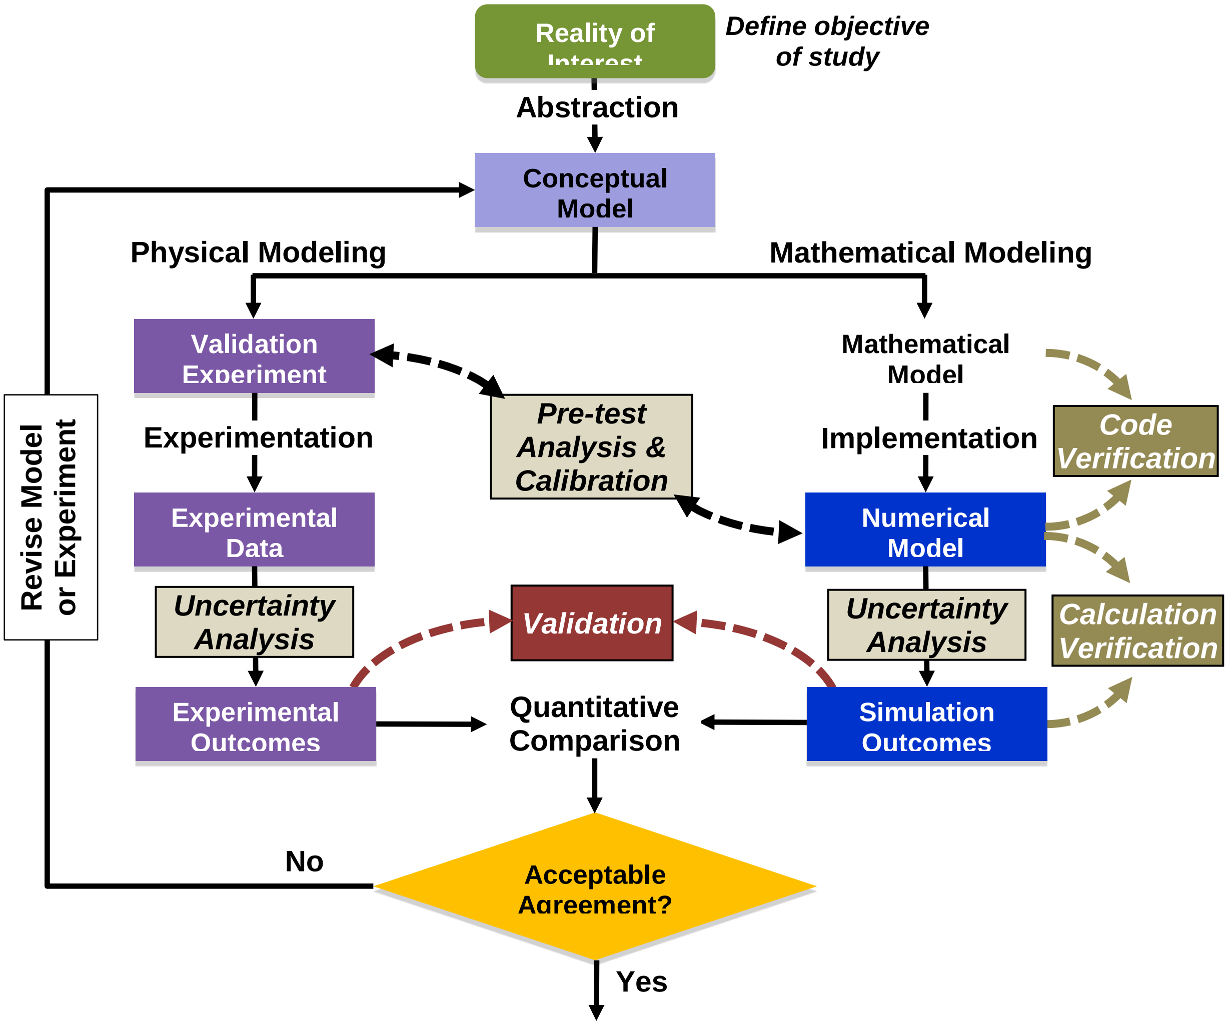

Figure 16: Representative validation process including a systematic verification and UQ component.[2]

D.2 Verification

All software contains errors. Potential sources of such error include: discretization errors from the discrete approximation of continuous functions; residual errors from iterative algorithms; finite-precision round-off errors; and programming errors (bugs). The purpose of verification in the context of simulation software is to find these errors (and eliminate bugs), to demonstrate that they behave as expected, and to estimate the size of errors in simulations results.

There is a well-accepted sub-division of verification effort into two activities: code verification and calculation (or solution) verification. The former addresses the questions of code correctness and order verification. The latter is an application of *a posteriori* error estimation techniques to make quantifiable error estimates for simulations with unknown solutions. Numerical error estimates in quantities of interest for computed solutions are needed to compare simulation results properly with experimental results.

D.2.1 Mission and Goals

The mission of the FSP verification activity is to produce correct codes with quantified error estimates in simulation results. The associated goals are to

- Find and eliminate programming errors in scientific software based on the discretization of mathematical models;

---

[2] Adapted from B. Thacker et al ., LANL Tech. Report LA-14167-MS, 2004.

- Demonstrate and document convergence to reference solutions for these mathematical models at the expected rates;
- Develop and provide the necessary tools and techniques required to provide reliable a posterior error estimates for simulation components and integrated simulations

The FSP must develop a verification strategy appropriate for individual components as well as for their integration. Key issues that the FSP verification effort must address include:

- Policies for verification: What are the scopes and levels of rigor of code and calculation verification required in the FSP? What are the acceptable techniques for code verification? How are the results recorded and distributed?
- Verification workflow: What are the processes for defining and conducting verification studies?
- Defining roles and responsibilities: What is the right level and form of independence in the verification approaches taken by ISA and advanced component efforts? To what degree are methods and techniques standardized across the FSP? Who has responsibility for executing verification activities and who evaluates these activities? Where does research and development of new verification techniques and tools occur within the FSP?
- Data management and documentation: How can data management methodologies be used for minimizing errors in the definition and execution of verification studies, and how do we best document and disseminate results of verification studies?

To develop this verification plan, the FSP planning effort undertook several activities. First, a small team with fusion, numerical analysis, and UQ expertise was tasked with writing a review document on common verification and UQ procedures. To determine the state of verification within the fusion community, input was sought from the six science driver efforts. The feedback was limited, indicating that at least some level of benchmarking, if not genuine code verification, occurs within fusion code projects. A proposal for an integrated verification and UQ effort was then written to provide a concrete basis for evaluation by verification and fusion simulation experts at the second FSP planning workshop. The findings included:

- Some level of code verification or at least benchmarking already occurs within most if not all of the fusion code projects, but the verification processes are not always systematic or extensive;
- The fusion community has not engaged in formal calculation verification activities, but there is a strong interest in doing so. Techniques such as the method of manufactured solutions could be powerful, but may be difficult to apply to problems with complex geometries.
- A systematic approach to documentation and archiving of the potentially vast quantities of data generated is important. A dynamic database for archival was proposed.
- The application of *a posteriori* numerical error estimation to fusion calculations for the purposes of calculation verification is an open research topic. New component development within the FSP should attempt to accommodate some form of a posteriori numerical error estimation.

Incorporating these findings, we present a strategy for each of the identified verification issues..

D.2.2 Policies for Verification

There are no verification procedures that will work for all applications. Verification procedures will in general be easier to apply to individual components than to integrated, multi-physics applications. We thus envision a tiered approach to the deployment of verification techniques and differentiate between policies and procedures for component development and for integrated application development. In addition, we in general favor policies that emphasize the required results over specific means to obtain these results.

D.2.3 Component development

The focus in component development will be on documentation, code verification through order verification, and calculation verification. It will be the policy of the FSP that component development projects will have the freedom to define and select appropriate test problems and methodologies, but component projects should attempt to apply the most rigorous techniques available.

Documentation must include the governing mathematical model; the discrete approximate model; a priori error estimates; reference solutions; quantities of interest relevant to major use cases; all numerical and model parameters; and all test results. Documentation also includes any scripts and input files used to generate the simulation results. We anticipate that the documentation will be saved and distributed within the framework established by the integrated data management policy outlined in Section 4.3.6.

It is recognized that rigorous code verification will be a challenge due to the complexity of the problem (in particular the coordinate system) and that benchmarking will continue to be a useful confidence-building tool. Nevertheless, the FSP must be diligent in its efforts to move to more rigorous code verification practices. It is preferred that code verification will be done rigorously through order verification, that is, a demonstration that the component converges to a reference solution at an expected convergence rate.[3] The most common technique for order verification is two-point Richardson extrapolation through grid refinement studies on reference problems. A suite of reference problems and their solutions will be created that can include known analytic, asymptotic, or manufactured solutions. The latter types of solutions are preferred, since these can be designed to exercise more complicated couplings in the code.[4] In addition, surrogate error tests based on intrinsic properties (invariants) of the discrete solution, e.g., conservation, where appropriate, should also be included in the component verification test suite. Well-defined code verification tests should be integrated into automated regression testing suites.

In the first few years of the FSP, calculation verification is expected to rely heavily on grid refinement and three-point Richardson extrapolation applied to sequences of four or more grids for *a posteriori* error estimation.[5] A suite of reference problems (without solutions) that demonstrate important behaviors of the component will be created; important quantities of interest for these reference problems will also be identified and defined. Component teams will be encouraged to incorporate more other, potentially more robust *a posteriori* error estimators, such as adjoint methods[6] or error transport methods[7]. Benchmarking (code-to-code comparison) and self-convergence (Richardson extrapolation using a fine-grid solution in place of the exact solution) will be used to build confidence, even if they are inadequate for proper calculation verification because they do not provide error estimates. In later years, the inclusion of *a posteriori* error estimators into new component development should be a goal.

D.2.4 Integrated application development

For integrated applications, rigorous verification becomes more difficult for several reasons. Fewer exact or approximate solutions are known, and manufactured solutions are more difficult to devise and implement. Robust *a posteriori* error estimators are more difficult to develop for integrated models. In addition, there are open research issues on how to develop composite error estimates from a configuration of components, which could possibly be dynamic.

As with component development efforts, model documentation (for the integrated application) and code and calculation verification, to the extent that these are possible, will be emphasized. Components to be coupled in an integrated application must have already undergone thorough code and calculation verification as described

---

[3] P. Knupp *et al* ., “Measuring progress in order-verification within software development projects,” *Engin. Comput*. **23**: 271 (2007).

[4] P. Roache, *Verification and Validation in Engineering*, Hermosa Publishers, Albuquerque, 1998, Chapter 3.

[5] *Ibid.*, Chapter 5.

[6] M. Giles and E. Süli, “Adjoint methods for PDEs: *a posteriori* error analysis and post processing by duality,” *Acta Numerica* **11**, 145 (2002).

[7] A. Hay and M. Visonneau, “Error estimation using the error transport equation for finite-volume methods and arbitrary meshes,” *Int. J. Comput. Fluid D.* **20**, 463 (2006).

above. Benchmarking and self-convergence will, in the early years of the FSP, play a key role in the absence of more rigorous code and calculation verification capabilities. In later years, it should be a goal of the FSP to include more robust *a posteriori* error estimators in integrated components for more rigorous code and calculation verification.

D.2.5 Workflow

The following outlines a best practice workflow for verification in the FSP:

- Define and document the test problem. This includes the mathematical model, the boundary and initial conditions, the geometry, and the solution, including an expression that bounds any errors for asymptotic expansions.
- Define and document the discrete problem. This includes the discretization, the solvers used, the discrete initial and boundary conditions, and the discrete mesh specification.
- Define and document the quantities of interest. These are generally functions of the solution that may be either problem-specific, e.g., a shock speed, or numerically motivated, e.g., a discrete norm of the solution. The methods for calculating such metrics should be documented as well.
- Identify suitable values for all numerical parameters. Numerical parameters include iteration parameters and any other parameters that depend on the discretization (mesh or timestep) parameters. Iteration parameters should be chosen to minimize iteration residuals. This step is non-trivial for certain types of algorithms, e.g., Arbitrary Lagrangian-Eulerian schemes.
- Execute the simulation. This is typically automated, and all runs should be done on the same machine with the same parallel decomposition.
- Analyze the results. In this step, the chosen verification methodology is applied.
- Document the results.

In this general form, the process applies equally well to both code and calculation verification. Code verification should always be treated as a pre-requisite to calculation verification. It is recommended that to promote accuracy and efficiency, documented workflow scripts or programmable workflow tools should be used to generate and drive any verification studies.

D.2.6 Roles and Responsibilities

There are three broad activities that fall under the scope verification:

- Development of methodologies
- Development of technologies
- Application of technologies and methodologies

Within the available resources of the FSP, the focus will be on the application of techniques, but there will be activity in the other two areas of research and development. For example, *a posteriori* error estimation for multi-physics simulations is an active research area.

Verification activities will take place throughout the FSP organization. A small, dedicated Software Quality (SQ) team will exist to focus on the development of technologies and, to a lesser extent, development of methodologies. This team will have the responsibilities to further specify program-wide policies and procedures and to provide expert guidance to the rest of the project. Each component development and ISA team will have an identified SQ point-of-contact and will define a verification plan appropriate to their task within the policies put forth by the SQ team. These strategies will be developed in consultation with the SQ team, who will also aid in the execution of the plans. The Software Integration team will work with the SQ team to adapt and develop program-wide infrastructure to support verification, such as a workflow tool that enables scripted study execution and allows for the incorporation of incorporates analysis plug-ins. Research into new methods suitable for FSP activities will take place within the SQ team. The relationships and responsibilities are summarized in Table 25.

For the purposes of verification, SQ point-of-contacts should:

- Have a strong grasp of numerical analysis and computation (running codes on HPC)
- Have a good grasp of the underlying physical theory and the nature of numerical errors
- Be able to work closely with SQ team
- Have experience with code verification
- Be able to implement a posteriori error estimators

For the purposes of verification, the SQ Team should have members with:

- A strong grasp of numerical analysis and computation (running codes on HPC)
- Experience with code and calculation verification
- Experience in the development of a posteriori error estimators
- The ability to work closely with members of application teams
- The ability to collaborate with software teams to develop verification infrastructure

Team	Responsibilities
SQ	● Define verification policies and procedures ● Work with component development and ISA teams and validation analysts in the selection and application of verification methodologies to their activities ● Evaluate the progress of verification efforts within the program ● Adapt and develop technologies, with the help of the Software Integration team, to provide program-wide verification infrastructure ● Evaluate external research results in verification for applicability to the FSP ● Develop new verification methodologies for the FSP as needed
Component or ISA	● Identify a point-of-contact to the SQ team ● Develop a code and calculation verification plan specific to their project in consultation with the SQ team ● Execute verification plans in collaboration with the SQ team as per the best practice workflow ● Provide feedback to the SQ team about needs and difficulties encountered in the execution of the verification plans
Software Integration	● Adapt and develop technologies, with the help of the SQ team, to provide program-wide verification infrastructure

Table 25: Verification software quality responsibilities within the FSP

D.2.7 Data Management and Documentation

As highlighted in the previous sections, the best practices for verification have implications for data management and documentation needs. In particular, tools are required to make the results of verification planning activities, exercises, conclusions available to all interested parties. The FSP public documentation should include the verification documentation as described above for individual components and integrated multi-component predictions for each science driver. The FSP data management plan is discussed in Section 4.3.6.

D.2.8 Research Opportunities

There are many open issues in the application of verification techniques to large-scale simulation codes. The following possible research topics have been identified

- Robust *a posteriori* error estimators
- Combination of error estimates between coupled components
- Determination of errors due to coupling
- Error estimation in the presence of models for unresolved physics
- Error estimation for multi-scale problems
- Error estimation for solution-driven model changes

Resources and therefore verification method research will be limited within the FSP, so the FSP will need to engage external research and development efforts (ASRC, ASC, NSF, etc.) for new tools and techniques.

D.3 Experimental Validation

D.3.1 Mission and goals of experimental validation

The Fusion Simulation Program (FSP) aims to provide the capability to confidently predict toroidal magnetic confinement fusion device behavior with comprehensive and targeted science-based simulations of nonlinearly-coupled phenomena in the core plasma, edge plasma, and wall region on time and space scales required for fusion energy production. The mission of the FSP experimental validation activity is to assess and improve physical and computational models by systematic, quantitative comparisons with experimental measurements. The associated goals are to:

- Develop and provide the necessary tools and documentation to allow FSP users/customers/stakeholders to determine what level of confidence they will give to predictions made by an FSP simulation
- Provide clear assessments of model and component physical fidelity to help guide their refinement and improvement.

The FSP must develop a validation strategy appropriate for individual components as well as their integration. Key issues the FSP validation effort must address include:

- Validation workflow: What is the process for assessing model and integration readiness for validation, identification of key phenomena/physics (including coupled physics) to be modeled, identification of key model sensitivities (also, the most sensitive parts of a model), readiness of the interface between codes and experiments i.e. diagnostics including synthetic diagnostics, and design and conduct of any needed validation experiments?
- Validation metric development: How are relevant primacy hierarchies identified, how should model-experiment (dis)agreement be quantified, and how do individual component (“unit problem”) metrics connect to integrated multi-component metrics?
- Defining roles and responsibilities of validation analysts: What is the right level and form of analyst independence, what modes of interaction with experimentalists and modelers should be used, and what are the analyst career development and support mechanisms?
- Data management and documentation: How can data management methodologies be used for minimizing errors at the code-experiment interface, and how do we best document and disseminate results of validation studies?

Based on lessons learned (http://www.psfc.mit.edu/FSP-Validation/index.php/Notes) from various validation studies and applications of predictive modeling within and outside of MFE research for each of the issues listed above, we come up with the following guideline for implementation of the FSP experimental validation.

D.3.2 Workflow best practices

A functional workflow for the "verification-validation-prediction" should have all the elements in Figure 4 above. The Verification and Validation (V&V) Planning Activities box is a crucial step to complete and check off before embarking on the Validation Experiment Execution. We have expanded that box to indicate five necessary components as suggested by the Phenomena Identification and Ranking Table (PIRT) approach identified in Oberkampf *et al*[8]. They strongly argue that the use of a PIRT is the most important tool in the validation planning activity for translating application requirements into prioritized V&V activities and ensuring the readiness to proceed to validation experiments. A complete PIRT should document

The key physics quantities and processes relevant for the specified application:

- Whether the conceptual model(s) to be tested is sufficient to describe the key physics of the specified application
- The verification and uncertainty quantification (UQ) methods and tests needed to ensure the accuracy and adequacy of the code solution(s)
- Needed critical experiments and associated diagnostic capabilities (including synthetic diagnostics) that will provide the key data for use in testing the model
- The quantitative metrics to be used in assessing the fidelity of the model results to the experimental validation measurements.

These points are expanded below using the pedestal ISA as an example.

D.3.2.1 Critical physics hierarchy

Critical physics is the essential phenomena that must be input to, and described by, the application code. Each phenomenon should be prioritized relative to others, and an explanation of this prioritization included in the PIRT document. In the course of answering these questions, a validation hierarchy should be constructed to help identify a range of experiments, possible separation of coupled physics, and levels of complexity relevant to the modeling of the ISA in question. A possible sample hierarchy for pedestal physics is shown in Figure 17.

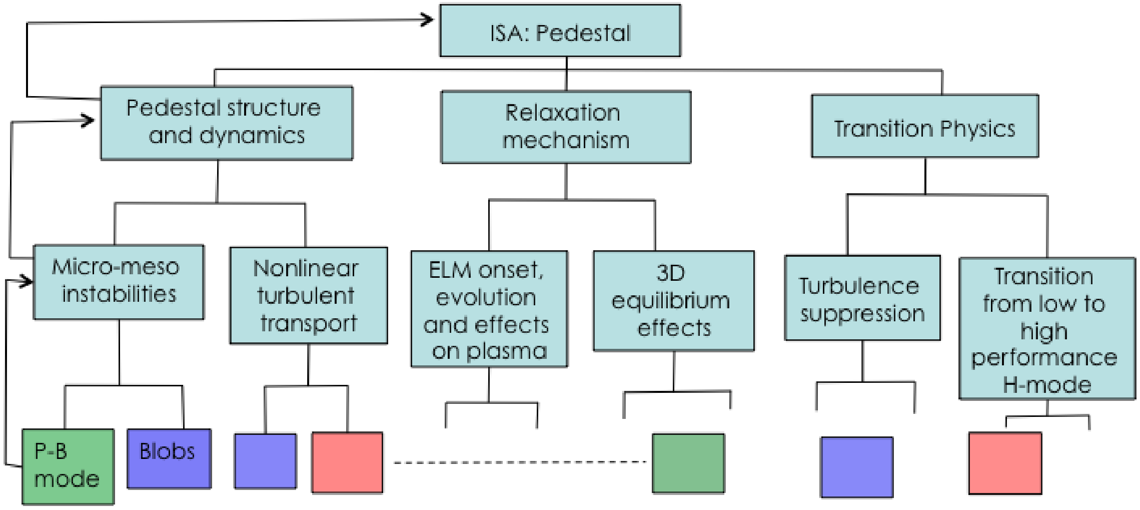

Figure 17: Example of a critical physics hierarchy for pedestal ISA*

**Uses information identified in the validation assessment tables (https://lce.txcorp.com/trac/2011_FspDefintionWorkshop).*

---

[8] W.L. Oberkampf, T.G. Trucano, C. Hirsch, "Verification, validation, and predictive capability in computational engineering and physics," *Applied Mechanics Reviews* **57** 345 (2004)

It should be emphasized that the critical physics hierarchy be extended to the finest granularity required for validation and understanding. In the example in Figure 17, the finest granularity is at the level of a single physics phenomenon, such as the peeling-ballooning (P-B) mode and density blobs. These are phenomena that can occur in isolation or in combination in an experiment. Hence it is logical to study them separately as a start. We also know when these phenomena appear together, they impact the plasma behavior differently. Advancing to the next level in the hierarchy to validate the coupled phenomena offers an opportunity to understand and confirm new physics at the micro-mesoscale. Adding other microinstabilities such as kinetic ballooning modes and electron temperature gradient modes, as well as their nonlinear evolution would move the validation to the next hierarchy where the pedestal structure and dynamics averaged over fast time would be the target. Finally combining the physics of different time scales, namely, the pedestal structure and dynamics, relaxation mechanisms, and transition physics would provide the full description of the pedestal behavior.

D.3.2.2 Code adequacy and readiness

Having developed the critical physics hierarchy, we need to similarly prepare the codes to describe the physics following the same hierarchy. At each step, we ask: Are the codes adequate and ready? This question entails several aspects. The first is conceptual model adequacy: Are the underlying conceptual models (e.g., ideal MHD, DF gyrokinetics, Branginskii equations, etc.) adequate for describing the physical phenomena of interested identified in (1)? As part of this process, one should identify the most sensitive or uncertain part of a model e.g., use of a quasilinear approximation, zero-orbit width, r/L <<1, equilibrium shape and accuracy. While this assessment is in some ways an output of the validation process, one should ensure that any phenomena under consideration do not violate assumptions made in the derivation and development of the model before beginning.

The second is code verification adequacy: Have the components and frameworks, which represent the numerical implementation of the conceptual models of interest, undergone adequate verification testing? One can similarly set up a flow chart and decision tree to evaluate the codes adequacy and readiness as illustrated in Figure 18. It shows for the same hierarchy of critical physics indicated in Figure 17, that one has to identify the application codes that theoretically model the physical phenomena in the regime of interest. Each application has to have gone through a rigorous process of code verification. Furthermore, key physical quantities computed by the applications that would be used to compare with experimental observables should be clearly identified.

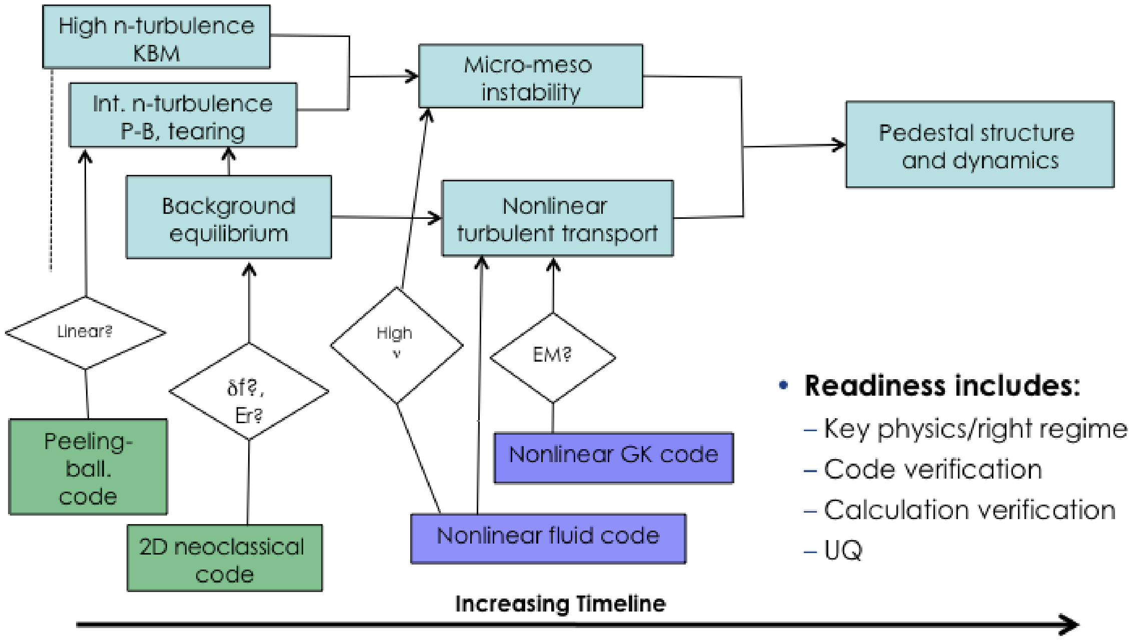

Figure 18: Logistics for determining adequacy and readiness of code applications

D.3.2.3 Best practices for designing validation experiments

In order to fully validate FSP components and frameworks, it is clear that new experiments will be required. These experiments should also fit in a hierarchical structure. The purpose of the validation hierarchy is to help identify a range of experiments, based on possible separation of coupled physics and level of complexity. A good hierarchical construction also enables selection of individual experiments in tiers that are practically attainable and capable of producing validation quality data, i.e. it allows precise and conclusive comparisons. The hierarchy should also help identify opportunities for small-scale experiments to test fundamental physics with enhanced controllability and characterization.

When designing these experiments, the following suggestions should be kept in mind:

- The experiments should be designed to test the most critical or uncertain parts of the physical models
- The physics assumptions should be well documented
- The experiments should help quantify region of parameter space for which conceptual model is acceptable for the application purposes
- The experiments should allow precise and conclusive comparisons of calculations with experimental data
- Determine what measurements (and diagnostics) are needed and at what accuracy and resolution
- Carry out code runs and experiments as independently as possible to minimize prejudicing the outcome
- Ensure synthetic diagnostics are available for apples-to-apples comparisons between code results and measurements
- Pay special attention to analysis of errors and uncertainties
- Document the entire process and results, including data reduction techniques and error analysis.

D.3.2.4 Validation metric adequacy

An essential component of the validation process is the development and application of metrics for quantifying model fidelity. Validation metrics should include the following properties:

- An explicit estimate and inclusion of numerical error in the model calculation, exclusion of such error if it can reasonably be shown to be small
- The metric should be a quantitative evaluation of the predictive accuracy of the system response quantity (SQR) of interest
- The metric should incorporate, or at least explicitly include, an estimate of the errors in the measurement(s) the model is to be compared against
- The metric itself should exclude any indications, either explicit or implicit, of the level of adequacy in agreement between computational and experimental results.

Metrics are fundamentally algorithms for quantifying agreement between a model and observations. They can be as simple as a chi-squared test of a prediction to a set of measurements, or highly complex and nonlinear functions that incorporate many model predictions and experimental data points. A commonly used terminology is to identify a “simple metric” as one that assesses the level of agreement between a single predicted field or observable and its corresponding measurement, whereas a “composite metric” is a (often weighted) combination of simple metrics. of particular note is with the adoption of the hierarchical approach for FSP validation, it is imperative that we squarely address the use of composite metrics at the integrated physics validation level. Take Figure 19 for instance; a prediction of the turbulent particle flux depends upon predictions of density and radial velocity fluctuation amplitudes, as well as the coherency and cross-phase of those fluctuations. Thus one might give more weight to comparisons of the individual components (e.g., fluctuation amplitudes, coherency, cross-phase) than the integrated product (the flux) when constructing a metric. However, these should be related in some way with justifications.

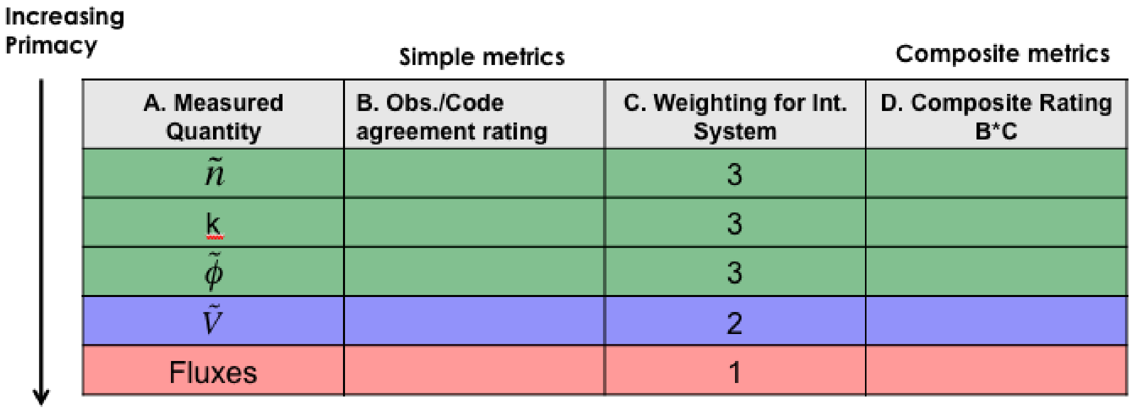

A. Measured Quantity	B. Obs./Code agreement rating	C. Weighting for Int. System	D. Composite Rating B*C
$\tilde{n}$		3	
k		3	
$\tilde{\phi}$		3	
$\tilde{V}$		2	
Fluxes		1	

Figure 19: Illustration of a composite metrics table relevant for primacy hierarchy validation

D.3.3 Collaboration with experimental facilities

The success of FSP validation hinges on building a true partnership between FSP and the experimental facilities. Both have to recognize there are significant benefits towards advancing the fusion energy goals by sharing resources and making a commitment to put experimental validation at a high priority programmatically. A document (https://ice.txcorp.com/trac/2011_FspDefinitionWorkshop) describing how the partnership should be implemented has been prepared based on ideas drawn from existing collaboration agreements used by the three major fusion experimental facilities, and their governing and planning processes. The facility management has reviewed this and provided constructive feedback. The document outlines:

- General principles for collaboration and IP sharing with major facilities
- Interactions with facilities on planning
- Roles for the FSP and the experimental team
- Lessons learned from experimental facilities for FSP in terms of organizing its own research efforts.

Two items deserves special mention. The first has to do with the roles for the FSP and the experimental teams in the partnership. It is envisioned that the FSP will provide

- Code suites and computer time subject to allocation process
- Help in understanding code capabilities and limitations
- Dedicated analysts
- Consideration of code developments, based on needs of the experiments

Experimental teams will provide

- Run time, subject to local planning processes
- Access to data
- Support for diagnostic data analysis
- Consideration of upgrades, based on the needs of the simulation program

The FSP and experimental teams will collaborate on

- Setting priorities
- Run planning
- Experiment analysis and interpretation

- Development of synthetic diagnostics
- Physics interpretation
- Preparation, presentation and publication of results

The second has to do with the roles of the analysts. Functionally FSP validation analysts are charged with

- Coordinating partnership with experimentalists, modelers, and theorists to identify conduct, document, and refine validation test case studies
- Working with modelers and theorists (including those not supported by FSP) to test new and improved models/theories against current FSP capabilities

Analysts are imbedded members of an ISA team

- While working closely with other team members, they will safeguard the independence of the code and experimental results, and the objectivity of the validation metric findings
- Watch out for broader implications of validation findings beyond a specific ISA

Analysts should have a broad background

- Strong grasp of underlying theory, computation (running codes on HPC), and experimental set up and measurements
- Able to work closely with verification & UQ
- Collaborate with software teams to develop synthetic diagnostics

The number of fusion scientists that meet the qualifications of the analyst is quite limited at present. The FSP will have to take on the responsibility of developing a new cadre of validation analysts and mentoring their professional growth.

Diagnostic	Simulation Code/Theory	Physics Application	Reference(s)
Electron Cyclotron Emission (ECE)	NOVA	Alfven eigenmode structure	M. A Van Zeeland *et al*[9]
Soft x-ray	GATO	MHD mode structure	J. S. Kim *et al*[10]
Fast ion Dα imaging	TRANSP(??)	Fast ion transport	M. A. Van Zeeland *et al*[11]
Visible Bremstrahlung imaging	Analytic Theory	NTM structure	M. A. Van Zeeland *et al*[12], J. H. Yu *et al*[13]
Phase Contrast Imaging (PCI)	TORIC	ICRF heating	Y. Lin *et al*[14]
Phase Contrast Imaging (PCI)	GS2, GYRO	Microturbulence (density fluctuation spectra)	D. R. Ernst *et al*[15], L. Lin *et al*[16]
Reflectometry (Doppler, Fast-sweeping)	GYRO, GYSELA	Microturbulence (density fluctuation spectra)	A. Casati *et al*[17]
Beam Emission Spectroscopy (BES)	GYRO	Microturbulence (density fluctuation spectra)	C. Holland *et al*[18]
Correlation Electron Cyclotron Emission Radiometry (CECE)	GYRO	Microturbulence (T_e fluctuation spectra)	C. Holland *et al*[xii]
CECE-reflectometry cross-correlation	GYRO	Microturbulence (n_e-T_e cross-phase)	A. E. White *et al*[19]
Doppler Backscattering (DBS)	GYRO	Intermediate-k drift-wave density fluctuation spectra	C. Holland *et al*[20]
Gas puff imaging (GPI)	SOLT	Edge turbulence	D. A Russell *et al*[21]

---

[9] M. A. Van Zeeland *et al* , Phys. Rev. Lett. **97** 135001 (2006); M.A. Van Zeeland *et al* , Phys. Plasmas **14**, 056102 (2007)

[10] J. S. Kim *et al* , Rev. Sci. Instrum. **80** 113503 (2009)

[11] M. A. Van Zeeland, W. W. Heidbrink, and J. H. Yu, Plasma Phys. Control. Fusion **51** 055001 (2009)

[12] M. A. Van Zeeland *et al* , Nuclear Fusion **48** 092002 (2008)

[13] J. H. Yu, M. A. Van Zeeland, and M. S. Chu, Rev. Sci. Instrum. **79** 10F516 (2009)

[14] Y. Lin *et al* , Plasma Phys. Control. Fusion **47** 1207 (2005)

[15] D. R. Ernst, N. Basse, W. Dorland, C. L. Fiore, L. Lin, A. Long, M. Porkolab, K. Zeller, and K. Zhurovich, "Identification of TEM turbulence through direct comparison of nonlinear gyrokinetic simulations with phase contrast imaging density fluctuation measurements," in IAEA Fusion Energy Conference, Chengdu, China, 16–21 October 2006, oral paper IAEA- CN-149/TH/1-3.

[16] L. Lin *et al* , Plasma Phys. Control. Fusion **51** 065006 (2009); L. Lin *et al* , Phys. Plasmas **16** 012502 (2009)

[17] A. Casati *et al* , Phys. Rev. Lett. **102** 165005 (2009)

[18] C. Holland, A. E. White, G. R. McKee, M. W. Shafer, J. Candy, R. E. Waltz, L. Schmitz, and G. R. Tynan, Phys. Plasmas **16** 052301 (2009)

[19] A. E. White *et al* , Phys. Plasmas **17** 056103 (2010)

[20] C. Holland *et al* , "Testing gyrokinetic simulations of electron turbulence," to be submitted to *Nuclear Fusion* (2011)

[21] D. A. Russell *et al* , "Comparison of scrape-off layer turbulence simulations with experiments using a synthetic gas puff imaging diagnostic," accepted for publication in *Physics of Plasmas*
[Batishchev 97] O.V. Batishchev *et al* ., Phys. Plasmas **4** (1997) 1672.
[Brooks 02] J.N. Brooks, Fusion Engineering and Design **60** (2002) 515.
[Chang 04] C.S. Chang *et al* ., Phys. Plasmas **11** (2004) 2649 .

Langmuir probe		Edge/pedestal ne, T_e, V_{float} fluctuations and profiles	
Mach probe		Edge/pedestal velocity fluctuations and profiles	

Table 26: Summary of existing synthetic diagnostic capabilities and gaps

D.3.4 Synthetic Diagnostics for Validation

In order for the quantitative comparisons of model predictions to experimental measurements that lie at the heart of validation to be meaningful, they must be “apples-to-apples” comparisons. In many cases, the quantities that are of interest and output by a computational or theoretical model are related, but not directly equivalent to what is measured experimentally. These differences can be related to spatiotemporal sensitivity or locality (e.g., the prediction of a local density or temperature profile by a model vs. a measured line-averaged value from a chord measurement), but can often be more complex (i.e. the translation of predicted plasma density, temperature, and electrostatic potential fields into the floating potential or ion saturation current measured by a Langmuir probe). Accounting for these differences in model validation is done via diagnostic and comparison-specific models termed synthetic diagnostics[22]. The complexity of a synthetic diagnostic model can be such that it requires verification and validation in and of itself.

The breadth of validated physics that the FSP aims to deliver will therefore require a substantial suite of synthetic diagnostics. These should likely be implemented as “standalone” components, not tied to a specific physics component (e.g., a given microturbulence or MHD code), and subject to the same software testing, verification and validation processes as other components. Fortunately, a number of existing synthetic diagnostics (listed in Table 26) have already been developed that should be used as starting points for the FSP synthetic diagnostic components. Noticeably missing are synthetic diagnostics for edge Langmuir and Mach probes, both of which will be essential for validating edge turbulence models. While there are differing levels of sophistication across these existing capabilities, they represent a highly useful starting point for FSP validation work. The key initial challenge for the FSP will be to translate these various existing models that have generally been implemented by individual researchers for use with single codes, into robust, standalone, code-independent components.

D.3.5 Documentation

As highlighted in the previous sections, the workflow, metrics, and analyst best practices all have implications for data management and documentation needs. In particular, there needs to be tools to make the results of validation planning activities, exercises, conclusions and lessons learned readily available to all interested parties. The FSP public documentation should include verification and validation metrics calculated for individual components and integrated multi-component predictions against relevant benchmarks and test cases for each science driver. A strong goal to aim for in the archiving of actual simulation results would be to build a database documenting the progression and evolution of FSP capabilities for these test cases, such that progress in physical fidelity (or lack therefore) can be clearly documented.

[Hassanein 02] A. Hassanein, Fusion Engineering Design **60** (2002) 527.
[Pigarov 09] A.Yu. Pigarov *et al* ., J. Nucl. Mat. **390-391** (2009) 192.
[Rognlien 05] T.D. Rognlien *et al* ., J. Nucl. Mat. **337-339** (2005) 327.
[Shestakov 03] A.I. Shestakov, R.H. Cohen *et al* ., J. Comput. Phys. **185** (2003) 399.
[Stotler 01] D.P. Stotler *et al* ., J. Nucl. Mat. **290-293** (2001) 967.

[22] R. V. Bravenec and W. M. Nevins, Rev. Sci. Instrum **77** 015101 (2006)

D.4 Uncertainty Quantification

Within the context of FSP, uncertainty quantification (UQ) refers to the quantitative characterization and reduction of uncertainties in simulation results. These uncertainties arise from uncertainty in input and model parameters, modeling errors, and unknown processes or mechanisms. Although the fields of sensitivity and uncertainty analysis are well-developed, the application of these ideas to large, multi-physics simulations for UQ is still a very active research area in applied mathematics and statistics. Thus, activities in this area must necessarily be dynamic and adapt to on-going developments.

D.4.1 Mission and Goals

The mission of the FSP UQ activity is to produce simulation results routinely with quantified uncertainties. The associated goals are to

- Identify and reduce sources of uncertainty in simulation results;
- Develop practical procedures for the routine quantification on uncertainty and integrate UQ into the standard practice of FSP simulation;
- Improve the rigor of validation activities and facilitate more productive collaborations between theory, computation, and experiment.

The FSP must develop a UQ strategy appropriate for individual components as well as for their integration. Key issues the FSP UQ effort must address include:

- Policies for UQ: What are the scopes and levels of rigor of UQ required in the FSP? What are the acceptable techniques for UQ? How is a UQ effort stood up in conjunction with the rest of the FSP?
- UQ workflow: What are the processes for UQ? How do they integrate with verification and validation activities?
- Defining roles and responsibilities: What is the right level and form of independence in the UQ approaches taken by ISA and advanced component efforts? To what degree are methods and techniques standardized across the FSP? Who has responsibility for executing UQ activities and who evaluates these activities? Where does research and development of new UQ techniques and tools occur within the FSP?
- Data management and documentation: How can data management methodologies be used aid in the documentation, data collection, and provenance association for of UQ activities, and how do we best document and disseminate results of UQ?

To develop this UQ plan, the FSP planning effort undertook several activities. First, a small team with fusion, numerical analysis, and UQ expertise was tasked with writing a review document on common verification and UQ procedures. To determine the state of UQ within the fusion community, input was sought from the six science driver efforts. Little feedback was obtained other than an interest in UQ. A proposal for an integrated verification and UQ effort was then written to provide a concrete basis for evaluation by UQ and fusion simulation experts at the second FSP planning workshop. The findings included:

- Systematic UQ has not really been applied in fusion energy simulation, but this subject area has the advantages of being data-rich and of having several experimental facilities.
- Documentation for UQ, especially of the sources of uncertainties and the assumptions used in analysis, is critical. A systematic approach to documentation and archiving of the potentially vast quantities of data generated in UQ studies is important.
- It is critical that precise, well-defined UQ questions relevant to the fusion community be formulated. An ongoing dialogue must be maintained between UQ practitioners and fusion energy scientists in order to define and to refine the UQ problems of greatest interest and in order to educate the fusion community about UQ and to educate UQ analysts about fusion energy simulation challenges and needs.

- Because application of UQ analysis techniques to a problem as difficult as the fusion simulation problem has not been done, the best initial approach will take incremental and iterative steps involving UQ experts, fusion scientists, computer scientists, and numerical analysts.
- It was recommended that a concrete sequence of problems be defined in the context of the two ISAs, and examples of investigation of hierarchies of components of increasing complexity from both the WDM/Disruptions and Edge/Pedestal areas were developed.
- The subject of disruptions was identified as the critical target moving forward, in the context of ITER, and this should be the initial focus of UQ investigations; as such work must be done to define the size of the parameter space for this problem.
- Uncertainties in input parameters must be nailed down by the fusion energy community.
- It was strongly suggested that not restriction to a particular class of UQ methods should be done at this time. All existing approaches to UQ should be investigated since there are no clear methods that will succeed, and different approaches can provide complementary information.
- New technologies will likely need to be developed, including methods of dimensional reduction and techniques to combine uncertainties and numerical errors.
- Massive amounts of computer time will be required, and it is unclear how this will be acquired.

Based on these findings, we present a high-level strategy for the integration of UQ activities into the FSP.

D.4.2 Policies for UQ

As stated earlier, there is no single agreed upon way to apply UQ to multi-physics simulations. This field is still young, and there are many outstanding research issues. In addition, there are several ways in which UQ can be used. We focus here on the forward sensitivity problem that produces uncertainty estimates for computed results due to uncertain inputs because that is the primary question for the validation use case. However, there other applications that will be of use in the FSP, such as the selection of the most likely model from a set of experimental data and the calibration of models given experimental data with known uncertainty. Resources will be limited, so the forward problem will be the emphasis with the pursuit of other applications as time and resources permit.

The SQ team will continually evaluate existing techniques and UQ infrastructures and periodically recommend best practices and tools for the FSP. There are numerous techniques being pursued in the research community including interval methods[23], deterministic forward sensitivity and adjoint methods[24], statistical methods[25], and stochastic PDEs[26]. The FSP, at least for the first few years of the program, will not restrict consideration to a single methods, but will instead rely on UQ expertise to select and evaluate several methodologies. Both deterministic techniques, like adjoints, and statistical techniques can be applied without modification to existing codes, so both approaches can be used, albeit with some degree of additional development. Infrastructure to support the large number of simulations common in UQ analysis should be adapted from any one of the number of existing tools.

---

[23] R. Kearfott and V. Kreinovich (Eds.), *Applications of Interval Computations*, Kluwer Academic Publishers, Dordrecht, The Netherlands, (1996).

[24] D. Cacuci, *Sensitivity and Uncertainty Analysis: Volume I Theory*, Chapman & Hall/CRC, New York, 2003; D. Cacuci *et al* ., *Sensitivity and Uncertainty Analysis: Volume II Applications to Large-Scale Systems*, Chapman & Hall/CRC, New York, 2005.

[25] J. C. Helton and F. J. Davis, “Latin Hypercube sampling and the propagation of uncertainty in analyses of complex systems,” *Reliability Engineering and System Safety* **81**, 23 (2003); J. C. Helton *et al .,* “Survey of sampling-based methods for uncertainty and sensitivity analysis,” *Reliability Engineering and System Safety* **91**, 1175 (2006); D. Cacuci *et al .,* *Sensitivity and Uncertainty Analysis: Volume II Applications to Large-Scale Systems,* Chapman & Hall/CRC, New York, 2005, Chapter I.A.

[26] R. Ghanem, “Ingredients for a general purpose stochastic finite elements implementation,” *Comput. Methods Appl. Mech. Engin*. **168**, 19 (1999); D. Xiu and G. Karniadakis, "Modeling uncertainty in flow simulations via generalized Polynomial Chaos," *J. Comput. Phys*. **187** (1). 137 (2003).

Because of the limitations on resources and the lack of prior work in the fusion simulation area, the UQ effort will, at least initially, be limited in scope. The focus will be on the ISAs and investigations will be limited to a subset of the problems in these science areas. We note that validation has and can be done in the absence of UQ analysis for the simulation output; the addition of UQ makes the validation stronger, more quantitative, and provides direction for addition improvements. Therefore, while it would be preferable to have UQ analysis involved in every validation study, given the state of UQ in fusion simulation and the challenges of the fusion problem, it is unrealistic to expect to have UQ fully integrated in the short term. In this sense, the UQ effort is in the near term more exploratory research than application. As successful UQ techniques are identified and developed, previous validation studies can be re-analyzed and re-interpreted.

Documentation is very important in UQ. A careful accounting of sources of uncertainty must be identified and documented so that appropriate models of the uncertainty can be constructed. There are many decision points and assumptions in the design of a UQ study that must likewise be recorded. All of the relevant parameters and inputs need to be identified, and this information should be recorded. Code development activities must also support code documentation activities, so that this information should be readily available. Documentation also includes any scripts and input files used to generate the simulation results as well as information about the simulations (provenance). We anticipate that the documentation will be saved and distributed within the framework established by the integrated data management policy outlined in Section 4.3.6.

D.4.3 Workflow

The following outlines a high-level workflow a typical FSP UQ campaign. It is assumed that the code (component, integrated application) has already undergone extensive code verification.

- For the given application, define and document the problem, including the models, initial and boundary conditions, problem geometry, etc.
- For the given application, identify and document all relevant quantities of interest. Define appropriate synthetic diagnostics for the results of the computations.
- Identify all the relevant physics and numerical parameters, and any other potential sources of error or uncertainty (e.g., model uncertainty, errors from the diagnostics, etc.). Quantify the uncertainty in the inputs, including these parameters and other input data.
- Define the parameter space to be explored for the computational sensitivity and error analysis. Use analysis, e.g., scaling analysis, and a priori knowledge to reduce the dimensionality of the parameter space to be explored as much as possible.
- Select an appropriate sensitivity analysis technique given the problem and available tools.
- Determine an efficient, rigorous, and practical evaluation (e.g., sampling) strategy for the numerical UQ campaign and define what data must be extracted from the simulation results for subsequent analysis. This includes a determination of how to assess and incorporate numerical error. Define what data must be archived for future use in analysis.
- Apply calculation verification methods to the numerical solution of the model equations. Adjust numerical parameters as necessary to ensure suitably accurate and converged results. If possible, build a surrogate of the error behavior over the parameter space to be explored. Otherwise, calculate an error estimate for each simulation in the study.
- Execute the sensitivity analysis: employ a UQ pipeline or other enabling technology to automate the examination of parameter space, the quantitative analysis of the results, and the archiving of data.
- Use the raw data from the simulation results and the derived results obtained from analysis to undertake validation against experimental data or design and predict new experiments. The validation step will involve the exercise of many UQ analysis tools.
- Coordinate these findings with experiment and theory, e.g., improve knowledge of parameters or improve models

- Document and archive important data and results.

Note that this process defines a routine uncertainty analysis where a reliable UQ methodology is known. In a research application, particularly for a multi-physics code, the process is similar, but in general iterative. A sequence of preliminary studies is followed to identify the inputs that produce the greatest sensitivity in the output. When much data is available, inverse problems can be solved to reduce the ranges of input parameters. Such studies are meant to reduce the size of the problem space that needs to be considered and to develop information about the structure of the response surface, which can limit the applicability of certain methods. In addition, correlations between parameters, possibly leading to the identification of more fundamental parameters (and thus reducing the problem space dimension) are sought.

Finally, it is anticipated that, until software infrastructure comes online, integrated application development will not be able to engage fully in sensitivity analysis and formal UQ. This lag will not pose a substantial problem because of the extensive documentation, component code and calculation verification activities, and UQ problem definition and study design that must also be undertaken before UQ simulations are begun.

D.4.4 Roles and Responsibilities

As with verification, there are three broad activities that fall under the scope of UQ:

- Development of methodologies
- Development of technologies
- Application of technologies and methodologies

Within the available resources of the FSP, the focus will be on the development and application of methodologies and the application of existing techniques, but there may be additional activity in the areas of research and development. Uncertainty Quantification for multi-physics simulations in particular is still a nascent and active area of research. Because a full-fledged UQ research and develop program is beyond the resources of the FSP, the FSP will seek to leverage existing DOE investments in UQ research and development, for example, from ASC-funded PSAAP centers, Laboratory Directed Research and Development efforts, and ASCR-funded projects.

UQ activities will take place throughout the FSP organization in much the same was as with verification. A small, dedicated Software Quality (SQ) team will exist to provide UQ expertise and focus on the development of methodologies. As before, this team will have the responsibilities to further specify program-wide policies and procedures and to provide expert guidance to the rest of the project. Each ISA team will have an identified SQ point-of-contact and will work with the UQ experts to define a UQ analysis appropriate to their task within the policies put forth by the SQ team. SQ team will also aid in the execution of the plans. The Software Integration team will work with the SQ team to adapt and develop program-wide infrastructure to support UQ, in particular a workflow tool that enables the generation and execution of many sensitivity analysis simulations and that provides a suite of statistical analysis tools. Several of such workflow tools already exist, , such as the DAKOTA/UQ [27], PSUADE [28], LLNL's UQ Pipeline, and the CalTech PSAAP Center's UQ Pipeline [29]. Thus, the FSP should leverage these existing technologies and adapt them for FSP use. Research into new UQ methods to incorporate into FSP activities will take place within the SQ team. The relationships and responsibilities are summarized in Table 27.

For the purposes of UQ, SQ point-of-contacts should

- Have a strong grasp of the underlying physical theory and computation (running codes on HPC)
- Have a good grasp of sensitivity analysis and UQ
- Be able to work closely with SQ team and with Validation Analysts

---

[27] http://dakota.sandia.gov

[28] https://computation.llnl.gov/casc/uncertainty_quantification

[29] http://www.psaap.caltech.edu/meetings/sitevisitoct10/posters/McKerns_PSAAP_poster.pdf

- Be able to develop analysis tools as needed

For the purposes of UQ, the SQ Team should have members with

- A strong grasp of uncertainty quantification, sensitivity analysis, statistics, and computation (running codes on HPC)
- Experience in the application and development of UQ analysis techniques
- The ability to work closely with members of application teams and Validation Analysts
- The ability to collaborate with software teams to develop verification infrastructure

Team	Responsibilities
SQ	• Define UQ policies and procedures • Work with ISA teams and validation analysts in the selection and application of UQ techniques to their activities • Evaluate the progress of UQ efforts within the program • Adapt and develop technologies, with the help of the Software Integration team, to provide program-wide UQ infrastructure • Evaluate external research results in UQ for applicability to the FSP • Develop new UQ methodologies for the FSP as needed
ISA	• Identify a point-of-contact to the SQ team • Develop a UQ plan specific to their project with guidance from the SQ team • Execute UQ plans in collaboration with the SQ team as per the best practice workflow • Provide feedback to the SQ team about needs and difficulties encountered in the execution of the UQ plans
Software Integration	• Adapt and develop technologies, with the help of the SQ team, to provide program-wide UQ infrastructure

Table 27: UQ responsibilities within the FSP

D.4.5 Data Management and Documentation

As highlighted in the previous sections, the best practices for UQ have tremendous implications for data management and documentation needs. In particular, tools are required to make the results of UQ planning activities, exercises, and conclusions available to all interested parties. UQ has the potential to generate tremendous amounts of data, and systematic approaches to the collection of this data and its provenance to facilitate easy retrieval are paramount. The FSP public documentation should include the UQ documentation as described above for individual components and integrated multi-component predictions for each science driver. The FSP data management plan is discussed in Section 4.3.6.

D.4.6 Research Opportunities

There are many open issues in the application of UQ techniques to large-scale simulation codes. The following possible research topics have been identified:

- Methods to ameliorate the curse of dimensionality
- Methods to propagate uncertainties through coupled components
- Dealing with instability, chaotic behavior, or lack of smoothness of response functions

- Efficient incorporation of deterministic error estimates with stochastic uncertainties
- UQ for multi-scale problems

As with verification, resources will be limited within the FSP, so the FSP will need to engage external research and development efforts (ASRC, ASC, NSF, etc.) for new UQ tools and techniques.

Appendix E: Configuration Management Plan

E.1 OVERVIEW

This *Fusion Simulation Program (FSP) Configuration Management Plan* is a standard auxiliary document to the primary FSP Program Execution Plan and its various project execution plans (PEPs) which describe the purpose and general plan for the relevant projects. Not all aspects of this standard plan may apply in a particular project undertaken, however. In that case, the project's PEP will describe how the configuration plan has been appropriately tailored.

Configuration management is the unique identification, controlled access, change control, and status reporting of selected intermediate and final work products, product components, and products during the life of a project or system. The management principle expressed in this is that once a plan has been approved, no deviation from it is permitted unless the deviation is formally identified, analyzed, approved, and reported to all stakeholders. All FSP projects will engage in rigorous configuration management processes to ensure that the project is completed successfully.

E.2 CONFIGURATION ITEMS

Configuration items (CIs) are those specific work products that are managed using configuration management processes and techniques. The CIs managed in the major FSP projects are listed in Table 28 below along with their abbreviated alternate names. Alternate names are useful as a quick reference or as the base of a computer file name. Smaller projects may not use all listed documents.

Configuration Item	Alternate Name
Integrated Project Team Charter	IPT Charter
Program/Project Execution Plan	PEP
Risk Management Plan	RMP
Quality Assurance Plan	QAP
Communications Plan	CommP
Configuration Management Plan (this document)	CMP
Cost and Schedule Baselines, i.e., a resource-loaded and scheduled Work Breakdown Structure	WBS
WBS Dictionary	
Software Application Codes	
Software Acceptance Test Plan	

Table 28. Configuration Items

E.2.1 Versioning

The first approved version of any CI will be Version 1, and subsequent versions will increment by +1. Subversioning will not be used except for software codes. All intermediate changes made to a planning document or baseline will be designated as a draft of the next version. Version numbers will not change until the draft is approved. The approval process is discussed in Section 3 below.

E.2.2 Storage, Handling, and Disposal

All CIs except for software codes, will be maintained by the FSP Project Management Office. All narrative CIs will be developed as either Microsoft Word or Excel documents. The WBS will be maintained in a commercial project management application such as Microsoft® Project. The latest approved versions of these will be made available on a FSP-managed website accessible to designated members of the FSP integrated project team (IPT) and the FSP project execution staff. The approved versions of these files will be write-protected or saved in a .pdf format. Prior versions of documents will be maintained in a separate file directory for historical reference.

The development of software codes will be managed through version control applications (to be determined).

There will be no other retrieval controls placed on these CIs. It will be a user's responsibility to make sure that any electronic or hard copy of any CI is consistent with the definitive versions described immediately above.

At program closeout and at each project closeout all approved versions of all CIs, including approved prior versions, will be electronically archived along with other relevant program and project materials.

E.3 CHANGE CONTROL

Any project stakeholder may submit a request to modify any CI. Change requests that might impact any project planning document (e.g., a new reporting chain or a different safety auditing procedure) will require a formal change request.[1] This includes a request to insert, delete, or modify any activity or milestone listed in the WBS.

Where appropriate, each FSP text-based (e.g., Microsoft Word) document has an approval page where required signatures are identified (by position). The signatures of the current position holders will be obtained for any document revisions. Changes to any planning document will be recorded within the document by the replacement of the approval signature page and also in a change control table located at the beginning of the document. The table will record the new version number, the date of the change, and a brief explanation of the change.

The approval process describing who authorizes changes to project performance baselines is discussed in the baseline change control section of the each project's PEP. These processes are not repeated here as they are tailored to the particular projects. The change decision and authorization process may involve the Department of Energy Program Offices of Fusion Energy Sciences and Advanced Scientific Computing Research, a member of the IPT, an executive from the host or partner institutions, the FSP Director, the Chair of the FSP Executive Management Committee, the head of the FSP Project Management Office (PMO), FSP Team Leaders, and/or any other stakeholders, depending on the nature of the change request.

All project change requests (PCRs) are to be submitted via a formal PCR form. The PCR form will be made available either in hard copy and/or as an online Microsoft Word of PDF form that may be completed and printed or e-mailed. The form requests identification and contact information; reason for the change; expected cost, schedule, and technical impacts; and any new or modified risks resulting from the change. An example of the current form is in Appendix A.

All PCR forms must be submitted to the FSP Director who, working with the PMO, will provide a preliminary evaluation of the PCR. If, after review, the PCR is accepted by the FSP project director, additional analysis and review and additional signatures may be required before the change is incorporated into the baseline. The hierarchy of change control authority for each type of change can be found in the appropriate PEP.

Changes to one document will often flow to other documents as well. A single PCR can cover all impacted documents.

Upon approval the PCR form will be updated to reflect the approving authority and approval date, as well as the final description of the change (which may differ from the original request.) Approved change request forms will be retained for the duration of the project and archived with other project documents at closeout.

For new projects some planning documents are generated as originals. In many cases, however, the project will refer to standard FSP plans that are relevant to all FSP projects. This configuration management plan is an example of a general FSP planning document. All of these general documents are reviewed for efficacy and currency as a matter of course before they are associated with each new project. As a result of the review, one

---

[1] Occasionally minor adjustments to the project planning documents (e.g., fixing typos, changes in named personnel or their contact information) are made without the formal submission of a project change request. Nonetheless, any resulting modifications to the documents must still go through a formal approval process.

or more documents may be updated to reflect changed conditions for the new project. In this case a formal PCR form is not used; however, as in all changes to CIs, all versioning and approval procedures are followed.

E.4 STATUS REPORTING

It is important that approved changes be announced and implemented. Depending on the nature of the change, the announcement and implementation may occur or commence immediately upon receipt of approval, but in any case, the announcement and implementation plan will be announced at the next regular weekly staff status meeting and at the next regular meeting of the IPT. Approved changes will be a standard content item for all monthly project status reports.

Appendix F: Risk Management

F.1 OVERVIEW

The *Fusion Simulation Program (FSP) Risk Management Plan* (RMP) is a standard auxiliary document to FSP program and project execution plans (PEPs), which describe the purpose and general plan for relevant FSP projects. Not all aspects of this standard plan may apply to a particular project undertaken; in which case the project's PEP will describe how this general risk plan has been appropriately tailored. Although this document discusses risk management in the context of project management, the FSP uses virtually the same processes to manage risks with respect to ongoing program operations.

A common understanding of risk management is that it is the process of identifying and analyzing risks and then taking appropriate steps to reduce risks to an acceptable level. The FSP RMP documents the processes employed to manage risk proactively as a component of effective management of the entire FSP. It is a management tool for mitigating the effects of events that may adversely impact the program.

- The FSP RMP describes the FSP processes for identifying, analyzing, tracking, and managing risk. The purpose of this plan is
- to document procedures for identifying and analyzing known risks to the program along with tactics and strategies to mitigate those risks;
- to serve program management as a basis for identifying alternatives to achieve cost, schedule, and performance goals; and
- to assist management in making decisions on budget and funding priorities by providing risk-related information for decisions.

While most people think of risk in its negative sense, the Project Management Institute's, *A Guide to the Project Management Body of Knowledge* (4th ed., Project Management Institute, 2008) defines risk as "an uncertain event or condition that, if it occurs, has a positive or negative effect on a project's objectives." Positive risk, i.e., an opportunity for improvement, is not ignored in FSP projects, and positive risk events are included in the risk register (an online software application maintained by the FSP) and are managed as important components of the project. It is, however, awkward to describe a plan for managing both kinds of risk at the same time. Therefore, since negative risk is the typical connotation, that aspect of risk management is covered in this RMP as if it were the primary concern of the project. Managing for positive risk is explained at the end of this document as a modification to general risk management processes and procedures.

F.2 RISK MANAGEMENT APPROACH

- A structured, disciplined approach for risk management has been developed and implemented for FSP using the Project Management Institute's best practices for risk management as a model. The risk management process must be ongoing and dynamic. The goals of risk management are to ensure
- that risk identification and analysis have the appropriate rigor;
- that risk issues are made visible early;
- that thorough, credible mitigation or alternative risk response plans are prepared/implemented; and
- that project budgets are maintained.

The FSP Program Director has overall responsibility for FSP risk management and the implementation of this RMP.

- The FSP Executive Committee and project control account managers[1] (CAMs)/milestone owners:

---

[1] Control Account Managers are those staff responsible for managing a particular scope of work and the financial account(s) (budgets) associated with that work.

- perform risk analysis, including identifying potential vulnerabilities/risks, likelihood of occurrence, and impact on the project;
- develop risk mitigation or other risk management strategies; and
- execute plans to accomplish risk reduction activities.
- The Head of the FSP Project Management Office:
- is responsible for the development of the risk management approach;
- schedules routine reviews of the risks;
- ensures that risk analysis results are documented and that risk mitigation plans are brought to closure;
- actively participates in the project's conduct of risk management, such as determination of mitigation plans, especially with interfacing risks between subprojects or activities; and
- collects, records, or provides budget estimates for risk management activities.

For the successful petascale project the FSP used an online risk management software program to track risks..

F.3 RISK MANAGEMENT PROCESS

Project risk management consists of a six-step process:

1. identifying potential vulnerabilities/risks;
2. determining their likelihood of occurring;
3. assessing their impact on the project scope, cost, and schedule baselines;
4. determining activities that would reduce/mitigate the risk;
5. executing a plan to accomplish these risk-reducing activities; and
6. reporting and tracking risk.

It is well known that the financial and project management benefits of risk management are less dependent on the specific formula used for quantitative assessment and more dependent on the frequency and rigor with which risk assessments are performed. The implications are that performing a risk assessment and basing it on a sound process are the most important aspects to achieving a positive impact.

F.3.1 Risk Identification

- FSP management evaluates project risk issues on a continuing basis. Various meetings, interviews, and other approaches are used for identifying project risks as well as for developing and tracking mitigation strategies and tactics. FSP looks at risk from two perspectives:
- An objectives- or activity-based perspective, wherein risks are identified that may impact specific project or operational objectives. For a project, these risks are generally identified by a bottom-up study of the work-breakdown structure (WBS).
- A scenario-based perspective, wherein risks are identified by analyzing situations in which potential risk events are not specific to a particular objective or that are complex or somewhat interdependent.

All reasonable risks that are identified are entered into the risk register software application, which provides the primary source of risk information for the project and from which various reports may be produced.

The primary technique used in an FSP project to identify risks is to hold a series of simple, separate interviews with CAMs and their staff wherein the WBS activities within their responsibility are examined one by one. Historical risk registers from previous projects are referenced to assist in the identification process. Once the individual interviews have taken place, another meeting with all CAMs and FSP management is held to go over all identified risks to determine whether there are gaps, interdependencies, or root causes. This subsequent meeting is also used to identify additional risks that are not associated with a particular WBS element but that may derive from or have impacts on broader situations. This set of meetings is held at the beginning of the project to establish a basic understanding of the scope of risk to the project and is held annually or whenever necessary thereafter to refresh that basic understanding. In the meantime, routine risk management meetings are held to refine the understanding and to keep it current.

F.3.2 Qualitative Risk Analysis

F.3.2.1 Process for Rating Risks

Risks are rated as high, moderate, or low as shown in Table 2. Two factors are combined to generate the overall rating: (1) the likelihood or probability of occurrence and (2) the impact or consequence to the scope (or technology), cost, schedule, and/or some other aspect of the project.

PROBABILITY	**CONSEQUENCE**		
	Marginal (M)	Significant (S)	Critical (C)
Very Likely (V)	Moderate	**High**	**High**
Likely (L)	Low	Moderate	**High**
Unlikely (U)	Low	Low	Moderate

Table 29: Risk Ratings Matrix

Likelihood is limited to three categories:

- **Very likely:** An event that is likely to occur with a probability ≥ 80%
- **Likely:** An event that is likely to occur with a probability ≥ 30% and < 80%
- **Unlikely:** An event with <30% probability of occurrence

These probability percentages are qualitative guides only and are not intended to represent absolute thresholds. It is important to refrain from inferring from these numbers a level of precision that is not really there.

“Consequence” or “impact” identifies the impact that any occurrence of an event will have on cost (amount increased), schedule (additional time), and/or technical scope (degradation from planned performance).[2] Each risk event will be evaluated on all three classifications. The highest of the three category values is used for the final rating. Some suggested impact thresholds are shown in Table 3. These and the likelihood thresholds mentioned above must be refined by the eventual FSP management team to fit their particular understanding of risk impacts.

CATEGORY	**IMPACT ON PROJECT**		
	Marginal (M)	Significant (S)	Critical (C)
Cost	< $250K	> $250K	> $500K
Schedule	< 1 month	> 1 month	> 3 months
Technical Scope (based on performance metrics)	< 10%	> 10%	> 20%

Table 30: Impact Categories and Thresholds

The risk rating in Table 2 is an indication of the perceived severity of the risk. It is derived from an evaluation of both likelihood and consequence levels.

---

[2] There is a fourth category of impact called “other,” which has so far been used infrequently in the FSP. It exists to handle situations that may arise when none of the other impact categories seems appropriate. “Other” adheres to the same likelihood scale that the other impacts do but has no defined impact thresholds. Impact intensity for “Other” is described in narrative form in the register’s description field.

The initial classification of a risk is conducted prior to any mitigation activities. A subsequent classification using the same criteria is performed when a mitigation action is planned and assumptions are made as to the probable effectiveness of the mitigation. The original and mitigated likelihood, impact, and rating values are maintained in the risk register with the risk.

F.3.2.2 Process for Ranking Risks

It is important to know the key or highest ranking risks to a project because they are likely to have the most impact on the project and on the stakeholder community. Therefore, they are the ones the project team should monitor closely. Risk ranking in an FSP project is accomplished by a simple, conventional comparative risk ranking exercise conducted periodically (no less than quarterly) or when a significant change has occurred to the overall risk environment of the project as determined by the FSP management team. The participants in the comparative risk ranking will be the FSP management team and project CAMs.

F.3.2.3 Quantitative Risk Analysis

Quantitative risk analysis produces cost and schedule estimates for activities that incorporate a consideration of risk. These activities may be in the current WBS or alternatives to WBS activities. The analysis techniques employed could range from simple educated guesses by experienced staff to running Monte Carlo simulations over the entire project.

An FSP project combines several techniques for determining contingency amounts. A general reserve of approximately 20% of estimated project costs and a 3–6 month buffer for each major project milestone (i.e., a Level 1–2 milestone) is recommended as a base contingency estimate. These amounts are derived from expert opinion and from experience with similar projects. This base is then validated by an analysis technique successfully used by Oak Ridge National Laboratory’s Spallation Neutron Source, whereby FSP staff provide weighted risk factors for cost, schedule, and technical risks for each WBS activity. This then provides information on what fraction of the conservative estimates staff developed for the baseline WBS should be set aside for a contingency fund. If the base and the calculated amounts are close, the base is validated. Additional validation involving Project (or Program) Evaluation and Review Technique (PERT) or Monte Carlo simulation analyses may be performed as well. If performed, the specifics will be described in or from the PEP.

Planned mitigation activities are integrated into the project WBS, and their progress and actual costs are tracked as any other project activity. If funds for the mitigation effort are external to the project, the mitigation effort would not be subject to project earned value analysis, however. Any additional costs associated with the remediation of accepted risks that do occur are also tracked and recorded in the risk register.

F.3.3 Risk Response Planning

- Once risks have been identified and characterized, they can be managed in several ways:
- Avoidance; i.e., taking prior action to eliminate the likelihood and/or the impact of a risk event. The use of a fixed-price contract with a vendor, is an example of avoiding the risk of price increases.
- Mitigation; i.e., taking prior action to reduce the likelihood and/or the impact of a risk before it happens. Close collaboration with vendor through on-site offices and periodic progress reviews are examples of actions taken to mitigate schedule and performance risks.
- Transference or sharing; i.e., taking prior action to assign some or all of the risk impact to another party; e.g., taking out an insurance policy.
- Acceptance; i.e., simply waiting for the risk event to occur, accepting the consequences, and/or finding a work-around.

Each registered risk will have a response assigned to it or developed for it. Most of the smaller risks will have a plan that simply states that the risk will be accepted and dealt with if and when it occurs. Larger risks will have more detailed response plans developed, and the larger the risk, the more detailed the response plan will be. A plan may include:

- mitigation tactics,
- detection mechanisms for early warnings of impending occurrence,
- trigger points or thresholds that initiate specific responses, and
- reporting paths for informing stakeholders that the risk has occurred and what is being done about it.

F.4 RISK MONITORING AND CONTROL

Once it has been identified, each risk is assigned to a risk owner, an individual who has the responsibility to track and report on that risk. That person is not responsible for the risk event happening or its consequences but is responsible for understanding all aspects of the risk in order to recognize when the risk environment changes or when thresholds or trigger points are reached. The risk tracker is responsible for reporting risk status to project management.

Risk monitoring and control are integrated with the identification process in many ways. The same meetings will be used to discuss existing risks and what needs to be done about them according to their response plans. If the impact or the likelihood has changed, the risk register is modified to reflect the new risk rating. If the response plan needs to be modified, that is done as well. If the risk event or its possibility has passed, the risk is retired. The FSP does not maintain a separate risk watch list. Highly rated risks are considered frequently and in depth at several regular meetings.

Risk management is intimately associated with change control (see the *FSP Configuration Management Plan*). Whenever a change request is submitted, the risk management process is engaged to help appreciate the full impact of the change. If the change is approved, the identified risks and the results of their qualitative and quantitative analysis, along with their response plans, are registered.

F.5 OPPORTUNITY MANAGEMENT

As stated in the first section, the Project Management Institute defines risk in having both a negative and a positive character. A positive risk is an opportunity to gain some advantage. Opportunities are incorporated into the initial planning and design phases of a project as a matter of course, and many new opportunities may develop after the project starts. Since frequent regular risk reviews are planned for all FSP projects, it will take little additional effort to include a component of opportunity identification along with the usual discussions of threats. Identifying these opportunities is important; however, identification is not always proactively attempted as a specific ongoing project management process.[3] Opportunities are recorded in the risk register. However, they are flagged so that may be analyzed and reported appropriately.

---

[3] Opportunity management is similar in concept to, and serves the same purpose as, Value Management or Value Engineering as discussed in Office of Management and Budget Circular A-131.

Appendix G: FSP Risk Register

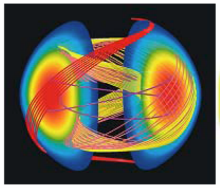

Fusion Simulation Program

Fusion Simulation Program (FSP)

Risk Rating Levels:
H = High,
M = Medium,
L = Low

Impact Horizon from Early Date:
Near = within 30 days,
Mid = 30 to 90 days out,
Far = greater than 90 days,
Current = early date passed

ACTIVE RISKS sorted by Exposure

IDz	Risk Title * Risk Description ** Management Plan *** Notes	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
2	Progress toward meeting delivery and other milestone schedules is slower than expected.	Bill Tang			Mitigate		Far

* There is a risk that the Program will experience delays meeting delivery schedules and other milestones because of research progress is slower than anticipated. It is the nature of any cutting-edge scientific research activity to have risks of this type.

** This risk can be mitigated in general by assuring that the Program teams are staffed with the most competent scientists available to it. In addition, negative impacts can be ameliorated somewhat by having the ability to react quickly and effectively by implementing timely and comprehensive cost and progress reporting procedures that provide Program management with early warning of schedule issues and by assuring the Directorate has the ability to redirect resources in an effective manner.

*** This is a general risk applicable throughout the all areas of the Program. If some area has a specific mitigation capability, a separate risk entry will be made.

Exposure

	Probability	Cost	Sched	Tech	Other
Original - 🟥 H 9	H	H	H	H	L
Residual - 🟥 H 6	M	M	H	H	L

IDz	Risk Title	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
11	Teams fail to collaborate effectively	Bill Tang			Mitigate		Far

* The distributed nature of the Progam effort means that some teams or team members may assume or affect a degree of independence that impacts negatively on overall Program progress. This risk includes situations where teams do not adopt FSP standards and common structures, tools, and procedures.

** This risk is mitigated by having: 1) a strong Directorate and Research Committee, both of which are committed to following the execution plan, 2) formal change management procedures that are rigorously enforced; and 3) by having appropriate management reporting and control procedures in place that will detect divergences from the plan in a timely manner and that allow for the reallocation of funding as necessary to keep to the plan.

Exposure

	Probability	Cost	Sched	Tech	Other
Original - 🟥 H 9	H	H	H	H	L
Residual - 🟩 L 1	L	L	L	L	L

IDz	Risk Title	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
7	Loss of key personnel	Bill Tang			Mitigate		Far

* The FSP execution plan is ambitious and depends very talented and experienced scientists to meet delivery schedules. The loss of key personnel (for whatever reason) could result in a significant schedule delay.

** As part of its Education, Outreach, and Training commitments, the FSP will require its management and team leaders to implement policies and procedures that accommodate succession planning and that provide subordinates reasonable leadership training opportunities. One of the selection criteria for leadership positions in the FSP will be the ability and willingness to mentor others.

Exposure

	Probability	Cost	Sched	Tech	Other
Original - 🟥 H 6	H	M	M	M	L
Residual - 🟨 M 3	H	L	L	L	L

FSP

ACTIVE RISKS sorted by Exposure

IDx	Risk Title * Risk Description ** Management Plan *** Notes	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
1	Program Directorate may not have requisit flexibility or authority to manage effectively. * Although a management structure and management procedures have been proposed in the Execution Plan, the specific mechanisms describing how funding will flow from DOE to the various participating institutions have not yet been decided. The risk is that the mechanisms eventually established by DoE will not provide sufficient capability to the FSP Directorate to react quickly and effectively to needed Program adjustments. ** It is critical to work cooperatively with the sponsoring DOE program office(s) to develop a funding mechanism that permits effective program management. *** Responsibility implies authority. If the FSP directorate will be held accountable for the success of the Program, it must have the ability to affect appropriate adjustments to the course of progress.	Bill Tang			Mitigate Exposure Original - Probability, Cost, Sched, Tech, Other H 6 M H H H L Residual - Probability, Cost, Sched, Tech, Other L 2 L M M M L		Far
8	Inability to find appropriate talent to fulfill program requirements or mean performance baselines. * The FSP execution plan is ambitious and requires very talented and experienced scientists to be successful. Competing interests may hinder recruitment to the Program. For example, because of the 'pause', experienced collaborators involved in the proto-FSP efforts may commit elsewhere. ** Provide potential staff with an attractive, well-managed, and well-supported program that will alow them to participate in some of the most important work in fusion science. *** Interest will be reduced if potential candidates perceive a lack of sustained interest in the Program by DOE.	Bill Tang			Mitigate Exposure Original - Probability, Cost, Sched, Tech, Other H 6 M L H H L Residual - Probability, Cost, Sched, Tech, Other L 1 L L L L L		Far
13	External dependencies may become unavailable. * The FSP will be dependent on the availability of several products and services (e.g., HDF5) such that progress will be impaired if they become unavailable or have support for them reduced. ** Continued monitoring of the support structures and community environment for these products or services will provide early warnings that problems may arise. As well FSP could provide funding to maintain critical parts of the software stack.	John Cary			Accept Exposure Original - Probability, Cost, Sched, Tech, Other H 6 H M M L L		Far

FSP

ACTIVE RISKS sorted by Exposure

IDx	Risk Title * Risk Description ** Management Plan *** Notes	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
15	Obsolescence of code from changing programming models. * New programming models will arise during the lifetime of the FSP and there will be different software stacks provided, just as in the past there were at least three different approaches (shmem, MPI, PVM) to parallel code writing. If the FSP picks an approach that ultimately is not supported, it will have to do software rewriting. ** A layered approach to software development during times of such decisions can allow rapid changing from one to another. *** re: PVM, shmem disappearing and being replaced by MPI, same for the current CUDA/OpenCL battle	John Cary			Accept Exposure Original - Probability, Cost, Sched, Tech, Other H 6 H M L L L		Far
3	Disruptions to funding. * Disruption to funding (e.g., a budget reduction or slow disbursement) to the Program is another high risk. These disruptions may result from Congressional (e.g., continuing budget resolutions, etc.) or DoE actions. The impacts to the Program can be significant causing schedule delays or loss of key personnel who may seek alternative financial coverage. ** The FSP Directorate may be able to ameliorate these impacts somewhat with management reserve funds.	Bill Tang			Accept Exposure Original - Probability, Cost, Sched, Tech, Other M 4 M M M M L		Far
14	Insufficient time on experimental facilities * Each validation task will require multi-layers of controlled experiments to the details of the physics models. This requirement translates to significant experimental time. The U.S. facilities are historically underfunded in experimental operation time, whereas experimental proposals typically outnumber experiments performed by more than a factor of three. Hence there is a real risk that FSP will not be able to have sufficient experimental time for the proposed WBS tasks. ** Work closely with experimental facilities as true partners in planning i.e. make validation an essential part of proposed experiments. Seek experimental time on international facilities	Vincent Chan			Accept Exposure Original - Probability, Cost, Sched, Tech, Other M 4 M L M M L		Far
16	Inadequate diagnostic capabilities * As FSP moves to validating the physics models at finer granularity e.g. 3D wall heat load, fluctuation correlations, etc., existing diagnostics on experimental facilities will not be adequate. New diagnostics funding is highly constrained and competes with experimental run time. Furthermore, it will require new ideas and time for development. ** Early planning with experimental facilities. Work closely with OFES experimental science/diagnostics. Solicit participation and contributions from international partners.	Vincent Chan			N/A Exposure Original - Probability, Cost, Sched, Tech, Other M 4 M L M M L		Far

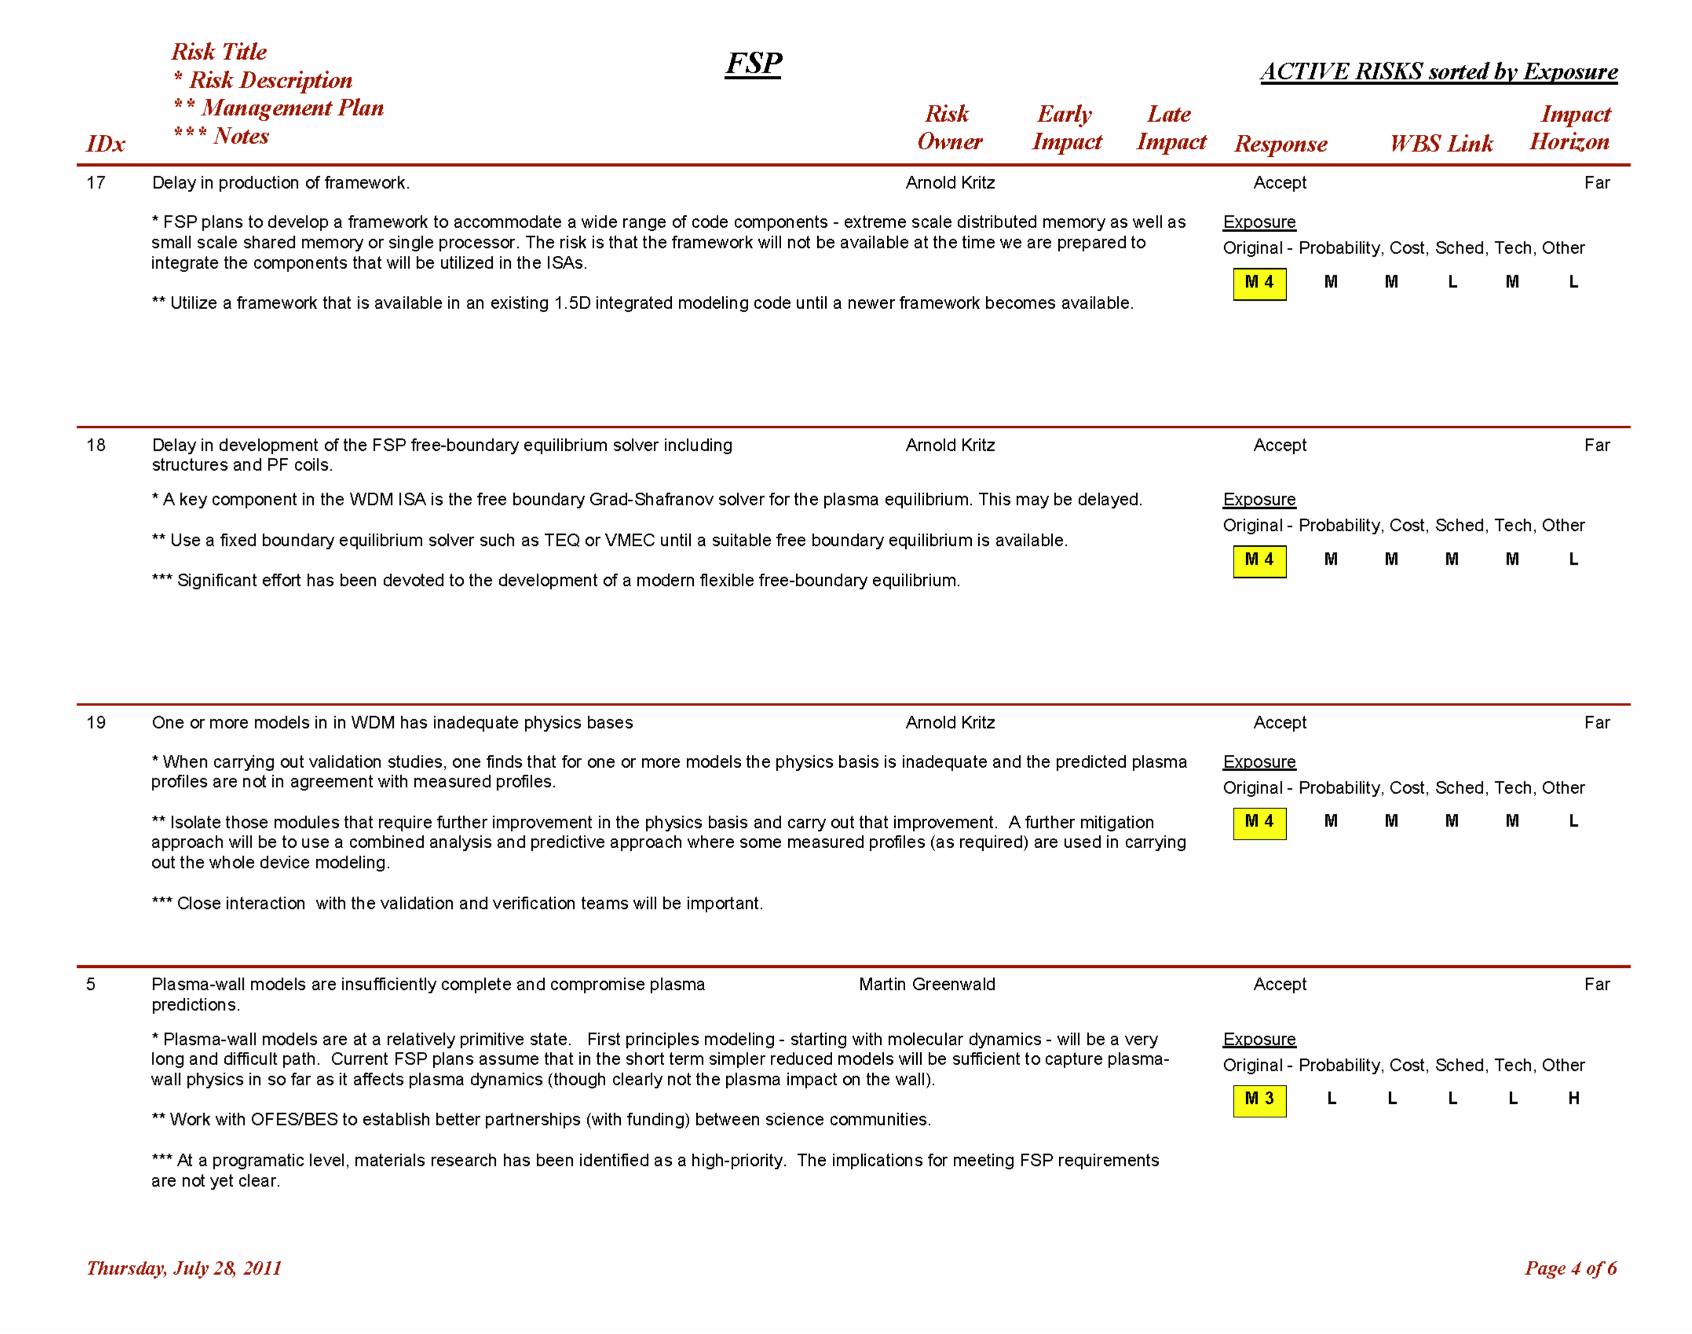

FSP

ACTIVE RISKS sorted by Exposure

IDx	Risk Title * Risk Description ** Management Plan *** Notes	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
17	Delay in production of framework.	Arnold Kritz			Accept		Far

* FSP plans to develop a framework to accommodate a wide range of code components - extreme scale distributed memory as well as small scale shared memory or single processor. The risk is that the framework will not be available at the time we are prepared to integrate the components that will be utilized in the ISAs.

** Utilize a framework that is available in an existing 1.5D integrated modeling code until a newer framework becomes available.

Exposure

Original -	Probability,	Cost,	Sched,	Tech,	Other
M 4	M	M	L	M	L

IDx	Risk Title	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
18	Delay in development of the FSP free-boundary equilibrium solver including structures and PF coils.	Arnold Kritz			Accept		Far

* A key component in the WDM ISA is the free boundary Grad-Shafranov solver for the plasma equilibrium. This may be delayed.

** Use a fixed boundary equilibrium solver such as TEQ or VMEC until a suitable free boundary equilibrium is available.

*** Significant effort has been devoted to the development of a modern flexible free-boundary equilibrium.

Exposure

Original -	Probability,	Cost,	Sched,	Tech,	Other
M 4	M	M	M	M	L

IDx	Risk Title	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
19	One or more models in in WDM has inadequate physics bases	Arnold Kritz			Accept		Far

* When carrying out validation studies, one finds that for one or more models the physics basis is inadequate and the predicted plasma profiles are not in agreement with measured profiles.

** Isolate those modules that require further improvement in the physics basis and carry out that improvement. A further mitigation approach will be to use a combined analysis and predictive approach where some measured profiles (as required) are used in carrying out the whole device modeling.

*** Close interaction with the validation and verification teams will be important.

Exposure

Original -	Probability,	Cost,	Sched,	Tech,	Other
M 4	M	M	M	M	L

IDx	Risk Title	Risk Owner	Early Impact	Late Impact	Response	WBS Link	Impact Horizon
5	Plasma-wall models are insufficiently complete and compromise plasma predictions.	Martin Greenwald			Accept		Far

* Plasma-wall models are at a relatively primitive state. First principles modeling - starting with molecular dynamics - will be a very long and difficult path. Current FSP plans assume that in the short term simpler reduced models will be sufficient to capture plasma-wall physics in so far as it affects plasma dynamics (though clearly not the plasma impact on the wall).

** Work with OFES/BES to establish better partnerships (with funding) between science communities.

*** At a programatic level, materials research has been identified as a high-priority. The implications for meeting FSP requirements are not yet clear.

Exposure

Original -	Probability,	Cost,	Sched,	Tech,	Other
M 3	L	L	L	L	H

Thursday, July 28, 2011 *Page 4 of 6*

FSP

ACTIVE RISKS sorted by Exposure

IDx	*Risk Title* ** Risk Description* *** Management Plan* **** Notes*	*Risk Owner*	*Early Impact*	*Late Impact*	*Response*	*WBS Link*	*Impact Horizon*
4	Theoretical formulation for edge/pedestal gyrokinetics may not be sufficiently complete and correct or computationally tractable.	Martin Greenwald			Mitigate		Far

* Conditions in the pedestal and SOL break the orderings (small fluctuation amplitude and gradient scale length much longer than gyro-radius) used to derive the gyrokinetic equation. Various approaches have been proposed, but we can not yet be certain that these are sufficiently faithful to the actual physics AND are computationally tractable.

** Work with OFES to target theory program and/or devote FSP resources to support theory development. To avoid schedule delays, needs to be proactive.

*** These challenges face the MFE program as a whole, not just the FSP. There should be opportunities to leverage these common interests.

Exposure

Original - Probability, Cost, Sched, Tech, Other

🟩 L 2	L	M	L	L	L

Residual - Probability, Cost, Sched, Tech, Other

🟩 L 1	L	L	L	L	L

IDx	*Risk Title*	*Risk Owner*	*Early Impact*	*Late Impact*	*Response*	*WBS Link*	*Impact Horizon*
9	Tools required for effective software quality assurance efforts are inadequate.	Jeff Hittinger			Accept		Far

* The FSP plans to minimize the resources required for ancillary software development, in particular for software testing systems, by leveraging existing tools. Many such generic tools exist. Neverthelesss, it may be the case that no freely-avaialble testing tools that are supported across all necessary platforms, for all computer relevant languages, and adhering to the remaining requirements of the ISA and component developers may exist.

** Develop required tools, either in-house or through contracts with external collaborators. This will require a redistribution of resources.

*** Complete requirements for software testing tools cannot be formulated until the ISAs and component efforts are more precisely defined. Even in this absence, the requirements of availability across development sites, cross-platform functionality and multiple language support are challenging constraints for some types of testing software.

Exposure

Original - Probability, Cost, Sched, Tech, Other

🟩 L 2	L	M	M	L	L

IDx	*Risk Title*	*Risk Owner*	*Early Impact*	*Late Impact*	*Response*	*WBS Link*	*Impact Horizon*
10	Not all required synthetic diagnostics are identified	Jeff Hittinger			Accept		Far

* Systematic validation of FSP software products cannot proceed without the ability to make apples-to-apples comparisions reliably and accurately between simulation and experimental results. Such comparison requires synthetic diagnostics that map simulation results to approximate experimental observables. Validation efforts will be delayed if all relevant synthetic diagnostics are not identified and thus developed in a timely manner

** Solicit broad input from theory, analysts, and experimentalists early. Any failure will require a redistribution of resources.

*** Early engagement between theorists, experimentalists, and analysists focused on experimental design (both in the laboratory and in silica) should identify all necessary diagnostics.

Exposure

Original - Probability, Cost, Sched, Tech, Other

🟩 L 2	M	L	L	L	L

FSP

ACTIVE RISKS sorted by Exposure

IDx	*Risk Title* ** Risk Description* *** Management Plan* **** Notes*	*Risk Owner*	*Early Impact*	*Late Impact*	*Response*	*WBS Link*	*Impact Horizon*
12	Existing SQ methods are inadequate for analysis	Jeff Hittinger			Accept		Far

* Verification, validation, and UQ techniques for multi-physics applications are active areas of research. Fusion computations are sufficiently complex that it is expected that current state of the art techniques will be inadequate. Resources have been explicitly allocated to research in these areas, but research progress is difficult to predict.

** Develop new methods, either in-house or through contracts with external collaborators. This will require a redistribution of resources.

*** While existing techniques may prove inadequate for rigorous and quantifiable confidence bounds, partial results can still provide qualitative confidence measures.

Exposure

Original -	Probability,	Cost,	Sched,	Tech,	Other
🟩 L 2	M	L	L	L	L

IDx	*Risk Title*	*Risk Owner*	*Early Impact*	*Late Impact*	*Response*	*WBS Link*	*Impact Horizon*
6	Insufficient computing cycles	Bill Tang			Accept		Far

* There is a risk that the developers, who need access to HPC computing resources, may not get what they require.

** Need to have a coordinated campaign for INCITE allocations on the LCFs and dedicated alliances for capacity resources at NERSC and other HPC facilities. Work with DoE to make available other computing resources. Apply to fusion community sources.

*** Most development work is done on local clusters.

Exposure

Original -	Probability,	Cost,	Sched,	Tech,	Other
🟩 L 1	L	L	L	L	L

End Page